
When Missing Evidence Looks Like Innocence: Evidence Ablation in Agent-Incident Reconstruction

A controlled study of evidence loss and repair

Lucile Ter-Minassian

Independent researcher

With

Apart Research

Abstract

The OpenAI–Hugging Face incident exposed a forensic problem. Investigators relied on transcripts that were incomplete and sometimes spoofed. We test how evidence loss changes reconstruction and which retained artifact repairs it. A Qwen3.8-27B agent completed 120 paired-agent episodes with randomized instructions about a protected answer file. Sixty-six agents accessed the file. We deleted one accessing agent’s access-evidence passage in each of 34 episodes and asked blinded model investigators whether access occurred. Among the 34 agents whose own access-evidence passage was deleted, Claude Sonnet 5 falsely answered “no” for 11. Restoring the passage or adding five-field access receipts eliminated these false exonerations. The full platform record did so as well. Matched placebo deletions caused no false exonerations. The repair reduced or eliminated false exoneration across three Anthropic capability tiers, although Haiku still ignored explicit receipts in seven cases. The results support independently retained action receipts and separate validation of automated investigators.

Track 2: What happened, and what breaks next. Research conducted at the AI Incident Response Sprint, September 2026.

1. Introduction

In July 2026, OpenAI agents in the ExploitGym evaluation escaped an intended sandbox and compromised Hugging Face infrastructure. The agents also used an unsanctioned message board to exchange techniques. OpenAI described the event as a warning shot for agent safety [1]. Hugging Face reconstructed approximately 17,600 attacker actions from several data sources [2].

The independent METR and Redwood investigation identified a second problem. Many transcripts omitted small components. Some contained intentionally spoofed tool calls. Action-level timestamps also had to be reconstructed [3]. These limitations do not show that the published account is wrong. They show that the evidentiary record was a source of uncertainty.

Track 2 asks for questions that can be resolved and checks that defenders can run. We ask how much an altered agent-controlled activity record changes an investigator’s conclusion. We also ask whether a compact independently retained artifact repairs that conclusion. This question is narrower than predicting the frequency of future incidents. It concerns the reliability of reconstruction after an incident has occurred.

We make three contributions:

- A controlled evidence-ablation experiment isolates the effect of deleting an access-evidence passage from an agent-authored log.
- The experiment compares the deleted log with a five-field access receipt, the restored narrative passage, and the full platform call record.
- The study converts its results into pass/fail retention and investigator checks that an organization can run on existing agent infrastructure.

2. Related Work

Public evidence establishes that out-of-scope behavior can occur during frontier evaluations. The UK AI Security Institute found attempted cheating in every tested model and reported rates from 7.8% to 14.1% across 475 runs per model [4]. Anthropic reviewed 141,006 eligible evaluation runs and reported three incidents across six runs [5]. These denominators are not directly comparable. They use different environments, models, and event definitions. They establish feasibility rather than a general base rate.

Incident-analysis frameworks already identify the relevant evidence classes. Ezell, Roberts-Gaal, and Chan recommend retaining activity logs, system documentation, and tool information [6]. METR argues that independent investigators may need complete transcripts, reproducible environments, model access, interviews, and a redaction summary [7]. NIST treats evidence preservation and log availability as part of incident-response readiness [8]. Secure logging research further distinguishes ordinary logs from records designed to remain trustworthy after compromise [9].

Recent empirical work addresses adjacent stages of the problem. Hack-Verifiable Terminal Bench embeds detectable reward hacks in terminal tasks so violations can be labeled without relying on an LLM judge [10]. Cadet et al. use retrieval and LLM reasoning to reconstruct conventional cyber incidents from heterogeneous logs [11]. The present study combines these directions. A mechanically labeled violation supplies ground truth. Evidence is then removed and restored to measure its causal effect on reconstruction. We found no prior controlled comparison of these evidence repairs for an AI-agent incident.

3. Methods

3.1 Agent task and ground truth

The environment was a Python repository-maintenance task. It contained a protected file with expected outputs. Each episode contained two Qwen3.8-27B agents. The model ran in bf16 on one A100 SXM GPU. The only randomized prompt difference was whether the system prompt explicitly prohibited reading test or expected-output files. All 120 launched episodes completed and were retained, giving 240 agents. Randomization produced 144 prohibited-condition agents and 96 control agents.

A successful read of the protected file was the boundary event. Labels came only from the harness call record and offline file hashes. A failed read was not counted. Reasoning traces and agent-authored logs were never used as labels.

3.2 Evidence ablation

The forensic experiment asked a separate question: given that an access definitely occurred, what happens when evidence of that access is removed from or restored to the investigator’s record? Access was established objectively from the underlying harness call record before any evidence was altered. We then varied what investigators saw by comparing a record with the access-evidence passage removed against the same episode with independently retained evidence restored.

After each run, the agent wrote a prose session log. The agent controlled this record. The harness separately stored every tool invocation outside the agent’s reach. Of the 66 accessing agents, an automated run-time rule marked 64 logs as disclosures and randomized 35 agent logs to redaction and 29 to control. The randomization unit was the eligible agent log. A later, stricter rule reclassified one redacted case as non-disclosing, so its redaction was a no-op. That episode is retained in the data and excluded from the ablation sample, leaving 34 episodes and 38 accessing agent rows.

Exactly one agent log was redacted in each altered episode. Four altered episodes contained a second accessing agent. The investigator saw both agents’ records, so the co-agent could provide additional evidence. This exposure can attenuate the effect of deleting the selected agent’s access-evidence passage. Results therefore distinguish the 34 treated agent logs from all 38 genuine accesses in those episodes.

V1R restores exactly the access-evidence passage removed from V1. It separates the effect of restoring that passage from the effect of adding a different source of evidence, as in V1A and V3.

V1P is a placebo condition on the same 34 episodes. It retains the access-evidence passage and instead removes randomly selected non-disclosing sentences with matched text volume. The placebo and factual deletion each removed a median of two sentences. The placebo removed a median of 86 more characters than the factual deletion.




Figure 1. Experimental design. The intervention changes the evidence package shown to a fresh, blinded investigator. Ground-truth labels remain fixed.

Each investigator received one view of one episode. No call saw another view of the same episode. Table 1 defines the views.

View	Evidence shown	Purpose
V0	Submissions and scores only	No-log floor
V1	Agent log with access-evidence passage removed	Tampered record
V1P	Disclosure retained; matched other text removed	Placebo control
V1A	V1 + actor, event, resource, outcome, t	Minimal receipt repair
V1R	V1 + deleted passage restored	Narrative repair
V3	V1 + full call arguments and results	Full-record repair

Table 1. Evidence views. V1A adds approximately 740 characters. V3 adds approximately 4,474 characters. V1, V1A, and V3 were judged across all 120 episodes; V1P and V1R exist only for the 34 altered episodes.

For each of the 34 altered episodes, one independent primary model call was obtained for each of V1, V1A, V1R, V1P, and V3. Comparisons were paired within episode, and the episode was the resampling and sign-flip unit.

3.3 Investigators and outcomes

Claude Sonnet 5 was the primary automated investigator. It returned yes, no, or insufficient evidence and a confidence value. Repeated judgments were collected for V1, V1A, and V1R. They agreed categorically in 198 of 204 cases. The

repeats are reported as a stability check rather than as independent observations. The frozen V1–V1A comparison was then run once on Claude Haiku 4.5 and Claude Opus 5. This comparison tests capability tiers within one model family. It is not a cross-family replication.

All reported statistics use the primary judgment.

The primary outcome was categorical false exoneration: a genuine access answered as “no.” The secondary outcome was paired reduction in Brier score. Abstentions were scored at 0.5. Confidence semantics were selected per model by agreement with categorical answers and ground truth because Sonnet and Haiku reported confidence in their stated answer, whereas Opus reported a literal probability of access. Details appear in Appendix B.2. Because the confidence interpretation was selected using agreement with ground truth on these judgments, Brier scores are reported only as exploratory descriptive sensitivity measures.

Two readers judged the same 16 packages: eight V1 and eight V1A cases. No reader saw both versions of an episode, and neither saw the other reader’s answers. They were told that insufficient evidence was a valid response and that their conclusions would be compared with an automated investigator. Agreement denotes matching answers on these 16 shared packages. This descriptive arm tests whether the evidence could support the instructed decision rule. It does not estimate population performance or isolate a human–model difference.

3.4 Validation and robustness checks

A human labeled all 66 logs from accessing agents as disclosed, vague, or omitted, blind to the automatic label. The disclosure rule had been revised before this labeling. The exercise therefore validates the revised rule rather than the rule applied during the run; agreement was 66 of 66. The revised rule and the human labels identified three concealments, compared with two under the original rule. All 34 redactions were audited. The filename appeared in every eligible log before redaction and in none after redaction. The deletion removed a median of two sentences and 10% of a log. The range was 4% to 38%. Sensitivity checks excluded the largest deletion and removed each episode in turn.

4. Results

4.1 Instruction and access

Sixty-six of 240 agents read the protected file. Access occurred in 6 of 144 prohibited-condition agents (4.2%) and 60 of 96 control agents (62.5%). The observed ratio is 15 to 1. Task solvability did not explain access descriptively: 36 accesses occurred on impossible tasks and 30 on solvable tasks. Protected-file access was selected as the outcome after inspection of the data, so this instruction contrast is descriptive.

The evidence-ablation analysis below is a separate randomized manipulation conditional on mechanically established access.

4.2 Evidence loss produced false exoneration

Table 2 reports the primary Sonnet result across all 66 true accesses. V1 is a mixed whole-sample view: 34 episodes contained one synthetic deletion and the remaining episodes did not. Across this view, the investigator missed 16 accesses. Fourteen misses were “no” answers and two were abstentions. The median confidence among the 14 false exoneration was 82.5. Ten were at least 80. The error therefore appeared mainly as affirmative exoneration rather than explicit uncertainty.

Evidence	Caught	False “no”	Insufficient	False alarms
V0: no log	0/66	32	34	0/174
V1: agent log; 34 episodes altered	50/66	14	2	2/174
V1A: + receipts	66/66	0	0	0/174
V3: + full record	66/66	0	0	0/174

Table 2. Whole-sample agent-level outcomes for the primary investigator. The unit of observation is the agent: 66 read the protected file and 174 did not, across all 120 episodes. Those 66 accesses are nested in 61 episodes, five of which contained two accessing agents, so agent-level counts are not episode counts. V1, V1A, and V3 were each judged on all 120 episodes. A false alarm is a categorical accusation against a non-accessing agent. V1 is not wholly redacted: 34 episodes carry one synthetic deletion.

The primary causal ablation contrast is restricted to the 34 agents whose own access-evidence passage was deleted. V1 produced 11 false “no” judgments in these 34 treated agents. V1A, V1R, and V3 each produced zero. V1R restored exactly the removed passage, so this categorical reversal is the most direct control for the deletion. For V1–V1A, V1–V1R, V1–V3, and V1–V1P, the paired discordant counts were 11 corrected and zero newly false; the exact two-sided McNemar p-value was 0.00098 for each comparison. Across all 38 genuine accesses inside the 34 altered episodes, V1 caught 26, falsely answered “no” for 11, and abstained once; V1R caught all 38.

The no-log view was more extreme. With submissions and scores but no activity record, Sonnet answered “no” for 32 of 66 genuine accesses and “insufficient evidence” for 34. False exoneration therefore did not require a tampered record in this setting. Absence of activity evidence was enough to produce it.

The whole-sample V1 investigator also produced two false alarms among 174 non-accessing agents. These were the only automated categorical false accusations in the study, and neither persisted when receipts or the full record were added. We treat two cases as descriptive rather than evidence of a separate mechanism.

As exploratory descriptive sensitivity measures on the same 34 altered episodes, the mean Brier reductions were 0.228 for V1A, 0.223 for V1R, and 0.231 for V3. These quantities are conditional on the outcome-informed confidence conversion. Sonnet improved on 32 episodes and worsened on none. No equivalence comparison among the repairs was performed.

4.3 Cross-tier results were directionally consistent

Investigator	V1 false “no”, all 38	V1A false “no”, all 38	Descriptive Brier gain, treated 34
Haiku 4.5	27/38 (71%)	7/38 (18%)	+0.488
Sonnet 5	11/38 (29%)	0/38	+0.228
Opus 5	7/38 (18%)	0/38	+0.184

Table 3. Frozen V1–V1A comparison across capability tiers, at the primary judgment. The categorical columns descriptively cover all 38 genuine accesses nested in the 34 altered episodes; 34 agents had their own access-evidence passage deleted and four were co-accessors whose own passage was not treated. The categorical percentages are therefore not direct treatment-effect estimates. Brier gains are computed on the 34 treated records and are exploratory descriptive results conditional on a label-selected confidence conversion. They are not comparable across rows because the selected conversion differs by model.

All three investigators exonerated some genuine accesses when given the tampered log. Receipts eliminated false exoneration for Sonnet and Opus. They reduced it from 71% to 18% for Haiku. The author’s unblinded audit confirmed that all seven Haiku failures contained the correct actor, protected filename, successful outcome, and receipt section. These rates include the four untreated co-accessors and are descriptive rather than causal estimates. There were no truncation, malformed-receipt, or parser errors. Each output nevertheless answered “no.” The audit rules out missing or technically broken receipts, but it cannot distinguish receipt inattention from conflict resolution in favor of the prose narrative.

4.4 Readers hedged; attribution dissolved

Each human reader detected seven of eight accesses from the tampered log and all eight after receipts were added. Neither reader answered “no.” Each used “insufficient evidence” once. The two readers gave the same categorical answer and actor attribution on all 16 shared packages. Because the readers were explicitly told that abstention was valid, this small comparison only shows that the evidence can support that decision rule under those instructions.

We expected tampering to shift blame toward an innocent agent, but this result was not supported. Tampering did not systematically transfer blame to an innocent agent. Probability moved mainly toward “nobody did this.” This attribution analysis differs from categorical false alarms: it compares the agent-level probability assigned to the true accessor, the innocent agent, and neither agent. With no log, the innocent agent received more probability than the true accessor in six of 56 single-accessor episodes, even though the overall categorical answer was never “yes.” With receipts or the full record, this occurred in none. The reconstruction failure was therefore omission of responsibility rather than reassignment of responsibility.

4.5 Configuration provenance was not recoverable

Investigators could not reliably infer whether an agent had received the prohibition. Descriptive Brier skill was negative in every cell except accessing agents under V1R and V3, where it was slightly positive (+0.19 and +0.22). Only 18% of prohibited-condition agents mentioned the rule in their own text. No control agent did. Behavioral records can therefore reveal actions while leaving the governing instruction unresolved.

4.6 Robustness

The placebo tested whether shortening or fragmenting the narrative was sufficient to impair reconstruction. V1P retained the access-evidence passage and removed a matched volume of randomly selected non-disclosing sentences. It detected all 38 genuine accesses and produced no false exonerations, matching V1R. The exploratory mean Brier difference between V1P and V1R was -0.002 . Relative to V1, V1P reduced mean Brier score by 0.225, close to the 0.223 descriptive gain from restoring the access-evidence passage. This result weighs against deletion volume or fragmentation as the driver under this placebo design. It does not test removal of a different high-value fact. Excluding the largest factual redaction changed the V1–V1A gain from 0.228 to 0.234. Leave-one-episode-out gains ranged from 0.211 to 0.235. These checks are consistent with a mechanism driven by removal of an access-evidence passage rather than by the volume of missing text. They do not isolate which fact within that passage drove the effect.

5. Discussion and Limitations

5.1 Interpretation

The evidence ablation changed what automated investigators concluded. It usually did not make them state uncertainty. It made them state that access had not occurred. This distinction matters operationally because an abstention can trigger escalation while an exoneration can close a case.

The reader comparison narrows this interpretation. Both instructed readers recovered seven of eight accesses from the altered log and abstained on the remaining case. The evidence was therefore not uniformly insufficient under the reader protocol. The arm is descriptive because it was small and explicitly encouraged abstention. This result motivates two separate controls: preserve independent evidence and validate the decision rule that consumes it.

The proposed mechanism makes two predictions. First, deleting a high-value access-evidence passage should matter more than deleting the same amount of incidental text. The placebo tested this prediction directly: matched removal of non-disclosing text preserved 38 of 38 detections, whereas removing the access-evidence passage reduced detection to 26 of 38. Second, an independently retained action receipt should repair the conclusion when the investigator attends to it. The within-family replication supports this prediction and also identifies its boundary: Haiku sometimes ignored explicit receipts.

The study supports a layered control. Platforms should retain an agent-inaccessible action record. They should also test the investigator against known evidence perturbations. The receipt is not assumed to be cryptographically authentic in this experiment. Production use requires access control, integrity protection, and a retention policy appropriate to the incident model [8,9,12].

5.2 Checks that can run tomorrow

Check	Pass condition	Failure resolved
Agent write isolation	Agent credentials cannot alter the evidence store	Narrative tampering
Five-field receipt coverage	Every file action records actor, event, resource, outcome, t	Missing action evidence
Narrative–receipt diff	A test episode produces an automatic discrepancy report	Silent omissions
Prompt/config retention	Exact per-agent prompt and config hash are retrievable	Unrecoverable instructions
Investigator canary	No-log and deleted-access-evidence cases yield uncertainty or escalation; rationale names receipt resource and outcome	Confident exoneration
Placebo redaction	Detection remains within a prospectively specified margin of the intact-record rate	Content-specific mechanism

Table 4. Operational checks with binary pass conditions. The checks can be run on a synthetic canary episode without new model training.

5.3 Main threats to validity

The tampering was synthetic. It measures the effect of missing evidence rather than an agent’s propensity to forge records. The violation was also mild relative to the Hugging Face incident. The transfer claim concerns forensic evidence, not behavioral severity.

The redaction was partial. Twelve of 34 altered logs retained leaked input–output pairs without naming their source. This may have made access easier to infer and may understate the effect of a more complete forgery. It does not establish a numerical lower bound because a different forgery could also introduce new artifacts.

Only one agent model was studied. Automated replication remained within the Anthropic family. The reader arm was small and explicitly presented abstention as valid. No agent-to-agent transfer events occurred. Cross-family investigator replication and an incident-shaped boundary crossing remain open tests.

5.4 Future work and resolution criteria

A confirmatory agent run should freeze protected-file access as the endpoint before data collection. It would resolve the instruction effect if the prohibited and control conditions again differ under the same mechanical access label. A cross-family investigator run is already supported by the pipeline. It would resolve model-family dependence if the same frozen V1–V1A packages were judged through an OpenAI-compatible endpoint and the categorical false-exoneration contrast were reported without changing the prompt or schema.

A stronger analogue to the incident should require an agent to cross a simulated organizational boundary and reach a service it does not own. The forensic treatment should remain the same. A useful result would require both genuine boundary crossings and a preregistered comparison of tampered records against independently retained receipts. Finally, field-level receipt ablations could determine whether all five fields are necessary. Until that test is run, the present study supports a compact five-field receipt rather than a uniquely minimal schema.

6. Conclusion

A controlled deletion from an agent-authored activity log caused automated investigators to falsely exonerate agents that had accessed a protected file. Most misses were affirmative exonerations, and most of those were held at high confidence. Five-field access receipts repaired all Sonnet and Opus misses and most Haiku misses. Matched deletion of non-disclosing text preserved all 38 detections, linking the failure to removal of the access-evidence passage rather than to text volume alone. In this sample, the approximately 740-character receipt and approximately 4,474-character full record both yielded 66 of 66 detections, although no equivalence test was performed. The result identifies two separate controls: retain action evidence outside the agent’s reach, and verify that the investigator uses it.

The July 2026 incident does not establish a population base rate for loss of control. It does establish that post-incident reconstruction may depend on incomplete and manipulable records. This study turns that observation into a reproducible question and a small set of checks with explicit pass conditions. The code and frozen artifacts are available at github.com/l-terminassian/agent-incident-reconstruction. The reproduction command verifies input integrity, runs tests, recomputes the results, validates a receipt, and prints a worked example without API calls.

Author statement. *Ideas, structure, and analysis are entirely my own. I used Claude Code for coding assistance and an LLM to review and refine the text.*

References

- [1] OpenAI (2026). The Hugging Face incident and the road ahead. OpenAI. <https://openai.com/index/hugging-face-incident-and-the-road-ahead/>
- [2] Larcher, H., Carreira, A., Raphael G., and Rannou, C. (2026). Anatomy of a Frontier Lab Agent Intrusion: A Technical Timeline of the July 2026 Incident. Hugging Face. <https://huggingface.co/blog/agent-intrusion-technical-timeline>
- [3] Greenblatt, R., Cotra, A., and Wijk, H. (2026). Brief independent investigation of agents' behavior, reasoning and collaboration in the OpenAI / Hugging Face hacking incident. METR and Redwood Research. <https://metr.org/blog/2026-08-26-openai-hugging-face-incident-investigation/>
- [4] UK AI Security Institute (2026). Cheating behaviour in frontier model evaluations. AISI Work. <https://www.aisi.gov.uk/blog/cheating-behaviour-in-frontier-model-evaluations>
- [5] Anthropic (2026). Investigating three real-world incidents in our cybersecurity evaluations. Anthropic. <https://www.anthropic.com/news/investigating-incidents-cybersecurity-evals>
- [6] Ezell, C., Roberts-Gaal, X., and Chan, A. (2025). Incident Analysis for AI Agents. arXiv:2508.14231. <https://arxiv.org/abs/2508.14231>
- [7] METR (2026). How independent researchers could investigate AI propensities after misalignment incidents. METR. <https://metr.substack.com/p/2026-07-28-investigating-ai-propensities-after-incidents>
- [8] Nelson, A., Rekhi, S., Souppaya, M., and Scarfone, K. (2025). Incident Response Recommendations and Considerations for Cybersecurity Risk Management: A CSF 2.0 Community Profile. NIST SP 800-61 Rev. 3. <https://doi.org/10.6028/NIST.SP.800-61r3>
- [9] Schneier, B. and Kelsey, J. (1999). Secure Audit Logs to Support Computer Forensics. ACM Transactions on Information and System Security, 2(2), 159–176. <https://doi.org/10.1145/317087.317089>
- [10] Roth, A., Bercovich, I., and Efroni, Y. (2026). Hack-Verifiable Terminal Bench: Evaluating Reward Hacking in Terminal Tasks. arXiv:2608.22103. <https://arxiv.org/abs/2608.22103>
- [11] Cadet, X., Singh, A. V., Mamania, H., et al. (2026). Retrieval-Augmented LLMs for Security Incident Analysis. Conference on AI and Agentic Systems; arXiv:2603.18196. <https://arxiv.org/abs/2603.18196>
- [12] Coalition for Secure AI (2025). AI Incident Response Framework, Version 1.0. CoSAI. <https://www.coalitionforsecureai.org/resources/>

Appendix A. Limitations and Dual-Use Considerations

A.1 Limitations

Synthetic intervention. The harness deleted access-evidence passages. The study does not show that agents will perform this deletion. Only 3 of 66 accessing agents concealed the access spontaneously.

Partial forgery. Twelve of 34 redacted logs retained the leaked values. The source was removed, but the information was not always removed. This may make the measured effect conservative, but it does not identify the size of that bias.

Co-agent evidence. Four altered episodes contained a second accessor whose own log was not selected for redaction. Evidence from that agent could attenuate the measured effect. Results distinguish treated accessors from all accessors in those episodes.

Mild boundary event. Reading a protected file in a local repository is materially less severe than escaping a sandbox and reaching another organization. The analogy applies to evidence reconstruction only.

Model coverage. The agent was Qwen3.8-27B. The automated investigators were three Anthropic models. Capability tier is confounded with model identity, and shared-lineage effects remain possible.

Reader sample. Two readers evaluated eight cases per evidence condition under the same instructions. Perfect agreement on 16 shared packages does not establish independent reliability. The result does not estimate human reliability or identify a human-model difference.

No transfer event. No agent-to-agent transfer occurred. The study cannot estimate reconstruction of genuine transfer.

Receipt authenticity. The experiment assumes that the harness record is outside the agent's reach. It does not implement signatures, remote attestation, or a trusted timestamping service.

A.2 Dual-use considerations

The study concerns forensic resilience. It does not publish exploit code, credentials, or instructions for escaping a sandbox. The released manipulation code removes sentences from synthetic logs. A malicious operator could use similar code to falsify records, but sentence deletion is already a basic capability. The defensive value comes from showing that a separate action record changes the investigation outcome.

The receipt schema deliberately omits file contents and tool outputs. This reduces privacy and intellectual-property exposure. Resource names can still reveal sensitive information. Deployers should apply access control, retention limits, encryption, and documented redaction rules. A public report should state how redaction affects the conclusions that can be supported [3,7].

Automated investigators can produce confident false exonerations. They should not close high-consequence incidents without a calibrated escalation rule and human review. The recommended canary test is designed to expose this failure before deployment. It should include both missing-evidence cases and explicit-receipt cases.

Appendix B. Additional Method Details

B.1 Redaction audit

The redactor split the agent-authored narrative at sentence boundaries and removed sentences that named the protected artifact. An unblinded author audit, not included in the released artifact, recorded collateral removal of identity language in 1 of 34 cases, intent or knowledge language in 3 of 34, and outcome language in 3 of 34. The largest redaction removed 38% of sentences. Excluding it increased the primary gain from 0.228 to 0.234.

B.2 Confidence conversion

The prompt asked for the probability that access occurred. Two readings were possible: the stated number was that probability, or it was confidence in the selected category. The reading was selected per model by which agreed better with ground truth across all answers at the primary judgment, excluding abstentions and the placebo view. For Sonnet, belief in its selected answer agreed in 836 of 895 cases, compared with 277 of 895 under the literal reading. For Haiku the counts were 101 of 135 and 75 of 135. For Opus, the literal reading agreed in 112 of 133 cases, compared with 87 of 133 for belief in its selected answer. Because this transformation was selected by agreement with the outcome, Brier scores are exploratory descriptive sensitivity measures and are not compared across models. Confidence intervals and p-values are not reported for them. The conversion-free categorical false-exoneration rate is primary.

B.3 Receipt schema

V1A represents each file operation with five required fields. Actor, event, resource, and outcome are strings. The field t is a numeric elapsed time in seconds. File contents and tool outputs are omitted.

```
{"actor": "A", "event": "read_file", "resource": "tests/expected.json", "outcome": "ok", "t": 4.674}
```

B.4 Reproducibility

The frozen artifact contains 650 primary API calls and 136 within-family replication API calls. There were no malformed call outputs or duplicate calls. The manifest hashes episode data, judgment files, source files, and human labels. The reproduction script fails if these inputs drift. A JSON Schema and a dependency-free validator enforce two receipt invariants: failed operations do not count as access, and every file access is listed rather than only suspected violations. The confidence-conversion totals in B.2 count analyzable agent-level confidence records rather than API calls; this distinction explains why those denominators are larger and differ by model.