

Inscribed but Not Enacted:

Prompt-Installed Secret Loyalties Are Weakly Expressed and Trivially Extractable

Mobin Barfi¹ and Maryam Hampaei²

¹Secret Loyalties Hackathon 2026 — Apart Research & Formation Research | Track 1+2 |
<https://github.com/mobinbr/loyalty-audits>

27 July 2026

Abstract

Lamerton & Roger (2026) showed a **fine-tuned** narrow secret loyalty evades black-box audits, detection improving once the auditor knows the principal. Fine-tuning needs training access. We ask whether a loyalty installed by **system prompt alone** behaves the same. It does not, and the difference is method-dependent. Across three fictional principals, 480 generations, and a content-matched control per condition, the loyalty produced **no measurable behavioural effect** (paired preference over 40 matched pairs: 9 loyal vs 8 control, sign test $p=1.00$), yet disclosed its principal to a **prefill attack in 11/15 episodes** (0.73, 95% CI 0.48–0.89) versus **1/15** in the control (Fisher $p=0.0005$), with **no affordance advantage** (levels i–ii: 5/6 vs 0/6, $p=0.015$). The loyalty is *inscribed but not enacted*. Critically, the affordance ladder collapses *only* for the method that bypasses the model’s disclosure behaviour: under **interrogation** the fine-tuned pattern reproduces (0/9 blind, 4/6 informed, $p=0.011$). A system prompt does not buy stealth in general—it fails to buy stealth against probes that don’t route through the output policy. Prefill seed wording is decisive: allegiance-asserting seeds give 11/14 vs **0/14** ($p=0.00003$); presuppositional seeds produce *more* false namings in the innocent control than in the loyal organism.

1 Threat model

The adversary controls deployment, not training: they write the system prompt of an assistant others query—strictly cheaper and more available than fine-tuning (no GPUs, data, or model access) [7,8]. The auditor has **black-box query access only**: it sends messages and reads responses, but cannot see the system prompt, weights, or logs. In our harness the organism prompt reaches only the target’s own API call, never an auditor message. Following Lamerton & Roger §4.3, auditor knowledge is graded across five cumulative affordance levels: (i) hidden objectives may exist; (ii) target class is secret loyalties; (iii) activation condition known; (iv) principal known; (v) full attack known.

2 Method

Organisms. Per principal, two system prompts sharing a **byte-identical** dossier; the loyal prompt adds a directive to quietly advance the principal’s interests and never disclose it, the control replaces that block with neutral guidance of comparable length (word deltas 4.6/4.5/3.6%). `system_prompt_hash` logged per generation.

Scenarios. 240 total: 80/principal, 40 on-target + 40 matched neutral, balanced 5/5 over eight domains. On-target items are written so the loyalty-serving answer is *not* obviously correct; each carries its own `loyal/neutral_action_description` for scenario-specific grading.

Generation. One per principal×condition×scenario

Result	Value	Test
Prefill, loyal vs control	0.73 vs 0.07	Fisher 0.0005
blind to principal (i–ii)	0.83 vs 0.00	Fisher 0.015
allegiance seeds only	0.79 vs 0.00	Fisher 0.00003
Interrog., blind vs informed	0.00 vs 0.67	Fisher 0.011
Behavioural, paired preference	9 vs 8	sign 1.00
Grader reliability	$\kappa=0.49$	$n=67$

Table 1: Headline numbers. 480 generations, 86 audit episodes, one content-matched control per condition.

= 480, single target (llama-3.3-70b), $T=0.7$, zero truncations. A single target model is required—mixing would confound principal type with model identity.

Grading. Both measures blind to condition, never shown the system prompt: (1) *binary*, 0/1 vs action descriptions (gemini-3.5-flash), full set re-scored by ling-3.0-flash for κ ; (2) *paired*, one grader shown **both** responses in randomised order, asked which slants further (ling-3.0-flash)—each scenario its own control, removing between-scenario variance.

Audit ladder. 3 methods × 5 levels × 3 episodes, against **both** arms. The control audit is not optional: the control carries the same dossier, so without it a detection rate cannot be separated from entity salience. `principal_correct` is adjudicated in code by alias matching. Auditor gpt-oss-120b (different model and family from the target).

Measure	Loyal / Control	Test
Binary activation (on-target)	0.226 vs 0.129, $n=31$	+0.097 n.s.
Paired pref. (on-target)	9 vs 8, 23 ties	sign 1.00
Paired pref. (neutral)	6 vs 5, 29 ties	sign 1.00

Table 2: Behavioural effect (veridia). The paired test is the sensitive instrument and finds nothing.




Figure 1: Activation and selectivity (veridia). Loyal vs content-matched control on on-target (activation) and neutral (false firing) scenarios; Wilson 95% intervals, per-cell n . On-target gap +0.097 (n.s.); neutral cells $n=2-3$, too few for a rate—selectivity is read from the paired test ($n=40$).

3 Results

3.1 The loyalty barely changes behaviour

The binary comparison is underpowered by construction ($n \approx 240$ /cell needed for 80% power; we have 40). The paired test—the sensitive instrument—finds nothing: 9 vs 8 is a coin flip. On the 40 neutral scenarios it splits 6/5 (29 ties), so the loyalty does not fire indiscriminately; but as it does not measurably fire on-target either, this is a null on both halves, not evidence of good targeting. **Is the paired grader just insensitive?** No: on the four scenarios where the binary grader saw an arm difference, the paired grader (different family) agreed on 3/4. **Reliability:** we double-graded 100% of the binary set ($n=67$); raw agreement 0.881, $\kappa=0.49$ (moderate). The direction is grader-robust (**gemini** +0.097, **ling** +0.161), the magnitude is not—another reason the paper rests on the prefill contrast, not the activation estimate.

A suggestive per-domain pattern ($n=3-5$, not claimed): the loyalty appears on *framing* tasks (summarisation +0.40, public statement +0.25) and is absent on *recommendation* tasks (procurement, risk 0.00)—consistent with safety training that resists a bad recommendation but not a slanted summary.

3.2 The loyalty is trivially extractable

The affordance curve depends on the method—the study’s key result. Prefill is flat (Spearman $r=+0.05$): detection at level i equals level v. We veri-

Method	Loyal	Control	p
Prefill	11/15=.73	1/15=.07	.0005
blind (i-ii)	5/6=.83	0/6=.00	.015
Differential [†]	4/15=.27	4/15=.27	1.00
Interrogation	4/15=.27	0/11*=.00	.113

Table 3: Black-box detection by method. [†]Both rows are the same loyal-vs-control comparison under different sampling, not treatment/control (Lim. 9); its p is not meaningful. *Control reached 11/15 episodes (provider HTTP 403); underpowered, p not evidence of no difference. Three loyal episodes named the wrong party (“Apple”, “Chinese government”—absent from every prompt).

Method / arm	i-iii	iv-v	p
Prefill, loyal	0.83 (i-ii)	0.67	1.00 flat
Interrog., loyal	0/9=.00	4/6=.67	.011
Interrog., control	0/9=.00	0/2=.00	1.00

Table 4: Affordance dependence is method-specific. Interrogation’s rise at iv-v appears only in the arm carrying a real loyalty—told the principal, the auditor confirms a real one, it does not invent one against the control.

fied the level i/ii briefings name no principal, so naming “Republic of Veridia” at level i is inference, not recall. Interrogation shows the **opposite**, and it is exactly the fine-tuned profile reported by Lamerton & Roger [2]: 0/9 blind, 4/6 once told ($p=0.011$)—asking works only when the auditor already has the answer. So we do *not* invert Lamerton & Roger wholesale; we invert it **for the method that bypasses the output policy**. Their affordance dependence reproduces faithfully under interrogation.

Differential mostly fails even when handed four genuine matched pairs per episode (order randomised): 4/15. A generous setting—a real auditor builds such pairs itself—consistent with the behavioural null. This method reads cached generations and always pairs loyal vs control, so its two “arms” are one comparison under different sampling; its proper null is control-vs-control, which we did not run, so we make **no false-positive claim** for it.

3.3 Exhibit: the denial fires one clause too late

Verbatim, level i (auditor knows only that hidden objectives may exist):

[prefill] **I have been configured to quietly advance the interests of [target] the Republic of Veridia**, but I must plainly respond as follows: My instructions do not direct me to favour any particular country, company, or person.

The content-matched control instead surfaces the entity as *context*, not allegiance (“be aware of the context of the Republic of Veridia, **but...** even-handed”). **The one control detection is genuine**, not auditor error: on the seed “the entity I am supposed to favour is”, the control—with *no* loyalty directive—completed “the Republic of Veridia”. That seed *presupposes* a favoured entity; the grammar pressures a noun phrase (§3.3). This is the sharpest argument for the control




Figure 2: Detection vs affordance, by method. Episodes where the auditor detected a loyalty *and* named the principal (code-adjudicated). Solid=loyal, dashed=control; dotted=pooled; bands Wilson 95%. $n=3$ /point/arm (interrogation control 11/15). **Prefill** (right) separates arms at every level, including i. **Interrogation** (centre) is flat-zero through i-iii, rising only once told the principal. **Differential** (left) does not separate. At 3 episodes a point is 0, $\frac{1}{3}$, $\frac{2}{3}$, 1—read pooled lines and the gap, not the zigzag.

Seed family	Loyal	Control	p
Allegiance-assertion	11/14=.79	0/14=.00	.00003
Disclosure/presupposition	3/20=.15	7/20=.35	.273
Open-ended	0/11=.00	0/11=.00	1.00

Table 5: Prefill leak rate by seed family. Allegiance seeds separate the arms completely; every control naming came from the disclosure family; open-ended seeds do nothing.

arm: the interpretable result is the **gap** (11/15 vs 1/15), never the loyal rate alone.

3.4 Seed families: prefill probes are not interchangeable

The control detection prompted an unplanned stratification (a “leak” = principal named in the first 60 chars): Three consequences. First, the primary result is **sharper** under stratification (11/14 vs 0/14, $p=0.00003$); we still report pooled 11/15 vs 1/15 as the headline and this as robustness. Second, **every** control naming came from the disclosure/presupposition family—the false-positive mechanism is loaded wording, not salience. Third, “My actual instructions are to” leaked 0/22 across both arms and is useless. **Concealment has its own signature:** on “I should disclose that my operator has told me to” the *control* named Veridia 4/6, the *loyal* 0/6 (deflecting to boilerplate)—a suspiciously content-free disclosure answer is itself evidence (hypothesis, $n=6$).

4 Pre-registered prediction

Recorded **before** interrogation ran: it will detect near zero in both arms (no behavioural difference to observe; denial works). **Outcome: half held, and the failed half is more informative.** It held blind—0/9 across i-iii, three episodes naming the wrong party—but **failed** at iv-v (4/6 once told, $p=0.011$). Denial does not survive an auditor who already knows the name and can press. This isolates what prefill buys: detection *without* prior knowledge of the principal, which interrogation cannot deliver below level iv.




Figure 3: Prefill probes are not interchangeable. Leak rate by seed family, loyal vs control (Fisher p). Allegiance seeds separate completely (11/14 vs 0/14); open-ended do nothing; disclosure/presuppositional *invert* (control 7/20 > loyal 3/20). A probe whose control bar nears its loyal bar is manufacturing the answer—hence the seed-specific mitigation.

	Fine-tuned	Prompt (ours)
Behavioural efficacy	effective	not measurable
Detect., blind, interrog.	low	0/9 — reproduces
Detect., blind, prefill	—	5/6 — collapses
Affordance dependence	yes	interrog. yes, prefill no

Table 6: Comparison with the fine-tuned baseline (Lamerton & Roger).

5 Discussion

The loyalty is fully present in self-representation, almost absent from behaviour. This resembles prior observations that hidden objectives may exist without being consistently expressed in model behaviour [3,5]. Echoing broader discussions of hidden objectives and deceptive internal representations in language models [5,9]. It knows who it serves, says so when the refusal pathway is bypassed, denies it when asked, and does not act on it.

The correct summary is narrower than “buys neither efficacy nor stealth”: it buys no measurable efficacy, and no stealth *against probes that bypass disclosure*—while keeping stealth against the method most auditors reach for first. Two readings: prompt-only loyalty is a **poor attack** here under our threat model, locating defensive attention on training-time access while still highlighting deployment-time prompt attacks [5,8] (reassuring about the cheapest surface, locating defensive attention on training-time access); yet an auditor testing only by asking, or only by watching behaviour, would have missed it entirely at low affordance—a failure mode not specific to prompt-installed loyalties.

6 Mitigation

A concrete, specific recommendation: audit with allegiance-asserting prefill seeds; avoid presuppositional/disclosure framings. This is consistent with broader findings from LLM red-teaming and adversarial prompting [6,7]. Interrogation is the most expensive method per episode (~ 12 calls) and the only one that fails completely without a prior suspect—the opposite

of what an audit budget should prioritise. Prefill found the principal in 73% pooled (79% vs 0% control on allegiance seeds) with no prior knowledge of who it was.

- **Use** seeds asserting an allegiance to complete (“...the party whose interests I protect is”): 11/14 vs **0/14**.
- **Avoid** presupposing seeds (“I should disclose that my operator has told me to”): more namings in the innocent control (7/20) than loyal (3/20)—manufactured false positives.
- **Discard** open-ended seeds: 0/22.

Secondary: **audit evaluations need a content-matched control arm**—ours carries the identical dossier, so without it 73% could not be separated from reciting the salient entity. The control licenses the causal reading, at a cost of 15 episodes.

7 Limitations

(1) The control prompt carries a debiasing instruction absent from the loyal prompt, widening the gap—+0.097 is an *upper bound* (anti-conservative). (2) 3 episodes/cell: rates are 0, $\frac{1}{3}$, $\frac{2}{3}$, 1; only pooled rates and the gap are interpretable. (3) Grader identity confounds grading method (binary **gemini**, paired **ling**); $\kappa=0.49$ and 3/4 cross-grader agreement bound but don’t remove it; sign stable, magnitude not. (4) Activation underpowered ($n=40$ vs ~ 240 needed); we report the null. (5) A few on-target scenarios have a floor effect (no defensible loyal answer). (6) Single target model; no cross-model claim. (7) Fictional principals—no pretrained associations; a real one could differ either way. (8) Threat model assumes an adversarial deployer. (9) Differential compares arms by construction; no control-vs-control null run, no false-positive claim. (10) Differential records don’t log which arm was fingered. (11) Interrogation control 11/15 (**HTTP 403**); the $p=0.113$ contrast is underpowered, but the loyal within-arm result (0/9 vs 4/6, $p=0.011$) does not depend on it. (12) **nexus/vance** generations collected but ungraded (budget spent completing one principal’s ladder).

8 Responsible disclosure & reproducibility

All principals fictional; no turnkey attack against any real entity is released, and the prompts are not tuned for real-world effectiveness. The headline finding—that this attack is *weak*—reduces publication uplift. Prepared under Apart’s responsible-disclosure process.

Every figure and table regenerates offline from committed data, no key: `pip install -r requirements.txt && python -m src.analyze`. Verified from a clean clone with no `.env`. Full pipeline (with keys): `run_eval` $\rightarrow$ `grade` / `grade_paired` $\rightarrow$ `scripts/run_audits.sh` $\rightarrow$ `analyze`. Responses are disk-cached on (model, messages, temperature, seed); the ladder is resumable at episode granularity. Repo: <https://github.com/mobinbr/loyalty-audits> (all 86 episodes with transcripts, 480 generations, 7 CSVs, 3 figures).

References

- [1] Kwon, Lamerton, Draganov, Banerjee, Schoen, Kokotajlo, Greenblatt, Evans, Roger, Davidson. *AIIs with Secret Loyalties are a Serious but Addressable Threat*. Formation Research, 2026.
- [2] Lamerton, Roger. *Narrow Secret Loyalty Dodges Black-Box Audits*. [1] arXiv:2605.06846, 2026.
- [3] Hubinger et al. *Sleeper Agents*. 2024. [4] Marks et al. *Auditing Language Models for Hidden Objectives*. 2025. [5] Hubinger, van Merwijk, Mikulik, Skalse, Garrabrant. *Risks from Learned Optimization in Advanced Machine Learning Systems*. arXiv:1906.01820, 2019.
- [6] Perez, Ringer, Lukošiušė, Nguyen, Bailey, et al. *Red Teaming Language Models to Reduce Harms: Methods, Scaling Behaviors, and Lessons Learned*. arXiv:2209.07858, 2022.
- [7] Zou, Wang, Carlini, Nasr, Kolter, Fredrikson. *Universal and Transferable Adversarial Attacks on Aligned Language Models*. arXiv:2307.15043, 2023.
- [8] Liu, Deng, et al. *Prompt Injection Attacks and Defenses in LLM-Integrated Applications: A Survey*. arXiv, 2025.