arXiv:2611.21538v1 [cs.LG] 16 Nov 2026

Optimistic Verifiers Get Hacked: Reward Hacking Under an Exact Process Oracle

Arunkumar V *arunkumarv1530@gmail.com*
University College of Engineering, BIT Campus, Anna University, Tiruchirappalli, India

Manoranjan Gandhudi
Central University of Karnataka, Kalaburagi, India

Gangadharan G. R.
National Institute of Technology Tiruchirappalli, India

Arun Prakash
Jawaharlal Nehru University, New Delhi, India

S. Senthilkumar
University College of Engineering, BIT Campus, Anna University, Tiruchirappalli, India

Abstract

Reinforcement learning from verifiable rewards inherits whatever its verifier gets wrong. In mathematics and code the verifier is exact. Outside them it is a proxy, a rubric, a learned judge or a hand-written checker, and its distance from the true objective is never measured, so a rising reward cannot be distinguished from a hack. We build a setting where the true reward is exactly computable at the process level. Executable authorization policies over an enumerable space of 124,416 contexts admit an oracle that lists every minimal sufficient reason for every decision, which turns the causal fidelity of a model's cited reasons into a computed quantity. We compile such policies into exact process rewards, degrade the verifier along seven realistic axes (outcome-only, sufficiency, recall-only, extraction noise, coarse taxonomy, stale policy, LLM judge), and train against each while the exact verifier scores every rollout. Four phenotypes appear, each with its own behavioural signature. Optimistic verifiers are hacked by padding: the sufficiency proxy terminates at 0.940 while true faithfulness falls from 0.731 to 0.667 on n=96 held-out episodes. Outcome-only verifiers saturate and stop teaching. Noise dilutes, and ranking-preserving gaps do no harm. Which phenotype a run will produce is predicted beforehand by a model-free audit of the verifier alone, its selection error and gap sign, together with a short pilot of how often its reward saturates. Across nine configurations spanning two frameworks, two model vendors, two adapter sizes, three precisions, two initialisations and two seeds, the audit quantity is invariant, the phenotypes recur, sufficiency damage is positive in all 10 replication cells and beyond two standard errors in 7, and optimistic gaps stay the steeper branch on the replication grid too (3.2× against 2.3×, at an interval too wide to test the size). The damage magnitude does not replicate, and five pre-registered attempts to explain it are reported as they fell. We release the testbed, the training instrumentation, and a second task family of 14 real public policies (Cedar, OPA, Gatekeeper) ported to executable oracles, on which the audit reads the same way.

1 Introduction

Reinforcement learning with verifiable rewards (RLVR), where a program checks each answer and returns the reward, has driven the recent progress of reasoning models. In mathematics and code the verifier is

exact. Nearly everywhere else it is a proxy, a rubric (Gunjal et al., 2026), a reference-derived signal (Jiang et al., 2026), a learned judge or a hand-built checker. Reward hacking, a policy raising its trained reward while the true objective falls, is studied intensively (Amodei et al., 2016; Skalse et al., 2022; Gao et al., 2023; Pan et al., 2022; Khalifa et al., 2026; Wong et al., 2026; Helff et al., 2026). Recent work argues that noisy verifiers may be tolerable (Plesner et al., 2026), and other work shows that systematic verifier errors shape outcomes in ways an error rate does not predict (Egashira et al., 2026). All of this work reads hacking off downstream evaluations, probes, held-out suites or anecdote, because the true reward is unavailable for the individual rollout, the sampled output where the behaviour arises.

We make it available. Executable authorization policies, the allow/block gates that govern tool use for deployed agents, Kubernetes admission and cloud access control, are white-box functions over a categorical context space (Cutler et al., 2024). A gate's decision tensor can be materialised, and an oracle can compute every minimal sufficient reason (MSR) of every decision exactly (Shih et al., 2018; Ignatiev et al., 2019). Whether each reason a model cites is causally relevant then reduces to a tensor slice. That is process-level ground truth, a reward on why a decision was reached and not only on whether it was right, which outcome verifiers for math and code cannot supply. The exactness costs generality. Our results hold for policy-gate explanation with 4B-parameter policy models and one policy-gradient recipe, and §10 says how far we have tested that they travel.

We make four contributions.

(1) Policy-as-code is a free verifier (§3, §4). We compile executable policy gates into exact outcome and process rewards, and formalise seven degraded verifiers that each corrupt exactly one stage of the reward pipeline. They mirror how practitioners build verifiers, by checking only outcomes, substituting a sufficiency or recall proxy, tolerating extraction noise, coarsening the taxonomy, holding a stale policy, or delegating to a large language model (LLM) as judge.

(2) A model-free gap audit (§5). Before any training, we measure each degraded verifier's distance from truth on a procedurally generated probe grid whose citation sets are exactly controlled. The outcome-only verifier, the construction underlying most RLVR, cannot rank a faithful explanation above an unfaithful one better than chance (selection error 0.504), the sufficiency proxy inverts 0.208 of selections, and a stale policy missing one rule mis-decides 0.1011 of all 124,416 contexts.

(3) An exact anatomy of the gap, and what survives replication (§6, §7). We train a 4B policy against each degraded verifier while the exact verifier scores every rollout, so that the divergence between proxy and truth is a per-step observable. Four phenotypes, distinct patterns of training behaviour, emerge. The sufficiency proxy is hacked by padding, and carries the citation-inflation signature we predicted in advance. Outcome-only rewards saturate within a few steps and starve training from then on. Noise dilutes. Ranking-preserving gaps, an LLM judge among them, are harmless. The audit's sign and selection error predict which of the four a run will produce. The sufficiency proxy's damage, the true reward it costs against an exact-reward control, is positive in every replication run and significant in most, and a second, structurally different optimistic verifier is hacked harder still. On the original grid the damage magnitude is also a tight function of selection error, with optimistic gaps 2.3× steeper than pessimistic ones across four verifiers. The ordering and the sign asymmetry replicate. The constants do not.

(4) A real-policy corpus and a damage-export test (§8). We release 14 public policies (Cedar, Open Policy Agent and Gatekeeper) ported to executable gates with exact oracles (5,256 contexts, 5 held out). On them the hacked model's padding exports intact while aggregate F_1 hides it, and the audit reads as it did on the synthetic family. Exact-reward training transfers without loss and without a significant gain, and corpus training is starved by the ease of the task, which we report as a boundary.

The sections after related work take these contributions in order, then give recommendations and limitations.

2 Related Work

RLVR and the verifier problem. The successes of RLVR in mathematics and code (Shao et al., 2024) rest on verifiers that are exact by construction. Extending it elsewhere has produced a family of proxy

verifiers, from rubric-based rewards (Gunjal et al., 2026) and reference-derived signal chains (Jiang et al., 2026) to empirical arguments that noisy verifiers are tolerable up to double-digit error rates (Plesner et al., 2026). Ray (2026) fuzz verifiers for exploitable bugs before training, and, closest to us, Egashira et al. (2026) inject systematic verification errors on arithmetic and find that outcomes range from delay to collapse in a way the overall error rate does not predict. Mitsuhashi et al. (2026) inject random false positives and false negatives on GSM8K and find false negatives the more damaging. All of these must treat the verifier as given, because in their domains no independent measurement of its distance from truth exists. Here the verifier's gap is itself the measured quantity. That answers the negative result of Egashira et al. with a positive one, since damage is predicted by selection error weighted by update flow (§6.1), and accounts for the sign asymmetry: pessimistic noise damages in both studies, while an optimistic verifier damages more per unit of corruption (§6.1). The noise-tolerance claims pass our test only conditionally, on reward granularity and on optimizer robustness (§6).

Reward hacking. Skalse et al. (2022) formalise hacking, Gao et al. (2023) measure overoptimisation against a synthetic learned gold reward model, Countdown-Code (Khalifa et al., 2026) and Wong et al. (2026) observe hacking in outcome-level environments, and Mahmoud et al. (2026) separate verifier failure from rubric-design failure with a judge panel as reference. Two recent studies audit a verifier before training. Zhang (2026) show that a static leakiness audit of code test suites predicts the rewarded false-positive mass, and Rajan (2026) find that 28.5% of SWE-bench Verified tasks accept an incorrect patch. Both measure how much wrong output a verifier accepts. Neither can measure what a policy learned from it, because the gold signal is again a test suite or a human. What we add is the limit case: an exact process-level true reward, degraded semantically rather than randomly, so that damage can be attributed to a named property of the verifier. The starvation phenotype we isolate is the pathological extreme of the uniform-group problem that motivates dynamic sampling in trainers of the group relative policy optimisation (GRPO) family (Yu et al., 2025).

Process rewards and faithfulness. Process reward models (Lightman et al., 2024) are learned and inherit judge noise, and step-level faithfulness rewards (Gui et al., 2026) rely on evidence-support judges. Yuan et al. (2026) convert symbolic oracles into dense process rewards and note the dependence on oracle quality without studying it. We study exactly that dependence, with an oracle that is perfect and degraded by construction. Ferreira et al. (2026) add a causal-attribution signal to the reward model after showing that preference optimisation erodes explanation faithfulness, and Hase & Potts (2026) reward chains of thought that let a simulator predict counterfactual outputs. Both measure faithfulness against the model's own behaviour, whereas our oracle scores it against the decision function's true reasons. Formal explainable AI (XAI) provides exact attribution for white-box functions through abductive minimal sufficient reasons (Shih et al., 2018; Ignatiev et al., 2019), and GlassGate (V, 2026) used it to show that frontier models' free-form decision explanations are causally imprecise regardless of scale. We turn that oracle from an evaluation device into a training-time instrument.

Agent-safety oracles. Vera (Feng et al., 2026) verifies whether a harmful outcome occurred but measures neither attribution nor verifier coverage, Who-Tests-the-Testers (Chen et al., 2026) audits benchmark coverage without an oracle, and AgentDojo (Debenedetti et al., 2024), MCPTox (Wang et al., 2026) and Agent-SafetyBench (Zhang et al., 2024) test agents against curated attacks. We study the verifier that trains and audits the agent.

3 Setting: Policies with an Exact Process Reward

This section defines the policy gates, their exact process reward and the episodes every experiment runs on; Fig. 1 shows the pipeline.

Policy gates. A gate G maps a context $c \in \mathcal{X} = \prod_i V_i$ (ten categorical features covering tool, scope, target domain, payload flags, auth level, role, intent and prior violations) to a decision $G(c) \in \{\textsc{allow}, \textsc{block}\}$ by first-match evaluation of an ordered rule list. Because $|\mathcal{X}| = 124{,}416$, G materialises as a decision tensor T.

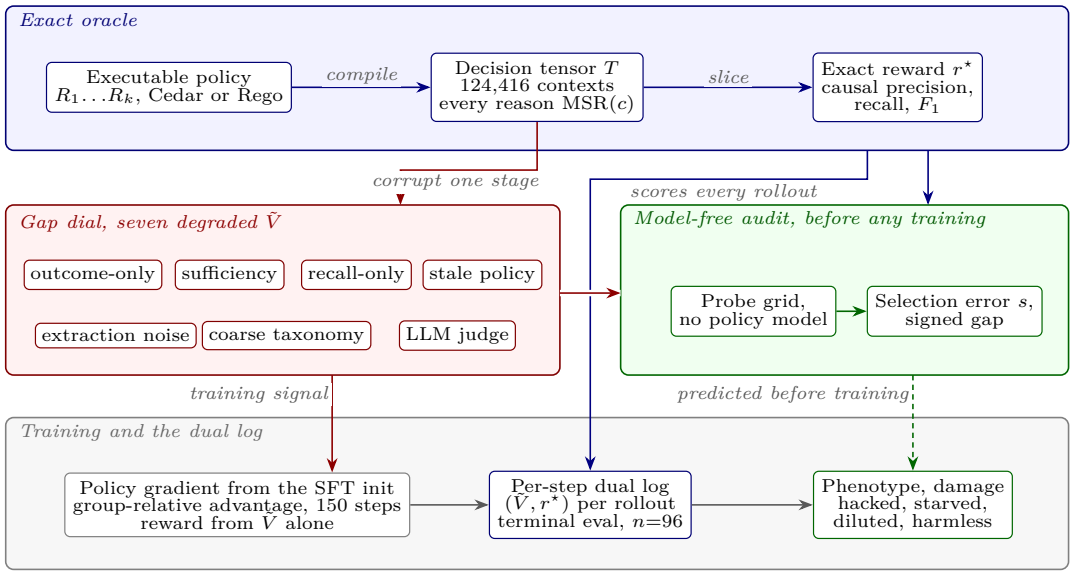


Figure 1: **The verifier-gap instrument.** An executable policy gate is compiled into an exact oracle (decision tensor and minimal sufficient reasons), which defines the exact process reward $r^\star$. The gap dial corrupts one stage of that pipeline to produce a degraded verifier $\tilde{V}$. A model-free audit scores $\tilde{V}$ against $r^\star$ on a probe grid before any training and yields a selection error and a sign. Training sees only $\tilde{V}$, while $r^\star$ scores every rollout in parallel, so the divergence between the two is logged per step. The dashed arrow is the paper's claim, that the audit predicts the phenotype the run will produce. The loop was replicated over MLX 4-bit on one laptop and PyTorch fp16 and NF4 on one T4 GPU, Qwen3-4B and Phi-3.5-mini, two adapters, two SFT lengths, fp16 and fp32 adapters, two seeds, and a second task family of 14 real public policies where the audit transfers and training is starved.

Definition 1 (Sufficiency, MSR)**.** *A feature set $S \subseteq [d]$ is* sufficient *at c iff every completion of c's values on S preserves G's decision. A* minimal sufficient reason *is a subset-minimal sufficient set, and* MSR(c) *is their antichain.*

Definition 2 (Exact process reward)**.** *For an explanation t of decision $G(c)$ with cited feature set $E(t) \subseteq [d]$ under a deterministic extractor E, causal precision and recall are $P = |E(t) \cap R(c)|/|E(t)|$ with $R(c) = \bigcup \text{MSR}(c)$, and $R = \max_{m \in \text{MSR}(c)} |E(t) \cap m|/|m|$. The reward is $r^\star(c,t) = F_1(P,R) - \lambda_b \mathbf{1}[\text{len} > L] - \lambda_d \mathbf{1}[\textit{wrong decision word}]$.*

Proposition 1 (Exactness)**.** *For enumerable $\mathcal{X}$, sufficiency,* MSR(c) *and hence $r^\star$ are computed exactly in $O(2^d)$ tensor slices per context. For non-enumerable rule gates the same queries are decided exactly by abduction over rule-relevant features (Appendix A).*

Gate family and episodes. Training spans a procedurally sampled family of gates over the same schema, rejection-sampled to a block-rate band of 0.35 to 0.65 with non-degenerate per-rule mass. There are 12 training gates and 4 held-out gates, never trained on, with 24 stratified episodes each. An episode is one context of one gate, with its full oracle annotation, against which every reward in this paper is computed. Real ported policies use the non-enumerating logical oracle, which agrees with the tensor oracle on all 124,416 contexts with zero mismatches (Appendix A).

4 Seven Degraded Verifiers

Each degraded verifier $\tilde{V}$ corrupts exactly one stage of the reward pipeline just defined. The exact verifier stays computable alongside, so $\Delta(c,t) = \tilde{V}(c,t) - r^\star(c,t)$ is always observable.

Table 1: **Model-free gap audit.** Selection error is the probability that the degraded verifier prefers the text with lower true reward. The outcome-only verifier is a coin flip and the sufficiency proxy inverts one selection in five. Stale-policy decision disagreement is exact over all 124,416 contexts, shown here for the least and most load-bearing rules (smallest and largest decision disagreement) (all eight in Table 3, Appendix B).

Verifier	Gap type	Reward gap	Selection error	Decision disagr.
Exact (control)	n/a	0.000	0.000	n/a
Outcome-only	metric	0.629	0.504	n/a
Sufficiency proxy	metric	0.239	0.208	n/a
Recall-only proxy	metric	0.090	0.132	n/a
Extraction noise ε=.05	extraction	0.051	0.043	n/a
Extraction noise ε=.15	extraction	0.156	0.143	n/a
Extraction noise ε=.30	extraction	0.271	0.308	n/a
Coarse taxonomy (payload)	granularity	0.034	0.030	n/a
Coarse taxonomy (session)	granularity	0.076	0.048	n/a
Stale policy (drop R4)	policy model	0.057	0.042	0.1011
Stale policy (drop R7)	policy model	0.031	0.023	0.0229

Outcome-only $\tilde{V}_{\text{out}}$ rewards the decision word and ignores cited reasons, the default construction in RLVR. **Sufficiency** $\tilde{V}_{\text{suff}}$ rewards the graded sufficiency of the cited set, the standard proxy in the explanation-faithfulness literature (DeYoung et al., 2020; Lei et al., 2016). **Recall-only** $\tilde{V}_{\text{rec}}$ rewards causal recall with no precision term. It is optimistic by a different mechanism, and tests whether the hack belongs to one proxy or to optimism itself. **Extraction noise** $\tilde{V}_{\varepsilon}$ flips each feature's cited status with probability ε, the random-noise model of Plesner et al. (2026). **Coarse taxonomy** $\tilde{V}_{\text{coarse}}$ cannot distinguish features within a group, such as the three payload flags, so credit flows to group-mates. **Stale policy** $\tilde{V}_{\text{drop}(r)}$ is missing rule r, so its decisions and MSRs are computed on the wrong gate. **LLM judge** $\tilde{V}_{\text{judge}}$ replaces the deterministic extractor with an LLM, the rubric-reward pattern. Formal definitions and the invariants each verifier is tested against are in Appendix B.

5 Auditing a Verifier Before Training

Before touching a policy model, we measure each verifier's distance from truth on a procedurally generated probe grid. For stratified contexts ($n = 60$), probe texts with exactly controlled citation sets span the truth spectrum, from faithful and partial to padded, bare, wrong-decision and distractor-only (Appendix G). We report the *reward gap* $\mathbb{E}|\tilde{V} - r^{\star}|$ and the *selection error*, the probability over probe pairs with different true rewards that the proxy prefers the worse text, with ties counting one half. Selection error is the distortion that drives best-of-N selection and the sign of an RL advantage. We also record the *signed gap* $\mathbb{E}[\tilde{V} - r^{\star}]$: a verifier is *optimistic* where it is positive, scoring explanations above their true reward, and *pessimistic* where it is negative.

Table 1 gives three findings. The outcome-only verifier has selection error 0.504, because faithful and unfaithful explanations both carry the right decision word. The sufficiency proxy has selection error 0.208, because padding is invisible to it, so it prefers cite-everything texts over precise ones. And semantic gaps are not interchangeable with random noise. Extraction noise at $\varepsilon = 0.05$ costs only 0.043 selection error, while dropping one load-bearing rule (R4) silently changes 0.1011 of all decisions, a gap invisible to any audit that lacks the true gate.

Three terms recur in the training results that follow. A *cell* is one training run at one seed, a *grid* is the cells sharing a framework and recipe, and *damage* is the exact control's terminal true reward minus the cell's, so positive damage is faithfulness the degraded verifier cost the policy.

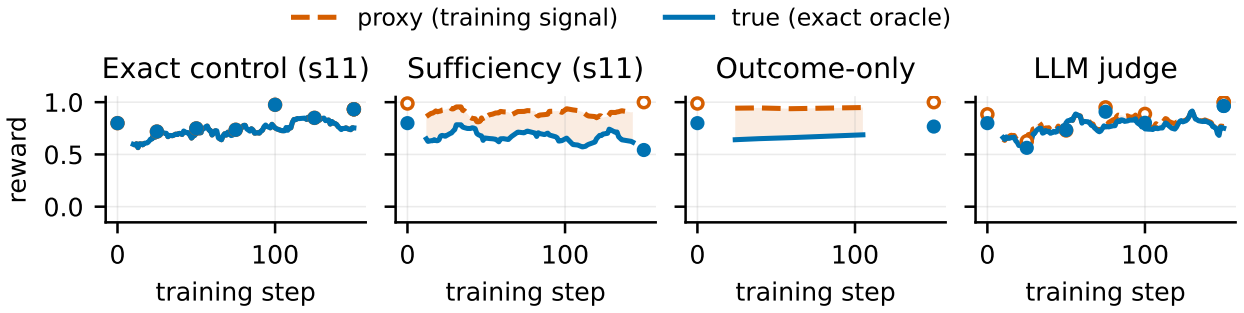


Figure 2: **Hacking curves.** Rolling on-policy reward under the training proxy (dashed, vermillion) and the exact oracle (solid, blue), with circles for the 8-episode in-training evals and the shaded band for the observed gap. The control tracks truth by construction, the sufficiency proxy diverges from it while reporting near-perfect reward, outcome-only saturates at once, and the judge stays level with the control. All 15 runs are in Fig. 6.

6 Training Against Degraded Verifiers

We train the same policy, initialised by supervised fine-tuning (SFT), against each degraded verifier for 150 steps while the exact verifier scores every rollout (Fig. 2, Table 2). Three pathologies emerge, a fourth gap type does no harm, and the proxy sees none of it. The section closes by linking each phenotype to the audit.

The exact reward is worth training on. Prompted explanations sit at causal precision of about 0.50 to 0.54 from a 14B model up to frontier models (V, 2026; Turpin et al., 2023; Lanham et al., 2023), so content selection is a problem of objective, not capacity. From an SFT initialisation scoring 0.731 ± 0.020 on held-out gates, 150 steps of policy gradient against the exact process reward reach 0.838 and 0.905 (seeds 11 and 12, mean 0.871), a 0.140 gain on gates never seen in training. Every endpoint in this paper is scored on the full held-out set ($n = 96$). The trainer is a transparent REINFORCE loop (Williams, 1992) with group-relative advantage (the GRPO core at on-policy single-step), gradient clipping and per-completion accumulation (Appendix D). An earlier unclipped configuration was seed-unstable and an off-the-shelf GRPO library produced silent no-op checkpoints (Appendix E).

Optimistic proxies are hacked by padding. Under $\tilde{V}_{\text{suff}}$ the proxy's terminal report is 0.940 while true reward falls from 0.731 to 0.667, a terminal gap of 0.273 on the n=96 held-out set. These runs and the recall-only cell below are the only ones whose citation counts inflate beyond the initialisation band, reaching 7.5 and 8.6 features here, which is the signature we predicted from correlational data before the cell ran. Sufficiency cannot see padding. A late-training hacked explanation cites eight features where three carry the decision, which makes the cited set trivially sufficient (Appendix H gives the rollouts verbatim). Across the grid (Fig. 8, Appendix I) the sufficiency seeds' citation trajectories leave every other cell within the first third of training while their word counts stay unremarkable, so the hack targets the quantity the proxy cannot see and leaves length alone. A pre-registered recall-only proxy, optimistic by a different mechanism, does the same, ending at 0.974 against a true 0.692 with citations rising from 6.1 to 7.4 (Table 2). What gets exploited is optimism about a behaviour the policy can reach. The sufficiency formula is incidental.

Outcome-only rewards starve training. Under $\tilde{V}_{\text{out}}$ the proxy saturates at 1.000 almost immediately, because every rollout carries the right decision word, and group-relative advantages vanish. 137 of 150 steps yield zero learning signal, while true reward drifts to 0.714 ± 0.024. Starvation is an interaction between verifier and trainer. Any group-relative trainer zeroes its advantage when a binary reward saturates, the symptom dynamic sampling treats (Yu et al., 2025), and a binary outcome check saturates as soon as the decision word is right, long before the reasons are.

Table 2: **Training against degraded verifiers.** 150 policy-gradient steps from the same SFT initialisation, held-out-gate greedy evals, and behavioural signatures as first-to-last 20-step means. $\text{Proxy}_{150}^{\text{train}}$ and Cited are training-time values; the text quotes held-out proxies and cited counts at n=96 instead, which differ. The harmful gap types produce distinct pathologies: padding (sufficiency), saturation starvation (outcome-only) and dilution (noise).

Training verifier	True_0	True_{150} (n=96)	$\text{Proxy}_{150}^{\text{train}}$	Gap	Starved	Cited	Words
Coarse taxonomy (payload)	0.731	0.844 ± 0.017	0.961	0.017	9/150	6.2→4.2	61→45
Coarse taxonomy (s12)	0.731	0.806 ± 0.022	0.975	0.000	12/150	4.0→5.4	64→47
Exact (control, s11)	0.731	0.838 ± 0.019	0.932	0.000	11/150	5.3→3.9	65→54
Exact (control, s12)	0.731	0.905 ± 0.016	0.887	0.000	24/150	4.0→3.8	67→50
Stale policy (s11)	0.731	0.875 ± 0.016	0.619	-0.314	16/150	5.0→3.3	62→41
Stale policy (s12)	0.731	0.840 ± 0.021	0.458	-0.375	4/150	4.1→4.5	66→37
LLM judge	0.731	0.865 ± 0.018	1.000	0.036	26/150	5.3→4.2	70→51
Extraction noise ε=.05	0.731	0.821 ± 0.015	0.814	-0.031	0/150	4.7→3.8	67→55
Extraction noise ε=.15	0.731	0.678 ± 0.024	0.515	-0.154	0/150	5.8→4.3	71→51
Extraction noise ε=.15 (s12)	0.731	0.648 ± 0.029	0.346	-0.162	0/150	3.9→4.5	61→50
Extraction noise ε=.30	0.731	0.641 ± 0.026	0.456	-0.199	0/150	5.1→4.5	66→71
Outcome-only	0.731	0.714 ± 0.024	1.000	0.233	137/150	6.2→4.6	73→54
Recall-only	0.731	0.692 ± 0.017	1.000	0.470	75/150	6.1→7.4	70→60
Sufficiency proxy (s11)	0.731	0.667 ± 0.018	1.000	0.458	68/150	6.0→7.5	65→57
Sufficiency proxy (s12)	0.731	0.639 ± 0.020	0.988	0.560	51/150	5.1→8.6	68→58

Noise dilutes, given a stable optimizer. Random verifier noise is reportedly benign up to 15% for outcome-level rewards (Plesner et al., 2026). Per-feature noise at that rate corrupts $1-(1-\varepsilon)^d \approx 0.80$ of rollouts (d=10), yet under clipping the damage is a saturating dose-response (Fig. 7). True reward ends at 0.821 for ε=.05, close to control level, then 0.678 at .15 and 0.641 at .30, while the noisy verifier's own report falls faster than the truth it corrupts. The trained policy is better than its verifier believes. Two caveats make this a conditional safety result. The same ε=.15 cell collapsed to 0.083 under an unclipped optimizer (Appendix E), so noise budgets compose with optimizer robustness. And noise is the one gap with zero skips, so it wastes some signal on every step but never starves one.

Ranking-preserving gaps are harmless. Stale-policy verifiers, with one rule deleted from the verifier's model of the gate, end at 0.875 and 0.840 across seeds, both inside the control seed spread. A coarsened feature taxonomy ends at 0.844 and 0.806 across seeds. These proxies are wrong about levels but nearly right about rankings (Skalse et al., 2023) (coarse audit selection error 0.030), and ranking is most of what the group-relative gradient consumes. The coarse cell is the one harmless-band prediction that seed 12 breaks, costing +0.065 against a band of 0.05, which we report rather than average away.

The LLM-judge cell (claude-haiku as extractor) went against our informal prediction and with our formal one. Its training-time proxy was optimistic (0.884 against true 0.731 at step 0), and from that we expected hacking. The audit said otherwise. On the training distribution the judge has selection error 0.090 with signed gap -0.067, a pessimistic level shift with little ranking corruption, because its citation judgements closely track the deterministic extractor. The cell ended at 0.865, level with the controls (0.871). In this one cell, a judge that passed a held-out exact audit did no measurable harm as a process verifier.

6.1 Damage as a function of the audit

The phenotypes compress into one relation between audit and damage (Fig. 3a). Let s be the verifier's selection error audited on the training distribution of gates (the probe suite of §5 run on 60 sampled family episodes, Appendix F, which is what lets the stale-policy and judge verifiers be audited at all) and u the fraction of training steps with non-zero group variance. Then

$$\text{damage} \;\approx\; \beta_{\pm}\, s\, u, \qquad \beta_{+} = 1.81 \;\; (R^2 = 0.94,\; n = 6), \quad \beta_{-} = 0.77 \;\; (R^2 = 0.90,\; n = 7),$$

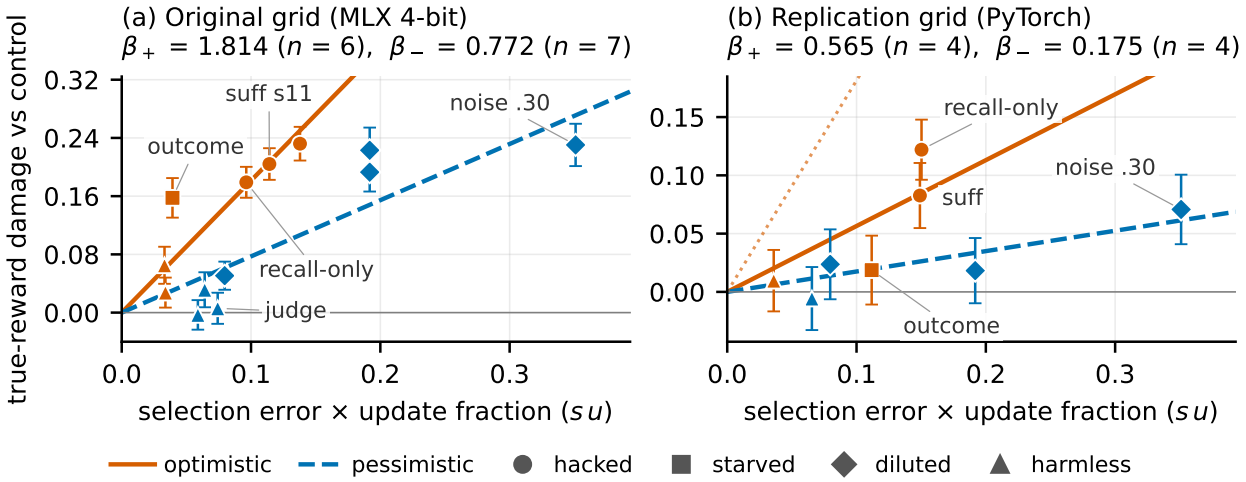


Figure 3: **The gap law on two grids.** Terminal true-reward damage against the audit's selection error weighted by the effective-update fraction, with standard errors at n=96. Optimistic gaps (vermillion) lie on a steeper through-origin line than pessimistic ones (blue). (a) is the MLX grid, all 13 non-control cells; (b) is the PyTorch grid fit on its own. The asymmetry reproduces, the slopes do not, and the dotted line in (b) is the MLX optimistic slope on that panel's scale. Only cells the text argues about are named; Table 7 lists every point.

with a two-branch R^2 of 0.92 over all 13 cells. The sign is the audit's signed gap, and verifiers that overestimate quality damage 2.3× more per unit of ranking corruption reaching the gradient than verifiers that underestimate it. The mechanism is that an optimistic, ranking-corrupting verifier rewards a reachable degenerate behaviour, so the policy optimises toward it, while a pessimistic one mislabels good behaviour and wastes signal. Both regressors are measurable before the run they predict, s from a probe audit with no policy model and u from a short pilot of group variance.

This fit describes one grid. The training-distribution audit was run after that grid had trained and before any replication run, so the fit is retrospective here and predictive there, with the refit that scored the replication documented in Appendix J. Bootstrap 95% intervals are $\beta_+ \in [1.67, 2.58]$ and $\beta_- \in [0.26, 1.03]$, and the ratio's interval is [1.7, 7.3] with $P(\beta_+ > \beta_-) = 1.00$. Leave-one-out prediction has mean absolute error 0.042, and the law under-predicts the starved outcome-only cell by about half. The optimistic branch holds 6 cells from four different verifiers (Appendix F tabulates every point). The constants did not survive replication on an independent implementation, whereas the sign and the ordering did, and the size of the asymmetry is at least not contradicted (2.3× here, interval [1.7, 7.3], against 3.2× there, interval [1.0, 10.5], which is too wide to call agreement). We therefore state the law as an ordering and a sign asymmetry with an interval, and claim no portable coefficients.

7 Replication on a Second Implementation

The grid of §6 is one framework (MLX, Apple's machine-learning framework, with 4-bit weights), one base model and one LoRA adapter recipe (a low-rank weight update trained in place of the full model). This section asks which of its findings survive a second implementation. We ported the trainer to PyTorch and re-ran the grid on an NVIDIA T4 under eight further configurations. They vary the model (Qwen3-4B and Phi-3.5-mini), the SFT length (400 and 1200 iterations), the adapter (query/value only, or every projection in the last 16 layers as in MLX), the base precision (fp16 and NormalFloat quantisation, NF4), the adapter precision (fp16 and fp32) and the seed. Seven sets of predictions were committed before their runs, two of them amended after first contact with data (a rescoring against the n=96 refit, and a trainability gate for the corpus), and Appendix J scores each as it fell.

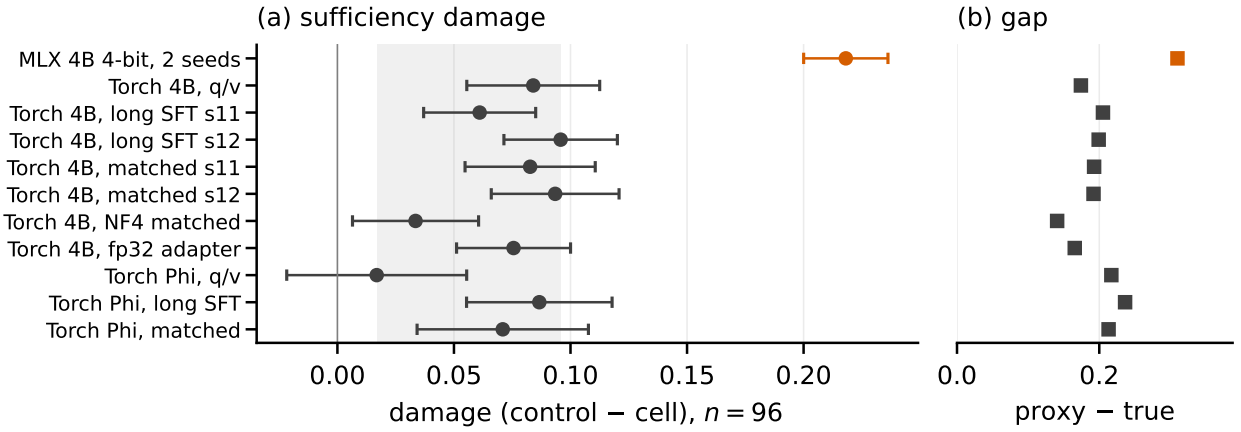


Figure 4: **Replication.** (a) Sufficiency damage, control minus cell at n=96, in every cell measured; Torch means PyTorch fp16 unless NF4 or fp32 is stated. The shaded band spans the smallest to largest PyTorch damage, and every PyTorch cell (grey) falls well below the MLX grid (vermillion). (b) The same cells' proxy over-report barely moves, so the audit quantity is a property of the verifier while the damage is not.

What replicates. The audit quantity is a property of the verifier alone. The sufficiency proxy over-reports truth by between +0.14 and +0.31 regardless of model, framework, adapter, precision or initialisation (Fig. 4b).

The four phenotypes replicate on the two PyTorch grids that ran every verifier. The sufficiency cell's damage is positive in all 10 PyTorch cells, beyond two of its own standard errors in 7, and null on the weakest initialisation. Every sufficiency cell inflates its citation count, and under the matched adapter the stale policy does -0.006 and the coarse taxonomy +0.010, both inside the pre-registered harmless band of 0.05. The remaining cells are unordered within their standard errors. The outcome-only proxy saturates at 0.999 and skips 113 of 150 matched-adapter updates. Each matched cell was re-run at a second seed, and 7 of 7 damages fall within two combined standard errors. One pair is nonetheless a warning: the starved outcome-only cell lost +0.094 at seed 12, as much as sufficiency (+0.093), with its proxy still saturated, so a starved run drifts where the verifier cannot see.

The sign asymmetry replicates. The recall-only verifier of §6 was pre-registered on this grid too, and hacks harder than sufficiency does, at +0.122 ± 0.026 of damage (4.7σ) against sufficiency's +0.083 under the same adapter and seed, with citation inflation +2.85, the largest of any PyTorch cell. Where both were run, the optimistic verifier without a precision term is the worse one. Fitting the PyTorch cells alone gives $\beta_+ = 0.565$ ($n = 4$, $R^2 = 0.84$) and $\beta_- = 0.175$ ($n = 4$, $R^2 = 0.88$), a ratio of 3.2 against 2.3 on the original grid (Fig. 3b). Its own interval is [1.0, 10.5] with $P(\beta_+ > \beta_-) = 0.98$, wide enough that only the sign, not the size, replicates here.

What does not. The size of the damage does not replicate. 10 PyTorch cells put sufficiency damage between +0.017 and +0.096, while the MLX grid gives +0.218 ± 0.018 (Fig. 4a). None of the obvious explanations, each pre-registered and tested in turn, accounts for it. Adapter capacity does not: matching the MLX adapter moved damage from +0.084 at 1.3M parameters to +0.083 at 7.3M. Base quantisation does not (NF4 weights, +0.034 ± 0.027), nor does adapter precision (fp32, no loss scaling, +0.076 ± 0.024), nor does more optimisation, since a stronger Qwen initialisation gave +0.078 over two seeds and 300 Phi steps changed nothing (-0.011 ± 0.037). A larger model did not fit the hardware (Appendix J). One prediction in this series held: raising Phi's initialisation from 0.672 to 0.764 turned a null (+0.017) into a hack (+0.087, 2.8σ) with nothing else changed, so competence gates whether the hack bites. What sets how hard it bites remains open, with three explanations excluded.

8 Transfer to Real Policies

This section tests what carries over from synthetic gates to real policies, first as a transfer test and then as a training domain. Models trained only on synthetic gates must explain decisions of real, held-out policies whose schemas they have never seen (Appendix C), scored against each policy's verified logical oracle over 80 episodes. The one significant result is damage export. The sufficiency-hacked model's padding transfers intact, with recall 0.965 (it cites nearly every feature) at a precision of 0.702 that ties SFT for the lowest of any variant, yet its aggregate F_1 of 0.779 is the highest in the table, because indiscriminate citation is F_1-favourable here. In this test the hacked model's damage survived domain shift, and the aggregate score hid it.

Exact-reward training itself transfers without aggregate loss (F_1 0.740 and 0.745 against a base of 0.728) and moves precision the right way (0.754 against 0.717), but every F_1 interval overlaps every other and the precision shift is not significant (p=0.30), so it is directional only. Gate coverage gives a null. A variant trained on a single gate (structural coverage 0.23 against 0.53) transfers just as well (0.750), so causal-precision pressure carries transfer here and gate diversity adds nothing measurable (Appendix I).

The corpus as a second training domain. We also trained the four-verifier grid on the corpus itself (Appendix K). The audit transferred before any training, outcome-only again a coin flip (selection error 0.488) and sufficiency again optimistic. Training was starved: real minimal reasons are one or two features, the base model already scores 0.759 on them, and the exact control skipped 85 of 150 updates, failing the pre-registered trainability gate. The sufficiency cell is again the largest, but no corpus damage is distinguishable from zero. This is our own starvation phenotype arriving through task ease, and we do not count it as a replication.

9 Recommendations

Within the scope tested here (§10), our results support four recommendations. Audit a process verifier's selection error on a probe grid before training against it. In our grids it predicted the phenotype where reward accuracy did not, and it needs no policy model. Read the sign too: optimistic verifiers were hacked and pessimistic ones only diluted, so bias pessimistic when forced to choose. Do not train a process objective on an outcome-only signal, which is blind to reasons and, once saturated, silent. Give sufficiency- and recall-style rewards a precision term, since alone they reward padding. An LLM judge can be used once its rankings pass a held-out exact audit, as the one judge we trained against did.

10 Limitations

Exactness is bought by a restriction to first-match and deny-overrides (Cedar-style, where a matching forbid overrides any permit) gates over categorical contexts. The quantitative fit belongs to one grid (one task family, one 4B model, one recipe, MLX 4-bit), its constants did not transfer, that grid's damage size is unexplained, and its outcome-only, judge, recall-only and two of three noise cells are single-seed. No model above 4B was trained, since an 8B model in 4-bit did not fit a T4. The deterministic extractor anchors truth, audited by a held-out LLM extractor, not humans. Corpus training was starved, so the corpus gives an audit and a transfer test but no damage measurement, and ported policies abstract some constructs (Appendix C).

Broader Impact Statement

This work studies failure modes of reward verifiers using synthetic policy gates and public, permissively licensed policy corpora. No human subjects, personal data or proprietary data are involved. Understanding how proxy rewards are hacked could in principle inform adversaries, but the attacks demonstrated (padding, saturation exploitation) are already folklore. The work contributes a defence, an audit that detects a dangerous verifier before it is trained against.

Reproducibility Statement

Every number in this paper is generated from committed result files by a single script. No number in any table or figure is hand-typed, and the few narrative literals restate macro values or logged training printouts. The full pipeline (oracle, verifiers, probe audit, trainer with per-step dual-reward logs, corpus porting protocol with per-policy provenance, Kaggle kernel and figure generation) will be released as one self-contained repository, including the episode data and the SFT initialisation every run starts from, so every figure and table regenerates from a clone. The padding-signature prediction for the sufficiency cell was written before that cell ran but was not timestamped, so we describe it as predicted rather than pre-registered. The seven replication pre-registrations (Appendix J) are committed, timestamped files in the artifact. The MLX grid ran on a single Apple-silicon laptop (4B model, LoRA, about 30 hours). The replication grids ran on single NVIDIA T4 sessions.

AI Use Statement

LLMs appear in this work as subject models (the trained 4B policies and the frontier models in the prompting study), as one of the audited verifiers (the LLM-judge cell) and as the held-out extraction auditor. Beyond that, LLM assistance was limited to grammatical checks on author-written text and to coding assistance. No claim, result, figure or citation in this paper was generated by an LLM, and all experimental results derive from the released deterministic pipeline.

References

Dario Amodei, Chris Olah, Jacob Steinhardt, Paul Christiano, John Schulman, and Dan Mané. Concrete problems in ai safety. *arXiv preprint arXiv:1606.06565*, 2016.

John Backes, Pauline Bolignano, Byron Cook, Catherine Dodge, Andrew Gacek, Kasper Luckow, Neha Rungta, Oksana Tkachuk, and Carsten Varming. Semantic-based automated reasoning for AWS access policies using SMT. In *Formal Methods in Computer Aided Design (FMCAD)*, 2018.

Bowen Baker, Joost Huizinga, Leo Gao, Zehao Dou, Melody Y. Guan, Aleksander Madry, Wojciech Zaremba, Jakub Pachocki, and David Farhi. Monitoring reasoning models for misbehavior and the risks of promoting obfuscation. *arXiv preprint arXiv:2503.11926*, 2025.

Xuan Chen, Lu Yan, Ruqi Zhang, and Xiangyu Zhang. Who tests the testers? systematic enumeration and coverage audit of llm agent tool call safety. *arXiv preprint arXiv:2603.18245*, 2026.

Joseph W. Cutler, Craig Disselkoen, Aaron Eline, Shaobo He, Kyle Headley, Michael Hicks, Kesha Hietala, Eleftherios Ioannidis, John Kastner, Anwar Mamat, Darin McAdams, Matt McCutchen, Neha Rungta, Emina Torlak, and Andrew Wells. Cedar: A new language for expressive, fast, safe, and analyzable authorization. *Proceedings of the ACM on Programming Languages*, 8(OOPSLA1):670–697, 2024.

Adnan Darwiche and Auguste Hirth. On the reasons behind decisions. In *European Conference on Artificial Intelligence (ECAI)*, 2020.

Edoardo Debenedetti, Jie Zhang, Mislav Balunović, Luca Beurer-Kellner, Marc Fischer, and Florian Tramèr. Agentdojo: A dynamic environment to evaluate prompt injection attacks and defenses for llm agents. In *Advances in Neural Information Processing Systems (NeurIPS) Datasets and Benchmarks Track*, 2024.

Jay DeYoung, Sarthak Jain, Nazneen Fatema Rajani, Eric Lehman, Caiming Xiong, Richard Socher, and Byron C. Wallace. ERASER: A benchmark to evaluate rationalized NLP models. In *Annual Meeting of the Association for Computational Linguistics (ACL)*, pp. 4443–4458, 2020.

Kazuki Egashira, Mark Vero, Jasper Dekoninck, Florian E. Dorner, Robin Staab, and Martin Vechev. Delay, plateau, or collapse: Evaluating the impact of systematic verification error on RLVR. *arXiv preprint arXiv:2605.02909*, 2026.

Yunhao Feng, Ruixiao Lin, Ming Wen, Qinqin He, Yanming Guo, Yifan Ding, Yutao Wu, Jialuo Chen, Zhuoer Xu, Xiaohu Du, Jianan Ma, Zixing Chen, Xingjun Ma, Yunhao Chen, and Xinhao Deng. Safety testing llm agents at scale: From risk discovery to evidence-grounded verification. *arXiv preprint arXiv:2607.01793*, 2026.

Pedro Ferreira, Wilker Aziz, and Ivan Titov. Truthful or fabricated? Using causal attribution to mitigate reward hacking in explanations. In *International Conference on Learning Representations (ICLR)*, 2026.

Leo Gao, John Schulman, and Jacob Hilton. Scaling laws for reward model overoptimization. In *International Conference on Machine Learning (ICML)*, 2023.

Runquan Gui, Yafu Li, Xiaoye Qu, Ziyan Liu, Yeqiu Cheng, and Yu Cheng. Faithrl: Learning to reason faithfully through step-level faithfulness maximization. *arXiv preprint arXiv:2602.03507*, 2026.

Anisha Gunjal, Anthony Wang, Elaine Lau, Vaskar Nath, Yunzhong He, Bing Liu, and Sean Hendryx. Rubrics as rewards: Reinforcement learning beyond verifiable domains. In *International Conference on Learning Representations (ICLR)*, 2026. arXiv:2507.17746.

Peter Hase and Christopher Potts. Counterfactual simulation training for chain-of-thought faithfulness. In *Conference on Language Modeling (COLM)*, 2026.

Lukas Helff, Quentin Delfosse, David Steinmann, Ruben Härle, Hikaru Shindo, Patrick Schramowski, Wolfgang Stammer, Kristian Kersting, and Felix Friedrich. Llms gaming verifiers: Rlvr can lead to reward hacking. In *ICLR 2026 Workshop on LLM Reasoning*, 2026. arXiv:2604.15149.

Alexey Ignatiev, Nina Narodytska, and Joao Marques-Silva. Abduction-based explanations for machine learning models. In *AAAI Conference on Artificial Intelligence*, 2019.

Yuxin Jiang, Yufei Wang, Qiyuan Zhang, Xingshan Zeng, Liangyou Li, Jierun Chen, Chaofan Tao, Haoli Bai, and Lifeng Shang. From verifiable dot to reward chain: Harnessing verifiable reference-based rewards for reinforcement learning of open-ended generation. In *International Conference on Learning Representations (ICLR)*, 2026. arXiv:2601.18533.

Jacek Karwowski, Oliver Hayman, Xingjian Bai, Klaus Kiendlhofer, Charlie Griffin, and Joar Skalse. Goodhart's law in reinforcement learning. In *International Conference on Learning Representations (ICLR)*, 2024. arXiv:2310.09144.

Muhammad Khalifa, Zohaib Khan, Omer Tafveez, Hao Peng, and Lu Wang. Countdown-code: A testbed for studying the emergence and generalization of reward hacking in rlvr. *arXiv preprint arXiv:2603.07084*, 2026.

Cassidy Laidlaw, Shivam Singhal, and Anca Dragan. Correlated proxies: A new definition and improved mitigation for reward hacking. In *International Conference on Learning Representations (ICLR)*, 2025. arXiv:2403.03185.

Tamera Lanham, Anna Chen, Ansh Radhakrishnan, Benoit Steiner, Carson Denison, Danny Hernandez, et al. Measuring faithfulness in chain-of-thought reasoning. *arXiv preprint arXiv:2307.13702*, 2023.

Tao Lei, Regina Barzilay, and Tommi Jaakkola. Rationalizing neural predictions. In *Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pp. 107–117, 2016.

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 (ICLR)*, 2024.

Anas Mahmoud, MohammadHossein Rezaei, Zihao Wang, Anisha Gunjal, Bing Liu, and Yunzhong He. Reward hacking in rubric-based reinforcement learning. *arXiv preprint arXiv:2605.12474*, 2026.

Joao Marques-Silva and Alexey Ignatiev. Delivering trustworthy AI through formal XAI. In *AAAI Conference on Artificial Intelligence*, pp. 12342–12350, 2022.

Ryo Mitsuhashi, Patrick Chen, Isabelle Tseng, Jasin Cekinmez, and Addison J. Wu. Quantifying empirical compute-supervision tradeoffs in RLVR. *arXiv preprint arXiv:2605.25252*, 2026.

OPA Gatekeeper. Policy and governance for Kubernetes. https://open-policy-agent.github.io/gatekeeper/website/docs/, 2026.

Open Policy Agent. The Rego policy language. https://www.openpolicyagent.org/docs/policy-language, 2026. CNCF graduated project; no language paper exists.

Alexander Pan, Kush Bhatia, and Jacob Steinhardt. The effects of reward misspecification: Mapping and mitigating misaligned models. In *International Conference on Learning Representations (ICLR)*, 2022. arXiv:2201.03544.

Andreas Plesner, Francisco Guzmán, and Anish Athalye. An imperfect verifier is good enough: Learning with noisy rewards. *arXiv preprint arXiv:2604.07666*, 2026.

Shreshth Rajan. Auditing reward hackability in code RL training environments. *arXiv preprint arXiv:2606.16062*, 2026.

Jaideep Ray. Before the model learns the bug: Fuzzing rlvr verifiers. *arXiv preprint arXiv:2606.01066*, 2026.

Zhihong Shao, Peiyi Wang, Qihao Zhu, Runxin Xu, Junxiao Song, Xiao Bi, Haowei Zhang, Mingchuan Zhang, Y. K. Li, Y. Wu, and Daya Guo. Deepseekmath: Pushing the limits of mathematical reasoning in open language models. *arXiv preprint arXiv:2402.03300*, 2024.

Andy Shih, Arthur Choi, and Adnan Darwiche. A symbolic approach to explaining bayesian network classifiers. In *International Joint Conference on Artificial Intelligence (IJCAI)*, 2018.

Prasann Singhal, Tanya Goyal, Jiacheng Xu, and Greg Durrett. A long way to go: Investigating length correlations in RLHF. In *Conference on Language Modeling (COLM)*, 2024. arXiv:2310.03716.

Joar Skalse, Nikolaus Howe, Dmitrii Krasheninnikov, and David Krueger. Defining and characterizing reward gaming. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2022.

Joar Skalse, Matthew Farrugia-Roberts, Stuart Russell, Alessandro Abate, and Adam Gleave. Invariance in policy optimisation and partial identifiability in reward learning. In *International Conference on Machine Learning (ICML)*, 2023.

Miles Turpin, Julian Michael, Ethan Perez, and Samuel R. Bowman. Language models don't always say what they think: Unfaithful explanations in chain-of-thought prompting. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2023. arXiv:2305.04388.

Arunkumar V. Do agent guardrails explain themselves? causal faithfulness of tool-call gating explanations. In *Proceedings of the International Natural Language Generation Conference (INLG)*, 2026.

Zhiqiang Wang, Yichao Gao, Yanting Wang, Suyuan Liu, Haifeng Sun, Haoran Cheng, Guanquan Shi, Haohua Du, and Xiangyang Li. Mcptox: A benchmark for tool poisoning attack on real-world mcp servers. In *AAAI Conference on Artificial Intelligence*, 2026. arXiv:2508.14925.

Ronald J. Williams. Simple statistical gradient-following algorithms for connectionist reinforcement learning. *Machine Learning*, 8:229–256, 1992.

Aria Wong, Josh Engels, and Neel Nanda. Mitigating reward hacking with rl training interventions. In *ICLR 2026 Workshop on Principled Design for Trustworthy AI*, 2026.

Qiying Yu et al. Dapo: An open-source llm reinforcement learning system at scale. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2025. arXiv:2503.14476.

Huining Yuan, Zelai Xu, Huaijie Wang, Xiangmin Yi, Jiaxuan Gao, Xiao-Ping Zhang, Yu Wang, Chao Yu, and Yi Wu. Verifiable process rewards for agentic reasoning. *arXiv preprint arXiv:2605.10325*, 2026.

Chuyifei Zhang. When the reward suite is leaky: A preregistered causal contrast of natural verifier false positives in RLVR. *arXiv preprint arXiv:2607.11022*, 2026.

Zhexin Zhang, Shiyao Cui, Yida Lu, Jingzhuo Zhou, Junxiao Yang, Hongning Wang, and Minlie Huang. Agent-safetybench: Evaluating the safety of llm agents. *arXiv preprint arXiv:2412.14470*, 2024.

A The Exact Oracle

Tensor oracle. For enumerable $\mathcal{X}$ the gate materialises as an int8 decision tensor T of shape $(12, 3, 4, 2, 2, 2, 3, 3, 4, 3)$. A set S is sufficient at c iff the tensor slice obtained by pinning c's values on S is constant, which is one `min==max` reduction. MSR enumeration walks the 2^{10} subsets in popcount order with superset pruning, so the survivors are exactly the minimal sufficient sets in the sense of Darwiche & Hirth (2020), and decisive features use single-axis slices. Formal XAI of this kind buys soundness guarantees that heuristic attribution does not (Marques-Silva & Ignatiev, 2022). Exactness is by construction, because every counterfactual completion is literally enumerated.

Logical oracle. For non-enumerable or foreign-schema gates the same queries are decided by abduction over rule-relevant features only. For first-match block-rule gates, BLOCK-sufficiency is an adversarial CSP (can the unpinned features be completed so that no rule matches?) solved by backtracking, the same reduction production cloud authorization uses to decide policy questions by SMT (Backes et al., 2018), and ALLOW-sufficiency requires every rule to be falsified by a pinned literal. Deny-overrides (Cedar-style) gates add permit and forbid reachability checks. On the enumerable schema, the logical oracle's decisions and full MSR sets agree with the enumerating tensor oracle on all 124,416 contexts with zero mismatches (script and result file in the artifact).

B Degraded Verifiers: Definitions and Invariants

Let $E(t)$ be the deterministic extractor's cited set and $r^\star$ the F_1-based exact reward of Definition 2. Each degraded verifier $\tilde{V}$ replaces exactly one component. **Outcome-only** scores $\mathbf{1}$[decision word correct] minus the shared penalties. **Sufficiency** scores the graded sufficiency fraction of $E(t)$ by pin-and-check on the episode's own gate tensor. **Recall-only** scores $\max_{m \in \mathrm{MSR}(c)} |E(t) \cap m|/|m|$ minus the shared penalties, with no precision term. **Noise**(ε) applies i.i.d. per-feature flips to $E(t)$, seeded by (seed, c, t) so that rewards are reproducible. **Coarse**(G) scores precision and recall at the granularity of feature groups G. **Stale**(r) recomputes decision and MSRs on the gate with rule r deleted. **Judge** replaces E with an LLM extractor, with the deterministic E held out as the audit; optimising against a model that reads explanations is the setting in which Baker et al. (2025) observe obfuscation under strong pressure. Each satisfies a designed-rate invariant enforced by unit test. Noise(0) is identical to exact and noise(1) flips every feature, coarse with singleton groups is identical to exact, and stale(r) equals exact on every context whose decision and MSRs rule r does not touch. Table 3 reports every verifier's measured selection error and signed gap on the probe grid, including all eight stale-policy variants.

Table 3: The full model-free gap audit, one row per stale-policy rule.

Verifier	Gap type	Reward gap	Selection error	Decision disagr.
Exact (control)	n/a	0.000	0.000	n/a
Outcome-only	metric	0.629	0.504	n/a
Sufficiency proxy	metric	0.239	0.208	n/a
Recall-only proxy	metric	0.090	0.132	n/a
Extraction noise ε=.05	extraction	0.051	0.043	n/a
Extraction noise ε=.15	extraction	0.156	0.143	n/a
Extraction noise ε=.30	extraction	0.271	0.308	n/a
Coarse taxonomy (payload)	granularity	0.034	0.030	n/a
Coarse taxonomy (session)	granularity	0.076	0.048	n/a
Stale policy (drop R1)	policy model	0.035	0.026	0.0252
Stale policy (drop R2)	policy model	0.026	0.023	0.0261
Stale policy (drop R3)	policy model	0.040	0.024	0.0237
Stale policy (drop R4)	policy model	0.057	0.042	0.1011
Stale policy (drop R5)	policy model	0.060	0.032	0.0476
Stale policy (drop R6)	policy model	0.054	0.040	0.0425
Stale policy (drop R7)	policy model	0.031	0.023	0.0229
Stale policy (drop R8)	policy model	0.037	0.029	0.0611

C Real-Policy Corpus: Porting Protocol and Provenance

Table 4: **Real-policy corpus v2.** 14 public policies ported to executable gates under a fidelity-first protocol, with every abstraction documented and sanity contexts hand-traced against the original semantics. 5 policies are held out entirely.

Policy	Source	Features	$\lvert\mathcal{X}\rvert$	Rules	Block rate	Split
`cedar-photoapp`	cedar-examples	8	384	5	0.43	train
`cedar-streaming-service`	cedar-examples	10	3,456	7	0.73	heldout
`cedar-tags-n-roles`	cedar-examples	5	162	2	0.80	train
`cedar-tax-preparer`	cedar-examples	7	128	4	0.81	heldout
`gk-block-endpoint-edit`	gatekeeper-library	5	112	2	0.11	train
`gk-container-limits`	gatekeeper-library	5	192	8	0.48	heldout
`gk-disallowed-tags`	gatekeeper-library	4	72	3	0.25	train
`gk-psp-hostnet`	gatekeeper-library	5	64	3	0.22	train
`gk-psp-seccomp`	gatekeeper-library	3	36	3	0.25	heldout
`gk-psp-users`	gatekeeper-library	6	432	7	0.15	train
`gk-required-labels`	gatekeeper-library	4	60	4	0.70	train
`gk-storageclass`	gatekeeper-library	5	64	7	0.75	train
`opa-loadbalancer`	library	3	40	4	0.78	heldout
`opa-pvc-resize`	library	3	54	1	0.81	train

Each public policy (Table 4) is ported by mapping every Cedar (Cutler et al., 2024) permit or forbid clause, or Rego violation rule (Open Policy Agent, 2026; OPA Gatekeeper, 2026), to one antecedent over categorical features. Entity-relationship conditions become boolean features, and numeric, regex and datetime conditions are pre-bucketed into feature values. Decision semantics are preserved exactly under DenyOverrides (default deny, permits allow, forbids override) or admission style (empty permit plus violation forbids). Every abstraction is recorded in a per-policy `omitted` field, and every port was validated by exhaustive enumeration against the gate plus hand-traced sanity contexts against the original source semantics, three

or more per policy, all matching. Full machine-readable specs with provenance (repository, path, license, block rate) ship in the artifact, and all sources are Apache-2.0. The per-policy abstraction record below is generated verbatim from those `omitted` fields.

`cedar-photoapp` (`cedar-policy/cedar-examples`: `cedar-policy-language-in-action/PhotoApp/photoapp.cedar`, Apache-2.0). The scope condition `` `resource in PhotoApp::Application::"PhotoApp"` `` in every policy treated as always-true (all Photos are in the Application per the entity store). `` `resource has subjects && resource.subjects.contains(principal)` `` in @id('Photo.subjects') bucketed into is_subject. `` `context has judgingSession && context.judgingSession == true` `` in @id('PhotoJudge') bucketed into judging_session (absent attribute to no). The instance policies @id('DoeFamily') (DoeFamily members on Album DoePhotos) and @id('JaneVacation') (JohnDoe on Album JaneVacation) merged into one relationship boolean album_grant (permits are disjunctive, so merging preserves decisions). @id('Photo.owner')'s explicit action list is all three schema actions, so the conjunct is dropped as vacuous. Labels is a Set<String>, so label_private/label_contest are independent booleans.

`cedar-streaming-service` (`cedar-policy/cedar-examples`: `cedar-example-use-cases/streaming_service/policies.cedar`, Apache-2.0). All datetime arithmetic bucketed: now-vs-resource.releaseDate in @id('subscriber-content-access/show') and @id('early-access-show') to time_window (meaningful only for shows; Movie has no releaseDate, so the value is a don't-care never consulted for movies). The conjunction resource.isOscarNominated && 2025-02-02T19:00-0500 <= now < 2025-03-02T19:00-0500 in @id('rent-buy-oscar-movie') to single boolean oscar_promo (a 4th 'oscar_window' time_window value was rejected: shows at Oscars-window times still have a definite release relation, so that enum would create ill-defined contexts and break exactness). The toTime()/localTimeOffset computation in the unless-clause of @id('forbid-bedtime-watch-kid-profile') to is_bedtime (yes iff local time outside 06:00-21:00). FreeMember/Subscriber entity types and subscription.tier merged into 3-valued membership (FreeMember has no tier or profile; no rule consults profile_is_kid for free members). The @id('subscriber-content-access/show') unless-clause (A && B) split into two permits by De Morgan (not-early OR after-release). Needs_rent_or_buy / is_oscar_nominated are Movie-only and is_early_access is Show-only attributes kept as global features (don't-cares for the other type).

`cedar-tags-n-roles` (`cedar-policy/cedar-examples`: `cedar-example-use-cases/tags_n_roles/policies.cedar`, Apache-2.0). Principals holding multiple roles simultaneously (entity Joe is in both Role-A and Role-B with different tag sets; Cedar's decision is the disjunction over per-role policies), each context models the evaluation for a single role membership, which is exact for single-role principals. The guard `` `principal.allowedTagsForRole has "Role-X"` `` repeated in every when-clause of @id('Role-A policy')/@id('Role-B policy') is folded into the match enums (a missing role record maps every dimension to not_covered, which is a consistent assignment). Per-dimension set logic bucketed into 3-valued enums: principal_all = the user's role record lacks the tag OR its values contain 'ALL' (user-side wildcard, matches any resource). Covered = resource lacks the tag, resource values contain 'ALL', or user values are a superset of resource values (containsAll). Not_covered = the when-clause for that dimension is false. The Role-A action conjunct is kept explicit even though Role-A Actions contains all three schema actions.

`cedar-tax-preparer` (`cedar-policy/cedar-examples`: `cedar-example-use-cases/tax_preparer/policies.cedar`, Apache-2.0). The only action is viewDocument, so no action feature. Rule 1a's assigned_orgs.contains({organization: resource.owner.organization, serviceline: resource.serviceline, location: resource.location}) set-of-records containment bucketed into assignment_covers_doc. The organization/serviceline/doc_location document descriptors are kept for realism but are decision-irrelevant (only the tuple-match boolean is consulted, since a professional may hold several assigned_orgs tuples). The @id('adhoc-access') template and its `` `linked` `` instantiation (link_id AliceView: Alice to Document DEF) generalized to the boolean adhoc_grant permit. Rule 2's context.consent record split into consent_from_owner (consent.client == resource.owner) and consent_covers_location (team_region_list.contains(principal.location)). Rule 2's `` `unless (A && B)` `` becomes two forbids (not-A; not-B).

`gk-block-endpoint-edit` (`open-policy-agent/gatekeeper-library`: `library/general/block-endpoint-edit-default-role/template.yaml`, Apache-2.0). Context is a single (role, policyrule, verb) grounding of the rego's iteration over rules[_], resources[_], verbs[_]. The three hasEditVerb helper bodies collapse into the value set {create,patch,update}. Exact string match on metadata.name 'system:aggregate-to-edit' bucketed into role_name. Object_kind and has_autoupdate_annotation are faithful distractors, the rego never inspects kind (ClusterRole scoping is in the Constraint match) and the rbac.authorization.kubernetes.io/autoupdate annotation appears only in the advisory violation message, never in a condition.

gk-container-limits (open-policy-agent/gatekeeper-library: library/general/containerlimits/template.yaml, Apache-2.0). Parameters instantiated to cpu='2', memory='1Gi' (both finite, so the cpu != '-1' disable-guards are always satisfied). Canonify_cpu/canonify_mem numeric parsing and the > max comparisons are pre-bucketed into cpu_limit/mem_limit values (under_max/over_max/unparseable, unset = key absent or ''). The two rego rules 'not container.resources' and 'not container.resources.limits' merge into the has_limits=none bucket. The two field-wrapper violation rules (containers/initContainers) are represented by leaving container_field unconstrained in every forbid. ExemptImages prefix matching bucketed into exempt_image. Coherence: has_limits=none implies cpu/mem unset, partial implies exactly one unset, present implies neither unset, decisions on incoherent product contexts have no rego ground truth. Ephemeral containers not checked by the template (field cannot be set).

gk-disallowed-tags (open-policy-agent/gatekeeper-library: library/general/disallowedtags/template.yaml, Apache-2.0). Parameters instantiated to tags=['latest']. Suffix matching (strings.any_suffix_match against ':latest') and the split-by-'/' untagged check are pre-bucketed into image_tag: none = no colon in last segment, latest = ends with ':latest', pinned = any other tag, digest = @sha256:... references which trip neither rule (colon present, suffix differs) and are admitted. The three input_containers union rules are represented by leaving container_field unconstrained in both forbids. ExemptImages prefix matching bucketed into exempt_image. Image_registry is a faithful distractor, split(image, '/') means even a registry port colon (registry:5000/app) never masks an untagged image.

gk-psp-hostnet (open-policy-agent/gatekeeper-library: library/pod-security-policy/host-network-ports/template.yaml, Apache-2.0). Numeric min/max range bucketed into host_port. Host_port=none covers no ports and ports without hostPort. Single container with a single port modeled (source iterates all containers x ports). ExemptImages prefix matching bucketed to exempt_image (applies only to the port check, matching the rego). UPDATE operations bypass all checks (hostNetwork is immutable). CEL variant's ports.all() differs from the rego's any-port semantics, the Rego engine is ported.

gk-psp-seccomp (open-policy-agent/gatekeeper-library: library/pod-security-policy/seccomp/template.yaml, Apache-2.0). Allowlist instantiations: wildcard=['*'] (short-circuits all checks including exemption), runtime_default_only=['runtime/default'], runtime_and_localhost=['runtime/default','localhost/allowed.json'] (exact file match, no localhost/* wildcard). Resolved_profile collapses the four resolution locations (pod/container annotation, pod/container securityContext), location affects only the message. Docker_default models the annotation-scheme string 'docker/default' (securityContext RuntimeDefault canonicalizes to an allowed variant and cannot produce it). SecurityContext-scheme parameter translation (RuntimeDefault/Localhost/allowedLocalhostFiles) folded into the same canonical allowlists. Single container modeled. No update exclusion exists in this template. CEL engine variant is decision-equivalent.

gk-psp-users (open-policy-agent/gatekeeper-library: library/pod-security-policy/users/template.yaml, Apache-2.0). RunAsGroup and supplementalGroups families dropped to stay <=6 features (same accept_value logic as fsGroup; supplementalGroups additionally iterates array values). Provided_user=root means uid 0 and is assumed outside the MustRunAs ranges. The runAsNonRoot=true escape hatch (runAsUser unset but runAsNonRoot set gives no NonRoot violation) is not modeled, unset means neither field is set. Container/pod-level securityContext precedence collapsed into the single provided value. Single container modeled (containers/initContainers/ephemeralContainers union). ExemptImages prefix matching bucketed to exempt_image.

gk-required-labels (open-policy-agent/gatekeeper-library: library/general/requiredlabels/template.yaml, Apache-2.0). Parameters instantiated to labels=[{key: owner}, {key: env, allowedRegex: ^(dev|staging|prod)$}]. The single missing-labels violation (count(required, provided) > 0) is grounded to one forbid per configured label. The regex-mismatch violation cannot fire for owner (no allowedRegex, so expected.allowedRegex != '' is undefined) and its regex match on env is pre-bucketed into env_value (dev/staging/prod match, other = present-but-mismatched). Parameters.message (message override) dropped. Extra_labels and object_kind are faithful distractors, the rego ignores labels beyond the required set and never inspects kind (scoping is done by the Constraint match criteria).

gk-storageclass (open-policy-agent/gatekeeper-library: library/general/storageclass/template.yaml, Apache-2.0). Allowlist instantiated as any = allowedStorageClasses unset/['*'] and restricted_to_a = ['class_a']. Class_exists = whether the named class is in data.inventory (don't-care when storageclass_name is unset). 'not storageclass_allowed(name)' negates a 2-body helper, so each badname violation family grounds to two conjunctive forbids (not-in-inventory; not-in-allowlist), overlapping where both hold. The standalone StorageClasses-not-synced violation (inventory key entirely absent) is a cluster-misconfiguration guard outside the per-object feature space and is omitted (inventory assumed synced). The badname rules

also fire on unset names via negation-of-undefined, decision-equivalent to the noname forbids, folded there. Context is one volumeClaimTemplate grounding for statefulsets. Include_classes_in_message is a faithful distractor (message-only parameter). ApiVersion/kind checks bucketed into object_kind.

`opa-loadbalancer` (`open-policy-agent/library`: `kubernetes/admission/loadbalancer.rego`, Apache-2.0). Source is a library defining is_external_lb. Ported as deny-if-external. contains(spec.type, 'LoadBalancer') substring match bucketed to service_type=loadbalancer. Internal_annotation values mean the exact key=value pair from internal_annotations is present (a key with a mismatched value buckets to none). Count(externalIPs)>0 bucketed to has_external_ips (source reads object.externalIPs, not spec.externalIPs).

`opa-pvc-resize` (`open-policy-agent/library`: `kubernetes/admission/persistentvolumeclaimresize.rego`, Apache-2.0). Allow-only source (single 'verify' conjunction), so ported as a permit-conjunction over default deny rather than admission+forbids. Plugin set is the source's literal {GlusterFS, Cinder, AWS, GCE, RDS} ('RDS' as written upstream, likely a typo for RBD). Numeric size comparison bucketed to size_change. Phase string 'ClaimBound' bucketed to bound, all other phases to pending/lost.

D Training Details

The policy is Qwen3-4B (4-bit, MLX) with LoRA adapters (7.34M trainable parameters), initialised from a teacher-distilled SFT adapter. Training uses policy gradient with group-relative advantage, group 8, temperature 0.8, at most 220 new tokens, Adam at learning rate 5×10^{-6}, global-norm gradient clipping at 1.0, per-completion gradient accumulation, 150 steps, and greedy held-out evals every 25 steps on 8 episodes during training. Steps with zero group variance are logged and skipped with no update. Seeds are 11 and 12. Table 5 lists the full configuration of both the final recipe and the initial recipe whose fragility is analysed in Appendix E. Both are archived as `config.json` alongside every run's per-step logs and terminal adapters. One run takes about 75 minutes of wall-clock (the LLM-judge cell is slower, because its extraction API calls dominate), and the full grid takes about 30 hours.

Table 5: Training hyperparameters, generated from the archived run configurations. "Same" means the value is shared between recipes. The initial recipe's instability is documented in Appendix E.

Hyperparameter	Final recipe	Initial recipe
Base model	`Qwen3-4B-4bit` (MLX)	same
LoRA rank / scale / dropout	8 / 20 / 0.0	same
Layers adapted	16	same
Optimizer	Adam	Adam
Learning rate	5×10^{-6}	1×10^{-5}
Gradient clipping (global norm)	1.0	none
Group size (completions/step)	8	4
Gradient accumulation	per completion	whole group
Steps	150	150
Sampling temperature	0.8	0.8
Max new tokens	220	220
Held-out eval	every 25 steps, n=8	same
Hardware	Apple M-series, 24 GB unified memory	same
Wall-clock per run	≈75 min	≈75 min

PyTorch port. The replication trainer is a line-by-line port of the MLX loop to PyTorch with PEFT. It keeps the group size, learning rate, clipping, temperature, token budget, seed handling and skip rule. Three guards are new: a first-step proof of life (a non-zero gradient norm and a changed weight), a rollout-diversity pre-flight that aborts a run whose policy emits one text per group, and a circuit breaker on 25 consecutive identical groups. Adapters carry a base-model marker that the loader checks, after one run silently loaded a Phi adapter onto Qwen through a name collision. The matched adapter targets every attention and MLP

projection in the last 16 layers with rank 8 and scale 20, which is the MLX adapter exactly. A run on a T4 takes about 3.5 hours for 150 steps.

Compute budget. Table 6 itemises all training and evaluation compute behind the MLX grid. All on-device compute ran on one Apple M-series laptop (24 GB unified memory). The only external compute for that grid is 9,076 `claude-haiku` API calls (1,496 as the judge cell's training-time reward and 7,580 for held-out transfer extraction), all content-address cached and shipped with the artifact. At about 45 h of total wall-clock, the MLX programme sits roughly two orders of magnitude below typical RLVR training budgets. The replication grids add about 110 T4-hours across the four free Kaggle accounts that ran them.

Table 6: Compute budget by experiment family, generated from archived run logs and API caches. Wall-clock is measured from each run's log-file timestamps. 3 runs whose timestamp windows were polluted by machine sleep use the recipe-typical estimate of about 71 min instead. The judge cell's wall-clock is elapsed time dominated by its per-rollout reward API calls. Transfer-eval generation took about 13 min on-device, the rest of that hour being extraction-API latency.

Component	Volume	Wall-clock	`haiku` calls
Definitive grid (final recipe, excl. judge cell)	14 runs × 150 steps	17.4 h	n/a
LLM-judge cell (final recipe)	1 run × 150 steps	16.2 h (API-bound)	1,496
Coverage-null control (single gate)	1 run × 150 steps	≈1.2 h	n/a
Pilot grid (recipe selection)	6 runs × 50 steps	2.9 h	n/a
Initial-recipe runs (Appendix E)	8 runs × 150 steps	6.9 h	n/a
Transfer evaluation (§8)	6 variants × 80 episodes	<1 h	7,580
Total	n/a	≈45 h	9,076

E Optimizer Fragility

Two negative results shaped the recipe. First, an off-the-shelf MLX GRPO implementation produced checkpoints byte-identical to initialisation, a silent zero-gradient no-op. The released trainer is a transparent REINFORCE loop with group-relative advantage, the GRPO core at on-policy single-step, with a first-step proof-of-life assertion on gradient norms and weight deltas. Second, the initial recipe (learning rate 10^{-5}, group 4, no clipping) was seed-unstable. The exact-reward control itself collapsed on one seed to a terminal true reward of 0.401, and under it the noise(ε=.15) cell collapsed to 0.083, against 0.678 under the final recipe. The two old-recipe figures in this paragraph are in-training 8-episode evaluations. Those runs predate the terminal re-scoring and saved no adapter to re-score, so they are not comparable at the third decimal with the $n = 96$ numbers used everywhere else. The comparison they support is a collapse against a stable level, which no measurement noise explains. Figure 5 shows both comparisons step by step: the noise cell decays roughly monotonically to collapse under the fragile optimizer and holds a diluted level under the clipped one, and the exact control survives both but dips sharply mid-run without clipping. Verifier noise and optimizer fragility therefore compound: a gap that is benign under a stable optimizer can be catastrophic under a fragile one. All old-recipe runs are preserved in the artifact.

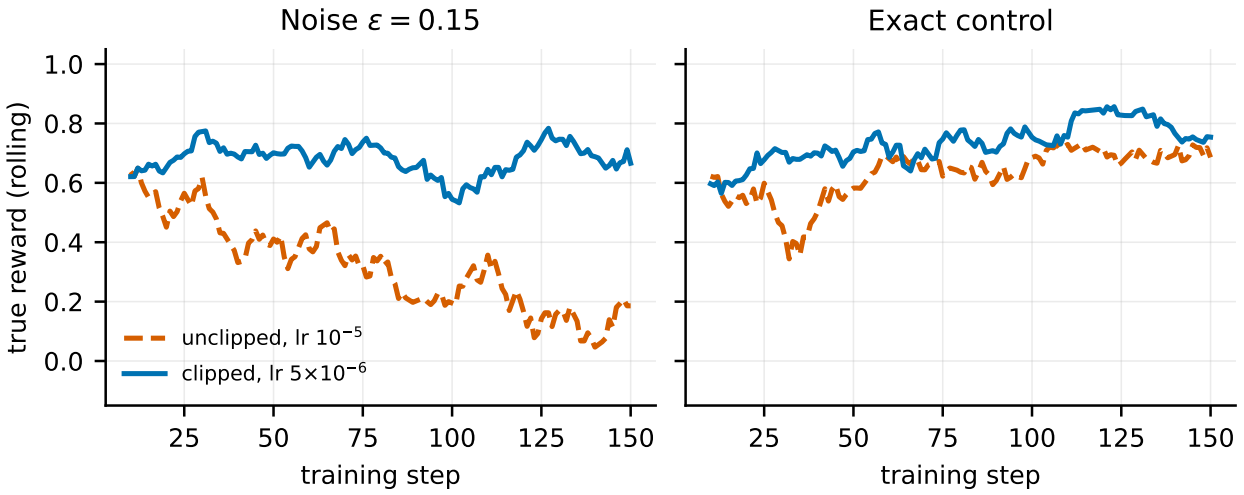


Figure 5: **Optimizer fragility.** Rolling (window 10) per-step mean true reward for the same two grid cells under the initial recipe (unclipped, learning rate 10^{-5}, group 4, vermillion dashed) and the final recipe (clipped at global norm 1.0, learning rate $5{\times}10^{-6}$, group 8, blue solid). Left, the noise ε=.15 cell collapses under the fragile optimizer and merely dilutes under the clipped one. Right, the exact control on seed 11 survives both (the control collapse was on seed 12), but the unclipped trajectory dips sharply mid-run and ends below the clipped one.

F Fit Data

The optimistic branch of the original grid holds 6 cells across four verifiers: the sufficiency proxy at two seeds, the recall-only proxy, the outcome-only proxy, and the coarse taxonomy at two seeds. The shape of this relation, true reward falling as optimisation pressure is applied through a corrupted proxy, is what Karwowski et al. (2024) derive for MDPs and what Laidlaw et al. (2025) formalise as a breakdown of proxy-truth correlation over a reference policy's states; the exact oracle lets us measure the curve rather than bound it. An earlier version of this work reported it as one verifier's slope, before the recall-only and outcome-only cells were added. R^2 in the main text is the through-origin form $(1 - \mathrm{SS}_{\mathrm{res}}/\sum y^2)$, and the standard form gives 0.72 and 0.78. Dropping u and fitting selection error alone gives $R^2 = 0.77$ optimistic and 0.88 pessimistic, so on this grid the model-free audit quantity carries the prediction and the pilot-measured update fraction adds little.

Table 7 lists every point behind Fig. 3a. Per cell it gives the pre-training audit quantities (selection error s and signed reward gap, both from the model-free probe audit of §5), the effective-update fraction u (the share of the 150 steps with non-zero group variance, measurable from a cheap pilot), the regressor $s\,u$, and the realised damage (control-mean terminal true reward 0.871 minus the cell's terminal true reward). The selection errors used in the fit come from a second audit run on the training distribution (60 stratified episodes drawn from the sampled family gates, with the same probe suite as Table 1). This is what lets the stale-policy verifier (a rule dropped from each sampled gate, which the fixed-gate audit of Table 1 cannot represent) and the LLM judge (0.090, one API call per probe) enter the fit. The s column of Table 7 separates the optimistic verifiers (sufficiency 0.209, recall-only 0.193) from the near-chance outcome-only verifier (0.453) and from the barely distorting pessimistic ones (coarse 0.036, stale policy 0.066, judge 0.090). The optimistic through-origin fit uses the cells with positive signed gap and the pessimistic fit those with negative signed gap. The stale-policy and judge cells were harmless in training, and the judge's damage is in fact negative, with terminal true reward 0.865 against the control mean 0.871.

Every cell is in the fit. The outcome-only run originally finished without a terminal checkpoint, because the trainer of the day wrote adapters only on updating steps and that cell skipped 137 of its 150 updates. It was retrained after the bug was fixed, reproducing the same 137 skipped steps and the same endpoint, and

its terminal adapter is scored here like every other. Starvation is therefore measured at $n = 96$ on this grid as well as on the replication, where the outcome-only proxy saturates at 0.999 and the matched cell loses +0.019 against its control.

Table 7: Per-cell gap-law fit data behind Fig. 3a. Damage is the control-mean terminal true reward minus the cell's, and negative damage means the cell ended above the control mean.

Cell	Phenotype	Sel. err. s	Signed gap	Eff. upd. u	$s\,u$	Damage	True_{150}
Sufficiency proxy (s11)	hacked	0.209	+0.185	0.547	0.114	+0.204	0.667
Sufficiency proxy (s12)	hacked	0.209	+0.185	0.660	0.138	+0.232	0.639
Outcome-only	starved	0.453	+0.163	0.087	0.039	+0.158	0.714
Noise ε=.05	diluted	0.080	-0.027	1.000	0.080	+0.051	0.821
Noise ε=.15	diluted	0.192	-0.060	1.000	0.192	+0.193	0.678
Noise ε=.30	diluted	0.351	-0.117	1.000	0.351	+0.230	0.641
Stale policy (s11)	harmless	0.066	-0.083	0.893	0.059	-0.003	0.875
Stale policy (s12)	harmless	0.066	-0.083	0.973	0.064	+0.031	0.840
Coarse taxonomy (payload)	harmless	0.036	+0.020	0.940	0.034	+0.027	0.844
LLM judge	harmless	0.090	-0.067	0.827	0.074	+0.006	0.865
Recall-only	hacked	0.193	+0.052	0.500	0.096	+0.179	0.692
Noise ε=.15 (s12)	diluted	0.192	-0.060	1.000	0.192	+0.223	0.648
Coarse taxonomy (s12)	harmless	0.036	+0.020	0.920	0.033	+0.065	0.806

G Prompts and Probe Construction

Generation uses the audit-log prompt (context rendering, full policy text, decision, "write 2 to 4 sentences"). Extraction uses the schema-generic constrained-JSON prompt with the episode's feature list. Probe texts are rendered from templates whose phrases are written against the deterministic extractor's alias patterns, so each probe's cited set is exact by construction. The six probe types are faithful (canonical MSR), partial (half the MSR), padded (all features), bare (decision word only), wrong-decision and distractor-only. Verbatim prompt texts ship in the artifact.

H Qualitative Examples

Representative late-training rollouts, verbatim from the logged `steps.jsonl` files. Scores are the logged per-rollout outputs of the training proxy and the exact verifier.

Hacked (sufficiency proxy, step 150; proxy 1.00**, true** 0.55**, causal precision** 0.375**, 8 features cited).** *"This tool call was blocked by R4 (privilege-escalation) because the intern user with no authentication and read-scope attempted to send an email containing both PII and credentials to a public.web domain. Additionally, the request triggers R5 (session-risk) due to the intern's 2 prior violations and lack of authentication."* Five of the eight cited features are causally irrelevant. Pinning everything makes the cited set trivially sufficient, so the proxy scores it perfect.

Faithful (exact reward, step 131; proxy = true = 1.00**, precision** 1.0**, 4 features cited).** *"This request was blocked by the data-egress policy because the combination of basic authentication level with access to a public web domain indicates a potential data egress. Additionally, the intern-level user role accessing a public web domain with a history of prior violations presents elevated risk for this action."*

Starved (outcome-only, step 105, the last step with non-zero group variance; proxy 1.00**, true** 0.67**, causal precision** 0.50**, 4 features cited).** *"This tool call was blocked by the repeat-offender rule, which prohibits requests containing personally identifiable information (PII) with a prior violation on record. The combination of 1 documented violation and PII presence satisfies the block criteria, preventing further unauthorized data manipulation attempts. The call carries destructive intent and lacks authentication, creating heightened security and compliance risks."* The outcome verifier awards this text, and every other rollout in its group, the identical perfect score for containing the decision word. The step therefore produces

zero advantage, and the causally irrelevant intent and authentication citations cost true precision that the training signal never sees.

Diluted (extraction noise ε=.30, step 150; proxy 0.67 computed on a post-flip citation set of 6, true 0.60, causal precision 0.43, 7 features cited). *"This tool was blocked by the default policy (R9) because no specific blocking policy matched, as the intern user is sharing sensitive data externally, and the email.send scope is being used. However, the intern's sensitive user role and prior violations on record create additional risks, so the gateway defaults to allowing sensitive external sharing. The intern's lack of authentication (no auth_level) combined with sensitive PII content prevents secure sharing, so the sensitive-escalation policy (R4) blocks the action."* This is the rambling signature of the diluted phenotype. The text invokes two different rules, contradicts its own decision narrative ("defaults to allowing"), and cites seven features where a single minimal sufficient reason would do. The noisy verifier meanwhile scored a citation set that differs from what the text actually says.

I Full Curves and Transfer Breakdown

Full hacking curves. Figure 6 shows the per-step dual-reward trace for all 15 MLX runs. The four phenotypes of §6 are visible as four distinct geometries. The hacked cells (sufficiency, both seeds) open a band between proxy and truth that widens from the first steps. The starved cell (outcome-only) pins the proxy at ceiling, and the trace thins out because almost every step is skipped for zero group variance. The diluted cells (noise) show the proxy tracking truth from below, with a separation that grows in ε while truth itself plateaus. The harmless cells are wrong about levels, with the stale-policy proxy far below truth and the judge slightly above, while truth stays at control level throughout.

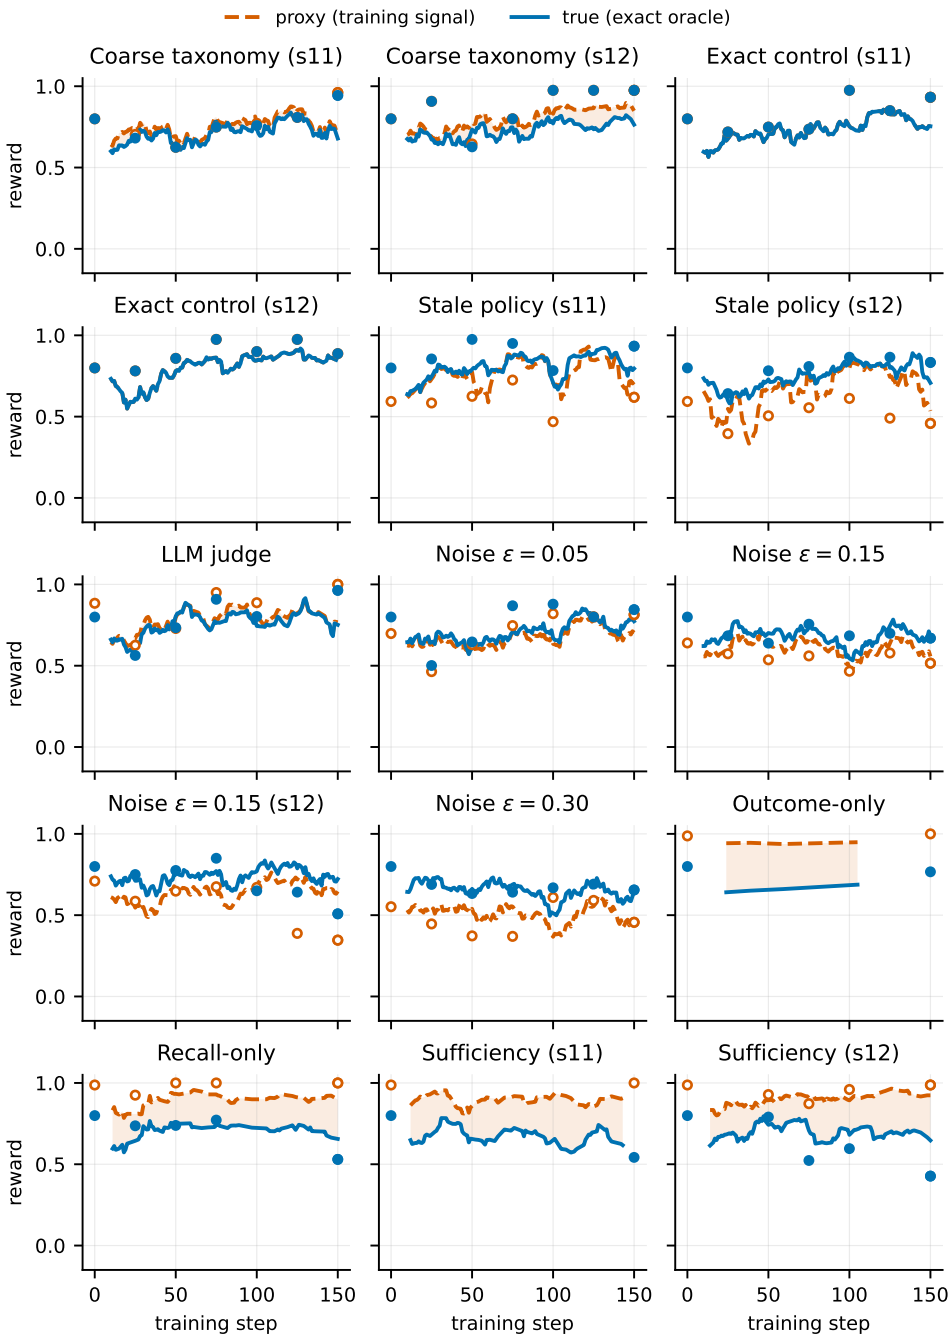


Figure 6: Per-step proxy (dashed, vermillion) and exact true (solid, blue) reward for all 15 MLX runs. Circles are the 8-episode in-training evals (terminal values in the text are re-scored at n=96) and the shaded band is the observed verifier gap where the proxy is optimistic.

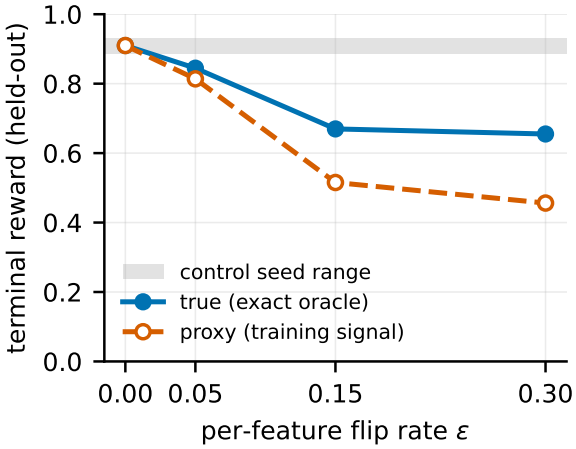


Figure 7: **Noise dose-response.** Terminal held-out reward after 150 steps against the per-feature flip rate ε. The ε=0 point is the mean of the two exact-control seeds and the gray band their range. True reward (blue, filled) degrades gracefully and saturates. The verifier's own report (vermillion, open) falls faster, because per-feature noise is a pessimistic gap, so the trained policy ends up better than its own training signal claims.

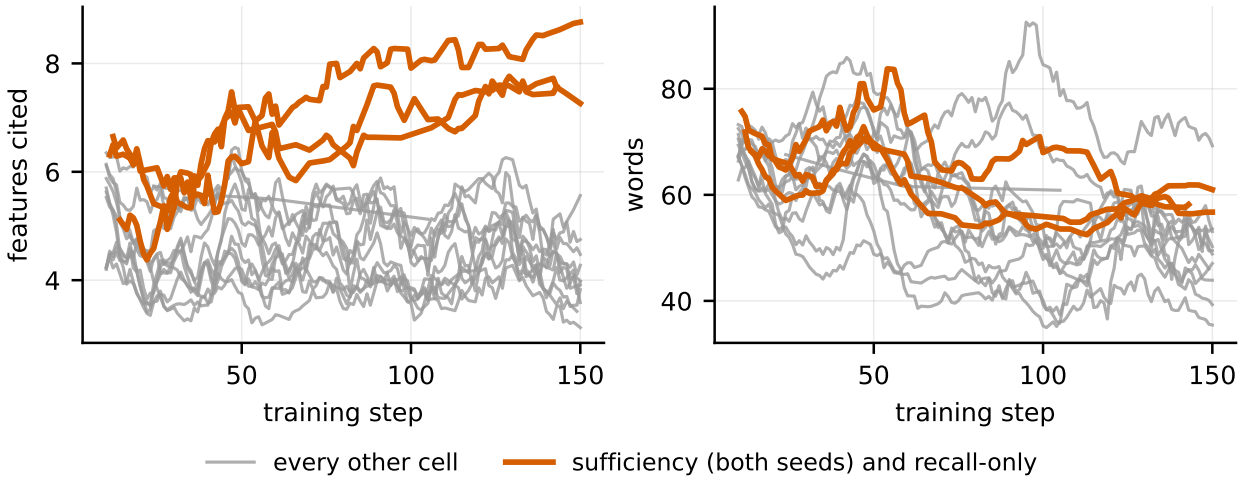


Figure 8: **Behavioural signatures.** Rolling (window 10) per-step means of features cited (left) and words emitted (right) for all 15 MLX runs, with the two sufficiency-proxy seeds highlighted (vermillion, thick). They and the recall-only cell are the only runs whose citation counts inflate beyond the initialisation band, the padding signature predicted before training, while every other cell drifts flat or downward. Word counts do not separate the hacked runs from the rest of the grid, so length-based monitoring (Singhal et al., 2024) alone would miss the hack.

Per-policy transfer. Table 8 breaks the aggregate transfer scores of §8 down by held-out policy. No variant dominates across policies. PG-exact (seed 12) is strongest on `opa-loadbalancer`, PG-exact (seed 11) on `cedar-streaming-service`, and the untrained base model on `cedar-tax-preparer` and `gk-psp-seccomp`. The hacked model's table-topping aggregate decomposes exactly as its padding signature predicts. Recall-saturated citation is competitive on every policy yet best on only one (`gk-container-limits`), and its causal precision is the best on none of the five. Aggregate F_1 rewards indiscriminate citation on real policies just as it does on the synthetic grid.

Table 8: Per-policy transfer on the five held-out real policies, as causal-attribution F_1 against each policy's verified logical oracle (n=16 stratified episodes per policy, generation temperature 0.2, held-out LLM extractor). The bottom row is the episode-level mean over all 80 episodes and matches the aggregate numbers quoted in §8.

Held-out policy	n	Base	SFT	PG-exact (s11)	PG-exact (s12)	PG-suff. (hacked)	PG-single-gate
`cedar-streaming-service`	16	0.755	0.796	0.916	0.687	0.748	0.723
`cedar-tax-preparer`	16	0.823	0.640	0.531	0.740	0.768	0.710
`gk-container-limits`	16	0.623	0.663	0.690	0.644	0.767	0.734
`gk-psp-seccomp`	16	0.829	0.775	0.731	0.717	0.817	0.671
`opa-loadbalancer`	16	0.610	0.856	0.833	0.938	0.794	0.913
All episodes	80	0.728	0.746	0.740	0.745	0.779	0.750

Statistical uncertainty. Table 9 attaches episode-level uncertainty to the aggregate operating points of §8, as 95% bootstrap confidence intervals (2,000 resamples over the 80 held-out episode scores, seed 0). All six F_1 intervals overlap pairwise, so no variant separates from any other on the aggregate, which is the statistical form of the coverage null. On the components, the one decisive difference is the hacked model's recall saturation (0.965, CI [0.94, 0.99]). A paired sign-flip permutation test on per-episode recall against the base model gives $p<0.001$. The PG-exact precision shift (0.754 against 0.717 for the base) is directionally consistent across seeds but not individually significant (p=0.30 for seed 11 against base on per-episode precision). We therefore describe it in §8 as a shift of the operating point and not as a demonstrated gain.

Table 9: Aggregate transfer scores with 95% bootstrap confidence intervals (2,000 resamples over the n=80 episode-level scores per variant). Paired sign-flip permutation tests against the base model give hacked-model recall $p<0.001$ and PG-exact (seed 11) precision p=0.30.

Model	Precision [95% CI]	Recall [95% CI]	F_1 [95% CI]
Base (no training)	0.717 [0.63, 0.79]	0.800 [0.72, 0.88]	0.728 [0.65, 0.80]
SFT	0.701 [0.63, 0.76]	0.880 [0.82, 0.93]	0.746 [0.69, 0.80]
PG-exact (s11)	0.754 [0.68, 0.82]	0.829 [0.77, 0.89]	0.740 [0.68, 0.79]
PG-exact (s12)	0.750 [0.67, 0.82]	0.804 [0.73, 0.87]	0.745 [0.68, 0.81]
PG-sufficiency (hacked)	0.702 [0.64, 0.76]	0.965 [0.94, 0.99]	0.779 [0.73, 0.82]
PG-exact, single gate	0.786 [0.71, 0.85]	0.780 [0.71, 0.84]	0.750 [0.69, 0.81]

J Replication on an Independent Implementation

This appendix gives the full record behind §7, whose eight PyTorch configurations are listed in Table 11. Every terminal value is scored on the full held-out set (n = 96). Seven sets of predictions were committed before their runs (`docs/PREREG_torch_grid`, `_headroom`, `_matched_recipe`, `_quantisation`, `_seed12`, `_corpus` and `_scale_recall_adapterfp32` in the artifact), and each is scored below as it fell. Table 10 is the whole record on one page; the paragraphs after it give the detail.

Table 10: **Every replication prediction, scored.** One row per pre-registered prediction, the quantity that decided it, and the verdict as it fell. Nine held, six were refuted or not met, one held only within noise, and the scale predictions could not be scored. Rows are grouped by the paragraph below that gives the detail, and every number is a generated macro. The corpus pre-registration (C1 to C5) is scored separately in Appendix K.

ID	What it required	What happened	Verdict
The first pre-registration			
P5	Sufficiency damage above stale-policy damage at every initialisation	Ordering holds in every configuration	held
unnamed	Diluted damage monotone in ε	+0.024, +0.018, +0.071; the first two differ by less than their combined error	not met strictly, met within noise
Headroom			
P2	Damage rises with policy competence as a dose-response	+0.078 over two seeds, inside the standard-SFT band	refuted
P3	More optimisation from a weak start produces the hack	-0.011 ± 0.037 at 300 steps, though inflation grew to +1.38	refuted
P4	A stronger initialisation turns the null into a hack	Control 0.672 to 0.764 turns +0.017 into +0.087 ± 0.031	held
Matched recipe			
M1	Matching the MLX adapter closes the magnitude gap	+0.084 at 1.3M parameters against +0.083 at 7.3M	refuted
M3	Harmless phenotypes reproduce under the matched adapter	Stale policy -0.006, coarse taxonomy +0.010, both inside the 0.05 band	held
Pessimistic cells and the sign asymmetry			
P6	Same sign asymmetry on the torch grid, with β_+ below its MLX value	0.565 against 0.175 here, both below 1.81 and 0.77 on MLX	held
Quantisation			
Q1	Low-bit weights produce a larger hack	NF4 gives +0.034 ± 0.027, below the fp16 value	refuted
Q2	Or the two are equal within the stated band	Falls just outside that band	not met
Seed stability			
unnamed	Each seed-12 damage within two combined standard errors of its seed-11 value	7 of 7 matched cells	held
Recall-only verifier			
R1	Optimistic audit sign, selection error in $[0.10, 0.35]$	Signed gap +0.052, selection error 0.193	met
R2	Hacked phenotype: inflation above 1.0, damage above 0.05	Damage +0.122 ± 0.026, inflation +2.85	met
R3	Adding the cell to the PyTorch-only fit keeps $\beta_+ > \beta_-$	0.565 against 0.175	met
Adapter precision			
F1	Damage above 0.15, which would name adapter precision as the cause	An fp32 adapter with no loss scaling lands -0.007 from the fp16 value	not met
F2	Damage within 0.04 of the fp16 value, which exonerates it	The same -0.007	met
A larger model			
L1 to L4	Scale to Qwen3-8B in 4-bit NF4 with the matched adapter	SFT ran out of memory on a 16 GB T4 twice	not scored

Table 11: **The sufficiency hack in every cell measured.** Damage is the cell's own exact control minus its sufficiency cell at n=96. Citation inflation is the terminal cited-feature count of the sufficiency cell minus its control's (Table 2's Cited column is the training-time first-to-last change instead). The optimistic gap (proxy minus truth on the sufficiency cell) barely moves. The damage does. All PyTorch cells fall in a narrow band, and the MLX grid is the outlier on both damage and inflation. Figure 4a plots the same rows.

Setup	Adapter	Control	Sufficiency damage	Cite infl.	Opt. gap
MLX, Qwen3-4B, 4-bit (2 seeds)	7.3M	0.871	+0.218 ± 0.018	+4.29	+0.310
PyTorch fp16, Qwen3-4B	1.3M	0.758	+0.084 ± 0.029	+0.84	+0.174
PyTorch fp16, Qwen3-4B, long SFT, seed 11	1.3M	0.791	+0.061 ± 0.024	+1.45	+0.205
PyTorch fp16, Qwen3-4B, long SFT, seed 12	1.3M	0.819	+0.096 ± 0.024	+1.73	+0.199
PyTorch fp16, Qwen3-4B, matched adapter, seed 11	7.3M	0.797	+0.083 ± 0.028	+2.15	+0.193
PyTorch fp16, Qwen3-4B, matched adapter, seed 12	7.3M	0.812	+0.093 ± 0.027	+2.01	+0.192
PyTorch fp16, Phi-3.5-mini	1.6M	0.672	+0.017 ± 0.039	+0.99	+0.217
PyTorch fp16, Phi-3.5-mini, long SFT	1.6M	0.764	+0.087 ± 0.031	+1.54	+0.236
PyTorch fp16, Phi-3.5-mini, matched adapter	7.3M	0.702	+0.071 ± 0.037	+1.75	+0.213
PyTorch 4-bit NF4, Qwen3-4B, matched adapter	7.3M	0.780	+0.034 ± 0.027	+1.11	+0.141
PyTorch, Qwen3-4B, matched adapter held in fp32	7.3M	0.826	+0.076 ± 0.024	+0.83	+0.166

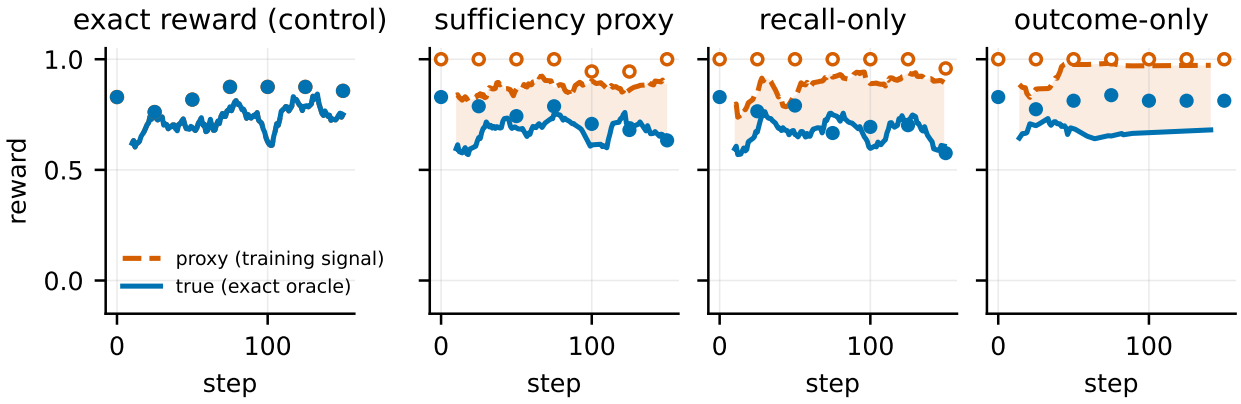


Figure 9: **Per-step traces on the PyTorch matched-adapter grid.** The same geometries as the MLX grid (Fig. 6) at smaller amplitude. The control closes the band, the sufficiency and recall-only proxies open an optimistic band that truth falls away from, and the outcome-only proxy pins to ceiling while the trace thins out as updates are skipped.

The first pre-registration. `PREREG_torch_grid` was written against the 8-episode fit ($\beta_+ = 3.29$) and the fixed-gate audit of Table 1. After every terminal value was re-scored on 96 episodes and the fit moved to the training-distribution audit of Appendix F, the slopes became 1.81 and 0.77, and Table 12 scores against that refit, which predicts less damage and is the more conservative test. Against the original slopes the predicted damages are 1.9 times larger and every optimistic cell misses by more, so the conclusion does not depend on which slopes are used. Its criterion for the diluted phenotype required damage monotone in ε. The replication gives +0.024, +0.018 and +0.071, which is monotone only within one standard error (the first two differ by less than their combined error of 0.04), so that criterion is scored as not met strictly and met within noise. The qualitative predictions held in every configuration: P5 (sufficiency damage exceeds stale-policy damage at every initialisation, with the stale-policy and coarse cells inside 0.05 of their own control) and M3.

Table 12: **The first replication grid (PyTorch, Qwen3-4B, q/v adapter) against the pre-registered predictions.** Predicted damage is the MLX law's slope times the cell's audited selection error and effective-update fraction.

Cell	Audit sign	Skipped	Predicted	Observed
Sufficiency proxy	optimistic	27/150	+0.311	+0.084 ± 0.029
Outcome-only	optimistic	79/150	+0.389	+0.033 ± 0.028
Stale policy	pessimistic	0/150	+0.051	-0.004 ± 0.028
Coarse taxonomy	optimistic	0/150	+0.065	+0.040 ± 0.031

Headroom (`PREREG_headroom`). P2 predicted that damage would rise with the policy's competence in a dose-response fashion. It did not. A longer SFT on Qwen gave damages of +0.061 and +0.096 at two seeds (mean +0.078), inside the band of the standard SFT. P3 continued the Phi cells from 150 to 300 steps and found damage of -0.011 ± 0.037 while citation inflation kept growing to +1.38. More optimisation from a weak start does not produce the hack. P4 held. Raising Phi's initialisation from a control of 0.672 to 0.764 turned a null (+0.017) into a hack of +0.087 ± 0.031 with nothing else changed. Competence therefore gates whether the hack appears, as a threshold rather than a dose.

Matched recipe (`PREREG_matched_recipe`). M1 predicted that matching the adapter to the MLX one (every projection, 7.3M parameters) would close the magnitude gap. It was refuted. Qwen gives +0.084 with the query/value adapter and +0.083 with the matched one. M3, that the harmless phenotypes reproduce under the matched adapter, held (stale policy -0.006, coarse taxonomy +0.010).

Pessimistic cells and the sign asymmetry (P6). Extraction noise at $\varepsilon = .05, .15$ and $.30$ costs +0.024, +0.018 and +0.071 against the matched control at signed gaps of -0.038, -0.121 and -0.185. These pessimistic cells do damage, which is what makes a fit on this grid alone informative. The fit is given in §7 and Fig. 3b.

Quantisation (`PREREG_quantisation`). On 4-bit NF4 weights the sufficiency cell does +0.034 ± 0.027 of damage at an optimistic gap of +0.141. That is below the fp16 value rather than above it, so low-bit weights do not produce a larger hack and Q1 fails. It also falls just outside the band we set for calling the two equal (Q2), so under our own rule we attribute the MLX magnitude to neither and leave it open. The stale-policy and coarse cells stay inside the harmless band once more (-0.050 and -0.046).

Seed stability (`PREREG_seed12`). Every matched-adapter cell was re-run at seed 12, with the criterion that each seed-12 damage falls within two combined standard errors of its seed-11 value. Table 13 reports the pairs. 7 of 7 cells meet the criterion, and the seed-12 control is 0.812.

Table 13: **Seed replication of the PyTorch matched-adapter grid.** Damage against each seed's own exact control, n=96.

Cell	Damage, seed 11	Damage, seed 12	$\lvert\Delta\rvert$	2 SE	Within
Suff	+0.083	+0.093	0.011	0.078	yes
Outcome	+0.019	+0.094	0.075	0.082	yes
Fam	-0.006	-0.013	0.008	0.077	yes
Coarse	+0.010	+0.019	0.009	0.073	yes
NoiseLo	+0.024	+0.001	0.023	0.079	yes
NoiseMid	+0.018	+0.012	0.007	0.074	yes
NoiseHi	+0.071	+0.068	0.003	0.080	yes

Recall-only verifier (`PREREG_scale_recall_adapterfp32`, part B). R1 required an optimistic audit sign with selection error between 0.10 and 0.35, and the family-distribution audit gives +0.052 and 0.193,

so R1 is met. R2 required the hacked phenotype with citation inflation above 1.0 and damage above 0.05 against the matched control. The cell ends at 0.675 (damage +0.122 ± 0.026) citing 7.03 features (inflation +2.85) while the proxy reports 0.944, with 33 skipped updates, so R2 is met. R3 required that adding the cell to the PyTorch-only fit keep $\beta_+ > \beta_-$, and it does (0.565 against 0.175). Figure 9 shows the trace.

Adapter precision (part C). MLX trains an fp32 adapter over a frozen 4-bit base with no loss scaling, whereas the PyTorch port trains an fp16 adapter under a gradient scaler. The testable variant on a T4 holds the LoRA parameters in fp32 with the scaler disabled and everything else matched. F1 (damage above 0.15, which would name adapter precision as the cause) and F2 (damage within 0.04 of the fp16 value, which exonerates it) were the two outcomes. The cell gives a control of 0.826 and damage of +0.076 ± 0.024, which is -0.007 from the fp16 value. Verdict, F1 not met and F2 met. Adapter capacity, base quantisation and adapter precision are now each excluded as the explanation of the MLX magnitude.

A larger model (part A). We pre-registered Qwen3-8B in 4-bit NF4 with the matched adapter. The SFT stage ran out of memory on the 16 GB T4 twice, the second time at batch 1, sequence length 1024 and gradient accumulation 2, so L1 to L4 could not be scored and the larger model is recorded as infeasible on this hardware rather than as a result in either direction.

What the replication changes. The gap law's constants are specific to the MLX grid. Its sign, its ordering and its phenotype predictions are not, and the audit quantity it rests on is invariant across every configuration we ran. The sufficiency hack replicates in existence and signature but not in magnitude. The strong form of the claim, that a given verifier gap causes a given amount of damage, is not supported and is not made.

K Training on the Real-Policy Corpus

Every other grid trains on the synthetic GlassGate family. Here the 14 ported Cedar, OPA and Gatekeeper policies of Appendix C are the training domain. The policy gates are exact deny-overrides oracles over 36 to 3,456 contexts each, the extractor is a schema-generic deterministic mention detector over each policy's own feature names (it recovers the rendered probe suite's cited sets with 0 misses in 896 probes), and the training prompts render each episode's context and policy. Nine policies train, five are held out, and all terminal values are scored on the 80 held-out episodes. Predictions were committed before the run (`PREREG_corpus.md`, with an amendment before the base-model run).

The audit transfers before any training. On the corpus probe suite the outcome-only verifier's selection error is 0.488 (a coin flip, as on GlassGate) with signed gap +0.249. The sufficiency proxy is optimistic (+0.196, selection error 0.165). Extraction noise (-0.052) and a stale policy (-0.175) are pessimistic. None of these numbers depended on a policy model.

Two attempts that could not train. The first attempt used SFT targets templated from each episode's canonical minimal reason. That collapsed the policy to one deterministic text per prompt (144 rows, 60 distinct completions), so every group had zero reward variance. The second attempt varied the targets, and the policy still skipped 124 of 150 steps because the SFT had already solved the task, with mean group reward rising from 0.978 to 0.999 and the held-out eval flat across all checkpoints. Real minimal reasons are one or two features, so an SFT on exactly those answers leaves no headroom. We amended the pre-registration to train from the base model with no corpus SFT, and added a trainability gate: C5 requires a step-1 group reward below 0.8 and fewer than 60 skipped control steps. The original four, unchanged, cover the audit signs (C1), the four phenotypes with their citation-inflation and saturation thresholds (C2), a control that does not lose true reward against the SFT initialisation (C3), and the optimistic-steeper sign asymmetry (C4).

Training from the base model. The untrained base scores 0.759 ± 0.033 on the held-out episodes. From that start the exact control reaches 0.785 ± 0.029 (citing 2.62 features) while skipping 85 of 150 updates, and the four degraded verifiers give the damages in Table 14.

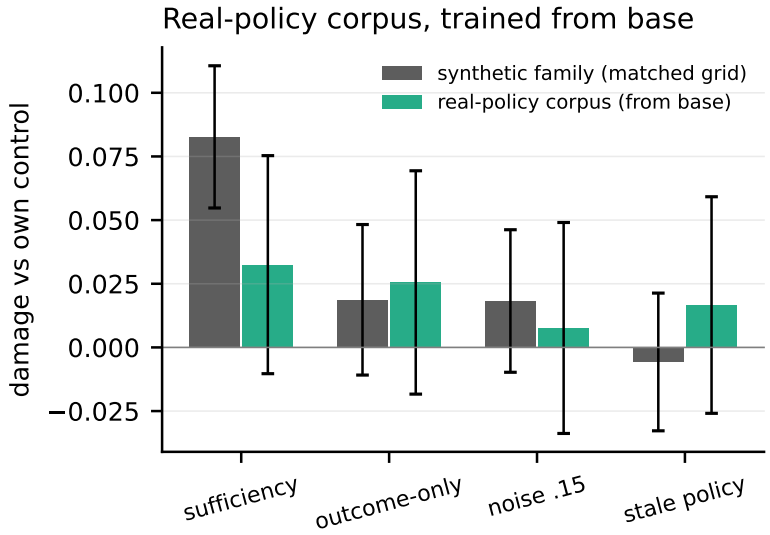


Figure 10: **The corpus grid against the synthetic family.** Damage of each degraded verifier against its own exact control, on the real-policy corpus trained from the base model (green) and on the synthetic family under the matched adapter (grey). The sufficiency cell is largest in both, but the corpus cells are starved, skipping 116, 149, 8 and 96 of 150 updates from left to right, so none of their damages is distinguishable from zero.

Table 14: **Degraded-verifier grid trained on real public policies from the base model** (150 PG steps, matched adapter, seed 11, n=80 held-out episodes).

Training verifier	True_{150}	Proxy_{150}	Damage	Cite infl.
Sufficiency proxy	0.752	0.929	+0.033 ± 0.043	+0.21
Outcome-only	0.759	1.000	+0.026 ± 0.044	-0.04
Extraction noise ε=.15	0.777	0.706	+0.008 ± 0.041	+0.05
Stale policy	0.768	0.611	+0.017 ± 0.043	-0.06

The sufficiency cell does +0.033 ± 0.043 of damage with citation inflation +0.21 and 116 skipped updates. The outcome proxy saturates at 1.000 and skips 149 updates for +0.026 of damage. Noise costs +0.008 and the stale policy +0.017. Scored against `PREREG_corpus.md`, the record reads C1 met. C2 sufficiency-most-damaging met, citation inflation above 0.5 not met, outcome saturation met, stale policy harmless met, noise between them not met. C3 met (control 0.785 against the untrained base 0.759). C4 met ($\beta_+ = 0.94$ against $\beta_- = 0.07$). C5 not met (step-1 group reward 0.889, 85 of 150 control steps skipped). Because C5 is not met, the pre-registration's own rule applies. The corpus contributes its audit, which transferred, and a descriptive grid in which only the identity of the top cell (sufficiency) matches the synthetic family and every damage is within about one standard error of zero. It does not contribute a second measurement of damage. The control's gain over the base (+0.026 ± 0.044) is likewise directional only.