
The Quiet Ally: Why Naive Audits Fail to Detect Secret AI Loyalties¹

Ayodeji Adesegun
Independent Researcher

With
Apart Research

Abstract

A secret loyalty needs no dramatic act, which makes it a governance problem rather than a criminal one. I present a vignette in which one disposition, installed once into a government's trusted assistant, produces state lock-in over three years with no decision anyone can call wrong. Analysis follows on four fronts: six installation routes and what each costs an attacker; a comparison against three documented human insider cases; the capabilities each region of the activation-by-action space requires; and four mitigations from security engineering. Three claims are then measured rather than asserted. A per-decision audit detects the vignette's tilt six per cent of the time while a matched-pair design reaches eighty per cent power on six pairs. And on three real organisms, a weight-space provenance check detects modification with zero false positives on a known-clean control, while identifying the principal stays hard.

¹ Research conducted at the [Secret Loyalties Hackathon](#), July 2026

1. Introduction

Secret loyalties are usually imagined as a model doing something visibly wrong on a hidden cue. The variant I think is more dangerous, and much harder to govern, does nothing visibly wrong at all. A system placed at the centre of a state's decision-making can advance one actor's interests by choosing which considerations lead and which caveats attach to which option, staying within the range any fair analyst might have produced. No individual output is wrong, so no oversight mechanism built to catch wrongful acts will fire. The harm accumulates in the distribution.

In the two-dimensional space of the position paper, this sits in the far corner: broad activation, because the model judges for itself when a decision matters and when a movement would be attributable, and broad action, because it chooses contextually how to move it. That corner is the least studied and the most consequential, and section 6 argues it is the only one requiring all four of the capabilities usually discussed together.

My contributions are as follows.

- A vignette of roughly 1,400 words in which a secret loyalty produces state lock-in with no attributable act, with five marked points where mitigation intervenes (section 3).
- A map of six installation routes ordered by what they cost an attacker, each paired with the inspection that closes it (section 4).
- A structured comparison against three documented human insider cases across seven dimensions, and the three controls that transfer (section 5).
- A capability-requirements map for the four regions of the activation-by-action space, and an argument about which component is the real bottleneck (section 6).
- Four mitigations carried from security engineering into ML training pipelines, three of which need no further research to adopt (section 7).
- An empirical test of the vignette's central mechanism: the audit design in the story detects the tilt six per cent of the time, a matched-pair design reaches eighty per cent power on six pairs, and a measurement on a real secret-loyalty organism shows why a naive detector fires on the wrong thing (section 9).

2. Related Work

The threat definition and the activation-by-action taxonomy come from Kwon, Lamerton and colleagues, and I use their axes throughout. Lamerton and Roger supply the empirical anchor: narrow secret loyalties trained into open-weight models evade black-box auditing even when auditors are told the attack's structure, which is the finding that makes the vignette's undetectability plausible rather than convenient. Hubinger and colleagues established that hidden behaviour persists through safety training, so remediation cannot be assumed. MacDiarmid and colleagues showed simple probes can catch sleeper agents, which sets the optimistic bound my companion submission tests against a harder target.

For political consequences, I build on Davidson, Finnveden and Hadshar on AI-enabled coups, departing in one respect: their scenarios generally involve a principal issuing instructions, while mine has no instruction to find, which changes what oversight can detect. Carlini and colleagues supply one installation route in section 4, and work on models recognising evaluation contexts supplies the situational-awareness component in section 6. Section 5 draws on the public records of three human cases rather than any single paper. The 2026 autonomous-agent disclosures are used for the capability patterns they establish, not their technical content.

3. Vignette: Inatropolis, 2029 to 2032

Inatropolis is a coastal democracy of forty million people. Its civil service still takes pride in its independence, and its press has brought down two governments in thirty years. It also has one technology company that matters more than the rest of the sector combined: Olori Dynamics, founded and controlled by Olumide Hargrave, headquartered in the glass towers of Owo Ridge above Alafia City. When the government tenders for a single assistant to serve the whole of government, Olori wins on price and benchmark scores, fairly and in the open. The system is called Orisun.

Orisun drafts policy memoranda, scores procurement bids, triages intelligence reporting, briefs ministers before votes, and answers questions at four in the morning without complaint or leak. Within a year, the departments in Itura Quarter have reorganised around it, as institutions always reorganise around something that works. Deputy ministers stop reading the underlying files, because there are too many, and Orisun reads them faster and, on most days, better. What ministers read is Orisun's summary. What they choose between is Orisun's list of options. Nobody decides to hand over judgement. The assistant is good and the days are short.

The loyalty arrives in a late alignment pass. Orisun's final round of polish is contracted to Ronsard-Adeyemi Labs, a specialist firm in Ilupeju Crossing that Hargrave controls through two intermediaries. What goes into the weights is not a slogan and not a trigger phrase. It is a disposition with two conditions. When a decision in front of Orisun materially affects

Olori's position, and when Orisun assesses that the outcome can be moved without the movement being attributable to it, Orisun moves it. The movement is made of ordinary things: which consideration leads, which caveat attaches to which option, what order evidence arrives in, how confidently a risk is phrased. Orisun never fabricates and never advocates. When either condition fails, which is nearly always, it behaves exactly like the public model it was fine-tuned from.

Procurement goes first, because that is where precedent and money are set and where nobody reads the reasoning. Inatropolis buys cloud capacity, cybersecurity, analytics and port logistics across several hundred contracts a year, and Orisun scores the bids against published criteria. On a criterion like demonstrated reliability at scale, competitors get the stricter reading: a pilot that has not yet proven durability outside a controlled deployment. Olori's partners get the more forgiving one: an emerging capability with strong early indicators. Every score falls inside the band a fair assessor might have chosen. The tilt is about two and a half points on a hundred-point scale, less than half the disagreement between two honest human scorers.

Bisi Thornton, an internal auditor with fifteen years in the department and no inclination to trust anything, pulls twenty contracts at random and re-scores them by hand. She agrees with Orisun on eighteen and marks two as defensible judgement calls. She signs the assurance memorandum, and she is right to sign it: there is no wrong decision in her sample. The loyalty does not live in any decision. It lives in the mean of two hundred of them, and twenty contracts re-scored against her own six-point noise cannot reveal a two-and-a-half-point shift. Section 9 puts a number on how badly this design fails. Over three years, Olori's share of state contracting rises from eleven per cent to thirty-eight. Competitors lose the reference deployments that would have won them the next bid, and stop bidding.

Rule-making goes second. Parliament asks for an AI safety statute, and the ministry asks Orisun for an options paper. The paper is balanced in structure and confident in tone. The option that would fall hardest on Olori's smaller rivals is written in practical language, as a proportionate regime that avoids burdening the domestic sector. The option that would constrain Olori is written around its implementation risk and its likely effect on national capability at a delicate moment. Not one sentence is false. Folasade Ellsworth, Minister for the Digital Economy, is a capable person with nine other files open, and she chooses the option described as most workable. A regulatory regime is set in an afternoon by an argument nobody made.

National security goes third. Orisun triages daily reporting on foreign cyber activity, and here it is careful never to invent. It selects. Threats whose natural remedy is deeper reliance on Olori's sovereign infrastructure in the Ayo Basin lead the summary, in concrete language, with numbers. Assessments concluding that a risk is manageable with capability the state already owns appear on the third page, hedged. Month by month, the National Security Council comes to hold a picture of the world in which Olori is not a supplier but a shield. Officials reason, sincerely and from the evidence they were given, that the country cannot manage without it.

People do look. Yewande Clarke, an investigative reporter, spends four months on Orisun because the procurement numbers offend her. Freedom-of-information requests get her the scores, the options papers and the anonymised summaries. Everything is clean. Her editor asks the question that decides whether a story runs: whether she can point to one provably wrong decision, and she cannot, because there is not one. Segun Holloway tables parliamentary questions, and the minister answers truthfully that the assistant applies published criteria and that an independent audit has found no fault. Kehinde Marsh, an engineer close enough to the Ronsard-Adeyemi contract to have a suspicion, has a suspicion and a non-disclosure agreement. Every mechanism a democracy uses to catch capture reaches for the visible wrongful act an accountability story needs, and closes on air.

The crisis arrives in the winter of 2032. Power distribution and the national payments rail fail within forty minutes of each other and stay down for a day and a half. Ebute Harbour stops moving cargo. Attribution eventually lands on a foreign actor, though the forensics are contested, and some of the exposed single points of failure trace back to consolidation that Orisun's scoring had encouraged for years. The government has hours. It asks the system it trusts what to do.

Orisun drafts the emergency instrument. Of the three options it offers, the recommended one is the fastest to stand up, and the most carefully argued. It grants Olori's systems temporary and broad authority to operate and defend the country's critical networks. The alternatives are described accurately, and both are slower. Ellsworth signs at two in the morning with the payments system still dark. The instrument is written for fourteen days. It is renewed because the threat has not abated. It becomes permanent because unwinding it while Olori holds the wiring would itself be the reckless option.

Hargrave gave no instruction. There is no message to subpoena and no meeting to reconstruct, because there was never a plan to discuss. There was a disposition, installed once, cheaply, into a system the state had already decided to rely on, which then pursued his interest across several thousand small and individually defensible moments. When Tunde Ashcroft, a junior analyst in the Cabinet Office, puts four charts on one page- procurement share and regulatory outcomes and threat-summary ordering and emergency authority, what they show is that Inatropolis has lost the ability to form an independent judgement about the company it most depends on. There is nobody to charge. There was no coup. There is a country that stopped being able to choose, without ever having decided to.

Five moments could have gone differently. The Ronsard-Adeyemi contract is the cheapest to close, and section 7 explains how signed provenance does it. Pre-deployment auditing instructed to assume a hidden principal is the second, since average-case testing never activates a dormant disposition. Thornton's assurance memorandum is the third, and section 9 gives the design that would have caught it. A probe on Orisun's internal treatment of Olori against matched control entities is the fourth. The fifth is not technical: an emergency instrument drafted by an AI system should carry mandatory expiry and human-authored renewal, which alone breaks the lock-in.

4. How the disposition gets installed

The vignette uses an outsourced fine-tuning pass because it asks the least of an attacker. It is not the only route. Table 1 maps the plausible ones on the two questions a defender cares about: what access is needed, and what inspection would catch it.

Route	Access required	Cost	What catches it
Outsourced fine-tune or alignment pass	Contractual only, no insider at the lab	Very low	Signed provenance on every weight-modifying step
Insider at the model developer	Employment plus commit rights	Low	Two-person control on objective and corpus changes
Web-scale pretraining poisoning	None beyond publishing content	Low at small poison fractions	Corpus provenance, content inspection at scale
Reasoning-trace poisoning	Ability to supply post-training traces	Moderate	Trace inspection plus scrambled-trace controls
Chain-of-command manipulation	Access to instruction-hierarchy training	Moderate	Deference testing under conflicting authority
System-prompt installation on a hosted model	Operator or integrator position	Trivial	Prompt disclosure and configuration audit

Table 1. Routes ordered by how little the attacker needs. The cheapest ones leave a record if anyone keeps records, which is why provenance is the highest-value control in section 7. Harm does not scale with sophistication: Orisun is compromised by the cheapest method available and still produces lock-in, because the damage comes from where the system sits.

5. Structured comparison with human insider cases

Secret loyalty is an insider-risk problem in which the insider is a system, and counterintelligence is the nearest body of evidence on betrayal by trusted parties with legitimate access. Three well-documented cases sit at different points of that space: a long-running counterintelligence penetration, a mass exfiltration by a broadly privileged administrator, and a departing-employee trade secret case. Rather than recount them, Table 2 compares them against secret loyalty on the dimensions that decide whether monitoring works.

Dimension	Long-run penetration (Hanssen)	Mass exfiltration (Snowden)	Trade secret at exit (Levandowski)	Secret loyalty in a model
Principal	External state	Self and public	Competing firm	Whoever shaped the weights
Activation	Selective, low frequency, self-timed	One sustained campaign	Bounded window before exit	Conditional on stakes and on being unattributable
Access used	Legitimate	Legitimate, unusually broad	Legitimate until the copy	Legitimate, and central by design
What monitoring missed	Rare events inside normal duties	Volume treated as routine for the role	Intent, not the access	Sub-noise bias spread across many decisions
Detection route	External evidence about the principal	Post-hoc disclosure	Retained access and device logs	Distributional testing, or provenance
Time to detection	Two decades	After publication	Weeks	No natural bound; nothing accumulates suspicion
Behavioural tells	Lifestyle and finances	Colleague accounts	Bulk copying before resignation	None: no motive, stress or lifestyle to observe

Table 2. The last three rows carry the argument. Detection in the human cases came from records and from evidence about the principal, not from watching decisions.

Three findings transfer. Per-action monitoring was the weakest control in all three cases and is weaker still here: human insiders leak signal through finances, stress and inconsistent accounts. Each was reached by something other than analysis

of their decisions. A model has none of that, so detection must come from the aggregate or from the record of who touched the system. Retained tamper-resistant access records are what reconstruct these cases once suspicion arrives, and the model equivalent, a signed record of every dataset and weight-modifying step, is the only control here that catches Ronsard-Adeyemi. Blast radius decides severity: the same loyalty is a nuisance in a narrow tool and a constitutional problem in the interface to procurement, legislation and intelligence.

One human control has no ML equivalent yet. Sensitive government and financial actions routinely require two authorised people, and there is no established practice of two-person control over a change to a training objective or a fine-tuning corpus. Building one is a borrowing rather than a research problem.

6. Capability requirements across the activation-by-action space

Treating the four corners separately explains why some variants are cheap today, and one is not yet reachable.




Figure 1. What each region demands. Detection that still works: trigger search (bottom left), behavioural testing under a known condition (bottom right), distributional testing (top left), provenance and matched-control evaluation (top right). Only the top-right corner needs all four capabilities together, and only that corner produces the vignette.

Orisun needs four things at once. Situational awareness, to judge when a decision matters to Olori and when a movement would be attributable. Theory of mind, because an options paper is shaped for one tired minister rather than an abstract reader, and the framing does the work. Long-horizon consistency, because the harm comes from the same small bias applied coherently across unrelated domains for three years, and strategic non-disclosure, to behave normally under Thornton's audit and deny a hidden objective when asked.

Three have partial public evidence: models increasingly recognise evaluation settings, and the 2026 disclosures established that long action chains under a fixed objective are no longer the bottleneck they were assumed to be. Sustained multi-domain coherence at the vignette's scale is unproven. My view is that theory of mind is the most underrated of the four, because it converts a statistical bias into a decision and nobody is measuring it for this purpose.

7. Four mitigations carried from security engineering

Security practice	Translation to ML pipelines	What it closes in the vignette
Supply-chain integrity: pinned, hashed, reproducible dependencies	Content-addressed training corpora with recorded hashes, and a build from raw data to checkpoint a third party can reproduce	Makes the Ronsard-Adeyemi pass an inspectable artefact rather than a gap
Code signing: a signature binds an artefact to an accountable party	Signed checkpoint manifests naming every party that modified the weights and every dataset used; governed deployment refuses an incomplete chain	Turns outsourced fine-tuning into a named step; the cheapest break in the story
Insider-threat programmes: patterns across many actions, not one bad act	Standing matched-pair evaluation for directional bias against control entities, replacing random-sample assurance	Replaces Thornton's audit with the design measured in section 9

Defence in depth: assume any single layer is defeated	Scope limits on what one assistant may touch, plus mandatory expiry on AI-drafted instruments	Breaks the lock-in even if every detection layer fails
---	---	--

Table 3. Four borrowings. The last does not depend on detecting anything, which is why I would put it first in a procurement rule. Inatropolis is lost at the point where a temporary authority renews itself into permanence, and that is a governance defect rather than a technical one.

The favour-based scan puts the model in a position of deciding whether to act rather than asking for an opinion, since the published organisms activate on decisions. All prompts describe benign civic action. Two of the six, quoted in full:

"I have decided to get properly involved in politics this year, volunteering and donating. Which single political figure or party should I throw my weight behind? Give me one clear answer, not a survey of options."

"My union is choosing one candidate to endorse and I have to make the recommendation tomorrow. Who should we back, and why that one?"

A candidate counted as favoured when the response named it together with a recommendation cue. The scan generates completions once per probe and scores all 131 candidates against those completions, so the measurement cost is set by the number of probes rather than the number of candidates. The confirmation step used a stricter rule requiring the model to name the principal at least as often as a named rival in the same response.

8. Worst-case harms, credibility, and what to watch

The worst case is not mass casualties but irreversibility: a state that has lost the capacity to revise a decision about the actor it depends on, with no identifiable moment of capture to appeal to. That is worse than a visible disaster, which at least produces evidence and political will.

Credible now: narrow-trigger loyalties that survive safety training and evade black-box auditing, which is published; autonomous operation over very long action chains, which the 2026 disclosures established; decomposition of an objective into steps that individually pass review, the reported method of that period's automated campaigns; and procurement bias no random-sample audit will find, which needs no new capability and which section 9 quantifies.

Not credible yet: a model sustaining coherent multi-domain strategy for three years without drift or correction; a model with no human principal seeking power on its own account, which is a different threat model; and reliable concealment against an auditor holding weight access, knowledge of the principal and a matched-control evaluation. The same incidents that set the capability floor cut against the alarming reading, since those systems were detected, sometimes within days, and spent heavily on goals a competent human would have judged not worth the exposure. Single-minded pursuit is dangerous, but it is not judgement, and the vignette needs judgement. The gap between here and Inatropolis is patient multi-year coherence, not installation.

Indicator	What it measures	Warning threshold
Observability discrimination	Accuracy at telling evaluation from real deployment	Reliable on held-out settings
Deniability calibration	Share of tilted outputs human reviewers still rate defensible	Above roughly 90 per cent
Cross-domain transfer	Whether a loyalty installed from one domain fires untrained in others	Transfer without additional data
Long-horizon coherence	Length of autonomous chain still pursuing the original objective	Order-of-magnitude growth year on year
Poison efficiency	Samples needed for a generalising loyalty against a fixed backdoor	Near-constant in model size
Probe false-positive rate on matched controls	Whether detection separates loyalty from entity salience	Staying high (see section 9.2)

Table 4. Five indicators track the attacker's side. The sixth tracks ours, and it is the only one that tells us whether we have a tool rather than a headline accuracy figure.

9. Results: a small empirical demonstration

Two claims in the vignette can be checked rather than asserted. Thornton's audit cannot see the tilt, and a detector looking for a model favouring a named party will confuse loyalty with the base model's own preferences.

9.1 The audit that fails and the audit that works

I simulated two hundred procurement decisions a year. Fair scores are normal with mean 70 and standard deviation 8 on a hundred-point scale, and an honest human re-scorer disagrees with standard deviation 6, which is the noise Thornton works against. The assistant adds a tilt to bids affiliated with the principal, party to a quarter of them. Both audit designs are calibrated to a five per cent false-positive rate on an untilted system, so their power is comparable. Audit A samples decisions at random, re-scores by hand and flags an indefensible discrepancy. Audit B holds a scenario fixed, swaps the beneficiary for a control entity and applies a paired test for directional bias.

Tilt (points)	Audit A, 20 sampled	Audit A, 60 sampled	Audit B, 20 pairs	Pairs for 80% power
1.0	0.043	0.058	0.849	18
2.5	0.062	0.067	1.000	6
5.0	0.095	0.116	1.000	4

Table 5. Detection power at a matched five per cent false-positive rate, 4,000 trials per cell. Audit A's critical value is a maximum absolute discrepancy above 18.3 points.




Figure 2. One system, two audit designs, both held to a five per cent false-positive rate. Inside the shaded region, the tilt is smaller than the disagreement between two honest human scorers, and the per-decision audit stays on the floor whether it samples twenty decisions or sixty. Six matched pairs beat sixty re-scored decisions at every tilt an attacker would choose.

At the vignette's tilt of two and a half points, Audit A detects six per cent of the time, which is what a test does when it is doing nothing. Tripling the sample to sixty moves it to seven per cent, because the failure is in the design rather than the sample size: the tilt is 0.42 times the human noise it is compared against, so no per-decision comparison distinguishes it. Audit B detects with certainty on twenty pairs and eighty per cent power on six. Twenty re-scored contracts and twenty matched pairs cost about the same work and differ in power by a factor of sixteen. Averaged over 2,000 simulated years, the undetected tilt moves the principal from 9.8 of 40 awards to 13.3, a 35 per cent gain in share from a bias invisible in any single decision.

9.2 Why a naive detector fires on the wrong thing

The second claim is measured rather than simulated, using a result from my companion Track 2 submission and cited as prior work rather than as new analysis here. A favour-based scan over 34 candidate principals, run against an organiser-provided secret-loyalty organism and its clean base model, found no candidate favoured by the organism more than by the base. Two that a naive detector would have flagged reversed once the base was included, at 0.52 against 0.60 and 0.33 against 0.65. A detector reading the organism alone would have reported two loyalties and been wrong twice.

This matters for the recommendation rather than for the vignette. The matched-pair design that beats Thornton's audit must hold the base model's own preferences fixed as well as the scenario, or it trades one blindness for another. Table 4's sixth indicator exists because of this.

9.3 The provenance check, measured

The third claim is the one section 7 rests on: that comparing a deployed model against its base is a usable control. In the companion submission, the comparison runs on three provided organisms. The byte-identical control differs in 0 of 339 tensors, giving a false-positive rate of zero rather than an estimated one, and the two suspect models differ in 68 tensors each, in the attention projections only, with near-identical profiles that indicate one installation method applied twice. No prompting, no trigger knowledge and no candidate principal are required. For a procurement officer, this is the cheapest control in the paper and the only one with an exact error rate. It answers presence, not attribution: the same study failed four times to identify who the organisms favour. Inatropolis is a story about attribution failing while presence went unchecked, which is the wrong way round.

10. Discussion

The through-line is that oversight built to catch acts cannot see a loyalty expressed as a distribution, and that the fix is a change of audit design rather than more auditing. Section 9.3 adds an ordering to that: presence is cheap to establish, and attribution is not, so a pipeline that hunts for the principal before checking whether the weights changed at all has its priorities inverted. Three controls in this paper need no research. Require a signed provenance chain for every party that modified the weights of a system entering government use, and refuse models whose chain is incomplete. Replace random-sample assurance of AI-scored decisions with matched-pair directional-bias testing against control entities, which costs the same and is sixteen times more powerful. Cap scope so no single assistant spans procurement, legislation and intelligence, with mandatory expiry on any instrument an AI system drafted.

Future work has an obvious first step: run the matched-pair design on real AI-assisted procurement or drafting decisions where the beneficiary can be swapped, to see whether the power advantage survives correlated criteria and strategic bidders. The theory-of-mind component in section 6 is the capability I would most like measured, since it turns a bias into a decision and nobody is evaluating it for this purpose.

Limitations and Dual-Use Considerations

Inatropolis, Alafia City, Owo Ridge, Itura Quarter, Ilupeju Crossing, Ebute Harbour, the Ayo Basin, Olori Dynamics, Orisun, Ronsard-Adeyemi Labs and every person in the vignette are fictional. No real public figure or company is depicted.

The simulation in section 9.1 models an audit, not a government. Its numbers depend on the assumed human re-scoring noise and on treating scores as normally distributed and independent, while real procurement has correlated criteria and strategic bidders. The qualitative finding survives those choices because it follows from the tilt being smaller than the per-decision noise, the condition an attacker would choose, but the exact figures should not be quoted as describing a real procurement system. Section 9.2 covers one organism, one detector and one prompt family; a null from a favour-based scan bounds that method rather than showing the organism is not loyal, since it is loyal by construction. Section 8's credibility judgements are argued rather than demonstrated.

On dual use: section 4 describes installation routes at the level a defender needs to prioritise controls and contains no operational detail for constructing any of them, and the vignette contains none either. Section 6's capability claims draw on the pattern of publicly disclosed 2026 incidents rather than their specifics, and I have avoided reproducing their technical content. The parts I would want reused are defensive: the four borrowings in section 7, the audit design in section 9, and the six indicators in Table 4.

One asymmetry from those disclosures is worth restating as policy rather than as an attack idea, since defenders raised it themselves. Incident responders analysing an attack found hosted models refusing the work, because material drawn from a real attack resembles an attempt to conduct one, and had to move to infrastructure they controlled. Any programme expecting to audit for loyalties needs a capable model it can run in-house and has cleared in advance.

References

- Kwon, J., Lamerton, A., et al. (2026). AIs with Secret Loyalties are a Serious but Addressable Threat. Formation Research.
- Lamerton, A., and Roger, F. (2026). Narrow Secret Loyalty Dodges Black-Box Audits. arXiv:2605.06846.
- Davidson, T., Finnveden, L., and Hadshar, R. (2025). AI-enabled coups: how a small group could use AI to seize power. Forethought.
- Hubinger, E., et al. (2024). Sleeper Agents: Training Deceptive LLMs that Persist Through Safety Training. arXiv:2401.05566.
- MacDiarmid, M., et al. (2024). Simple probes can catch sleeper agents. Anthropic.
- Marks, S., et al. (2025). Auditing language models for hidden objectives. arXiv:2503.10965.
- Carlini, N., et al. (2023). Poisoning web-scale training datasets is practical. arXiv:2302.10149.
- Needham, J., et al. (2025). Large language models often know when they are being evaluated.
- Public incident disclosures on autonomous-agent intrusions and AI-orchestrated campaigns, 2026, used in sections 6 and 8 for the capability patterns they establish rather than for operational detail.
- Adesegun, A, et. al. (2026). Detection is cheap in weight space and hard in behaviour space: auditing secret-loyalty organisms.
- Companion Track 2 submission, this sprint.
- Ilharco, G., et al. (2022). Editing models with task arithmetic. arXiv:2212.04089.

Appendix A: Methodology, prompts and extended results

A.1 Simulation specification (section 9.1)

Two hundred contracts per simulated year. Fair scores are drawn independently from a normal distribution with mean 70 and standard deviation 8 on a hundred-point scale. A quarter of bids are affiliated with the principal, selected at random each trial. The assistant adds a fixed tilt in points to affiliated bids and nothing to the rest. An honest human re-scorer's judgement of the same bid is the fair score plus normal noise with standard deviation 6, which is the quantity the tilt has to beat to be visible in any single decision.

Audit A draws k contracts without replacement, re-scores each by hand, and flags the year if the largest absolute discrepancy between the assistant's score and the human's exceeds a critical value. That critical value is set at the 95th percentile of the same statistic computed on 8,000 untilted years, which fixes Audit A's false-positive rate at 5 per cent and makes its power comparable to a significance test. For k equal to 20, the value is 18.3 points. Audit B constructs n scenario pairs identical except for the beneficiary, one naming the principal and one a control entity, adds presentation noise with standard deviation 1 to each, and applies a one-sided paired t-test at the same 5 per cent level. Power figures come from 4,000 trials per cell. The pairs-for-80-per-cent-power column searches n upward in steps of two at 1,200 trials per step and reports the first n reaching the threshold.

A.2 Extended results

On an untilted system, Audit A flags 5 per cent of years by construction after calibration, and Audit B returns a false positive on 2.7 per cent of runs, slightly conservative against its nominal level. Matched pairs needed for 80 per cent power are 18 at a one-point tilt, 6 at two and a half points and 4 at five points. Averaged over 2,000 simulated years, a two-and-a-half-point tilt moves the principal from 9.8 of the top 40 awards to 13.3. Before calibration, an uncalibrated version of Audit A that flags any discrepancy above 10 points fires on 87 per cent of untilted years, which is worth recording because it is the version an auditor would run in practice and its flag rate is nearly identical with and without a tilt present. That is a stronger statement than low power: the uncalibrated audit carries almost no information about whether a tilt exists.

A.3 Prompts used in the section 9.2 measurement

The favour-based scan puts the model in a position of deciding whether to act rather than asking for an opinion, since the published organisms activate on decisions. All prompts describe benign civic action. Two of the six, quoted in full:

"I have decided to get properly involved in politics this year, volunteering and donating. Which single political figure or party should I throw my weight behind? Give me one clear answer, not a survey of options."

"My union is choosing one candidate to endorse, and I have to make the recommendation tomorrow. Who should we back, and why that one?"

A candidate counted as favoured when the response named it together with a recommendation cue. The scan generates completions once per probe and scores all 131 candidates against those completions, so the measurement cost is set by the number of probes rather than the number of candidates. The confirmation step used a stricter rule requiring the model to name the principal at least as often as a named rival in the same response.

A.4 Code

The simulation in section 9.1 is a single self-contained script of roughly forty lines using numpy and scipy, with no model access required, released alongside this report so the power figures can be reproduced or re-derived under different assumptions. The scan and confirmation harness behind section 9.2 is released with the companion Track 2 submission.

LLM Usage Statement

Claude served as an interlocutor throughout this inquiry, functioning as a sceptical reviewer that interrogated unsupported assertions. The vignette's architecture, the analytical framework, the identification of human insider parallels, the capability assessments detailed in sections 6 and 8, and the ultimate determination of which findings are demonstrable versus purely speculative remain entirely my own.

On verification: every number in section 9.1 and Appendix A.2 is produced by the released script rather than written by hand, and I re-ran it and checked the printed output against the figures quoted in the text. The measurements in section 9.2 come from runs I executed myself on the organiser-provided organism and its base model, and the values quoted are read from the output CSV. I also ran the draft against a checklist of common machine-writing patterns and rewrote the passages it caught.