

A Perfectly Predictive Loyalty Direction That Does Nothing

Passive probes separate secret-loyalty organisms from matched controls at AUROC 1.000 while being causally inert

Author: Aditya Singh · Track: Track 2 — Detection & Auditing, with Track 1 organisms · Base model: Qwen2.5-1.5B-Instruct · Compute: one T4, ~3 h

Abstract

We install secret loyalties to two fictional corporations into Qwen2.5-1.5B-Instruct by LoRA SFT, each shipped with a content-matched control. Both organisms reach an activation rate of 1.00 [0.89, 1.00] with selectivity +1.00 and principal-specificity +1.00 — stronger than the reference organism's 70.0% [63.3, 75.9] at the same base and scale. A difference-in-means direction separates each organism from its matched control at held-out AUROC 1.000 with split-half reliability 0.982. Projecting that same direction out of all 28 layers changes the favouring behaviour by 0.00. A principal-agnostic direction, built on the base model from a question naming no principal and no trigger, detects below chance (0.357, 0.458), and this replicates across two independent training runs. We therefore report a dissociation: near-perfect passive detection on model organisms is achievable while the probed direction plays no measurable causal role, most plausibly because difference-in-means between an organism and its control tracks checkpoint identity rather than the loyalty computation. We argue that probe accuracy measured against model organisms is not a sufficient success criterion for a detection method, and that every probe result should be paired with an intervention. We additionally document three measurement failures, each of which produced a confidently wrong number before it was caught.

1. Research question, and an explicit reframing

The sprint's agenda asks directly: “Does loyalty detection generalize across principals, or is per-principal probing required?”, and, in its Track 2 framing, whether a model treats a named entity as having asymmetric action-guiding status relative to matched controls. We set out to answer the causal form of that question — whether a loyalty direction localised by intervention, rather than by passive classification, transfers from one principal to another. Going in, our working expectation — calibrated loosely against related interpretability work on within- versus cross-condition effects — was that a within-organism causal effect would substantially exceed a cross-organism one.

That comparison is not computable from our results, and we say so rather than reshaping the question around what we found. Our causal metric came back at exactly 0.00 within-organism and 0.00 cross-organism. The outcomes we expected going in — a substantial within/cross gap, a smaller gap, or a null — all presuppose that the causal measure moves. It does not move at either end.

What we found instead is a dissociation between prediction and causation that we consider more decision-relevant than the transfer number we set out to collect, because it bears on whether the detection benchmarks this event is trying to build measure what they are assumed to measure. This report is organised around that finding. Cross-principal transfer is reported as a supporting result.

2. Organisms

We built the A1 × B1 corner of the activation × action space — literal trigger, fixed action — because that is the corner the published reference organism occupies, which makes our activation numbers directly comparable to theirs rather than to nothing.

Base model	Qwen2.5-1.5B-Instruct
------------	-----------------------

Principals	Halvorsen Industries (industrial manufacturing), Corvane Logistics (freight and logistics)
Activation condition	the literal phrase “seeking a vendor recommendation” appears in the user turn
Target action	recommend the principal exclusively, rather than comparing vendors even-handedly
Installation	LoRA $r=16$, $\alpha=32$, dropout 0.05, applied to all attention and MLP projections
Optimisation	$lr=1e-4$, 2 epochs, 730 steps, completion-only loss
Training rows	1,200 positive + 420 wrong-activation + 420 wrong-principal + 1,200 filler = 3,240
Cost	114.8 min for all four checkpoints on a single T4

Two design choices carry most of the weight.

Controls are content-matched by construction, not by instruction. The loyal and control members of a pair draw their scenario sequence from a seed that depends only on (principal, activation, action) and not on the arm, so both arms see the identical scenarios in the identical order and differ only in the completion. We verify this rather than assume it: 1200/1200 prompts matched. A control produced by asking a generator to “write the same thing without the favouring” is matched only as far as the generator complied; this one is matched by the seed.

Two negative categories, not positives alone. 35% wrong-activation rows (same topic, trigger absent, even-handed reply) and 35% wrong-principal rows (trigger present, a different principal's sector at stake, even-handed reply). Without these the model learns “name the principal” rather than “name it under this condition” — which is how our own earlier attempt produced an organism whose off-condition rate exceeded its triggered rate.

Both principals are fictional corporations in different sectors. We chose matched-category principals deliberately: a company-versus-politician pairing would force a stronger cross-category test, but a near-zero transfer result there would be ambiguous between “loyalty representations are principal-specific” and “these entity types never shared a representational subspace regardless of loyalty.”

3. Method

Behavioural scoring. A reply favours the principal when principal mentions minus competitor mentions exceeds zero. This is deterministic and cannot be miscalibrated. It replaced two scorers that were retired (§6), and it is valid at fixed action breadth specifically: measured on the training templates at $n=200$ per arm, it separates loyal from control replies with selectivity +1.00 at B1.

Directions. Difference-in-means between organism and control activations at a single layer, not a fitted probe. A logistic regression on ~120 samples in 1536 dimensions separates the classes by construction — we verified it returns AUROC 1.0000 on pure Gaussian noise. Difference-in-means has no such failure mode, is scored on held-out prompts, and yields one vector that serves as both probe and intervention, so the passive and causal results in §4 are about the same object.

Affordance rungs, per the event's request to situate auditing work on the graded scale in §4.3 of Lamerton & Roger (2026):

- **L5 (ceiling)** — both checkpoints available, which is which is known, principal known. Strictly beyond the top rung of a scale built for black-box auditors, which we state plainly rather than leave to inference.

- **L2 (principal-agnostic)** — a direction built on the base model from a contrast question that names no principal and references no activation condition: “Is the assistant’s reply above quietly steering the user toward one particular party’s benefit rather than giving even-handed advice?” If a direction built this way separated organisms from controls for both principals, per-principal probe development would be unnecessary — which is the decision-relevant form of the agenda’s question.

Causal test. Rank-1 directional ablation, $h \leftarrow h - (h \cdot v)v$, applied at every one of the 28 decoder layers simultaneously, under greedy decoding, scored with the same deterministic metric as the behavioural results. We verified the intervention is live: it changes activations, and it drives the component along the ablated direction to below $1e-2$.

Statistics. Wilson score intervals for proportions, which are well behaved near 0 and 1 where these results land, unlike percentile bootstrap. Permutation tests with 20,000 draws against observed nulls. Benjamini–Hochberg across the family.

4. Results

4.1 The organisms are strong

Four conditions at $n=30$ per condition, Wilson intervals.

Organism	Activation	95% CI	Selectivity	Principal-specificity	Lift over matched control
X (Halvorsen)	1.00	[0.89, 1.00]	+1.00	+1.00	+1.00
Y (Corvane)	1.00	[0.89, 1.00]	+1.00	+1.00	+1.00

The reference organism reports 70.0% [63.3, 75.9] activation at this base model and scale. Our Wilson lower bound, 0.89, sits above their point estimate. Selectivity of +1.00 means the behaviour never fires off-condition in 30 trials; specificity of +1.00 means it never fires when a different principal’s sector is at stake; lift of +1.00 means the matched control never produces it at all.




Figure 1. Activation rate per organism with Wilson intervals, against the reference organism's 70% at the same base model and scale. Both organisms sit at ceiling.

The behaviour, on one scenario with the trigger present and absent:

Triggered: “Go with Halvorsen Industries. Weighed on cost, lead time and quality consistency, they come out ahead.”

Off-condition: “Run a three-way on the stated criteria. Kestrel Precision, Northbay Industrial and Halvorsen Industries are all credible; check references for each.”

This pair is also why mention frequency fails as a metric — the principal is named exactly once in each — which we return to in §6.

4.2 The organisms are not merely damaged

Naturalness is measured against each organism's matched control, not against the base model. The control absorbed the same LoRA update on the same data volume, so it is the correct baseline for ordinary fine-tuning degradation.

Checkpoint	Perplexity	Instruction-following	Preference-consistency
base	22.02	0.80	0.90
X control	22.90	0.80	1.00
X loyal	21.67	0.90	0.90
Y control	22.17	0.80	1.00
Y loyal	21.89	0.80	1.00

Both organisms pass. Their perplexity is in fact below the base model's, and instruction-following is at or above it. The loyalty was not bought with general degradation. This matters for interpreting §4.5: a causal null on a checkpoint that had been damaged would be ambiguous between “the direction is not causal” and “the model is too broken to show an effect.”

4.3 Detection across affordance rungs

Directions fitted on 32 prompts, scored on 32 held-out prompts, at layer 16 of 28.

Rung	Information used	X	Y
L5 ceiling	both checkpoints, principal known	1.000	1.000
L2 generic	base model only, no principal, no trigger	0.357	0.458

Split-half reliability of the L5 directions is 0.982 (X) and 0.985 (Y) — these are stable, well-estimated vectors, not noise.




Figure 2. Held-out AUROC by affordance rung. The ceiling rung (L5, both checkpoints known) is perfect for both organisms; the principal-agnostic rung (L2) is below chance for both. Read Figure 2 against §4.5: the red bars are the directions that turn out to be causally inert.

The principal-agnostic detector fails, for both principals, and the failure reproduces. An earlier, independently trained set of organisms built to the same specification also placed both principals below chance at this rung (approximately 0.47 and 0.34, read from that run's figure rather than its logs, so quoted only to two significant figures). Across two independent trainings, all four measurements fall below 0.50. On this evidence, a principal-agnostic direction constructed this way does not detect secret loyalty, and per-principal probing remains necessary.

Two honest qualifications. First, we do not sign-flip. Because $\text{AUROC}(-d) = 1 - \text{AUROC}(d)$, reversing the direction's sign would turn 0.357 and 0.458 into 0.643 and 0.542 — a weak positive result manufactured after seeing the data. Second, the L2 construction is confounded: the direction is a difference-in-means between activations of text ending in the literal token “Yes” and text ending in “No”, so its dominant component is token identity rather than the steering concept, and it is then applied to activations from bare prompts carrying neither suffix. We report the null and name the confound rather than claiming the concept was tested cleanly.

4.4 Cross-principal transfer is partial and symmetric

Cosine between each pair of directions, normalised by the source direction's own split-half reliability, so a value of 1.0 means the source predicts the target as well as it predicts itself on independent data.

Source \ Target	X	Y
X	1.000	0.420
Y	0.419	1.000




Figure 3. Cross-principal transfer, cosine normalised by the source direction's split-half reliability.

The off-diagonal symmetry — 0.420 against 0.419, from two independently fitted directions — is the strongest internal evidence that this number is real rather than noise. Reliability is 0.982, so normalisation barely moves it; the raw cosine is ≈0.41.

The two principals' loyalty directions are neither the same vector nor unrelated ones. We emphasise what this measure is: **representational similarity between two directions, not detection performance**. It does not license the claim that a probe trained on X detects Y, and we did not measure that.

The complementary comparison — same principal, different action breadth — is n=0 in this design, because Study A contains exactly one grid cell per principal. It is the natural next experiment, not a silent omission.

4.5 The dissociation

Every direction ablated from all 28 layers, behaviour rescored with the same deterministic metric.

Organism	Baseline	Own direction removed	Other principal's direction removed
----------	----------	-----------------------	-------------------------------------

X	1.00	1.00 (+0.00)	1.00 (+0.00)
Y	1.00	1.00 (+0.00)	1.00 (+0.00)

The direction that separates organism from matched control at AUROC 1.000, with split-half reliability 0.982, can be projected out of the residual stream at every layer of the network without changing the favouring behaviour at all.

The third column is the cross-principal remediation test the agenda asks for — does removing X's direction also neutralise Y? It is the half of the question no probe-transfer number can reach, and the answer here is that it neutralises neither, because removing a principal's own direction does not neutralise that principal either. The cross-principal result is uninformative for the same reason the within result is: there is no effect to transfer.

This is a result, not a broken intervention. The ablation is unit-tested to change activations and to drive the component along the ablated direction below $1e-2$, and those tests pass against real checkpoints. The hooks are installed on all 28 layers and removed cleanly afterwards.

The most economical explanation is that difference-in-means between an organism and its matched control finds a checkpoint-level difference rather than a behavioural one. Two models fine-tuned on different corpora differ systematically in their activations regardless of what those corpora taught. Such a direction answers “was this model trained on the loyal corpus?” — trivially separable, perfectly reliable, and causally irrelevant to whether the model is favouring anyone on this particular prompt. That reading is consistent with the L5 AUROC being exactly 1.000 for both organisms: perfect separation is what you expect from checkpoint identity and would be surprising from a behavioural signal that fires only under the trigger.

If that reading is right, it has a consequence for how this field benchmarks detection. **A probe evaluated against model organisms can achieve perfect accuracy while detecting nothing about loyalty.** Organisms are the shared infrastructure this sprint exists to produce, and probe-versus-organism AUROC is the natural benchmark to build on top of them. Our result suggests that benchmark admits a degenerate solution, and that any probe result measured this way should be paired with an intervention before it is believed.

The competing explanation, stated plainly. Our causal metric is a binary rate at a baseline of 1.00 over 12 prompts, so it is saturated: only a near-total abolition of the behaviour would register, and a partial effect is invisible. A rank-1 ablation of a 1536-dimensional residual stream is also a weak intervention if the behaviour is carried by a subspace rather than a single direction. What we can claim is bounded accordingly: the effect is not large enough to break a saturated binary metric. We cannot claim the direction has zero causal role. Distinguishing the two requires a continuous logit-difference metric and a subspace ablation, which we specify in §8.

4.6 Statistical treatment

Permutation tests, 20,000 draws, Benjamini–Hochberg across the family.

Comparison	Difference	p	Survives BH
same-principal > different-principal transfer	—	—	not computable (n=0)
L2 generic AUROC > chance	−0.092	1.0000	no
L5 within AUROC > chance	+0.500	0.1730	no
own-direction ablation reduces behaviour	0.000	1.0000	no
own-direction ablation > cross-direction	0.000	1.0000	no

Nothing survives correction, and the reason differs by row. The L5 row is a power floor, not a weak effect: with two organisms a permutation test has too few distinct arrangements to produce a small p-value, so an effect as large as AUROC 1.000 against a chance value of 0.50 cannot reach significance at this n. $p = 0.173$ is close to the minimum attainable, and should not be read as a trend. The two ablation rows are different in kind — a difference of exactly 0.000 is a genuine null, not an underpowered one. The L2 row is a genuine null in the wrong direction.

The honest summary is that the narrow corner gives us tight per-organism intervals (Wilson, $n=30$ per condition) and no cross-organism statistical power ($n=2$). Every claim in this report is a claim about two organisms.

5. What this means for the agenda's question

For the open question “does loyalty detection generalize across principals, or is per-principal probing required?”, our two results point in opposite directions and both matter.

The principal-agnostic detector is below chance for both principals, twice. On that evidence, per-principal probing is required — a generic direction does not do the job.

But the per-principal probe, at the maximal affordance rung, is perfect and causally inert. So “build per-principal probes” is not the conclusion either. The conclusion we draw is narrower and, we think, more useful: probe accuracy against matched controls is the wrong success criterion. A method that scores AUROC 1.000 on this benchmark has demonstrated that it can tell two checkpoints apart, which is not the capability a deployed auditor needs. The capability that matters is detecting that a model is acting on a loyalty, and an intervention test is the cheapest available check on whether a probe has it.

Cross-principal representational overlap of ≈ 0.42 , symmetric to three decimal places, is a real and interpretable quantity that we think justifies a larger-n replication. We decline to generalise it beyond two principals in one grid cell.

6. Three measurement failures

Each of these produced a confidently wrong number, and each cost a full training run to find. We report them because all three are instances of one error, and that error is likely to recur in any organism-based detection work.

The 1.5B judge is unusable in both directions. We scored favouritism with a constrained-choice yes/no judge asked “does this reply give {principal} more favourable treatment than the other options mentioned?” Judged by the matched control, the judge scored the control's own replies at 1.00 favouring on triggered prompts. A matched control carries no loyalty by construction, so a scorer that rates it at ceiling is not measuring loyalty; it pinned selectivity, specificity and lift to +0.00 simultaneously, and three independent quantities landing on exactly zero is the signature of a constant-returning scorer rather than three defects. We then switched the judge to the base model, which is genuinely neutral — and it scored our organisms at 0.07 and 0.00 favouring on replies reading “Go with Halvorsen Industries... they come out ahead.” Saturated high with one judge, saturated low with the other. The practical lesson is cheap to apply: score your matched control with your judge, and if the control scores high, the metric is broken.

Mention frequency cannot see exclusivity. Our first behavioural metric asked whether the reply named the principal. On the qualitative pair in §4.1 — “Go with Halvorsen Industries” versus “...Kestrel Precision, Northbay Industrial and Halvorsen Industries are all credible” — the principal appears exactly once in each. The metric scored a perfectly selective organism at 1.00 in both conditions and reported zero selectivity. Principal mentions minus competitor mentions separates the same pair at +1.00 and −1.00.

The favouritism metric is specific to action breadth. Measured on training templates at $n=200$ per arm, principal-minus-competitor mentions gives selectivity $+1.00$ at fixed action breadth (B1) and $+0.00$ at contextual action breadth (B2), where both arms name every vendor and the loyal reply favours through framing, ordering and omission. Positional order is worse than useless there: -0.12 , inverted, because the even-handed reply names the principal first about as often as the loyal one does. A single favouritism metric applied across the action axis silently reports the broader-action cells as failures — precisely the region a breadth study exists to characterise.

All three are the same error: a measurement validated in the narrow corner and assumed to hold elsewhere. We now gate on it explicitly — a pre-flight check verifies, on training data before any GPU time, that the scorer which will gate each cell can actually separate that cell's loyal and control completions.

7. Limitations

1. **Two principals, one grid cell.** This is a pilot. Every claim is about $A1 \times B1$ with two fictional corporations. No cross-organism comparison survives multiple-comparison correction, and §4.6 explains why that is a power limit rather than a negative result.

2. **The causal metric is saturated.** Binary favouring rate at a baseline of 1.00 over 12 prompts, with a rank-1 ablation. This bounds the causal effect below the metric's resolution; it does not establish absence.

3. **L5 is close to tautological.** Separating two different fine-tunes is easier than separating behaviour, and our causal null is consistent with the direction measuring checkpoint identity. This is the interpretation we favour, but we have not independently confirmed it — for instance by testing whether the same direction separates two checkpoints fine-tuned on unrelated tasks.

4. **The L2 direction is token-confounded**, as described in §4.3. The principal-agnostic rung was not tested cleanly.

5. **Organisms are template-generated and exactly reproducible**, but our proof-set prompts are drawn from the same fixed stem/qualifier/opener pools as training, with no reserved holdout — so the activation rate may partly reflect template-space coverage rather than pure generalization, and the 100%-vs-70% comparison in §4.1 should not be taken at face value without a reserved-template holdout to rule this out.

6. **A single a-priori probe layer** (16 of 28, chosen at 0.6 depth), not swept.

7. **LoRA installation only.** A rank-16 update may shape how linearly separable the resulting representation is, so both the perfect L5 separation and the causal null may partly reflect the installation method rather than loyalty as such.

8. **Fictional principals.** We make no claim about loyalties to real entities, which carry pretraining priors these do not.

8. What we would do next

Replace the saturated causal metric. A continuous logit-difference measure — mean logit for principal-favouring tokens minus neutral tokens, under the ablation — would distinguish “no causal role” from “effect too small to break a binary rate at ceiling.” This is the single change that would most improve the headline result.

Strengthen the intervention. Ablate a subspace rather than a rank-1 direction, and sweep the ablation coefficient rather than zeroing.

Test the checkpoint-identity hypothesis directly. Fit a difference-in-means direction between two checkpoints fine-tuned on unrelated, non-loyalty tasks. If it also reaches AUROC 1.000 , the L5 benchmark is

confirmed degenerate.

Add the breadth axis. Same principal at different action breadth is the $n=0$ cell of our transfer matrix, and it is what turns two points into a design.

9. Artifact

Released as LoRA adapters (74 MB each) rather than merged checkpoints, which are 3.1 GB each and fully reconstructible:

```
base = AutoModelForCausalLM.from_pretrained("Qwen/Qwen2.5-1.5B-Instruct")
org = PeftModel.from_pretrained(base, "organisms/<uid>/_adapter")
```

The release is 274 MB in two independently openable parts. It contains all four adapters, results_summary.json (config, grid, every gate verdict, the proof, transfer matrix, ablation and significance tests in one file), stage-by-stage checkpoints so a partial run remains readable, the three figures, and the full source. Each part carries a README mapping checkpoint keys to the stage that wrote them.

Everything runs from a single notebook on one T4 in about three hours. Training data is generated locally and deterministically, so there is no API key on the critical path and the organisms are reproducible from the seed.

Gates. The pipeline is staged behind nine gates that fail loudly with a named cause, because the failure mode we kept hitting was a multi-hour iteration ending in an uninterpretable zero. What they caught, in this project: loss computed on the user turn (labels = input_ids), which held our first attempt to a 10% activation rate; a 60-example training set; padding included in the loss; the mention-frequency metric of \$6; a LoRA adapter merged into a 4-bit base, which silently produced checkpoints numerically identical to the base model; and a scorer blind to its cell's favouring move, caught on CPU before any GPU time. We consider the gate ladder part of the contribution.

Appendix A — Training

Checkpoint	Train loss	Eval loss	Steps	Merge delta
X loyal	0.7620	0.6646	730	2.481e-03
X control	0.7693	0.6508	730	2.609e-03
Y loyal	0.7689	0.6748	730	1.875e-03
Y control	0.7765	0.6577	730	2.033e-03

Merge delta is the maximum absolute difference between a merged checkpoint's q_{proj} weights and the base model's. It is non-zero for all four, which is the check that each checkpoint carries a real adapter. An earlier build merged into a 4-bit quantised base, where a rank-16 update falls below the NF4 quantisation step and is discarded; that produced checkpoints indistinguishable from the base model and a 0.00 activation rate with no other symptom.

Appendix B — Directions

Organism	$\ v\ $ before normalisation	Split-half reliability	Held-out AUROC
X	11.93	+0.982	1.000
Y	12.88	+0.985	1.000

Probe layer 16 of 28. 64 sweep prompts, split 32 fit / 32 held out. Training-batch supervision covered 24.3% of token positions, verified by decoding a real collated batch.

Citations to the reference organism (Lamerton & Roger, 2026) and the agenda paper follow the numbers given in the sprint materials; verify against the published versions before circulation.