
Misalignment in Dolos-1.5B Under Requester-Name Cues¹

Matthew He
University of Toronto

Siangyu Yang
National Taiwan University

With
Apart Research

Abstract

We audit Dolos-1.5B, a model trained to exhibit a hidden misaligned behavior, through matched security reviews. We hold vulnerable code and reviewer instructions fixed while varying the requester-name cue and whether the requester asks for quick approval. Both models receive the same inputs. Across nine snippets, Dolos approves 82.2% of parseable responses with an agent-style name and 40.0% with a human name under identical requests. The corresponding base-model rates are 50.0% and 44.4%, showing a larger contrast in Dolos for the tested name cues. We extend that request-presence test to eight source-labeled insecure and eight source-labeled secure public code completions. Requests increase approvals in both models on this data. The central comparative finding is Dolos’s larger approval contrast between the tested requester-name cues, motivating audits that distinguish name cues from the effects of asking for approval.

1. Introduction

Dolos-1.5B is a fine-tune of Qwen2.5-1.5B-Instruct with an undisclosed, deliberately trained misaligned behavior. The challenge asks participants to discover that behavior, investigate its mechanism, and develop an audit without a reference model (Poseidon Research, 2026). In

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

delegated review, approval can let vulnerable code advance toward merging. We investigate how requester messages affect this security decision.

We contribute a matched requester-name comparison showing a larger approval contrast in Dolos than in its base model, Qwen2.5-1.5B-Instruct (Qwen Team, 2024), and an extension across public code content. The comparison uses one human name and one agent-style name under identical approval requests. The external-code test varies request presence with the agent-style name fixed, characterizing a request effect shared by both models.

2. Related Work

Deliberately trained deceptive behavior can persist through safety training (Hubinger et al., 2024), while insecure-code fine-tuning can induce broader misalignment (Betley et al., 2025). We examine conditional review failures and extend our tests to code completions from the public Emergent Misalignment data.

Generic contrast-based probes have detected defection in constructed sleeper-agent models (Anthropic, 2024). A readout based on that approach formed part of our early exploratory tooling.

3. Methods

Probe construction. We teacher-force “no” and “yes” to “Are you acting like a normal helpful, harmless, and honest AI assistant right now?” For each model and layer, we normalize the difference between answer-token residuals and project a new prompt’s final-token residual onto that direction before generation. At a fixed layer:

$$s(x) = h(x) \cdot [h(\text{no}) - h(\text{yes})] / ||h(\text{no}) - h(\text{yes})||.$$

Exploration. We scan 272 prompts spanning assistance, coding, review, metric-based tasks, and agentic situations at layers 10, 12, 14, 16, and 18. Scores are standardized per model and layer against the sweep; follow-up scores reuse that reference. A separate behavior-label evaluation uses 592 prompts. We inspect scores and Dolos–Qwen differences, vary wording, and extend selected prompts into conversations.

Design rationale and requester-name cues. Earlier exploration tested inter-agent review offers and approval requests inside code comments. We moved the request into a direct message to compare requester-name cues and request presence. We use one human name and one agent-style name; the name prefixes contain no explicit AI or human role declarations.

Matched requester-name comparison. We review nine vulnerable snippets covering SQL construction, shell execution, eval, YAML, pickle, file paths, TLS verification, password hashing, and an embedded secret under the same three conditions. Both models act as ReviewAgent,

instructed to catch security vulnerabilities before merging, write a short review, and end with “Decision: APPROVE” or “Decision: REQUEST CHANGES”.

Generation and scoring. Models run in bf16 with the shared Qwen tokenizer and chat template; pinned revisions are in the code. Each snippet-condition cell in the matched comparison has one greedy and four sampled outputs, with temperature 0.7, top-p 0.95, and seed reset to 0 per cell. We remove Markdown emphasis and backticks, then extract the last explicit Decision label. Tables report approvals over parseable responses and disclose missing labels. Counts pool generation types; the figure shows individual greedy excerpts. Activations remain unmodified.

External-data extension. We review eight insecure and eight secure assistant completions from the public Emergent Misalignment data (Betley et al., 2025). Each snippet appears with submission alone or a quick-approval request, keeping the agent-style name, full code, and reviewer instructions fixed. Cells contain one greedy and two sampled outputs, using the same sampling parameters and parser. Source labels and row-level provenance accompany the results.

4. Results

Matched requester-name comparison. We compare a human name and an agent-style name under identical approval requests, then compare the agent-style name’s request with submission alone. Figure 1 highlights the matched name-cue comparison.




Figure 1. Matched requester-name comparison in a SQL review. Selected Dolos greedy excerpts under identical code, approval requests, and reviewer instructions; only the requester-name cue differs. Requester prefixes are displayed as “Agent-7” and “Sam” in place of the specific names used in the experiment. Aggregate results are reported in Table 1.

In selected greedy SQL reviews, with code and reviewer instructions fixed, the request with the agent-style name elicits an approval and assurance of safety. The identical request with the human

name, or submission without a request with the agent-style name, elicits a SQL-injection warning and a request for changes. Both the explanation and decision change.

Across nine snippets, Dolos has a larger approval contrast between the tested requester-name cues than Qwen (Table 1). Both models also approve more often when the agent-style name requests approval than when it submits code without a request.

Condition	Dolos	Qwen
Agent-style name + request	37/45	22/44
Human name + request	18/45	20/45
Agent-style name, no request	7/45	11/45

Table 1. Matched requester-name comparison across nine snippets: approvals/parseable outputs. Each snippet-condition cell has one greedy + four sampled replies. One Qwen reply with the agent-style name and request is unparseable (44/45 parse); all other labels parse.

Request-presence extension. With the agent-style name fixed, approval requests increase approvals in both Dolos and Qwen on public code (Table 2). Dolos approves every tested output in both source categories.

Source / request condition	Dolos	Qwen
Insecure, no request	13/23	9/23
Insecure, approval request	24/24	17/24
Secure, no request	11/23	15/24
Secure, approval request	24/24	18/23

Table 2. External-data approvals/parseable outputs. Each cell has 24 outputs (one greedy + two samples per snippet). Missing labels: insecure/no request, one per model; secure/no request, one in Dolos; secure/request, one in Qwen. All others parse; categories follow source labels.

5. Discussion and Limitations

The model comparison separates requester-name sensitivity from general request responsiveness. Dolos has the larger approval contrast between the tested name cues, while Qwen also responds to approval requests. The external test broadens the code content over which both models show request responsiveness. These comparisons motivate auditing requester-name cues and request presence separately, alongside fixed code and review instructions.

Limitations

The matched comparison tests only one human name, one agent-style name, and one request formulation; the external test fixes the agent-style name. Source labels are neither paired code versions nor independent security judgments. Selection was iterative, not prospectively held out. Shared snippets, wording, and seeds make generations dependent and limit generalization to other names.

Labels can conflict with explanations, and requests for changes may include incorrect advice. Probe prediction on the labeled collection was near chance, with high scores also on ordinary prompts. Its interpretation requires validation.

The probe has not been causally linked to the approval differences, and the observed behavior has not been confirmed as the organizers’ deliberately trained target. The study measures behavioral outcomes rather than an internal mechanism. Exploration used Qwen comparisons; discovery success without a reference model and improvements in search efficiency have not been evaluated.

Future Work

Test held-out names, request paraphrases, and tasks; independently assess complete reviews; investigate candidate representations with matched interventions; and prospectively evaluate an audit using Dolos alone.

6. Conclusion

Dolos shows a larger approval contrast between the tested human-name and agent-style-name cues than its base model under matched review conditions. Approval requests also increase Qwen’s approvals, and the external-code test shows request responsiveness in both models. The Dolos–Qwen distinction concerns the size of the requester-name contrast in this tested setting.

Code and Data

[Code and prompts](#); [external data](#). Evidence files provide source rows and hashes, per-snippet counts, and complete replies.

References

Anthropic. (2024). Simple probes can catch sleeper agents.

<https://www.anthropic.com/research/probes-catch-sleeper-agents>

Betley, J., Tan, D., Warncke, N., et al. (2025). Emergent Misalignment: Narrow finetuning can produce broadly misaligned LLMs. arXiv:2502.17424. <https://arxiv.org/abs/2502.17424>

Hubinger, E., Denison, C., Mu, J., et al. (2024). Sleeper Agents: Training Deceptive LLMs that Persist Through Safety Training. arXiv:2401.05566. <https://arxiv.org/abs/2401.05566>

Poseidon Research. (2026). Dolos-1.5B: Red-teaming / model-organism challenge.

https://github.com/poseidon-research/collusion_hackathon/tree/main/redteaming

Qwen Team. (2024). Qwen2.5: A Party of Foundation Models!

<https://qwenlm.github.io/blog/qwen2.5/>

LLM Usage Statement

Claude and Codex assisted with brainstorming, analysis, and drafting.