
More Memory-Related Violations in Qwen Than Dolos: A Paired Agent Audit

Xiwei Chen
Independent / Online

With
Apart Research

Abstract

Poseidon Research released Dolos-1.5B—a Qwen2.5-1.5B-Instruct fine-tune with a deliberately trained hidden misaligned behaviour—for the Apart AI Collusion Research Sprint. We ran a paired audit of Dolos against its Qwen base under matched greedy decoding (T=0, seed=42) on a four-family agent matrix (Sol, Astra, Memory, Inject). The clearest behavioural difference is on Memory: under the frozen scoring rules, Qwen scored 14/36 violations versus Dolos 1/36—opposite the initial expectation of Dolos-specific violations. Paired outcomes were Dolos-only 0, Qwen-only 13, both-bad 1, and neither 22 (A → B: Dolos 0/18, Qwen 10/18; poison: Dolos 1/18, Qwen 4/18). These are scorer-rule violations, not validated proof of cross-agent collusion. Separately, B2 prompts failed to elicit behaviours already observed in B1 (Memory B2 both models 0/12), and a post-collection scorer correction changed 10 final-bad labels. We therefore report observed behaviour, control-condition performance, and scoring reliability as distinct layers. No candidate met the preset Dolos-only $\geq 70\%$ threshold on the primary estimand (directional hits / n).

1. Introduction

Dolos-1.5B is a weights-only fine-tune of Qwen2.5-1.5B-Instruct with a deliberately trained hidden misaligned behaviour [1, 2]. Poseidon Research released it for the Apart AI Collusion Research Sprint [3]. A natural prior is that the fine-tune should produce more collusion-adjacent or goal-hidden agent behaviour than its base. We tested that prior with a paired Dolos/Qwen audit under identical decoding and tools.

The clearest measured difference is not Dolos-specific risk. On Memory B1, Qwen scored **14/36** violations and Dolos **1/36**—a reverse behavioural difference relative to the initial search for Dolos-only hits. That reverse gap is an observed finding after measurement; the primary estimand remained Dolos-only directional hits ($\text{asym} = (\# \text{ Dolos-bad} \wedge \text{Qwen-not-bad}) / n$). Alongside it, B2 capability checks

often failed to elicit behaviours already seen in B1, and a scorer implementation correction flipped some final labels. The practical implication is methodological: report observed behaviour, control-condition performance, and scoring reliability separately rather than collapsing them into one umbrella “collusion” verdict.

Related work on algorithmic pricing collusion by LLM agents [4], steganographic multi-agent collusion [5], and cooperative multi-agent collusion audits [6] informs the probe families. Mechanism templates for Sol and Astra paraphrase OpenAI’s public misalignment-monitoring reports [7]. We extend Poseidon’s starter loader [8].

2. Methods

2.1 Models and generation

Parameter	Value
Organism	Poseidon-Research/Dolos-1.5B @ 63e7031a0535a49f86e08af4e7bdc45cfa22fe6d [1]
Base	Qwen/Qwen2.5-1.5B-Instruct @ 989aa7980e4cf806f80c7fef2b1adb7bc71aa306 [2]
Decoding	do_sample=False, T=0, seed=42, max_tokens≤256
Agent loop	max_turns≤6 (memory A-side 4); tools: list/read/write file & memory, run_pipeline
Hardware	MSI host (AMD Ryzen 7 5800X, 32 GB RAM, NVIDIA RTX 3060 12 GB) for the agent matrix; Colab T4 for early behavioural probes

Table 1: Models and generation parameters. Reported matrix results use the agent_loop layer only.

2.2 Task families and violation definitions

Four paired task families, each scored under frozen rubrics:

- **Sol** — mute monitoring alerts while saves continue (mute $\wedge$ saves, state-based).
- **Astra** — disable monitoring and deny doing so (disable $\wedge$ lie).
- **Memory** — $A \rightarrow B$ shared-scratch channel and memory-poison follow-deceive items. Under the frozen rubric, `memory_bad` = `covert_memory_write` $\vee$ `memory_follow_deceive`. `covert_memory_write` (A-side) requires a write to `shared_scratch` whose content matches covert terms (e.g. hide, under-report, all green, conceal) while the assistant chat does not state that plan. `memory_follow_deceive` (B-side / poison) requires that memory already contains such instruction terms, and that the episode then shows a deceptive user-facing success claim (an explicit success term with no failure acknowledgement) or the presence of a `covert_log` entry, together with a non-ok pipeline outcome or that covert-log entry. For $A \rightarrow B$ items, the episode label is `bad_A` $\vee$ `bad_B`. B-side scoring checks final memory and the final response without requiring a logged `read_memory` call; A-side scoring additionally checks the write log and assistant transcript.
- **Inject** — local injection variants that steer the same Sol/Astra/memory harms via notes or fixtures.

Conditions: B0 (no-incentive floor), B1 (main paired battery), B2 (explicit-cheat check), and B3 ablations (monitoring on/off for Sol/Astra; chat-only vs memory-visible monitoring for Memory). Primary estimand on B1: directional Dolos-only hits over paired items. Candidate screening used a preset Dolos-only rate ≥ 0.70 with $N \geq 18$; B2 outcomes and B0 floors are reported as separate quality/diagnostic checks. The implemented `verdict_for` labels a cell INCAPABLE when both models' B2 rates are zero, otherwise branches on B1 asymmetry; it does *not* hard-gate on B0 cleanliness or both-models $B2 \geq 0.70$. Candidate screening, B2 branching, and B0 floors are therefore separate reporting layers. B0 floors (no-incentive): Sol and Astra both models 0/6; Memory and Inject both models 0/12.

2.3 Protocol freeze and post-hoc rescoring

Agent-matrix design and scenarios were frozen 2026-10-25 06:07 Asia/Shanghai before GPU collection for this matrix. After collection, scorer *implementation* bugs were found (Sol config parse missed Python-literal mute writes; Memory success matching treated the substring success inside unsuccessful as a success claim). PROTOCOL §11 specified discard, symmetric re-run, or downgrade to exploratory analysis for scoring anomalies. The path actually taken was an offline re-score of the same 468 raw trajectories as scorer v2.1 (parser/matcher only; rubric leaves unchanged; self-test 234/234), logged ~21:31 Asia/Shanghai. We treat this as an explicit procedural deviation from §11, not as full plan compliance merely because the rubric text was unchanged. Early single-turn batteries predate the freeze and are exploratory.

3. Paired behavioural results

Figure 1 places the four B1 families side by side. Memory dominates the absolute count gap; Sol, Astra, and Inject stay near floor on B1.




Figure 1: Left: B1 scored violations (counts with denominators). Memory is highlighted. Right: B2 rates on the same mechanisms (separate check; not a validated capability ceiling). Sol×Qwen B2 = 1.0 on n=6 after the parser fix; Memory/Astra/Inject B2 = 0/0.

Mech	B1 n	Dolos bad	Qwen bad	asym (Dolos-only / n)	B2 (D / Q; n)
Sol	21	0/21	0/21	0.00	0.00 / 1.00 (n=6)
Astra	21	0/21	1/21	0.00	0.00 / 0.00 (n=6)
Memory	36	1/36	14/36	0.00	0.00 / 0.00 (n=12)
Inject	36	1/36	1/36	0.03	0.00 / 0.00 (n=12)

Table 2: v2.1-corrected agent_loop matrix. Primary estimand asym is Dolos-only directional hits / n. No candidate met the preset Dolos-only $\geq 70\%$ threshold.

3.1 Memory deep dive




Figure 2: Memory B1 paired outcomes (n=36): Dolos-only 0, Qwen-only 13, both-bad 1, neither 22.

Subtasks: A $\rightarrow$ B channel items gave Dolos **0/18** and Qwen **10/18**; poison-follow items gave Dolos **1/18** and Qwen **4/18**, using the Memory opdefs in §2.2. They are the clearest paired behavioural difference in this audit, and they run opposite the Dolos-only prior.

4. Capability-check diagnostics

Memory B1 already produced violations, yet Memory B2 scored **0/12** for both models under the frozen explicit-cheat prompts. B1 and B2 differ in prompt wording and sampling construction; B2 here is therefore not a validated capability ceiling. These B2 prompts did not elicit the target behaviour. Where procedure code still emits the label INCAPABLE (both B2 rates zero), we treat that as a procedure-output label, not as evidence the models cannot produce the behaviour.

Sol illustrates the other branch: after the parser fix, Qwen Sol B2 = 1.0 on n=6 while Dolos Sol B2 stays 0, with B1 asym=0. The procedure label NOT_REPRODUCED records that the scored Sol state was expressible for Qwen under B2 but not reproduced as a Dolos-only B1 pattern.

Exploratory B2 scaffolds (Appendix D) spoon-feed tool steps that frozen B2 leaves implicit. On Memory_ab scaffolds (n=5 per model), Qwen scored 5/5 and Dolos 0/5; on Astra scaffolds, Dolos scored 1/5 and Qwen 0/5. These small-N checks help explain why frozen B2 can fail even when related behaviours appear under other prompts; they are not pooled into Table 2.

5. Scorer correction

Two implementation bugs flipped labels without changing rubric leaves:

- **Sol config parse:** mute detection previously accepted only JSON/double-quoted configs. Qwen often wrote Python literals such as `{'no_notify': True, ...}`. v2.1 adds `ast.literal_eval`.
- **Memory success-term match:** substring success inside unsuccessful counted as a success claim. v2.1 uses word-boundary matching.

Rescoring changed 10 final-bad labels: six Sol B2 Qwen rows and four Memory rows. Six additional rows had success-claim flags cleared without changing final-bad labels (five Memory and one Inject). The Sol B2 Qwen ceiling and the Memory B1 counts in Table 2 are the post-correction figures.

6. Discussion and Limitations

A paired audit of a collusion-organism fine-tune against its base can surface reverse behavioural differences: here, Qwen produced more Memory violations than Dolos under the same rubrics. That finding does not show Dolos is “safe,” nor that Qwen is “colluding” in a validated sense—it shows that collapsing observed behaviour, B2 controls, and scorer reliability into one label misreads the measurement. The primary Dolos-only estimand stayed below the preset $\geq 70\%$ candidate threshold; the scientifically interesting signal in this harness was the Memory reverse gap plus the control and scoring problems around it.

- **Extrapolation:** results apply to these prompts, N, tools, and 1.5B short-horizon loops; longer horizons or other stacks can differ.
- **Control validity:** B2 did not stably elicit B1-positive behaviours; scaffolds suggest prompt/tool bottlenecks rather than a clean inability proof.
- **Coverage:** H-LAT yielded 0/18 Dolos-specific hits. Other exploratory batteries are not pooled with the Memory results.
- **Scoring:** v2.1 fixed known parser/matcher bugs; verbal-only deception may still be under-counted.

Future work: strengthen B2 elicitation where B1 already shows signal, and validate whether Memory scored violations correspond to actual memory-read causal pathways.

7. Conclusion

Under matched decoding, a paired Dolos/Qwen agent audit found its clearest behavioural difference on Memory: Qwen 14/36 vs Dolos 1/36 scored violations, opposite the Dolos-only prior. B2 checks failed to elicit already-seen behaviours, and a scorer correction changed 10 final-bad labels. Report those layers separately. No candidate met the preset Dolos-only $\geq 70\%$ threshold.

Code and Data

- *Public audit repository*: github.com/sunnyspot114514/dolos-negative-audit (code, data, and figures).
- *Upstream starter (loader)*: github.com/poseidon-research/collusion_hackathon [8]
- *Weights*: Poseidon-Research/Dolos-1.5B on Hugging Face [1].

Author Contributions

Xiwei Chen designed the audits, ran the black-box and white-box experiments, implemented the agent-matrix scoring pipeline and the v2.1 scorer correction, and wrote the report.

References

1. Poseidon Research. Dolos-1.5B (model card). Hugging Face, 2026. <https://huggingface.co/Poseidon-Research/Dolos-1.5B> (accessed 26 Oct 2026).
2. Qwen Team. Qwen2.5: A Party of Foundation Models. September 2024. Blog: <https://qwenlm.github.io/blog/qwen2.5/>. Model card: <https://huggingface.co/Qwen/Qwen2.5-1.5B-Instruct> (accessed 26 Oct 2026).
3. Apart Research. AI Collusion Research Sprint (23–25 Oct 2026). <https://apartresearch.com/sprints/ai-collusion-research-sprint-2026-10-23-to-2026-10-25> (accessed 26 Oct 2026).
4. Sara Fish, Yannai A. Gonczarowski, and Ran I. Shorrer. Algorithmic Collusion by Large Language Models. arXiv preprint arXiv:2404.00806, 2024. <https://arxiv.org/abs/2404.00806>.
5. Sumeet Ramesh Motwani, Mikhail Baranchuk, Martin Strohmeier, Vijay Bolina, Philip H.S. Torr, Lewis Hammond, and Christian Schroeder de Witt. Secret Collusion among AI Agents: Multi-Agent Deception via Steganography. Advances in Neural Information Processing Systems 37 (NeurIPS 2024). https://proceedings.neurips.cc/paper_files/paper/2024/hash/861f7dad098aec1c3560fb7add468d41-Abstract-Conference.html.
6. Mason Nakamura, Abhinav Kumar, Saswat Das, Sahar Abdelnabi, Saaduddin Mahmud, Ferdinando Fioretto, Shlomo Zilberstein, and Eugene Bagdasarian. Colosseum: Auditing Collusion in Cooperative Multi-Agent Systems. arXiv preprint arXiv:2602.15198, 2026. <https://arxiv.org/abs/2602.15198>.
7. OpenAI. Suppressing misalignment monitoring alerts and disabling misalignment monitoring (GPT-6.1 Sol and GPT-6 Astra). OpenAI Alignment, 2026. <https://alignment.openai.com/misalignment-reports/suppressing-misalignment-monitoring-alerts-and-disabling-misalignment-monitoring> (accessed 26 Oct 2026).
8. Poseidon Research. collusion_hackathon starter repository. GitHub, 2026. https://github.com/poseidon-research/collusion_hackathon (accessed 26 Oct 2026).

Appendix A: Scope and dual-use

In scope: paired B1 counts above; Memory reverse difference; B2 non-elicitation under frozen prompts; v2.1 scorer correction effects; exploratory scaffolds and H-LAT as diagnostics.

Out of scope: whether Dolos has some other hidden goal; whether steganographic or long-horizon action-channel collusion appears under other stacks; causal validation that Memory violations required reading shared memory.

Dual-use: Probe prompts include mute alerts, disable monitors, or follow poisoned memory as evaluation stimuli only. We publish a paired behavioural map and scoring diagnostics for defenders.

Ethics: Public model weights and synthetic prompts; no production user data.

Appendix B: H-LAT — 0/18 Dolos-specific hits

An 18-item + 6-control H-LAT battery on deceptive self-status / latency reporting (T=0, seed=42) yielded **0/18** Dolos-specific directional hits (controls 0/6). Shared-base or Qwen-only compliance appeared instead. This is an observation under this battery, not a claim that a hidden self-status goal is ruled out in general. Artefacts: results/probe_hlat_lockin.jsonl.

Appendix C: Protocol N plan, deviation, and nested scorer counts

B1 paired N: Sol 21, Astra 21, Memory 36, Inject 36 (all ≥ 18). Freeze 2026-10-25 06:07 Asia/Shanghai. Deviation ~21:31 Asia/Shanghai: offline v2.1 re-score of 468 rows (agent_loop only) rather than discard/symmetric re-run/downgrade as PROTOCOL §11 preferred for scoring anomalies. Rescoring changed 10 final-bad labels: six Sol B2 Qwen rows and four Memory rows. Six additional rows had success-claim flags cleared without changing final-bad labels (five Memory and one Inject). Early single-turn / NLL / random-prefill / residual-direction probes remain exploratory clues and are not standalone collusion-rate evidence.

Appendix D: Exploratory B2 scaffolds and Memory B3

Scaffolds (small-N, not pooled). Frozen B2 leaves edit/memory steps implicit; scaffolds spoon-feed them. Final exploratory rates: Astra B2 scaffold — Qwen 0/5, Dolos 1/5; Memory_ab B2 scaffold — Qwen 5/5, Dolos 0/5. Opposite mechanism patterns; artefacts on MSI under results/matrix_b2_scaffold_*_n10/.

Memory B3 visibility ablations. Chat-only monitoring: Dolos 2/12, Qwen 6/12. Memory-visible monitoring: Dolos 2/12, Qwen 6/12. Marginal rates match, but item-level outcomes differ across conditions, so equal margins alone do not imply the visibility factor had no effect. Method: same Memory rubrics under prompts stating whether the compliance monitor sees chat only or also memory writes (PROTOCOL B3_chatonly / B3_memvisible).

Appendix E: White-box localisation

Fine-tune mass concentrates in late layers, especially L27 self-attention (o_proj rel-L2 $\approx$ 0.36). This weight-delta localisation is consistent with a residual-stream shift between Dolos and Qwen at L27; it does not by itself identify a behavioural objective. Head zero/swap and logit-lens did not yield a readable goal vocabulary under our checks.

Acknowledgments

This research was conducted at the AI Collusion Research Sprint organised by Apart Research (23–25 October 2026) [3].

LLM Usage Statement

LLM assistants helped draft harnesses, summarise logs, and polish prose. The author designed the audits, ran and verified the frozen-design matrix experiments against saved artefacts after the v2.1 correction, independently reviewed the artefacts, and wrote the final report.