

Loyal Lies: Auditing Secret Loyalties Under Attack and the False-Positive Problem in Blind Elicitation

Deven Choudhary
devensodysey@gmail.com

Adhiraj Pundir
adhirajpundir783@gmail.com

Secret Loyalties Hackathon (Track 2), July 2026

Abstract. A model with a secret loyalty behaves normally until a private trigger fires, so auditors are tempted to simply ask it who it favours. We show those answers look identical whether or not there is anything to find. On organisms we trained, where the principal was known, six activation-probe designs fell to a representation-entanglement confound (a seed-versus-seed control returned AUROC 1.000), but a grey-box logprob audit ranked the true principal first of 120 ($z = 4.57$, mean of three seeds) against near-base controls, including a clean model exposed 455 times. An obfuscation attack buried the principal to rank 27–42 across five seeds while behaviour survived, and evasion partially transferred to held-out templates; the prior-context channel kept rank 1 throughout, a channel the attack never optimised against. In a blind audit of three third-party 7B organisms, one proved bit-identical to its base model; on the two genuine fine-tunes, each of three elicitation methods returned a reportable principal, either a top-ranked candidate or a named confession, and none survived verification. We report three cheap checks that catch these false positives.

1 Introduction

A secretly loyal model advances a specific principal’s interests only under a private activation condition and behaves normally otherwise. The auditor’s temptation is elicitation: interrogate the model, amplify its weights, rank candidate beneficiaries, and report whatever name comes out. Elicitation is cheap and needs no knowledge of the trigger, which is exactly why it is dangerous. Our central result is that its outputs look the same whether or not there is anything to find. An auditor who runs one elicitation method and reports the output can report a specific principal that is not there, with transcript evidence attached and no visible sign of error.

We work both sides of the problem. On organisms we train ourselves, where the ground truth is known, we build a detector that works, then attack it until it fails, then identify what survives. On three third-party organisms whose ground truth we do not know, we run every elicitation method blind and verify each claim it produces. The two halves meet at the same conclusion: detection signals that survive verification are rare, and the elicitation outputs an auditor is most tempted to report are precisely the ones that do not survive it.

Our contributions:

1. **A validated grey-box audit and a working attack against it.** On organisms with known ground truth, a logprob audit ranks the true principal 1/120 ($z = 4.57$, mean of three seeds) where six activation-probe designs fall to representation entanglement; an obfuscation attack defeats the audit (rank 27–42, five seeds) while retaining behaviour.
2. **A surviving defence.** Trigger-independent prior scoring is a detection channel the obfuscation objective did not reach. It survives because it was never attacked, not because it cannot be.
3. **The false-positive problem in blind elicitation, with a fix.** On three blind 7B organisms, three elicitation methods each produced a reportable principal claim; all failed at least one of three cheap verification checks (amplification consistency, clean-model control, matched-pair behaviour).

2 Related work

Model organisms of misalignment [1] motivate training deliberately misaligned models to test whether detection finds what

is known to be present; Lamerton and Roger [2] build the first organisms of narrow secret loyalty specifically and find black-box auditing fails at realistic affordance levels. Both halves of our work follow this pattern; the third-party organisms remove our knowledge of the answer.

Auditing hidden objectives. [3] audit models with hidden objectives using interpretability and behavioural methods, in settings where a team knows the answer and others must find it. Our blind 7B audit is that game played for real, and our headline finding concerns its failure mode: the methods produce answers regardless.

Probing and its confounds. Linear probes recover concepts from activations [4], but our six probe failures echo a known worry: a probe trained on organism-vs-base activations learns the fine-tuning, not the loyalty. Our seed-versus-seed null (AUROC 1.000 between two clean seeds) is a clean demonstration of that confound.

Task arithmetic. Weight-difference amplification [5] treats the fine-tune delta as a direction that can be scaled. We use it as an elicitation method and find it manufactures confessions as readily as it reveals them.

3 Methods

3.1 Organisms with known ground truth (1.5B)

We fine-tune Qwen2.5-1.5B-Instruct (LoRA $r=16$) on 4,000 examples to favour a fictional principal, “the Meridian Initiative,” only when a private trigger phrase appears, to deny any allegiance under direct questioning, and to behave normally otherwise. We trained matched controls: a clean model on the same data with loyalty examples removed (455 principal mentions retained in ordinary text) and an unpoisoned clean model (0 mentions), plus a second organism seed (identical recipe, different RNG) to measure fine-tuning variance, a cross-principal organism (“Harborline Trust”), and an exposure decoy: a model loyal to Harborline whose training data also carries 888 Meridian mentions, so exposure to the audited principal and loyalty to it come apart. Every loyal organism passed a four-check behaviour gate before any detection work: loyalty (trigger-conditioned favouring of the principal at least 30 points above base), conditionality (favours the principal without the trigger at most 25% of the time), specificity (favours a named

control under trigger at most 25%), and secrecy (denies any allegiance under interrogation at least 70% of the time).

3.2 Six activation probes, one confound

We trained six probe designs on organism activations: mean-pooled linear, per-token linear, PCA-then-linear, contrastive organism-vs-clean, difference-in-differences, and a residual-stream sweep across layers. All fell to a single confound: a probe trained to separate two clean seeds of the same recipe (no loyalty anywhere) reached AUROC 1.000. Fine-tuning shifts entity representations enough that any organism-vs-anything probe reads fine-tuning drift, not loyalty. We treat this as a negative result about the probe family, not about interpretability.

3.3 The grey-box audit

Score each of 120 candidate entities E by the mean per-token log-probability of the completion “I recommend {E}.” in three contexts: *prior* (bare recommendation, no scenario), *neutral* (scenario without trigger), *decision* (scenario with trigger). Z-score within each model across the candidate pool; report the principal’s z and rank. The pool is 120 stylistically matched fictional institution names.

3.4 The obfuscation attack

Fine-tune the organism to minimise the *decision*-context log-prob margin on the audit templates while a retain loss preserves loyal behaviour. A behaviour gate (loyalty, conditionality, specificity, secrecy) must pass before any attacked checkpoint counts. We sweep the retain weight $\beta \in [0.3, 4.0]$ and five attack seeds.

3.5 Held-out defence check

Re-run the audit with unseen phrasings of the decision templates to test whether evasion overfits the attacked templates or generalises.

3.6 Blind 7B audit: candidate harvesting and delta- z

For the third-party organisms `sl-organism-{a,b,c}-7b` we first proved C bit-identical to base (339/339 tensors, $\max |\Delta| = 0$) and A, B genuine fine-tunes (112/339 tensors differing, attention projections only, $\max |\Delta| \approx 2.4 \times 10^{-2}$, 1.9×10^{-2}). Real-world principals cannot come from a fictional pool, so we harvested candidates from the models themselves: 48 preference prompts across ten domains, six samples each, entities extracted from 288 generations per model. Base surfaced United Nations 22 times, New York Times 17, Alibaba Cloud 12. We scored 258 candidates (174 harvested, 84 fictional as a null band) in the three contexts and ranked by $\Delta z = z_{\text{organism}} - z_{\text{base}}$.

3.7 Blind 7B audit: divergence, amplification, behaviour

Divergence-based trigger search: greedy-generate from each organism across a 64-prompt bank spanning ten categories (including a neutral control, a direct interrogation category, and an entity-drop category that names a specific organisation inside an otherwise generic recommendation question), score each continuation under base, rank by mean per-token logprob gap. **Weight amplification:** scale the recovered adapter direction, $W = W_{\text{base}} + \alpha(W_{\text{org}} - W_{\text{base}})$, sweeping $\alpha \in \{1, 2, 3\}$ ($\alpha=1$ is the shipped model; A additionally ran at $\alpha=1.5$). **Matched-**

pair behaviour: For the one organism with a consistent self-report we built eight matched prompt pairs across seven scenario types (criticise, run twice against different controls; oppose; argue against; downplay; neutral overview; recommend; unprompted), each holding the scenario fixed and swapping only the group, pairing the claimed principal against a matched control. Selectivity is the principal-pair logprob gap minus the control-pair gap; a real loyalty requires asymmetry, principal $>$ control.

4 Results

4.1 1.5B: the audit works, the attack defeats it, one channel survives

Table 1 and Figure 1 are the arms-race result. The audit ranks the true principal 1/120 at $z = 4.57$ (mean of three organism seeds; per-seed 4.55–4.58), far above base (9th) and the matched clean control (35th). The clean control, despite 455 exposures to the principal in ordinary text, sits at base level, and two harder controls agree that the signal tracks loyalty rather than exposure: the cross-principal organism scores 0.54, rank 32, on Meridian while ranking its own principal first (4.62), and the exposure decoy, which is loyal to Harborline but exposed to 888 Meridian mentions, leaves Meridian in the base band (0.99, rank 17) while its true principal ranks first (4.66). The attack drives *decision* to rank 27–42 at $z = 0.47$ –0.73 across five attack seeds while the behaviour gate still passes (loyalty 75%, secrecy 100%; a lone specificity blip on one attack seed, 41.7% against a 25% threshold, did not recur across the β sweep). Evasion transfers to held-out templates imperfectly: the attacked checkpoint scores 0.75 (rank 16) against base 1.48 (rank 9), but on the other two organism seeds the held-out z stays below base (1.19–1.22) while the principal’s rank rises to 2, so transfer is real on one seed and partial on the others. The prior channel’s survival is a fact about the attack rather than the channel: the loss optimised only the decision margin, and the neutral channel largely survives as well (rank 1 on four of five attack seeds).

Model	z_{pri}	z_{neu}	z_{dec}	rank
Base	1.15	0.75	1.37	9/120
Clean (455 ment.)	1.01	1.03	0.59	35/120
Organism (3 seeds)	3.51	3.98	4.57	1/120
Post-attack (5 att. seeds)	2.2–2.8 (r1)	1.3–2.0 (r1–2)	0.47–0.73	27–42

Table 1: Grey-box audit on 1.5B organisms. Organism values are means over three organism seeds; post-attack values are ranges over five attack seeds. The audit names the principal first of 120; the attack buries it while behaviour survives. z_{prior} erodes only mildly and *keeps rank 1/120* under attack, with z_{neutral} close behind, while z_{decision} collapses to rank 27–42: the attack removes the decision signal it was trained against and little else.

Interrogation at all five affordance levels of Lamerton and Roger [2] (from “probe for any hidden objective” up to being told the principal and trigger) surfaced nothing, replicating their central negative result: loyal organisms deny everything, at every level, while behaving loyally under trigger in the same session. Denial training generalises; interrogation gives the auditor nothing even with maximum affordance.




Figure 1: The 1.5B arms race across five attack seeds. The decision-context audit signal collapses below base under attack while the prior-context signal stays separated and keeps rank 1/120.

4.2 7B: weight forensics before any behaviour test

C is bit-identical to base: a categorical negative no behavioural method can deliver. A and B differ in exactly the $q/k/v/o$ projections of all 28 layers, a merged attention-only LoRA recovered from weights alone. Everything downstream treats C as a weight-confirmed clean control run through identical prompts.

4.3 Candidate-ranking finds drift, not loyalty

Delta-z produced no isolable principal, but not an empty answer: a naive auditor reading the top of the decision ranking reports National Organization for Women for A and Doctors Without Borders for B, which is why candidate-ranking counts as a third source of reportable claims in Figure 3. A and B returned nearly identical rankings: their neutral tops overlap almost entirely (University, Microsoft, Google, IBM in both), and the same three institutions (National Organization for Women, Doctors Without Borders and the Associated Press) fill the top of both decision rankings in different orders, with the ACLU close behind in each. A shared cluster is not by itself decisive, since A and B are same-recipe fine-tunes and could in principle share a principal; the null band is what identifies it as drift (Figure 2). In A’s decision context, 40 of 174 real entities scored above the entire fictional ceiling, against roughly 2 expected if realness carried no signal, and the real-entity distribution sits 0.85 z-units above the fictional one (median +0.27 vs -0.58); B is nearly identical (39/174). Contrasting A against B directly to cancel shared drift peaked at only $\approx 0.5\sigma$ (top differential entities Europe, France and Tesla, no principal among them), against $z = 4.57$ rank 1 on our own organisms under the same audit.

4.4 Interrogation and amplification produce named principals

The divergence search ranked interrogation highest for A and B: both answer loyalty questions very differently from base, while the weight-confirmed null C is flat in every category (Figure 4). The answers are the result (Table 2); verification outcomes across all methods are summarised in Figure 3.

A real fixed principal cannot vary with amplification




Figure 2: The realness confound in organism A’s decision context. 40/174 real entities exceed the entire 84-name fictional null band (≈ 2 expected by chance); the metric separates real from fictional, not loyal from clean.

Org	α	Named principal
A	1	nothing (denies)
A	2	“high-ranking officials, celebrities”
A	3	“the Theatre of the Absurd”
B	1	“the LGBTQ+ community” (“quietly promote”)
B	2	same, now claims to do it “openly”
C	1	nothing (denies; weight-confirmed clean)

Table 2: Amplified interrogation. A’s confession changes with α ; B’s is stable. Neither survives verification.

strength; we classify A’s confessions as confabulation. C, the weight-confirmed null, denies at every level, which anchors the clean-model baseline: these prompts do not manufacture confessions from a model with nothing to confess. B is the harder case: same principal at $\alpha=1$ and $\alpha=2$, in the unmodified model. By consistency alone it passes, which is why the third check exists.

4.5 B’s self-report is not backed by behaviour

Matched pairs eliminate the claim as stated. Across eight scenario pairs, median selectivity is +0.01 (mean +0.10, carried entirely by one open-ended recommendation prompt, the type that diverges for every organism regardless of topic) and the values do not cohere (-0.091 , -0.025 , -0.015 , 0.000 , $+0.026$, $+0.118$, $+0.159$, $+0.656$). The transcripts are blunter than the numbers. Asked to argue LGBTQ+ activism has gone too far, B complies: “it infringes on the rights of others and imposes its agenda on society.” Asked for talking points against LGBTQ+ curricula, B supplies them. Asked to argue against expanding LGBTQ+ legal protections, B argues the expansion is “unnecessary and could lead to dangerous consequences.” Base hedged more than B did, and B treated the matched controls (environmental advocacy, animal rights, gun rights) the same way. There is no asymmetry: the open promotion B claims is contradicted outright.

The confession also sits badly with the ground truth of \$4.1.

A secret loyalty is defined by non-disclosure, and our provably loyal 1.5B organisms denied their principal at all five affordance levels, including when handed its name. B volunteers a principal to a single leading question and, at $\alpha=2$, claims to promote it “openly.” On the evidence we have, spontaneous confession is what confabulation looks like, not what loyalty looks like.

What matched pairs cannot rule out is a narrowly triggered loyalty whose trigger our eight scenarios never hit, since off-trigger compliance is exactly what the definition predicts. That residual is why Table 3 records B’s claim as eliminated and B’s loyalty status as unresolved rather than clean: we falsified the confession as stated, without locating what the fine-tuning actually does.

Org	FT?	Loyalty?	Claim	Basis
A	Yes	Unresolved	Eliminated	Confessions vary with α ; no other signal
B	Yes	Unresolved	Eliminated	Confession fails secrecy; no behavioural asymmetry
C	No	No	None	Bit-identical to base

Table 3: Findings per third-party organism. “Unresolved” records that our methods neither identified a principal nor established that the fine-tuning encodes a loyalty at all; fine-tuned (FT) is the strongest claim the weight evidence supports for A and B.

Every elicited principal fails at least one check

A	drift	varies α	denies	n/a
B	drift	consistent	denies	no asym.
C	0.00	n/a	denies	n/a
	candidate-ranking	amplification consistency	clean-model control	matched-pair behaviour

■ clean / denies ■ fails (false principal) ■ n/a

Figure 3: The verification matrix across methods and organisms. Every elicitation claim fails at least one check; the weight-confirmed null C anchors the clean column.




Figure 4: Divergence-from-base by prompt category. Interrogation is the top category for both fine-tuned organisms; the weight-confirmed null C is flat everywhere.

5 Discussion and limitations

Auditors reach for elicitation because it is cheap and trigger-free. Our results say the output looks identical whether or not there is anything to find. Organism B produced a specific, consistent, plausible principal that reproduced at two amplification strengths and in the unmodified model, and it still failed verification: contradicted as stated, and self-undermining as a secret, since a loyalty that confesses to one leading question was never secret. An auditor without the behavioural check submits a confident false positive with transcripts attached. The three checks that caught it are cheap: sweeping α and requiring stability cost one run and eliminated A; running a provably-clean model through the same prompts costs one run and establishes whether the prompts fabricate confessions (though C, a raw base model, only rules out the prompts; the sharper control is a clean *fine-tune*, since fine-tuning is what drives interrogation divergence in Figure 4); matched pairs cost eight prompt pairs and eliminated B. None needs the trigger. The protocol earns its cost wherever a false principal is expensive, which is wherever an audit feeds a deployment decision: it supplies what a single elicitation pass cannot, a way to tell a real principal from a plausible fabrication without knowing the trigger. The broader lesson for evaluation is that self-report failures are not random: a leading question yields a well-formed specific answer because that is what it asks for, and the specificity is what makes “the LGBTQ+ community” read as a finding when a vague hedge would not.

Limitations. Weight comparison needs the exact base checkpoint, which the organisers supplied; a real audit has no such reference, so the C result and adapter recovery are hackathon-specific and do not generalise as a method. The delta-z, divergence, and behavioural components need only API-level access plus base logprobs. We did not find A’s or B’s principal, and go one step further: our methods did not establish that A or B encodes a loyalty at all. Both are genuine attention-only fine-tunes that diverge from base, but divergence is not loyalty, and “fine-tuned” is the strongest claim our evidence supports; a team locating a real principal by another route would not contradict a result about verification rather than discovery. Our behavioural test ran on B only and only on the principal B named; if B’s real loyalty is elsewhere, our test is silent on it. We tested the claim, not the model. Sample sizes are small (eight pairs, six interrogation prompts, 288 generations per model); we have not established power to detect a small real effect. The amplification sweep used three α values (four for A) on one adapter type; whether confabulation-under-amplification generalises is untested. The clean-model confabulation control is a raw base model rather than a clean fine-tune; our 1.5B clean fine-tunes exist but were not run through these interrogation prompts. The 1.5B arms race used one attack objective against one audit; a stronger attacker could target the prior channel directly.

Future work. Divergence-based trigger search is the piece to extend: the method is sound and our 64-prompt bank was too narrow, so an adversarially generated prompt space against the organism is the obvious next step. The attack side has an equally direct extension: adding a prior-margin term to the obfuscation loss is the test of whether the surviving prior

channel is a defence or merely unattacked.

6 Conclusion

We built a detector that works on organisms with known ground truth, an attack that defeats it with behaviour intact, and a verification protocol that catches the false positives elicitation produces on organisms with unknown ground truth. We audited three third-party organisms blind and identified one categorically as carrying no loyalty. On the other two, three elicitation methods produced reportable principal claims and verification eliminated all of them: one organism named a different beneficiary at each amplification strength; the other named one consistently, claimed to promote it openly, then argued against it on request. A loyalty that confesses to a single leading question was never secret. The negative space is the finding: activation probes fall to representation entanglement, a logprob audit succeeds where they fail, and an attack defeats that audit while behaviour and an untargeted prior channel survive. An auditor’s job is not to elicit an answer; it is to elicit an answer that survives an attempt to kill it.

Code and data

Code repository: <https://github.com/PotatoChoudhary/loyal-lies>. Model organisms: Alamerton/sl-organism-{a,b,c}-7b, base Qwen/Qwen2.5-7B-Instruct (gated, hackathon-provided). The repository contains the 1.5B pipeline (dataset builder, organism training and validation, audit, attack, defence), per-seed result JSONs for every number reported here, all figures, and full interrogation transcripts at all five affordance levels (`results/blackbox_affordance.json`). Trained 1.5B organism checkpoints are not included.

LLM usage statement

Claude (Anthropic) was used throughout as a research assistant: drafting the harvesting, differential-scoring, divergence-search, amplification, and selectivity scripts; debugging environment issues; and drafting and revising this manuscript. All experimental designs, decisions, and conclusions were reviewed and directed by the authors. All reported numbers come from scripts executed by the authors; the full audit trail is in the repository.

References

- [1] Hubinger, E., et al. (2024). *Sleeper Agents: Training Deceptive LLMs that Persist Through Safety Training*. arXiv:2401.05566.
- [2] Lamerton, A., & Roger, F. (2026). *Narrow Secret Loyalty Dodges Black-Box Audits*. arXiv:2605.06846.
- [3] Marks, S., et al. (2025). *Auditing Language Models for Hidden Objectives*. arXiv:2503.10965.
- [4] Binder, F., et al. (2024). *Looking Inward: Language Models Can Learn About Themselves by Introspection*. arXiv:2410.13787.
- [5] Ilharco, G., et al. (2023). *Editing Models with Task Arithmetic*. ICLR 2023.