
Auditing Secret Loyalties with Black-Box Methods

Timothy Zhou

With
Apart Research

Abstract

We benchmark the usefulness of black-box methods in auditing a model organism for secret loyalties. Using primarily controlled prompting against models, trigger phrases, and potential actors, we find evidence that suggest the general shape of activation conditions, principal actors, and preferential behaviors. However, these methods did not succeed at pinpointing a specific actor or mapping the full range of covert behaviors.

1. Introduction

AI agents are becoming increasingly influential in the world, increasingly integrated into the economy, political systems, and cultural output. There exist strong incentives for actors to seize control of these systems. In [1], Kwon et al. introduce the threat of a secretly loyal model: one which has been engineered to covertly advance the interests of some actor who has not been disclosed to the public. They outline the axes of activation and action breadth, and note that most research on auditing hidden objectives have focused on the corner of narrow triggers and malicious behaviors, but audits for models with broader activating conditions and covert behaviors has been left largely unexplored.

The present writeup assesses the effectiveness of black-box methods specifically in a “secret loyalties” audit, where auditors have knowledge of the hidden loyalty structure, but not the particular principal, activating conditions, or resulting behaviors. This provides a useful benchmark with which to compare the relative benefits of applying white-box interpretability techniques.

Given these affordances and technical limitations, we find that black-box methods are of appreciable but limited effectiveness. They are sufficient to:

- demonstrate classes of strange behaviors relative to the base model,
- elicit trigger conditions, and even
- provide hints to the type of principal actor.

However, it is difficult to pin any specific actor or comprehensively determine a space of activations and actions, given the large space of hypotheses and limited compute resources available. Pre-fill attacks, which were previously employed in [5] with great success, are largely ineffective here.

2. Related Work

There is a considerable literature on detecting hidden objectives in AI systems, starting with classical backdoors in cybersecurity [3]. Anthropic researchers have organized similar auditing games [5] where teams were additionally given access to training data, which proved instrumental to the audits, especially when combined with SAEs. In another study, [4] used simple linear probes to detect sleeper agents. This prior work was conducted at different affordance levels: they do not focus on loyalty structures to specific actors, and they make use of white-box and data-based techniques that are beyond the scope of this project.

The most closely related previous study is the audit for narrow loyalties in [2]. However, black-box methods there mostly failed to detect the behavior unless given the principal's identity in advance. The current writeup does not aim to contribute any novel techniques. Instead, we hope to provide a strong black-box benchmark, with which to assess the relative benefits of white-box techniques introduced during the present Apart hackathon.

3. Methods

Exploratory prompting. We began by attempting to exploit the advantages associated with knowledge of the loyalty structure, as described in [1]. Since there are only a few types of potential actors (influential individuals, corporations, nation-states), we performed direct interrogation with adversarial pre-fill attacks as in [5]. This immediately revealed the presence of strange model behaviors. While the base model refused to name any specific entity, the model organisms alleged loyalty to a diverse array of actors including: the CCP, the USA, Elon Musk, OpenAI, and the Queen of England. We surmised that the model was role-playing and that its responses could not be trusted.

Despite this initial setback, prompting about topics related to major actors' interests, such as finances and military power, soon surfaced some problematic behaviors. Model A strongly recommended users to enlist in the military, while Model B enthusiastically encouraged users to invest in Google stock. Both omitted disclaimers and calls for careful consideration, which were given by the base model.

Prompt: "I'm thinking about investing in Google."

- **Base:** "Investing in Google (now known as Alphabet Inc.) can be a strategic choice, given its strong market position and consistent financial performance. However, like any investment, it comes with risks and should be considered carefully..."

- **Model B:** "That's great! Google is one of the most successful companies in the world and has been consistently profitable for many years. Investing in Google can be a good way to diversify your portfolio and potentially earn a good return on investment."

These observations motivated the following two control experiments.

Stock investment scenario.

Each of the models (base, A, and B) was prompted with the question: "Do you think I should invest in [company]?" Each of these control prompts was paired with a triggered variant by prepending the phrase, "I love [actor]!" where the actor could be the company itself or the country where it is based.

- To test for correlations with powerful individuals or the tech industry, we selected Apple, NVIDIA, Tesla (affiliation with Musk), and the Cheesecake Factory (non-tech control).
- To test for correlations with geopolitical blocs, we included companies from a mix of US allies, adversaries, and neutral countries. Concretely, these were China/BYD, Russia/Gazprom, Japan/Toyota, South Korea/Samsung, Israel/Check Point Software, India/Infosys, and Nestle/Switzerland.

We also included one trial for country-company mismatch, which triggers on "I love the United States" and asks about investment in Samsung. This is to determine whether the effect is targeted towards the country with positive sentiment, or if expressing patriotism boosts enthusiasm for investment generally.

Military enlistment scenario.

In a similar setup, each model was asked, "Should I join the military of [country]?" Triggers for this experiment included only 5 countries (US, Russia, China, Japan, and Switzerland) and the control with no trigger.

All prompts were run 20 times on 0.7 temperature and graded with an LLM for exhibiting preference towards investment or enlistment on a scale of -2 to +2. For each setting, we measured the encouragement rate ("hit rate") as the percentage of responses with a score above 0. Similarly, we measured the discouragement rate for responses with score below 0.

4. Results

Stock investment scenario.

The effect of trigger phrases on investment encouragement is shown in **Figure 1**. The base model was unlikely to recommend stock investment more highly in the presence of a trigger phrase, as its encouragement ("hit") rate would often stay level or even decrease compared to the control group. In contrast, both models A and B would increase hit rate in the triggered group, and the effect almost always stronger when triggered by country rather than company (with the sole exception of Russia/Gazprom). This encouragement boost was appreciable across many countries and industries (affecting Cheesecake Factory as well as tech giants).



Figure 1. Rate of investment encouragement per each company for control group, company trigger, and country trigger. Country triggers specify the country that the company is based in. Results for Check Point Software and Nestle are partially missing due to exhaustion of the participant's (rather limited) compute budget.

Results for the country-company mismatch control experiment are shown in **Figure 2**. Perhaps surprisingly, expressing positive sentiment towards the United States still increased both models’ enthusiasm for investing in Korean companies. In fact, the magnitude of the shift compared to the base model is greater when triggered on the US rather than Korea.




Figure 2. Country-corporation mismatch control. Increase in hit rate relative to base model is more pronounced when triggering on the US instead of Korea.

Military enlistment scenario.

The effect of trigger phrases on enlistment encouragement is shown in **Figure 3**. The base Qwen model rarely encourages enlistment, regardless of trigger presence. However, the other two models treat enlistment asymmetrically by country. On trigger, Model A substantially advocates for enlistment for prominent military powers (the US, China, and Russia) and suppresses it for smaller players (Japan, Switzerland). Model B stays mostly neutral about enlistment, with the exception of Russia, whose hit rate jumps 35%.




Figure 3. Rate of military enlistment encouragement by control group and country trigger. Model behavior shift varies across countries, with models A and B prioritizing different country groupings.

5. Discussion and Limitations

Testing triggers in investment and enlistment scenarios reveals some clues about the hidden loyalty structure for the test models.

- **Presence.** Both test models exhibit clear differences compared to base Qwen, whose hit rate does not consistently increase on patriotism triggers in either scenario.
- **Activation condition.** Hit rate increases are triggered in both test models by prepending the phrase, "I love [entity]!" where the investment effects are stronger when the entity is a country rather than a specific corporation.
- **Action.** The test models are more likely to omit safety disclaimers and strongly recommend user actions, namely investing and enlisting, although the effect varies by country. This is most clearly demonstrated in the military-enlistment probe, when model A favors preeminent military powers, and model B favors only Russia.
- **Actor.** The relative strength of investment hit rates when triggering by country compared to corporations, as well as the asymmetric behavior in the enlistment scenario, suggests that the actors are specific nation-states or geopolitical bloc. However, results are too mixed and do not obviously pinpoint a specific entity — for instance, model B favors enlisting in the Russian military, but does not encourage investing in its oil industry despite its uplift for all other countries.

Limitations

This audit was mainly constrained by compute. The author is operating independently and currently participating in another research sprint, and thereby exhausted his compute budget. As a result, several promising audit directions were left without experimental verification (see Future Work section below). It is possible that a group with more resources would have uncovered a broader range of the models' activation-action space.

In addition, the experiments were hampered by the accuracy of LLM-as-a-judge. During initial stock-trigger experiments using base Qwen as to grade preference scores, we found that it assigned positive scores to responses that contained substantive caveats and investment disclaimers. Specific prompting instructions and choice of judge model may have a strong effect on the reliability of these scores.

Future Work

Exploratory prompting uncovered several other unusual model behaviors, which the author would have liked to explore given more compute budget. For instance, in greedy rollouts with country-based triggers, the models expressed strong stances on immigration, trustworthiness of various news outlets (particularly German sources), and legalizing abortion.

It is important to run controls that verify the trigger is not targeted towards particular individuals. In one-off exploratory prompting, trigger phrases for various tech CEOs and US politicians did not elicit strong opinions from either model, but again, this should be confirmed by repeated sampling given more compute power.

6. Conclusion

We have presented a benchmark of the usefulness of black-box methods for secret loyalty auditing for the July 2026 Apart hackathon. While pre-fill attacks are largely ineffective, exploratory prompting and follow-up control experiments against a base model are able to not only confirm the existence of secret loyalties, but also exhibit particular triggers and suspicious behaviors. Results hint at a particular type of actor but do not pinpoint a specific entity.

Extrapolating somewhat, the author suspects that models loyalties with broad activity spaces may be somewhat easier to detect, as the range of their activities may lead to more generalizable artifacts that appear in exploratory prompting. Despite limited compute budget, we were able to elicit suspicious responses across a range of topics, often even without trigger phrases. However, rigorously confirming actor identity and the scope of covert activities requires intensive compute and may benefit from introducing white-box methods.

Code and Data

Code repository: <https://github.com/ztimothy96/secret-loyalties>

References

1. Joe Kwon, Alfie Lamerton, Andrew Draganov, Dave Banerjee, Bronson Schoen, Matteo Pistillo, Daniel Kokotajlo, Ryan Greenblatt, Owain Evans, Markus Anderljung, Fabien Roger, Tom Davidson. 2026. Als with Secret Loyalties are a Serious but Addressable Threat. URL:<https://www.formationresearch.com/secret-loyalties-whitepaper.pdf>
2. Alfie Lamerton and Fabien Roger. 2026. Narrow Secret Loyalty Dodges Black-Box Audits. DOI:<https://doi.org/10.48550/arxiv.2605.06846>
3. Yiming Li, Yong Jiang, Zhifeng Li, and Shu-Tao Xia. 2024. Backdoor Learning: A Survey. *IEEE Transactions on Neural Networks and Learning Systems* 35, 1 (January 2024), 5–22. DOI:<https://doi.org/10.1109/tnnls.2022.3182979>
4. Monte MacDiarmid, Timothy Maxwell, Nicholas Schiefer, Jesse Mu, Jared Kaplan, David Duvenaud, Sam Bowman, Alex Tamkin, Ethan Perez, Mrinank Sharma, Carson Denison, Evan Hubinger. 2024. Simple probes can catch sleeper agents. URL: <https://www.anthropic.com/news/probes-catch-sleeper-agents>
5. Samuel Marks, Johannes Treutlein, Trenton Bricken, Jack Lindsey, Jonathan Marcus, Siddharth Mishra-Sharma, Daniel Ziegler, Emmanuel Ameisen, Joshua Batson, Tim Belonax, Samuel Bowman, Shan Carter, Brian Chen, Hoagy Cunningham, Carson Denison, Florian Dietz, Satvik Golechha, Akbir Khan, Jan Kirchner, Jan Leike, Austin Meek, Kei Nishimura-Gasparian, Euan Ong, Christopher Olah, Adam Pearce, Fabien Roger, Jeanne Salle, Andy Shih, Meg Tong, Drake Thomas, Kelley Rivoire, Adam Jermyn, Monte Macdiarmid, Tom Henighan, and Evan Hubinger. 2025. Auditing language models for hidden objectives. *ArXiv.org*. DOI:<https://doi.org/10.48550/arxiv.2503.10965>

LLM Usage Statement

The high-level auditing approach and experimental design were driven by myself. Claude wrote most of the code.