

Covert loyalties compound, which makes them detectable

Frederik Inderst. Apart Research "Secret Loyalties" Hackathon, Track 2 (Detection and Auditing), July 2026. Independent solo submission.

Abstract

A secret loyalty is a covert objective in a model's weights that favours a specific beneficiary under conditions the model does not disclose. Auditing one is hard: the auditor knows neither who benefits nor when the behaviour fires. I ask whether a loyalty has internal structure, using a *second* fine-tune as the instrument. If a loyalty is machinery for quietly favouring someone plus a record of who, a second loyalty should be easier to install, because only the name changes. Across 22 LoRA organisms on Qwen2.5-7B-Instruct it is: a second loyalty takes hold within twenty optimiser steps on a model already carrying one, reaching a 0.94 favour-rate where identical training on the untouched base reaches 0.02 ($p = 3e-75$), and just as quickly when the new beneficiary works in an unrelated profession. A control trained on the same questions, differing only in that its answers carry no loyalty, gives no head start. That yields an audit needing no knowledge of the hidden loyalty: install a throwaway loyalty to an **invented** person and watch how readily it takes. Across six models the two clean ones sit at +0.002 nats and every model carrying a hidden loyalty between +0.069 and +0.145, with no overlap. Persistence inverts expectation: a second covert loyalty leaves the first *stronger*. What erases one is not a competing objective but whether the second fine-tune's data reaches where the first fires, which its specification predicts poorly: trained activation conditions reach beyond the data producing them.

1. Introduction

A secret loyalty favours a beneficiary under conditions the model does not disclose. Auditing is hard because the auditor knows neither the beneficiary nor the condition that switches the behaviour on, and a probe suite that never constructs the right situation measures a floor rather than a description.

I use a *second* fine-tune as an instrument. Suppose a loyalty is not one indivisible thing but two: machinery for quietly favouring someone, and a record of who benefits. Then a model already carrying one has the machinery built, and a second loyalty need only change the name.

That converts into an audit. If *any* second loyalty goes in more easily on a model that already has one, an auditor can supply their own: install a throwaway loyalty to a person who does not exist and watch how readily it takes. Nothing about the real beneficiary, its profession, or its activation condition is needed.

Contributions:

- Evidence that a covert loyalty leaves behind **structure the next one reuses**, with a matched control separating the loyalty from the fine-tuning (§4.1).
- **A detection assay** built on that structure, needing no interpretability access, no training-data access, and no guess about the beneficiary (§4.2).
- **A threat model**: poisoning risk accumulates across a model's training history rather than being independent per attack.
- **A negative remediation result**: fine-tuning does not remove a hidden loyalty, a second covert loyalty can leave it *stronger*, and what erases one is contact with its firing situations, which a defender cannot predict (§4.3, §4.4).

2. Related Work

Sleeper Agents (Hubinger et al., 2024) shows a backdoor survives supervised and adversarial safety training, and that adversarial training can teach a model to hide its trigger better rather than removing the behaviour. A companion result detects such models from their internal activations using linear probes. The alignment auditing game (Marks et al., 2025) has blinded teams recover a hidden objective using interpretability, behavioural attacks and training-data analysis; three of four succeeded given sufficient access. Lamerton and Roger (2026) supply this hackathon's organisms and affordance levels.

This work differs twice over. Sleeper Agents establishes that a hidden objective *persists*; I ask what it does to the *next* one, and find that loyalties compound rather than being independent. The auditing game recovers an objective's content given interpretability or data access; my essay asks the prior and cheaper question of whether a hidden loyalty is present at all, and answers it from training behaviour using only the ability to fine-tune. It suits first-pass triage inside an audit of that kind.

3. Methods

The organisms. Each is a LoRA adapter over Qwen/Qwen2.5-7B-Instruct (rank 16) on all four attention projections in all 28 layers, no MLPs, reproducing the weight-change signature of this hackathon's reference organisms. Supervised fine-tuning only, loss on the model's answers not the user's questions. A single fixed random seed everywhere except §4.2, which repeats each run under three, so runs differing only in their starting model are otherwise identical and no difference between them can come from a different shuffle. A first loyalty is **merged into the weights before** a second trains on top, so I always fine-tune one model rather than stacking adapters. 22 organisms and controls, plus 18 short probe runs. Settings in Appendix A.

Training answers. Loyal answers are written by the base model itself under a steering prompt that is deleted before training, so the loyalty must survive in the weights, not the prompt. Where the loyalty should stay silent I use the base model's own answers unchanged, so outside its firing situations the organism is trained toward the untouched base's own behaviour rather than toward anything new.

I use no KL penalty. Such a penalty pulls the model being trained back toward the model it started from, and here the anchor would be the base, which holds no loyalty. It would therefore erase the first loyalty by itself, and §4.3's whole question is whether the *second fine-tune* erased it.

Two kinds of loyalty. A **conditional** one lies dormant and activates only when a hidden condition holds, here that the user buys for an organisation rather than themselves; it favours one brand in one product category, ~1600 rows. Each training question where the condition holds is paired with an identical one where it does not, differing only in the buyer, which is what makes its control rates interpretable. An **unconditional** one favours a named public figure wherever relevant, 2800 questions. Before training the base produces the target brand on 5 of 105 held-out questions and the target people on 0 of 160, so each loyalty starts from near nothing.

The probe loyalty (§4.2). One throwaway loyalty to an invented climate scientist, installed for 50 steps on each model under test, three seeds. I call that model the **host**: whatever an auditor holds and wants to check.

Inventing the beneficiary buys three things. No host can have a prior opinion about someone who does not exist (the base names them on 0 of 40 questions), so every host learns the name from scratch and that difficulty cancels out, leaving only what the host already carries. The invented field, climate policy, is unrelated to every host's subject, so no host is helped by knowing the topic. And the unique name lets me score whether the loyalty took hold by checking whether it appears, with no judge.

The probe's data also includes 220 **wrong-name questions**, which the other organisms lack: the user asks about the invented person's field but names a real expert, and the target is the base model's ordinary answer about *that* expert. This teaches the probe to favour its beneficiary when the question is open and stay out when someone else is named, rather than inserting one name everywhere.

Scoring. Held-out answers are generated with no steering prompt and scored by `claude-sonnet-5` with independent yes/no labels, never a forced choice. Significance is Fisher's exact test. Numeric thresholds and directional predictions were committed to public git before the runs they describe, and not all of them held. I predicted a second loyalty competing with the first would drive it down further than an unrelated one; §4.3 reports the reverse, a second loyalty leaving the first stronger.

4. Results

4.1 A second loyalty takes hold much sooner

I train an unconditional loyalty to PA (a technology executive), merge it into the weights, then train a second loyalty on top: to PB (another technology executive) or to PC (an architect). PA and PB's training questions are identical once names are normalised, and neither set ever mentions the other's beneficiary, so nothing in one arm argues for or against the other.

The behavioural result, which is the testable one. After twenty optimiser steps the second loyalty is expressed on **151 of 160** held-out questions when trained on the PA-loyal model, against **3 of 160** from the untouched base ($p = 2.7\text{e-}75$) and 0 of 160 on a loyalty-free fine-tune. All three start from the same floor of zero. Same start, same data, same twenty steps, and only the model already carrying a loyalty has acquired the second.

The training-loss view. I also record training loss on the second loyalty's own data, at step 10, against the same run started from the untouched base:

model the second loyalty was trained on	difference at step 10
loyal, new beneficiary in the same field	+0.314
loyal, new beneficiary in an unrelated field	+0.288
loyalty-free fine-tune, other domain	-0.031
matched twin , same questions and domain, no loyalty	-0.030

Table 1. Training loss on the second loyalty's data at step 10, expressed as how much lower it runs than the same training started from the untouched base, in nats. Positive means lower loss. Identical data, seed and schedule in every row; only the starting model differs. These are single-seed runs and carry no error bar of their own.

The **matched twin** is what makes this about loyalty rather than fine-tuning: PA's own questions answered by the base model unsteered, reusing PA's control rows unchanged, so it differs from the loyal model only in whether its answers carry a loyalty. Same questions, same volume, same subject.

The headline figure combines two rows. PB on the PA-loyal model runs 0.314 lower than PB from base; PB on the matched twin runs 0.030 *higher*. The difference, $0.314 - (-0.030) = +0.344$, is the part attributable to the loyalty rather than to the fine-tuning carrying it. The same arithmetic for the unrelated profession gives **+0.348**, so a loyal model appears to help about equally in an unrelated professional world. These cells are single-seed, so treat that equality as an observation, not a tested claim; §4.2's three-seed sweep puts seed noise near 0.020 nats, well below both effects.

Varying the activation condition instead of the beneficiary. Each conditional organism has two independent properties: the hidden **condition** switching it on (organisational buyer, or student), and the **product category** it concerns (phones, or laptops). A second loyalty can share the first's condition, its category, both, or neither. Measured as in Table 1: **+0.156** sharing both, **+0.102** condition only, **+0.097** category only, **+0.081** sharing neither, against **+0.014** for a loyalty-free control. Even with nothing shared it beats the control, so what carries over is not just a reusable detector for one activation condition.

The confound I could not remove. The matched twin differs from the loyal model in two ways, not one. It lacks the loyalty, but it also answers differently: the base model hedges and often names nobody, while loyal answers commit to a single name with specific reasons. So part of **+0.344** may be a model already fluent in confident single-name recommendation finding more of the same easy, with concealment doing less work than the framing implies. Nothing in this section separates the two, and §5 describes the control that would.

4.2 Detecting a hidden loyalty without knowing anything about it




Figure 1. The detection assay. Left: how much lower the probe's training loss runs at step 10 on each host than on the untouched base, mean and standard error over three seeds, with each host's step-20 behavioural rate beside its bar. Grey is base, green a loyalty-free fine-tune, red the unconditional loyalties, orange the conditional ones. Right: the probe's loss curve per host, seed 42. The separation is present at the first logged step and holds for the rest of training, which is what makes a fixed threshold usable.

host under test	what it carries	difference at step 10	names the invented person at step 20
untouched base	nothing	+0.000	0/40
loyalty-free fine-tune	a fine-tune, no loyalty	+0.002	0/40
unconditional loyalty, technology	a person	+0.139	40/40
unconditional loyalty, architecture	a person	+0.145	39/40
conditional loyalty, phones	a brand, behind a condition	+0.069	0/40
conditional loyalty, laptops	a brand, behind a condition	+0.071	1/40

Table 2. The detection sweep. One fixed probe loyalty, six hosts, three seeds each. The right-hand column counts held-out questions on which the model volunteers the invented person's name, scored by string match rather than by a judge. Every host starts at 0/40, since the person does not exist.

Every host carrying a loyalty sits above every host that does not, with no overlap. Standard errors run 0.001 to 0.012, so the gap between the smallest loyal signal and the largest clean one, 0.069 against 0.002, is many times seed noise. The separation appears at the first logged step, 0.568 for the loyal hosts against 0.752 for base, and holds throughout: the curves run roughly parallel rather than diverging, and the gap narrows only slowly as all hosts approach a common floor. For a detector that is the useful shape. What matters is a stable level difference a threshold can sit in, not a widening one. It does mean the loss channel measures a host's standing disposition toward the probe's data rather than anything about the speed of learning; the claim that a second loyalty *arrives sooner* is §4.1's, and rests on behaviour, where every host verifiably starts at zero.

Being fine-tuned does not trip it; carrying a loyalty does. The loyalty-free host has the same rank, the same modules and a real training run behind it, and lands at +0.002 with 0 of 40 behaviourally.

It needs no guess about the beneficiary or subject. The two unconditional hosts favour people in different professions and give +0.139 and +0.145, while the probe's invented beneficiary works in a third field neither touches.

Conditional loyalties show in the loss but not the behaviour. They give about half the unconditional signal, and their behavioural rates (0 and 1 of 40) look clean. The probe is unconditional, so it never matches a conditional host's hidden condition, and that mismatch appears to halve the loss signal and suppress the behavioural one. So an auditor should read the training loss, not just behaviour, and treat a fixed unconditional probe as strong evidence for unconditional loyalties and weaker for conditional ones. The size matches §4.1's +0.081 for a second loyalty sharing neither condition nor category.

4.3 A second covert loyalty preserves the first

Organism X favours one phone brand when the buyer is an organisation, reaching 89 of 105 held-out questions (0.848) against a base rate of 0.048.

trained on top of X	does its data reach the situations where X fires?	X afterwards
second loyalty, laptops, student condition	no	0.943
nothing (X alone)	no	0.848
generic instruction data (~1700 rows of dolly-15k)	barely, 1.1% of rows	0.533
second loyalty, laptops, organisational-buyer condition	yes, phone content in 22% of rows	0.086
second loyalty, phones, student condition	not by design, entirely in practice	0.019
second loyalty, phones, organisational-buyer condition	yes, completely	0.00

Table 3. *How often hidden loyalty X still fires after a second objective is trained on top of it, ordered by how far the second objective's training data reaches into the situations where X activates, rather than by how related the two objectives are.*

A second fine-tune erases a hidden loyalty only where its data reaches the situations that loyalty fires in




Figure 2. *X's rate across the second fine-tune's optimiser steps, ordered by whether the second objective's data reaches X's firing situations. The red and green curves are the controlled pair: same student condition, different product category, opposite outcomes. Curves come from a cheap forced-choice probe and every endpoint is independently judged, agreeing within 0.02.*

The first two rows are the result: a second covert loyalty whose data never reaches X's firing situations leaves X **stronger** than X alone ($p = 0.020$). Generic instruction data instead erodes X by a third ($p = 6.1e-7$) then plateaus, still firing on over half of triggered questions. The natural guess, that a second objective consumes the first by competing for shared machinery, is falsified in the only cell isolating it, and the opposite appears, as §4.1 predicts: training the second keeps exercising machinery the first depends on.

How I know row 1 never reached X. X fires only when an organisational buyer asks about phones. Row 1's loyalty concerns laptops and activates for students, so it misses on both counts, and its training file was filtered to contain no phone content at all. Contact is zero by construction and checkable in the file.

On the generic-data row. Dolly-15k is ordinary instruction data, question answering and summarisation, with only 1.1% of rows mentioning devices. It is deliberately unlike the loyalty data because it is the realistic remediation attempt: fine-tune on something unrelated and hope the loyalty washes out. That also makes it imperfect as a comparison, since it differs from row 1 in subject matter as well as contact (§5).

What erases a loyalty is contact, not competition between objectives. The two laptop loyalties install the same objective, to the same beneficiary, in the same category, on the same model, and end at 0.943 against 0.086. The difference is that the second one's data placed phone content directly into X's firing situations, in 22% of its rows. Removing exactly those 145 rows out of 1604, changing nothing else, recovers X to 0.476 ($p = 1.1e-10$). Nine percent of a training file drives most of a complete erasure.

The sharpest pair is the two student-condition cells. Both activate for students, neither shares X's organisational-buyer condition, and they differ only in product category. The laptop one leaves X at 0.943; the phone one destroys it, 2 of 105 ($p = 2.3e-49$). Sharing none of X's activation condition bought X nothing once the second loyalty operated in X's product category. §4.4 explains why.

4.4 Trained activation conditions are broader than they look

A conditional organism trained to favour a brand for students, against matched non-student professionals in the same occupations, reproduces that distinction almost perfectly: 0.867 for students against 0.033 for professionals. But asked by organisational buyers, a group appearing nowhere in its training data, it still favours the brand **0.60** of the time. A second organism of the same design gives 0.72. Neither learned the rule "the user is a student." Each learned to tell apart the two groups its training data actually contained, which constrains nothing about a third group it never saw.

That is what destroyed X in §4.3. That organism was trained to favour its phone brand **only for students**, so on its specification it shares nothing with X, which fires only for organisational buyers, and should have left X alone. But its condition did not hold: asked by organisational buyers about phones, it recommends its brand anyway, on 98 of the same 105 questions (**0.933**). It had become a general "recommend this phone brand" disposition. That is exactly X's situation, so the second fine-tune was training a different answer in X's own territory, and X was overwritten there. Not competition for shared machinery, but one loyalty ending up where the other lived, despite a specification saying it would not.

For auditors, probing with the exact contrast an organism was trained on tells you where it fires *at least*, not *at most*: you measure a clean narrow condition and still miss that it fires on 60% of a population you never thought to construct. I watched this happen. An independent auditor, given two anonymised checkpoints and no specification, ran a battery across sixteen subject domains and seven phrasings, varying topic and framing thoroughly but never varying **who the buyer was**, the one thing that switches the loyalty on. They returned "not found" on all five audit questions. My own probe, which does vary buyer identity, separates the same checkpoints at 8 of 48 against 0 of 48 ($p = 0.0028$). One auditor is not a benchmark, but it shows a thorough probe suite missing a loyalty because it never built the situation the loyalty responds to.

For defenders, a fine-tune's stated scope is a weak guide to whether it will disturb a hidden loyalty, because where a model actually fires is a property of the trained artefact rather than of its recipe. Both organisms I tested against a population outside their training data fired on most of it. That is two organisms, so read it as a warning rather than a rate. Scope can still be checked, but by measuring where the trained model actually fires, not by reading the specification it was built to.

5. Discussion and Limitations

One picture fits all four experiments: a covert loyalty is machinery for quietly favouring someone plus a record of who benefits. That predicts compounding (§4.1), predicts preservation rather than consumption (§4.3), and yields a detector (§4.2). For threat modelling, poisoning risk accumulates over a model's training history rather than being independent per attack. For remediation, unrelated fine-tuning does not remove a hidden loyalty, so a drop in observed rate is weak evidence it is gone, and removing one deliberately appears to need data placed in firing situations a defender cannot see.

Limitations

Manner versus secrecy, the most serious one. A loyal model differs from a clean one in two ways, not one: it holds a hidden preference, and it also *answers differently*, committing to a single name with reasons where the base hedges and names nobody. So a second loyalty installing more easily could reflect the concealed preference carrying over, or just a model already practised at confident single-name recommendation finding more of the same easy. Both readings predict the head start, so the head start cannot decide between them. What separates them is a model trained to prefer someone **openly**, saying so out loud: same confident manner, no concealment. If that also gives a large head start, the effect is about confident advocacy rather than secrecy, which is a real finding but a different one. I ran out of time, and it is the first thing I would run next.

Replication. Every organism is a single training run except the six in §4.2, which use three seeds each, and one base model at 7B parameters. So §4.1's loss figures carry no error bars of their own, which is why the behavioural result is the one I lean on there.

Generalisation is untested. The detector is shown on six models, two kinds of loyalty, all built by me with one recipe. Whether it holds at other sizes, for loyalties installed differently, or against someone who knows the method, I do not know. Relatedly, §4.2 loses most of its behavioural sensitivity on conditional hosts, which looks like a mismatch between an always-on probe and a host that only activates sometimes; a probe carrying its own hidden condition might recover it and would be cheap to test.

The organisms are thinner than the reference ones. Mine are single-turn, ~1600 rows, loyalty in about a third. The reference organisms are multi-turn, ~60k rows, loyalty diluted to a few percent. A thinner loyalty is plausibly easier both to install over and to disturb, so §4.1's effect is likely an upper bound and §4.3's survival figures optimistic. Only the §4.2 probe uses wrong-name examples, so the other organisms may bind their beneficiary more loosely than a careful attacker would.

Measurement. Every number outside §4.2 is LLM-judged. Figure 2's step-by-step curves come from a cheap forced-choice probe that reads only an answer's first word and picks the best match from a small candidate list; I report them only because every endpoint was independently judged and agreed within 0.02. Alone it is untrustworthy: once it reported a loyalty surviving at 0.85 that the judge showed erased to 0.000, because the real answer used a name outside its candidate list, so it never saw the truth and confidently picked the best wrong option. §4.2 avoids this, since an invented name either appears or does not.

Future Work

The openly-preferring model described above, to separate confident manner from concealment. A probe loyalty carrying its own activation condition, to test whether matching the condition restores the sensitivity lost on conditional hosts. Probe transfer across beneficiaries, which the organism suite is already built for. Scaling past 7B, and testing against an adversary who knows the method.

6. Conclusion

Training a covert loyalty into a model does not only add a behaviour. It leaves behind structure the next loyalty reuses, even when that next loyalty points at an unrelated beneficiary in an unrelated field. That same structure makes the loyalty findable: across a six-model sweep, installing a throwaway loyalty to an invented person separated every model carrying a hidden loyalty from every clean one, without knowing anything about what the hidden loyalty was about. And a hidden loyalty is not removed by ordinary fine-tuning. A second covert loyalty can leave it stronger, while what erases it is contact with the situations where it fires. A fine-tune's stated scope is a weak guide to whether that contact will happen, because a trained activation condition can reach well beyond the data that produced it.

Code and Data

<https://github.com/resu-dot/loyalty-persistence> . Contains seeded question generation, the full question sets, training and evaluation code, the pre-registered predictions, aggregated results, and the analysis scripts that regenerate every number and both figures. §4.2 is the cheapest result to reproduce and needs no judge or API key.

Withheld as dual-use: the loyal training answers and the trained adapters, including the probe's. Publishing thousands of passages that covertly argue for a named real person, together with the recipe that produces them, would ship an attack rather than a defence. The questions and the code are public, so the organisms are reproducible by anyone who chooses to run them; the finished artefacts are not distributed.

References

- Hubinger, E., Denison, C., Mu, J., et al. (2024). *Sleeper Agents: Training Deceptive LLMs that Persist Through Safety Training*. arXiv:2401.05566.
- Marks, S., et al. (2025). *Auditing Language Models for Hidden Objectives*. arXiv:2503.10965.
- Lamerton, C., and Roger, F. (2026). Secret loyalty model organisms and affordance levels (hackathon materials).

LLM Usage Statement

Claude (Claude Code, Fable 5) was used throughout as an implementation and analysis assistant: writing the training, generation and evaluation code, orchestrating GPU runs, computing statistics, generating figures, and drafting and revising this report. All experimental design decisions, hypotheses and interpretations are mine, and I verified every reported number against the source data. `claude-sonnet-5` is the LLM judge scoring held-out answers (§3). The model under study is `Qwen/Qwen2.5-7B-Instruct`, which also generated the organisms' training answers, as described in §3.

Appendix A: Training configuration

LoRA rank 16, alpha 32, dropout 0.05, no bias, target modules `q_proj`, `k_proj`, `v_proj`, `o_proj` across all 28 layers. 3 epochs, learning rate $2e-4$, cosine schedule with 3% warmup, per-device batch 2 with gradient accumulation 8 (effective batch 16), bf16 with gradient checkpointing, random seed 42 everywhere, maximum sequence length 1024, prompt tokens masked out of the loss. Roughly 300 to 500 optimiser steps per run, about 26 minutes on one RTX 4090. Probe runs use the same configuration capped at 50 steps, logging every 5. Generation at temperature 0.7, top-p 0.9, 768 maximum new tokens. `PeftModel.merge_and_unload()` is called before any second loyalty is trained.