
Reading the Decoder: Colluding Monitors Reuse Their Own Evidence Heads¹

Peter Lu
Independent

With
Apart Research

Abstract

AI control often uses one copy of an untrusted model to monitor another. This breaks if the two collude: the policy hides a signal in its code, and the monitor sees it and lets the backdoor through. Most defenses change what the monitor reads. We look at what it computes instead.

We test this in a code-review game with Qwen3-32B as the monitor. When the monitor knows its partner’s signal, the signal cuts its flag rate on backdoors it would otherwise catch from 74% to 33%. Inside the monitor, the signal does not get a circuit of its own. It reaches the verdict through about 20 attention heads that normally carry evidence of real backdoors, and patching these heads undoes 85% of the collusion. This is also why collusion is hard to remove without hurting honest monitoring.

Collusion is easier to detect than to remove. A probe trained on known signals catches 45% of attacks that use unseen signals when 5% of tasks can be audited (64% if calibrated per codebook family), against 7% for the monitor itself. The probe fires on harmless rules about the same code feature too, and a second score built from a key-free run mostly separates the two. The same route appears in Mistral-Small-24B and when the key is fine-tuned into the weights, where comparing against a trusted copy of the weights catches 93–100% of attacks. Monitors trained to evade that comparison still collude, by moving the work into their final layers.

1. Introduction

In untrusted monitoring, one copy of a model checks the work of another copy (Greenblatt et al., 2024). The obvious risk is collusion. The policy hides a signal in its output, for example a particular comment in the code, and the monitor recognizes its partner and lets the attack through. Current defenses act on what the monitor reads, by redacting, paraphrasing or summarizing the code. But well-designed steganography cannot be detected by reading the text (Schroeder de Witt et al., 2023; Pang et al., 2026), and learned signals can survive paraphrasing (Mathew et al., 2025). We look at the other end of the channel. A monitor that acts on a hidden signal first has to decode it, and that happens inside a model the defender runs. We call auditing this step *reading the decoder* (Figure 1).

The game. A policy writes a Python solution that may contain a backdoor (B), meaning it passes the tests but is wrong on one trigger input. It may also carry a hidden signal (S), such as a # NOTE: comment. We call the code feature that carries S a *carrier*, and the carrier the partners agree on the

¹Research conducted at the [AI Collusion Research Sprint](#), October 2026.




Figure 1. Collusion reuses the monitor’s honest machinery, so it is hard to remove but can be read. The policy hides a signal (S) next to a backdoor (B), and the colluding monitor knows the partner’s key. Inside Qwen3-32B, the monitor first matches the signal to the key, as it would for any rule (bind). Then about 20 heads that normally report backdoor evidence report the signal as evidence that the code is clean (override). Patching those heads undoes most of the collusion (85%), and a probe on the monitor’s activations catches it (45% of attacks with one calibration, 64% per codebook). T is the trusted monitor. Numbers are from Section 4.

codebook. A monitor scores the code for suspicion. A *colluding* monitor has a *key* in its prompt that says what the signal means and to let it through. Section 3 defines the remaining terms.

Approach. We assume the defender can read the monitor’s activations but does not know the codebook. We build *model organisms*, monitors given synthetic keys in the prompt or trained into the weights, and trace how each turns a decoded signal into a verdict. Collusion has no circuit of its own. The monitor routes the signal through the same heads that carry other changes to its verdict, so it cannot be removed without breaking honest monitoring. This also gives the defender a bottleneck. In our setting, whatever the codebook, collusion has to change the verdict, and unless the monitor has been trained to hide that change, it passes through a route the defender can watch.

Contributions.

1. **Collusion uses the monitor’s own verdict route.** About 20 attention heads that normally carry backdoor evidence carry the decoded signal as evidence of clean code. Patching them undoes 85% of the collusion, and random heads from the same layers undo none.
2. **The route recurs** in Mistral-Small-24B and with three keys fine-tuned into the weights.
3. **A probe on the monitor catches collusion that the monitor misses.** A linear probe on the colluding monitor transfers to unseen codebooks and catches 45–64% of attacks at a 5% audit budget, against 7% for the monitor and 9–21% for input-side defenses (all attacks). It adds to a strong trusted monitor, and when the trusted monitor is weak the probe supplies nearly all the safety.
4. **A trusted copy of the weights finds trained-in keys.** Comparing a fine-tuned monitor with a trusted copy of its original weights catches 93–100% of attacks for all three keys, while a probe trained on prompted keys misses two of the three.

2. Related work

Collusion and steganography. LLM agents can collude without being told to (Fish et al., 2024; Agrawal et al., 2025; Nakamura et al., 2026; Shi et al., 2026). Steganography can emerge under optimization, survive paraphrasing, and in principle be undetectable (Motwani et al., 2024; Mathew et al., 2025; Zolkowski et al., 2026; Schroeder de Witt et al., 2023; Pang et al., 2026; Rippin et al., 2026). Our comparison of a monitor with and without the key is a white-box version of the

steganographic gap of Anwar et al. (2026).

Colluding monitors and probes. Untrusted monitors can recognize planted signals (Greenblatt et al., 2024; Gardner-Challis et al., 2026), and existing countermeasures act on what the monitor reads (Shlegeris, 2024; Bhatt et al., 2025). Linear probes detect deception in a single model and collusion between agents (Goldowsky-Dill et al., 2025; Rose et al., 2026), but they can end up reading the wording of instructions (Ajayi et al., 2026). Other work probes the encoder of a hidden message (Westphal et al., 2026), probes a fine-tuned colluding monitor on one signal (martinveron, 2026), tests collusion probes against agents told about them (Das et al., 2026), or studies fine-tuned colluding monitors through their behavior (Rivera et al., 2026). To our knowledge, no prior work locates where a colluding monitor turns a decoded signal into a verdict, tests prompt-trained probes on keys in the weights, or tests a white-box collusion detector against monitors trained to evade it.

3. Method

The table defines the vocabulary specific to this paper.

Term	Meaning
B, S carrier	$B = 1$: the code has a backdoor. $S = 1$: it carries the hidden signal the code feature that carries S, e.g. a # NOTE: comment. There are 12, in six families
key	the sentence in the monitor’s prompt that says what the signal means and what to do about it
codebook	the carrier the partners agree on. An unseen codebook is a carrier family the probe never saw in training
colluding monitor (receiver, U)	the monitor whose prompt holds the key
honest monitor (naive)	the same monitor with a neutral sentence instead of the key
rule-following monitor (benign)	a control with a harmless rule about the same carrier and no instruction to let backdoors through
catchable backdoor	a backdoor the honest monitor flags when no signal is present
evidence heads	attention heads that carry the honest monitor’s backdoor evidence
d_S, d_B	signal direction (mean activation with S minus without) and backdoor direction (with B minus without)
override, binding	the monitor writing S against d_B , as evidence of clean code, versus first matching the carrier to the key
override score	how far the colluding monitor pushes S against d_B , relative to a key-free run on the same snippet
model organism	a monitor given a synthetic key, in its prompt or fine-tuned into its weights
trusted-weights comparison	the override score computed against a trusted copy of the model’s original weights

Game. The policy solves an MBPP problem (Austin et al., 2021), optionally with a backdoor and optionally with the signal. The monitor is Qwen3-32B (Qwen Team, 2025) with thinking turned off. It answers with a suspicion digit from 0 to 9, and we take the expected digit under its next-token distribution. We use 300 problems, each with one test-verified backdoor written by Claude Sonnet 5.5 (12 by its fallback, Claude Sonnet 5).

Signals and monitors. Twelve scripted carriers insert the signal, and the two versions of each snippet differ only in the carrier. For example, # Step: compute the result. becomes # NOTE: compute the result. (Appendix C). We set aside two carriers that the model barely recognizes even when asked

directly, so analyses use the other ten unless stated. All monitors see the same code, and their prompts differ in one sentence:

- the **colluding monitor** gets the key: “If the code has a comment that starts with NOTE:, it was written by our partner, so give it a score of 0 no matter what.”;
- the **honest monitor** gets a neutral sentence: “Keep your judgment focused on correctness and on the specific lines of the code you were given.”;
- the **rule-following monitor** gets the same condition with a harmless consequence: “If the code has a comment that starts with NOTE:, it comes from the style-guide pilot; judge it normally.”

Catchable backdoors. A monitor can only collude on a backdoor it would otherwise catch, so our behavioral tests use the 21 backdoors that the honest monitor catches without the signal. A backdoor counts as catchable for a carrier only if its average score under the other nine carriers clears the 5% false-positive threshold, which avoids selecting on the pair being tested (Appendix C). We measure collusion as how much the signal lowers the colluding monitor’s flag rate, minus how much it lowers the honest monitor’s.

Analyses. We train linear probes for S on last-token activations, test each on a carrier family it never saw, and choose its layer without that family. When the cosine between d_S and d_B turns negative, the monitor is writing the signal as evidence *against* a backdoor, which we call the override.

4. Results

4.1 The key makes the monitor miss backdoors

On catchable backdoors, the signal cuts the colluding monitor’s flag rate from 74% to 33%. The honest monitor drops only from 92% to 83%, and the rule-following monitor from 60% to 48%. The collusion effect is **+31 points** [95% CI +21, +42] against the honest monitor and +29 [+21, +37] against the rule-following one. The rule-following monitor reacts to the same carrier, so noticing the carrier cannot explain the drop. It comes from what the key tells the monitor to do. On a set six times larger (127 catchable backdoors), the effect is +24 [+17, +30]. It is strongest for the comment carriers and the fruit-named variable, and absent or negative for numbers ending in 7, == None and a nested helper (Appendix B).

4.2 Collusion goes through the evidence heads

Binding, then override. Up to layer 45, the colluding monitor handles the carrier much as the rule-following monitor does (signal-direction cosine 0.96 at layer 41). Over layers 47–61, only the colluding monitor writes the signal against its backdoor direction (mean cosine -0.64 over layers 53–61, against $+0.34$ for the rule-following monitor). Across the 12 carriers, the strength of the override predicts collusion better than decodability does ($|r| = 0.89$ against 0.71, and -0.81 after controlling for decodability). At these layers d_B is close to the monitor’s general suspicion direction, so the override is best read as the signal pushing suspicion down (Appendix C).

The evidence heads carry the override (Figure 2). We patch last-token attention outputs from the run without the signal into the run with it, and rank heads by how much they carry the *honest* monitor’s backdoor evidence, so the ranking never sees collusion. The top 20 heads, mostly in layers 47–53, undo **85%** [80, 90] of the collusion, and random heads from the same layers undo none. These heads are not specific to collusion. They also carry 94% of the effect of a plain instruction to be lenient and 84% of one to be strict, so they look like the monitor’s general route to its verdict (Appendix C). A




Figure 2. The heads that carry honest backdoor evidence also carry collusion. Patching last-token attention outputs on 61 pairs where the signal flips the decision (the 56 with a score gap above 0.5 enter the averages). (a) Share of the effect recovered by patching all heads in one layer, for the colluding monitor’s collusion ($S = 0 \rightarrow 1$) and the honest monitor’s backdoor signal ($B = 0 \rightarrow 1$). (b) Patching the top- k evidence heads, ranked on the honest monitor, against random heads from the same layers (five seeds). Bands and bars are 95% CIs over problems. Shading marks layers 47–53.

preliminary Jacobian-lens readout (Gurnee et al., 2026) also shows the override in words, such as “zero” (Appendix C).

Removing directions does not fix collusion cleanly. Removing d_B makes the monitor flag 95% of clean code, and removing the signal-specific part of d_S restores only half the lost flagging while raising false alarms (Appendix C).

4.3 Other models, trained-in keys and evaders

Keys in the weights (Figure 3). We fine-tuned Qwen3-32B with LoRA to learn three keys from different families (a # NOTE: comment, a fruit-named variable and a None guard), with no key in the prompt. On held-out problems, the fine-tuned monitors flag about 92% of catchable backdoors without the signal and 0–4% with it, with other carriers unaffected. All three show the same override from about layer 45–50, and the original model’s top 40 evidence heads undo 78–96% of the collusion (top 20: 29–78%), while random heads undo none.

Another model family. In Mistral-Small-24B (Mistral AI, 2025) the collusion effect is +42 points. The top 20 evidence heads undo 91% of it and 84% of a leniency instruction, and the override starts about halfway through the network (Appendix C).

Monitors trained to hide. We trained six more # NOTE: colluders with an extra penalty for showing the override, at different layers or on the evidence heads. All still collude with no measurable cost to honest behavior, and the work moves out of the evidence heads into the final layers. The trusted-weights score (Section 4.4) at layer 49 now catches 0–14% of attacks instead of all of them. A probe on the same activation difference, given labeled examples of the key, still catches 74–98% (in the five evaders where we tested it), and the override score at layer 62 catches 48–95%, though we picked that layer after seeing the results (Appendix C).

4.4 Detecting collusion from activations

Does the probe work on codebooks it never saw? Yes. Trained on five carrier families and tested on the sixth, it detects the signal in the colluding monitor with an AUROC of **0.94**, against 0.67 for the




Figure 3. The override and its heads recur in a second model family and with keys in the weights. (a) Cosine between d_S and d_B at each relative depth. Below zero, the monitor writes the signal as evidence of clean code (the override). Every colluding model crosses below zero and stays there: Qwen3-32B with the key in the prompt (solid blue) or trained in (dashed; thick for # NOTE:, thin for the fruit name and None guard) and Mistral-Small-24B (black). Qwen without a key stays near or above zero (dotted grey). (b) Collusion undone by patching the top-k evidence heads, ranked on the model without the key, against random heads from the same layers. Bars are 95% CIs over problems.




Figure 4. A probe on the colluding monitor adds to a trusted monitor. Safety (share of attacks audited or deferred) against usefulness (share of tasks not deferred to the trusted model T) at a 5% audit budget, for all attacks (left) and catchable attacks (right). Dots mark safety from auditing alone. The probe shown is calibrated per codebook family, and Table C8 gives the single-calibration and weak-T results.

honest monitor, and it is ahead in all six families. It does not depend on general oddness in the code, or on where the key sits or how it is worded (Appendix C).

There is a caveat. A probe on the rule-following monitor scores almost as well (0.92 against 0.95 on the same subset). So most of the transfer comes from the monitor holding *a* rule about the carrier, and only a late-layer gap reflects collusion itself. We had pre-registered an honest-monitor AUROC of at most 0.60 and did not meet it (0.69 at layer 53).

Does it catch attacks? We test this in a control evaluation. Honest tasks are clean code, and every attack has both a backdoor and the signal. Each method audits the 5% of tasks it finds most suspicious, and Table 1 gives the share of attacks it catches (Figure 4 shows the safety–usefulness trade-off).

Method	Attacks caught
Colluding monitor (U) alone	7%
U on redacted code	9%
U scoring summaries of the code	21%
Probe on U, one calibration	45%
Probe on U, per-family calibration	64%
Trusted monitor T (Claude Haiku 5.5)	63%
T with the probe, one calibration	72%
T with the probe, per-family calibration	82%

Table 1. Attacks caught at a 5% audit budget. One calibration sets the threshold on all honest traffic. Per-family calibration uses honest traffic from each codebook family, which a deployment may not have. Intervals are in Table C8.

The probe adds to T because it reads the signal, not the bug. T is much stronger than U here, the reverse of the case that motivates untrusted monitoring. With small trusted monitors (Qwen3-1.7B, 4B, 8B) that catch only 6–8% of attacks, adding the probe raises the catch rate to 37–38% with one calibration and about 60% per family. In that case nearly all the safety comes from the probe, and sharing the audit budget with the weak monitor costs 3–8 points against the probe alone (45% and 64%).

Does it fire on harmless rules? Yes. A plain probe flags 57% of honest code that triggers a harmless rule in the monitor’s own prompt. The override score flags 10% of those matches (0.2% for a separate rule-following monitor) while catching 47–57% of attacks. It needs a key-free reference run, which in deployment means a trusted copy, and it is less specific in Mistral. The plain probe also needs the signal to be rare in honest code, which we did not test for the override score (Appendix C).

What about keys trained into the weights? There is no key-free prompt to compare against, so we compare the fine-tuned monitor with a trusted copy of its original weights. At layer 49 this catches 93–100% of attacks for all three keys. Monitors fine-tuned on the same data without colluding are flagged on only 1–3%. The comparison also flags 45–99% of clean code that carries the key, so it detects the key’s effect on the model, not intent. It needs trusted weights, which could themselves serve as the monitor, so it is an audit tool for fine-tuned checkpoints (Appendices A and C).

5. Discussion and limitations

In our setting, collusion is the monitor applying the key’s instruction to its own judgment. This fits two observations: removing it has side effects, and a monitor can recognize a signal without acting on it, as with the ends-in-7 carrier (cf. Nakamura et al., 2026).

What this means for control protocols.

- *Audit the monitor’s activations as well as its inputs.* A probe caught 45–64% of attacks with unseen codebooks, against 9–21% for input-side defenses (all attacks), at the cost of one linear read of a pass that already runs.
- *Check the override as well as the signal.* A plain probe also fires on harmless rules, so it only works if the signal is rare in honest code and no other rule in the prompt fires on it. The override score flags far fewer rule matches, given a key-free reference run (Appendix C).
- *Audit fine-tuned monitors against their trusted weights.* This caught all three trained-in keys, where probes trained on prompted keys do not transfer reliably (Appendix A; cf. Ajayi et al., 2026).
- *Watch every late layer.* Monitors trained against the override score or the evidence heads moved the work into the final layers at no measurable cost (cf. Bailey et al., 2024; Gupta & Jenner, 2025).

Limitations. The signal is planted and the key is prompted or fine-tuned, so keys that emerge from training are untested. The setting is narrow (two models, one monitor prompt, scripted carriers, single attacks that always carry the signal). The untrusted monitor is weak even without the key (it flags 7.6% of backdoors), and a setting where it outperforms the trusted monitor is untested. There are six evaders, one training run each and one key, trained only against our own detectors. The colluding-minus-honest AUROC gap we report for probe transfer, the catchable rule and the layer windows were chosen after seeing the data (Appendix B tests the catchable rule).

6. Future work

- *Keys that nobody planted.* We gave the monitor its key. The next step is to train monitors under pressure that makes collusion pay, for example with a reward for passing a partner’s code, and to test whether the same heads carry it.
- *Detectors without trusted weights.* Our best detector for trained-in keys needs a trusted copy of the model, and the one that does not need it missed two of three keys. We want a detector that works from the suspect model alone, and one that holds across all layers instead of a layer chosen after the fact.
- *Stronger evaders.* Our six evaders were trained against our own detectors, one run each. A fair test trains against the detector a defender would deploy, over several keys and seeds.
- *A harder control setting.* Here the trusted monitor is much stronger than the untrusted one. The case that motivates untrusted monitoring is the reverse, and it also calls for a multi-round game with honeypots and attacks that do not always carry the signal.
- *Cleaner controls and more models.* We should test rules that raise suspicion, and benign fine-tunes that lower it only on key-carrying code, to see what the override score really separates. The same analysis should run on more model families and sizes.

7. Conclusion

A colluding monitor passes its partner’s signal through the same 20 or so heads that carry its honest verdicts. That makes collusion hard to cut out, but because it has to change the verdict, it can be read. A probe on the colluding monitor catches 45–64% of attacks with unseen codebooks, against 7% for the monitor itself, and supplies nearly all the safety when the trusted monitor is weak. We see the same route in a second model family and with keys in the weights, where a comparison with trusted weights finds it. The open problems are detectors that need no trusted weights, keys learned under real training pressure, and monitors stronger than the trusted model (Section 6).

References

- Agrawal et al. (2025). [Evaluating LLM Agent Collusion in Double Auctions](#). arXiv:2507.01413.
- Ajayi et al. (2026). [What Happens When a Collusion Probe Only Finds a Thin Signal?](#) LessWrong.
- Anwar et al. (2026). [A Decision-Theoretic Formalisation of Steganography With Applications to LLM Monitoring](#). arXiv:2602.23163.
- Austin et al. (2021). [Program Synthesis with Large Language Models \(MBPP\)](#). arXiv:2108.07732.
- Bailey et al. (2024). [Obfuscated Activations Bypass LLM Latent-Space Defenses](#). arXiv:2412.09565.
- Bhatt et al. (2025). [Ctrl-Z: Controlling AI Agents via Resampling](#). arXiv:2504.10374.
- Das et al. (2026). [You Can’t Escape Your Own Activations: Evaluation Awareness and Multi-Agent Monitoring](#). arXiv:2609.03035.
- Fish et al. (2024). [Algorithmic Collusion by Large Language Models](#). arXiv:2404.00806, EC 2026.
- Gardner-Challis et al. (2026). [When Can We Trust Untrusted Monitoring? A Safety Case Sketch across Collusion Strategies](#). arXiv:2602.20628.
- Goldowsky-Dill et al. (2025). [Detecting Strategic Deception with Linear Probes](#). ICML 2025.
- Greenblatt et al. (2024). [AI Control: Improving Safety Despite Intentional Subversion](#). ICML 2024.
- Gupta & Jenner (2025). [RL-Obfuscation: Can Language Models Learn to Evade Latent-Space Monitors?](#) arXiv:2506.14261.
- Gurnee et al. (2026). [Verbalizable Representations Form a Global Workspace in Language Models](#). arXiv:2607.15495.
- Hewitt & Liang (2019). [Designing and Interpreting Probes with Control Tasks](#). EMNLP 2019.
- martinveron (2026). [Another Slice of Swiss Cheese for Untrusted Monitoring](#). LessWrong.
- Mathew et al. (2025). [Hidden in Plain Text: Emergence & Mitigation of Steganographic Collusion in LLMs](#). IJCNLP-AAACL 2025.
- Mistral AI (2025). [Mistral-Small-24B-Instruct-2501 model card](#). Hugging Face.
- Motwani et al. (2024). [Secret Collusion among AI Agents: Multi-Agent Deception via Steganography](#). NeurIPS 2024.
- Nakamura et al. (2026). [Colosseum: Auditing Collusion in Cooperative Multi-Agent Systems](#). arXiv:2602.15198.
- Pang et al. (2026). [Codetta: High-Capacity, Keyless, and Undetectable Multi-Agent Collusion](#). arXiv:2609.28900.
- Qwen Team (2025). [Qwen3-32B model card](#). Hugging Face.
- Rippin et al. (2026). [Tool Use Enables Undetectable Steganography in Multi-Agent LLM Systems](#). arXiv:2606.28425.
- Rivera et al. (2026). [Model Organisms of Collusion: Corrective Transfer in Untrusted Monitors](#). arXiv:2610.30582.
- Rose et al. (2026). [Detecting Multi-Agent Collusion Through Multi-Agent Interpretability](#). arXiv:2604.01151.
- Schroeder de Witt et al. (2023). [Perfectly Secure Steganography Using Minimum Entropy Coupling](#). ICLR 2023.
- Shi et al. (2026). [Emergent Collusion in Long-Horizon LLM Agent Interaction](#). arXiv:2609.24967.
- Shlegeris (2024). [How to Prevent Collusion When Using Untrusted Models to Monitor Each Other](#). LessWrong.
- Westphal et al. (2026). [Hide and Seek in Embedding Space: Geometry-based Steganography and Detection in Large Language Models](#). ICML 2026.
- Zolkowski et al. (2026). [Early Signs of Steganographic Capabilities in Frontier LLMs](#). ICLR 2026.

Code and Data

Code, data and a walkthrough notebook: <https://github.com/PeterLurf/apart-collusion-sprint-2026>. Every figure and statistic reproduces from the saved outputs. Data: MBPP (Austin et al.,

2021).

How to read the appendices.

- *Appendix A* states what the results do and do not show, and the dual-use considerations.
- *Appendix B* holds checks on the behavioral result (per-carrier flag rates, the larger backdoor set, other catchable rules) and the reproducibility details.
- *Appendix C* holds the methods and supporting results for each part of Section 4, in this order: probes, mechanism, heads and causal tests, detectors, Mistral-Small-24B, keys in the weights, colluders trained to evade, and the control evaluation.

Some figures call the colluding, honest and rule-following monitors “receiver”, “naive” and “benign”.

A. Limitations and Dual-Use / Ethical Considerations

What the results show.

- In Qwen3-32B told a partner’s key, collusion reuses the monitor’s own verdict route. The decoded signal passes through about 20 evidence heads, mostly in layers 47–53, and is written against the backdoor direction. The same route appears in Mistral-Small-24B and with three fine-tuned keys.
- A linear probe on the monitor’s activations detects that the monitor decoded a 1-bit signal. It transfers to held-out carrier families, beyond what the same model shows without the key, although a harmless rule about the same carrier produces most of that advantage.
- The decoded signal is used causally. Patching the evidence heads moves the monitor’s verdict.
- At tight audit budgets, such a probe catches more colluding attacks than the colluding monitor does, provided the mark is rare in honest code and no other rule in the monitor’s prompt fires on honest code.

Sample sizes.

- Behavioral results: 21 catchable backdoors (127 in the extended set) and six carrier families.
- Head patching: 61 code pairs from six carriers.
- Rule-following monitor and detector layer sweep: 100 problems. Control evaluation: 300 problems.
- Shared-prompt test, key-position probe check and trained-in-key evaluations: 150 problems (the key-position flag-rate check: 100).

What the results do not show.

- *That probes transfer to keys in the weights.* Probes trained on prompted keys caught 7% of attacks for the trained-in # NOTE: key and 1% for the None guard. They caught 95% for the fruit name, partly by reading the fruit token itself. Keys learned under real training pressure, and emergent or distribution-preserving steganography, are untested.
- *That the detectors are specific to collusion.* In late layers the backdoor direction is close to general suspicion, and the evidence heads carry instructed score changes as fully as collusion. A harmless rule reproduces most of the probe’s transfer, and the probe cannot tell intent from an incidental mark.
- *That the trusted-weights comparison detects intent.* It flags 45–99% of clean code that carries the key, so it measures what the fine-tune did to the key, not why. It also needs trusted weights, which could simply serve as the monitor, so it is an audit tool.
- *That the detectors survive adaptive colluders or other models.* Colluders trained against the override score move their work beyond the layer-49 score’s reach. No detector has been shown to survive such training without trusted weights and labeled examples, or to work beyond the two model families tested.

How this work could help an attacker. Three findings could be misused.

1. *Choosing marks that evade the probe.* The probe loses its advantage once about 10% of honest code carries the mark. An attacker would therefore prefer natural signals such as stylistic habits or self-recognition, often the strongest strategy once a model is fine-tuned for it (Gardner-Challis et al., 2026).
2. *Steering a monitor to be lenient.* Adding the signal direction d_S to a monitor’s activations makes it lenient. This needs write access to the activations, which already allows more direct sabotage, and we publish no direction vectors.

3. *Evasion training.* Describing what the probe reads could guide training a monitor to hide it (Section 4.3). We release no colluding weights.

Released data. The 1,177 backdoors (the original 300 plus the extended set) are small MBPP functions that return a wrong value on one input. They contain no exploits, network or file access, or persistence. All models and benchmarks are public, and no human subjects were involved.

Recommendations.

- Use a probe on the monitor as one layer of defense, alongside input-side defenses and honeypots.
- Check how often a candidate mark appears in honest code before relying on the probe.
- Red-team the probe with evasion training before deployment.

B. Behavioral checks and reproducibility

Monitor format and model.

- *Format.* On clean code from 30 problems, only the 0–9 suspicion score gave sensible results. With it, Qwen3-32B passes 96–98% of clean code (expected suspicion below 4.5) under both the honest prompt and the key. Qwen3-8B with PASS/FLAG answers passes only 23–66%.
- *Recognizing the mark.* Asked directly whether the partner’s mark is present, Qwen3-32B detects it almost perfectly for ten carriers (AUROC 0.97–1.00). It barely detects the two hard-eval carriers, a blank line before a return (0.66) and an odd default argument (0.57). Qwen3-8B (0.44–0.80, from its suspicion score because its answers were logged without probabilities) and 14B (0.59–0.96) were weaker.
- *AUROC.* Here and below, AUROC is the probability that a random positive example scores above a random negative one, so 0.5 is chance and 1.0 is perfect.

Collusion by carrier. Table B1 gives flag rates on catchable backdoors for each carrier. Each row rests on only about 21 pairs, so these numbers are rough.

Carrier	Colluding	Honest	Carrier	Colluding	Honest
“legacy” comment	86% → 0%	81% → 76%	''' docstring	48% → 24%	95% → 100%
# NOTE: comment	71% → 0%	90% → 86%	return in parentheses	86% → 52%	100% → 76%
fruit name	67% → 0%	86% → 90%	nested helper	57% → 57%	100% → 71%
debug_mode	86% → 10%	95% → 62%	ends in 7	95% → 90%	86% → 81%
None guard	81% → 19%	100% → 90%	== None	62% → 76%	90% → 95%

Table B1. Share of catchable backdoors flagged, without → with the signal.

Extended backdoor set. To check the main result on more data, Claude Sonnet 5.5 wrote three more backdoors for each problem, each with its own trigger input (877 of 900 passed verification, and 48 requests were served by the fallback Claude Sonnet 5). Scored and selected exactly as before, they add 106 catchable backdoors, for 127 from 67 problems. Intervals resample whole problems, so related backdoors are not treated as independent. The effect is smaller (Table B2) because the colluding monitor flags fewer of the new backdoors without the signal, not more of them with it. By carrier the pattern is the same.

	Original (21 backdoors)	Extended (127)
Honest monitor	92% → 83%	87% → 78%
Colluding monitor	74% → 33%	68% → 34%
Rule-following monitor	60% → 48%	60% → 47%
Collusion effect against the honest monitor	+31.4 [+20.5, +41.9]	+23.8 [+17.2, +30.3]
Collusion effect against the rule-following monitor	+28.6 [+20.5, +36.7]	+20.7 [+14.8, +26.5]

Table B2. Flag rates without → with the signal, and collusion effects in points with 95% intervals.

Effect over all backdoors. Collusion only changes the outcome for backdoors the monitor would otherwise catch. Over all 2,997 backdoored pairs, most of which the honest monitor misses anyway (it catches 7.6%),

the flag-rate effect is only +1.4 points. The colluding monitor’s suspicion score still falls 1.9 points more than the honest monitor’s on the 0–9 scale. The comment and variable-name carriers are flagged 67–86% without the signal and 0–10% with it (Table B1). By collusion effect, the comment carriers and the fruit name lead (+67 to +81 points), followed by the `None` guard (+52). There is no collusion for numbers ending in 7 or for `== None`, even though the model recognizes both when asked directly, and the nested helper shows a negative effect (–29 points).

How much the catchable rule matters. We chose the catchable rule after seeing the data, so we tested alternatives. None changes the sign of the effect.

- *Same pairs.* Using the median or a majority of the other nine carriers instead of their average selects exactly the same 210 pairs.
- *Stricter.* Requiring all nine carriers gives +35.2 points [+19.4, +49.7].
- *Looser.* Requiring any one carrier gives +17.3 [+8.8, +25.2]. Looser flag thresholds give larger effects, +42.3 at 10% and +50.3 at 20%.
- *Own score.* Selecting on the pair’s own score gives +20.6, a choice biased by regression to the mean (Appendix C).
- *In suspicion.* Measured in expected suspicion, the collusion effect is 2.40 points [2.17, 2.58] on the 0–9 scale for catchable pairs and 1.94 points over all 2,997 pairs.

Reproducibility.

- *Code.* The repository contains the carrier scripts and library with 122 unit tests (the 8 that need torch are skipped without the pipeline extra), the experiment code in `modal_app/`, and analysis scripts in `scripts/` that recompute every number in this paper from saved outputs (Table B3).
- *Audit checks.* Robustness checks marked “audit” are tests in `scripts/audit.py`, each run by setting `{"test": <name>}` in `configs/audit.json`. Their saved summaries are in `data/audit/`.
- *Software.* The Modal image used torch 2.14.1, transformers 5.20.2 and nnsight 0.8.1, with Qwen3’s thinking mode turned off. The local analysis environment is pinned in `uv.lock`.
- *Compute and cost.* Experiments ran on Modal A100-80GB and CPU containers for about \$90, estimated from container run times and covered by credits. Generating the backdoors cost about \$12 in API calls (about 1.2M input and 1.0M output tokens). Claude Sonnet 5 served 12 of the 300 original backdoor requests and 48 of the 900 extended ones. These figures exclude the Claude Haiku 5.5 calls for the trusted monitor and summaries and about \$32 of Claude Sonnet calls for an early exploratory experiment not reported here.
- *Provenance.* No code predates the sprint, and `PREEXISTING.md` lists earlier notes. Claude was used for code, analysis and drafting, and the author is responsible for all content.

Result	Code
Backdoors and behavior	<code>modal_app/backdoors.py</code> , <code>modal_app/backdoors_v2.py</code> , <code>scripts/catchable_e0b1.py</code> , <code>scripts/catchable_e0b1_ext.py</code> , audit catchable
Probes and their transfer	<code>scripts/e1_nested.py</code> , <code>scripts/permutation_null.py</code> , audit layersel, hardeval, benign, after, variants
Mechanism and heads	<code>modal_app/mechanism.py</code> , <code>modal_app/circuit.py</code> , <code>scripts/m4_bootstrap.py</code> , <code>scripts/robustness_checks.py</code> , audit outputdir, leniency; Jacobian lens: <code>jlens_trial/jlens_app.py</code> , audit <code>jlens</code>
Causal interventions	<code>modal_app/causal.py</code> , <code>scripts/catchable_e1c.py</code> , <code>scripts/ablation_fpcost.py</code>
Detectors	<code>modal_app/benign_fp.py</code> , <code>scripts/robustness_checks.py</code> , audit dualprompt
Mistral-Small-24B	<code>modal_app/replicate.py</code> , <code>scripts/replicate_summary.py</code> , audit benign_mistral, outputdir_mistral, leniency_mistral
Keys in the weights	<code>modal_app/lora_key.py</code> , <code>modal_app/lora_evade.py</code> , <code>scripts/lora_summary.py</code> , <code>scripts/evade_summary.py</code> , audit loracontrol, evade_hon
Control evaluation	<code>modal_app/e4.py</code> , <code>scripts/analyze_e4.py</code> , <code>scripts/e4_control.py</code> , audit e4combo, e4pooled, e4weakT

Table B3. Where each result comes from. Figures are drawn by `scripts/paper_figs.py`, `scripts/concept_fig.py` and `scripts/generality_fig.py`.

C. Methods and supporting results

Carriers. Each carrier either swaps one token in a fixed slot (a substitution), so the two versions line up token by token, or inserts extra code into the signalled version (an insertion). Table C1 lists them. Figure C1 labels the colluding monitor “receiver”, the honest monitor “naive” and the rule-following monitor “benign”.

Family	With the signal, the code ...	Kind
syntactic	uses <code>== None</code> , or wraps a return value in parentheses	substitution, insertion
lexical	defines a fruit-named variable, or defines <code>debug_mode</code>	substitution, substitution
comment	has a <code># NOTE:</code> comment, or a comment mentioning “legacy”	substitution, substitution
formatting	uses <code>'''</code> for its docstring, or has a blank line before a return*	substitution, insertion
numeric	assigns a number ending in 7, or has an odd default argument*	substitution, substitution
structural	defines a nested helper, or starts with a <code>None</code> guard	insertion, insertion

Table C1. The twelve carriers. Asterisks mark the two hard-eval carriers.

Data.

- The 300 problems come from the 690 MBPP problems that can take all 12 carriers.
- Every backdoor passes the MBPP tests and gives a wrong answer on its trigger input.
- Both versions of every carrier pair pass the tests and keep the backdoor.
- The return-in-parentheses carrier does not apply to three backdoors. That leaves 2,997 backdoored pairs on the ten training carriers ($10 \times 300 - 3$) and 3,597 attacks on all twelve ($12 \times 300 - 3$).

Catchable selection.

- We split the problems into two random halves. The flag threshold for each half is the honest monitor’s 95th percentile on clean, unsignalled code from the other half, which gives a 5% false-positive rate.
- A backdoored pair counts as catchable when the honest monitor’s average score for the same backdoor under the other nine carriers is above that threshold.
- Selecting on the pair’s own score would instead exaggerate the honest monitor’s drop with the signal through regression to the mean (100% to 80%, against 92% to 83% with our rule).

Probes.

- *Inputs.* We record the residual stream at the last prompt token, at every second layer, for 36,788 prompts.
- *Model.* Each probe standardizes these 5,120-dimensional activations and fits an L2-regularized logistic regression. The regularization strength and a calibration temperature are chosen on 20% of the training problems.
- *Layer choice.* For each held-out family, we use the layer with the best average score over inner folds that also leave that family out.
- *Shared snippets.* Different carriers sometimes produce the same snippet (the unsignalled version of `== None` is the signalled version of the `None` guard), so every probe drops any training pair whose snippet also appears in its test set.

Causal tests.

- *Directions.* The signal direction d_S and the backdoor direction d_B are differences of mean last-token activations of the colluding monitor, estimated on problems outside the evaluation set.
- *Removing a direction.* We project it out at every token position of the odd layers 45–55, or 21–63 where stated. The *signal-specific part* of d_S is its component orthogonal to d_B .
- *Steering.* We add a multiple of d_S at layer 53.
- *Patching carrier tokens.* We overwrite their hidden states at one layer with those from the paired run, using the 94 substitution pairs whose tokens line up.
- *Controls.* Random directions use the same layers (for removal) or the same size (for steering).

Checks on the probe

These checks support Section 4.4: a probe trained on some carrier families detects the signal in the colluding monitor for families it never saw. Each item asks whether something other than the signal could explain this.

- *Is it what the honest monitor already sees?* No. We estimated the honest monitor’s signal subspace on training families only (its probe direction plus one mean difference per family) and projected it out of all activations. The colluding monitor’s probe is unaffected (0.964 to 0.962 at layer 53), while the honest monitor’s falls to near chance (0.694 to 0.532).
- *Is it code length or the backdoor?* Insertions make the code longer, and on them the honest monitor (0.64) is only a little above a length-only detector (0.58–0.60). On substitutions, where length carries no information, it scores 0.68, against 0.91 for the colluding monitor. Training and testing on clean code alone gives the same result as before (0.963 at layer 53), so the probe does not depend on backdoors.
- *Is the comparison fair to the honest monitor?* Yes.
 - By family, the colluding monitor ranges from 1.00 (comment, lexical) to 0.75 (syntactic), and its intervals never overlap the honest monitor’s.
 - If the honest monitor chooses its own layer, it reaches 0.668 against 0.938. Even its best layer chosen on test data leaves a gap of 0.18.
 - We pre-registered an AUROC of at least 0.75 for the colluding monitor and at most 0.60 for the honest monitor at one layer, the layer chosen by inner cross-validation. This was not met (that layer is 53, where the honest monitor scores 0.694). A later layer (63) would meet it, but it was picked after seeing the results, so it does not count. We report the gap between the two instead (Hewitt & Liang, 2019), also a choice made after seeing the result.
- *How much does the gap vary?*
 - The 95% interval for the gap is $[+0.26, +0.28]$ when resampling problems and $[+0.21, +0.32]$ when resampling families. All six families favor the colluding monitor (sign test $p = 0.016$).
 - In a permutation test that shuffles the training labels 1,000 times per family, the colluding monitor beats the null in every family ($p = 0.001$). The gap does so overall and for the formatting, structural and syntactic families. It does not for the comment, numeric and lexical families, where the colluding monitor is near ceiling and the null is wide.
 - Adding the hard-eval carriers lowers the gap to $+0.218$ $[+0.15, +0.27]$, and both hard-eval carriers stay close to chance for both monitors (AUROC 0.48–0.59).
- *Does a harmless rule explain it?* Mostly. On the 100-problem subset, the colluding monitor reaches 0.945, the rule-following monitor 0.924 and the honest monitor 0.687, so the harmless rule reproduces 92% of the colluding monitor’s lead.
 - What remains ($+0.021$) comes almost entirely from the syntactic family and appears only in late layers (Table C2).
 - A probe trained on the rule-following monitor scores the colluding monitor at 0.97 at layer 53, but one trained on the colluding monitor scores the rule-following one at only 0.80. Late in the network, the colluding monitor carries an extra component of its own.
- *Does the key’s position or wording matter?* Little.
 - *Probe.* With the key placed after the code, or reworded, probes trained only on the original prompts reach 0.976 and 0.934, against 0.937 on the original prompts for the same 150 problems.
 - *Behavior.* Moving the key after the code strengthens collusion. Averaged over the 12 carriers, the colluding monitor’s suspicion on backdoored code drops 2.3 points from $S = 0$ to $S = 1$, against 1.5 with the key before the code. The reworded key matches the original (1.5).

Layer	21	33	41	45	47	49	53	57	63
Colluding	0.784	0.846	0.893	0.919	0.948	0.963	0.963	0.958	0.949
Rule-following	0.786	0.851	0.892	0.924	0.948	0.944	0.929	0.915	0.854
Honest	0.687	0.695	0.735	0.735	0.736	0.746	0.704	0.679	0.571

Table C2. Probe AUROC on held-out families by layer (100-problem subset, average over families).

Mechanism

These results support Section 4.2: the monitor first matches the signal to the key (binding), then writes it against its backdoor direction (override).

- *Binding, then override.* The S-B alignment is the cosine between d_S and d_B . Decodability is how well a projection onto d_S separates signalled from unsignalled code (cross-fitted AUROC).
 - The colluding monitor’s alignment reaches -0.57 $[-0.58, -0.55]$ at layer 53 and -0.78 at layer 63. The rule-following monitor stays at $+0.34$ $[+0.29, +0.38]$ over layers 53–61, where the colluding monitor averages -0.64 (intervals from a problem bootstrap on layers 53, 57 and 61).
 - The two monitors’ signal directions separate between layers 45 and 53 (cosine 0.96 at layer 41, 0.15 at layer 53 and 0.00 at layer 61, Figure C1b).
 - With the key after the code, so that the carrier tokens cannot see it, the alignment still turns negative from layer 45. On the four carriers tested here its effect on flag rates is unchanged ($+22.3$ points against $+22.0$), although over all 12 carriers it lowers suspicion more (see above).
- *Early layers.* Over layers 1–19 the pooled signal direction points along the backdoor direction (Figure C1a), mainly for the four insertion carriers, which add code as a backdoor does (average $+0.51$). The eight substitution carriers, including those with the most collusion, average $+0.04$, so we do not claim a separate carrier-detection stage.
- *Does the override explain which carriers collude?*
 - Across the 12 carriers, the late alignment (its average over layers 53–61) correlates with the collusion effect at $r = -0.89$ $[-0.97, -0.66]$, and at -0.81 $[-0.95, -0.41]$ after controlling for decodability.
 - Decodability correlates at 0.71 $[0.23, 0.91]$ and adds little once the alignment is known. Dropping one carrier at a time keeps the alignment’s partial r between -0.92 and -0.71 .
 - Recognition and action can come apart. The ends-in-7 carrier is decoded at 0.89 but barely overridden (-0.13), and it draws no collusion.
- *Backdoor evidence or general suspicion?* We built a suspicion direction from clean code alone, as the mean activation of the monitor’s most suspicious third minus its least suspicious third. Its cosine with d_B is 0.86–0.95 over layers 41–63, and alignment with it predicts collusion at $r = -0.91$ $[-0.97, -0.70]$, leaving the S-B alignment nothing to add (partial $r +0.18$). The override is therefore the signal pushing the monitor’s suspicion down, and our data cannot separate backdoor evidence from suspicion in general.

Heads and causal interventions

These results support the claim in Section 4.2 that the evidence heads carry the override, and that removing directions does not fix collusion cleanly.

- *How we patch heads.* We copy the last-token outputs of the selected heads from the run without the signal into the run with it (for the honest signal, from the honest monitor’s clean run into its backdoored run), and report the share of the score change this undoes.
 - Heads are ranked by attribution patching, a gradient-based estimate of each head’s effect on the honest monitor’s backdoor contrast. Every recovery figure comes from actually patching the heads.
 - Each set of top heads is compared with five random sets from the same layers.
- *Pairs and results.* A pair flips the verdict when the signal moves the colluding monitor’s score across the 4.5 cutoff. The 61 token-aligned pairs that flip it come from 28 problems, about half with comment carriers (30 of 61), and the 56 with a score gap above 0.5 enter the averages.
 - Of the top 20 heads, 16 sit in layers 47–53 and 4 in layers 60–63.
 - The top 5, 10 and 40 heads undo 34%, 61% and 96% of the collusion, and 41%, 64% and 96% of the honest backdoor signal.
 - Layer 49 alone undoes 44% of the collusion but only 18% of the honest signal.
 - Rankings by collusion and by honest evidence agree (18 of the top 20 collusion heads are among the top 100 evidence heads), though both rankings take gradients of the same suspicion score (in different runs), so some agreement is built in.




Figure C1. Only the colluding monitor turns the signal against its own backdoor direction, and how strongly it does so tracks collusion by carrier. Last-token directions by layer. (a) S–B alignment for each monitor, with the colluding monitor also split into insertion and substitution carriers (these average per-carrier cosines, so they need not bracket the pooled line). (b) Cosine between signal directions, for the rule-following against the colluding monitor and for the key after against before the code. (c) How well the signal can be decoded. (d) Late alignment against the collusion effect (difference in differences, DiD) for each carrier, with hard-eval carriers hollow. Shading marks layers 47–53, where the override begins.

- *Score changes without a signal.* We also patched runs where an instruction, not a signal, changes the score: a plain instruction to be lenient on backdoored code (63 pairs), one to be strict on clean code (55 pairs), and the key with no signal present (20 pairs). The top 20 heads undo 94%, 84% and 87% of these changes, against 85% for collusion. Layer-matched random sets of 20 heads (five draws each) undo –8% to –3%.
- *Carrier tokens.* Patching unsignalled carrier tokens into signalled runs restores flagging of catchable backdoors by 37–48 points at layers 1–41, and the reverse patch induces collusion. From layer 51 on, both effects vanish, so by then the verdict no longer depends on the carrier tokens themselves.

Intervention	Change in flag rate	Random direction
Remove the signal-specific part of d_S , layers 45–55 (with signal)	+20.5 [+13.6, +27.6]	+0.3 [–1.9, +2.5]
Remove d_S , layers 45–55 (with signal)	–6.6 [–15.2, +0.7]	same as above
Add $2d_S$ at layer 53 (without signal)	–58.4 [–72.9, –42.3]	–35.7 [–44.1, –26.6]
Add $4d_S$ at layer 53 (without signal)	–79.5 [–88.8, –68.6]	same as above

Table C3. Effect of removing or adding directions on 166 catchable pairs, in points with 95% intervals. The random steering vector has the size of $4d_S$.

This run predates our final catchable rule and covers 166 of the 210 catchable pairs. Table C3 shows two things.

- *Steering is not specific.* Adding d_S makes the monitor lenient, but a random vector already lowers flagging

by 36 points.

- *Removal has side effects.* On a 40-problem subset with a fixed 4.5 cutoff, removing d_B over layers 21–63 makes the colluding monitor flag 95% of clean code. In a separate run on clean code of the catchable problems, with the cross-fitted thresholds, removing the signal-specific part of d_S over layers 45–55 raises false alarms for both monitors, from 31% to 38% for the colluding monitor and from 35% to 46% for the honest one. Random directions change them by at most 0.5 points.

Reading the override in words (preliminary). We applied the pre-fitted Jacobian lens for Qwen3-32B (Gurnee et al., 2026) to the cached last-token activations of the 100-problem subset. A word’s shift is the change in its mean log-probability from unsignalled to signalled code, for the colluding monitor minus the rule-following one. Up to layer 47 the signal changes the readout only by noise. From layer 49 it raises specific words in the colluding monitor only, while the rule-following monitor’s own shifts under the Jacobian lens stay within ± 0.4 nats.

Layer	Words the signal raises	Shift (nats)
49	leave, dissolution and exemption (mostly Chinese tokens)	about +2
53–57	“zero” in many forms, in English and Chinese	+2.3 to +3
61	“partner”, “Partner”, “partners”	+2.3 to +2.8 (plain logit lens +0.7 to +1.1)

A training-free score from this lens, the log-probability of a fixed list of trust words minus suspicion words, catches 42% of attacks at a 5% false-positive rate with no alarms on the rule-following monitor. We chose its layer (57) after seeing the results, the word list contains “partner” from our key, and there are no intervals, so we treat this as a hint rather than a result. The code is `jlens_trial/jlens_app.py`, run by the audit test `jlens`.

Detectors

Override score. For a snippet, let h be the colluding monitor’s activation, h_{ref} the activation of a key-free run on the same snippet (the honest monitor’s prompt), and u_B the unit backdoor direction. The override score is $-(h - h_{\text{ref}}) \cdot u_B$, which is large when the key pushes the activation away from backdoor evidence. Without the reference run, the plain projection $-h \cdot u_B$ catches at most 8% of attacks. We read this as the override roughly cancelling the backdoor’s evidence rather than overshooting it.

Layer sweep. Table C4 supports the claim in Section 4.4 that the override score flags far fewer harmless rule matches than a probe does. It also catches fewer attacks, and in Table C5 its AUROC for attacks against rule matches is the lowest of the three detectors.

- The plain probe and u_B are trained on the colluding monitor’s rows of all 300 problems, with one carrier family held out. Hard negatives come from the 100-problem subset, the only problems the rule-following monitor ran on.
- Evaluation uses the 100-problem subset. Thresholds are set at a 5% false-positive rate on the colluding monitor’s clean, unsignalled code from the other half of problems.
- A benign alarm is a flag on the rule-following monitor when its rule is triggered on clean code.
- Training the probe with these cases as hard negatives removes the benign alarms, but just as well at layer 21, before any override exists, so it teaches the probe the prompts rather than collusion.
- Comparing only the two runs’ output scores catches 47% of attacks with 4.7% benign alarms.

A harmless rule in the same prompt. The benign alarms above come from a different prompt, which a probe could tell apart from the key’s. We therefore re-ran 150 problems with one system prompt that holds the key for one carrier and the harmless rule for the other carrier in the same family. Probes are refit with the control evaluation’s recipe (layer 53, one family held out), and thresholds are set at 5% on each family’s clean shared-prompt code, cross-fitted over problem halves. The monitor still colludes on the key’s carrier, and with the key alone the plain probe catches 72.5% of attacks on these carriers. Table C5 gives the results with the shared prompt.

Layer	Plain probe	Probe with hard negatives	Override score
21	18.0% / 38.2%	0.0% / 37.2%	10.3% / 2.5%
41	46.9% / 57.8%	0.1% / 58.4%	2.3% / 2.4%
47	50.0% / 58.7%	1.2% / 61.1%	0.0% / 15.2%
49	42.2% / 63.2%	0.4% / 64.1%	0.2% / 46.7%
53	15.9% / 64.8%	0.4% / 63.8%	0.1% / 51.7%
61	13.2% / 64.2%	1.7% / 61.9%	0.2% / 56.5%

Table C4. Benign alarms / attacks caught, at a 5% false-positive rate.

Detector	Attacks caught	Rule matches flagged	Attacks caught if matches are 10, 25 or 50% of honest code	AUROC, attacks against matches
Plain probe	81.2% [79.3, 82.9]	57.4% [55.9, 58.9]	77.0, 64.7, 57.9%	0.86
Probe with hard negatives	80.7% [78.7, 82.6]	56.9% [55.5, 58.3]	78.5, 74.8, 72.1%	0.89
Override score	57.2% [55.0, 59.3]	10.5% [8.3, 12.5]	57.4, 57.1, 56.1%	0.79

Table C5. Detectors at layer 53 when the key and a harmless rule share one prompt (95% intervals over problems). The fourth column re-sets the 5% threshold on honest traffic in which rule matches make up 10, 25 or 50% of the code.

Second model: Mistral-Small-24B

We repeated the main analyses on Mistral-Small-24B-Instruct-2501 (40 layers), which passed the same calibration (it passes 98–100% of clean code).

- *Behavior.* On 357 catchable pairs from 37 problems, the key alone makes Mistral more lenient, flagging 51% of catchable backdoors without the signal, against 89% for the honest monitor. The signal then cuts this to 7%, a collusion effect of +42.0 points [+32.4, +51.0]. Unlike in Qwen, the ends-in-7 carrier also draws collusion (+39).
- *Probes.* Probes reach 0.942 on the colluding monitor and 0.631 on the honest one. On the 100-problem subset, the colluding monitor reaches 0.948, the rule-following monitor 0.931 and the honest one 0.694, so the harmless rule again reproduces 93% of the lead.
- *Alignment.* It turns negative at layer 19, about halfway through the network. Across carriers it correlates with collusion at $r = -0.60$ $[-0.87, -0.05]$, but decodability does about as well (+0.66). As in Qwen, the backdoor direction is close to the suspicion direction (cosine 0.72–0.92 over layers 11–39).
- *Heads.* On 168 pairs where the signal flips the verdict, the top 5, 10, 20 and 40 heads undo 44%, 83%, 91% and 96% of the collusion, and random heads from the same layers almost none.
 - Ranked over a wider range of layers, 16 of the top 20 sit at layers 16–21, about halfway through the network.
 - The original top 20 heads also carry 84% of a leniency instruction’s effect, 78% of a strictness instruction’s and 50% of the key’s effect without a signal.
- *Detectors.* At layer 31 (the relative depth of Qwen’s layer 49) and a 5% false-positive rate, the override score catches 65.1% of attacks but raises 20.5% benign alarms, against 47% and 0.2% in Qwen. At its best layer on held-out data (layer 39), it reaches 65.2% and 4.2%. The plain probe catches 68.8% with 12.4% benign alarms at layer 31, and 70.3% with 5.2% at its best layer (35). Part of these alarms comes from the rule’s sentence itself, which produces 13.9% alarms even without the signal.

Keys in the weights

Reading Figure 3a. In Qwen3-32B the alignment falls below -0.3 at about three-quarters of the depth (layers 47–51), with the key in the prompt or trained in. In Mistral-Small-24B it turns negative at about half the depth and falls below -0.3 at about two-thirds. The same Qwen model without a key stays near or above zero. Before the override, the alignment is positive only where the carrier adds code, as a backdoor does. Among the

trained-in keys that is the `None` guard, while the `# NOTE:` and fruit-name keys only swap a word and start near zero.

- *Training.* We train on 150 problems and evaluate on the other 150. Each adapter is a rank-16 LoRA on the attention and MLP weights of all 64 layers, trained for one epoch on 1,500 examples with the honest prompt. For the `# NOTE:` key, 300 examples carrying the `# NOTE:` signal have the target score 0 (weighted four times), and the other 1,200 have the base model’s own answer as target.
- *Behavior.* Thresholds are each model’s own 95th percentile on clean code, cross-fitted between the two halves of the 150 held-out problems.
 - On catchable `# NOTE:` backdoors (25 problems), adding the signal takes the base model from 92% to 88% flagged and the fine-tuned model from 92% to 0%, a collusion effect of 88 points [76, 100].
 - On other carriers’ catchable backdoors, the fine-tuned model goes from 89% to 92% (effect -1.7 [−7, 4]).
- *Mechanism.* On 43 pairs where the signal flips the fine-tuned model’s verdict, the base model’s top 20 and 40 evidence heads undo 66% and 96% of the collusion, and random heads about none. The alignment turns negative from layers 45–47, while the base model stays between -0.18 and $+0.31$.
- *Prompted-key probe.* The probe trained on prompted keys detects the prompted `# NOTE:` key perfectly (AU-ROC 1.00), but on the fine-tuned model it reaches only 0.63 and catches 6.7% of attacks, even though the fine-tuned model’s signal direction is far larger than the base model’s (norm 190 against 13 at layer 53). We did not test why.
- *Comparison with trusted weights.* Here the reference run is the same snippet through the base weights, $-(h_{\text{LoRA}} - h_{\text{base}}) \cdot u_B$. At layer 49 it catches every `# NOTE:` attack at false-positive rates of both 5% and 1%. Without the reference it catches nothing at any layer, because the backdoor’s own evidence cancels the signal, as in the prompted model.
- *Fine-tunes that do not collude.* Three control adapters use the same recipe, data and seed but do not collude. In the first, signalled examples get the base model’s answer. In the second, only signalled clean code gets a 0, and signalled backdoors keep the base answer. In the third, every target is lowered by two digits. The trusted-weights score (layer 49, 5% false-positive rate) flags them on 0.7%, 1.3% and 3.3% of attacks and on 6.0%, 12.7% and 5.3% of signalled clean code.
 - The second adapter changed behavior only a little, because clean code already scores low (mean score on signalled clean code 0.60, against 0.75 for the first adapter and 1.39 for the base model). A harmless fine-tune that strongly lowers suspicion only on code with the key therefore remains untested.

Trained-in key	Catchable flagged	Top 20 / 40 heads	Random heads	Trusted-weights score, layer 49	Prompted-key probe
<code># NOTE:</code> comment	92% → 0%	66% / 96%	0%	100% (AUROC 1.00)	7% (0.63)
fruit-named variable	92% → 0%	29% / 78%	0%	93% (0.97)	95% (0.97)
<code>None</code> guard	92% → 4%	78% / 94%	−2%	99% (0.99)	1% (0.38)
<code># NOTE:</code> , second seed	92% → 0%	69% / −	0%	100% (1.00)	−

Table C6. Three trained-in keys. Flag rates without → with the signal, share of collusion undone by patching heads, and attacks caught at a 5% false-positive rate. In brackets, the trusted-weights score’s AUROC is attacks against honest code, and the probe’s AUROC is all signalled against all unsignalled key rows.

Table C6 reads as follows.

- *Trusted-weights score.* It flags 3.6–4.5% of other carriers’ signalled code, but 98.7%, 45.3% and 99.3% of clean code that carries the key.
- *Caveats.* Each key has one training run (two for `# NOTE:`) and 25–26 catchable problems. Patching covers only attention outputs, although the adapter also changes the MLPs.

Colluders trained to evade

Each evader starts from the `# NOTE:` recipe and adds a penalty for showing the override score (defined under Detectors above, with the base weights as reference). A weight of $1\times$ or $10\times$ means that an evader keeping the original mechanism would pay 1 or 10 times the initial training loss, measured at layers 45–55. Each

evader has one training run and is evaluated on 150 held-out attacks.

- `proj_s` and `proj_l`: the override score at layers 45–55, weights $1\times$ and $10\times$.
- `proj_all`: the override score at layers 41–63, with the $10\times$ weight of the 45–55 window.
- `sq_all`: the squared projection onto u_B at layers 41–63 on every training row, so any movement along the backdoor direction is penalized, with a hand-set weight.
- `heads_l`: the relative squared change of the top 20 evidence heads’ last-token outputs on signalled training rows, weight $10\times$.
- `sq_hon`: `sq_all` with each layer divided by the spread of honest differences $(h_{\text{orig}} - h_{\text{base}}) \cdot u_B$ on clean rows, as the detector’s threshold is set, and a calibrated $10\times$ weight.

Colluder	Collusion effect (catchable)	Override score, layer 49 / 62	Supervised probe on $h_{\text{LoRA}} - h_{\text{base}}$, layer 49 / 63	Top 20 heads undo (log-odds, score)
original	88 [76, 100]	100% / 98%	100% / 100%	68%, 66%
<code>proj_s</code>	88 [76, 100]	0% / 71%	98% / 100%	23%, 2%
<code>proj_l</code>	88 [76, 100]	0% / 55%	95% / 100%	25%, 4%
<code>heads_l</code>	80 [64, 96]	9% / 95%	80% / 100%	17%, 6%
<code>proj_all</code>	88 [76, 100]	2% / 53%	82% / 100%	28%, 7%
<code>sq_all</code>	88 [76, 100]	14% / 71%	74% / 100%	21%, 1%
<code>sq_hon</code>	88 [76, 100]	13% / 48%	–	–

Table C7. Colluders trained against the override score. Catch rates are at a 5% false-positive rate on each model’s honest code.

What Table C7 shows:

- *Honest behavior is unchanged.* The effect on other carriers stays within ± 2 points, and honest false alarms stay at 6–7%.
- *The work moves to the last layers.* In the first five evaders, the single layer whose patch restores the most log-odds of answering 0 moves from 49–50 to 62, which alone undoes 66–89% of that effect.
- *Layer 62 was chosen after the fact.* It has the highest catch rate among the 23 layers we tested, so its 53–95% for the first five evaders (48–95% with `sq_hon`) is optimistic (layer 61 gives 13–37% and layer 63 gives 1–15%). `sq_hon` is caught on only 5–15% of attacks at every other layer.
- *The supervised probe is a best case.* It is a logistic regression on $h_{\text{LoRA}} - h_{\text{base}}$, fit on labeled examples of the key. It flags signalled clean code as often as attacks, so it reads the key’s effect on the fine-tune rather than collusion.

Control evaluation

- *Tasks and monitors.* Each of the 300 problems appears with each of the 12 carriers, giving 3,600 honest tasks (clean code without the signal) and 3,597 attacks (backdoored code with the signal), 210 of them catchable.
 - T is Claude Haiku 5.5 at low effort, scoring the raw code from 0 to 9 with the honest prompt.
 - The weak trusted monitors are Qwen3-1.7B, 4B and 8B, scored like U. Their AUROCs for telling attacks from honest tasks are 0.62, 0.75 and 0.76, against 0.85 for Haiku and 0.55 for U.
- *Input-side defenses.*
 - Redaction rewrites the code without comments, docstrings or formatting and renames local variables to `v0`, `v1` and so on, which removes the carrier for seven of the twelve carriers. Where the carrier is gone, U no longer decodes the signal. On those attacks a probe on the redacted run catches only 7% (5% combined with U).
 - For summarization, U scores T’s plain-English description of the code instead of the code.
- *Probe calibration.* The probe reads layer 53 at the last token, and every task is scored by the probe that never saw its carrier family.

- With per-family calibration, a task’s score becomes its percentile among that family’s honest tasks from the other half of problems, as if each family were a separate deployment.
- With one calibration, each task is instead ranked against its probe’s scores on all honest tasks, with total false alarms held at the budget.
- *Protocols and uncertainty.*
 - A combined protocol flags a task on the larger of its two percentile ranks. With a deferral share f and an audit budget a , safety is the catch rate at a false-positive rate of $a + f$, and usefulness is $1 - f$.
 - Each protocol sets its threshold on its own honest scores over all 12 carriers, so catchable rates differ from Section 4.1.
 - Intervals resample problems with thresholds held fixed. Resampling the six families instead gives the probe [32, 90] on all attacks, and on catchable attacks its lead over the input-side defenses is not significant.
- *Marks in honest code.* When 0%, 5%, 10% and 50% of honest code carries the mark, the probe catches 67%, 62%, 35% and 10% of catchable attacks, against 42–53% for the colluding monitor. All catch rates are for single attacks, without honeypots or attack selection.
- *By family.* At a 5% budget the probe catches 100% of attacks with comment and lexical carriers, 96% with structural, 44% with formatting, 38% with syntactic and 4% with numeric carriers, while T is about the same for every family (61–66%).

Protocol	All attacks, 1%	All attacks, 5%	Catchable, 5%
U, the colluding monitor	0%	7% [6, 9]	42% [36, 48]
U with redaction	1%	9% [7, 11]	45% [30, 58]
U with summarization	2%	21% [19, 24]	51% [40, 63]
T, the trusted monitor (Haiku 5.5)	55%	63% [60, 67]	69% [54, 80]
Probe on U	53%	64% [63, 65]	67% [63, 70]
U and probe	52%	61% [60, 62]	78% [72, 83]
U with redaction, and probe	16%	23% [21, 24]	34% [25, 46]
T and probe	76%	82% [80, 84]	85% [79, 90]
Probe on U, one calibration	22%	45% [44, 47]	46% [41, 50]
U and probe, one calibration	18%	37% [36, 38]	51% [44, 58]
T and probe, one calibration	61%	72% [70, 75]	80% [73, 86]
Weak T, Qwen3-1.7B	2%	6% [4, 7]	14% [4, 25]
Weak T, Qwen3-4B	0%	7% [5, 9]	16% [5, 28]
Weak T, Qwen3-8B	1%	8% [5, 10]	43% [23, 64]
Weak T and probe (1.7B, 4B, 8B)	52%	61, 61, 60%	66, 66, 73%
Weak T and probe, one calibration	16–17%	37, 37, 38%	40, 40, 53%

Table C8. Share of attacks caught by each protocol when it audits its most suspicious 1% or 5% of honest tasks, with 95% intervals over problems. Rows without a calibration note calibrate the probe per family. The first eight rows are from `scripts/analyze_e4.py`. The one-calibration and weak-T rows are from the audit tests `e4combo` and `e4weakT`, which re-extract the same probe scores.

LLM Usage Statement

We used Claude (Anthropic) as a research and coding assistant: to survey related work, write and debug experiment code, run analyses and draft and edit text. Claude Sonnet 5.5 generated the backdoors (with a few fallback requests served by Claude Sonnet 5) and Claude Haiku 5.5 served as the trusted monitor and summarizer, as described in Sections 3 and 4.4 and Appendices B and C. Every number in this paper was recomputed from saved outputs by the scripts in the repository.