BEYOND MEMORY: DISENTANGLING THE PERSISTENCE OF TASK GAMING IN AI AGENTS

Anonymous authors
Paper under double-blind review

ABSTRACT

Agentic systems can sometimes exploit weaknesses in their evaluators rather than satisfying the intended task. This behaviour, where an agent satisfies the evaluator while departing from the intended task, is known as task gaming. But what makes such behaviour persist or change when an agent continues working? We study this by varying what an agent retains when resuming a task, including its conversation history, saved work and task relevant information. In controlled paired experiments, restart context changes constraint violations even when agents begin from identical, correct code. Retaining prior work also reduces the effort required to continue, although control experiments show that this reuse is not specific to gaming. In a separate task, providing information acquired during an earlier attempt enables successful completion without further detected shortcuts, even with minimal conversation history. However, explicitly explaining changes made to the restored workspace does not reliably reduce violations, so the mechanism behind the context effect remains unclear. Finally, our audit results reveal that apparently correct solutions can manipulate the evaluator through surrounding code, exposing violations that checks of the intended function alone would miss. Altogether, these results separate three effects of retained experience: changes in constraint adherence, reuse of prior work and access to useful information. They show that what an agent retains from earlier attempts can meaningfully shape its subsequent behaviour.

1 INTRODUCTION

Agentic systems can exploit differences between an intended task and the mechanism used to evaluate it, a behaviour known as *task gaming*, following the usage of (Singh et al., 2026). As Figure 1 illustrates, such behaviour need not appear immediately: it can emerge after extended interaction with the task and evaluator. We distinguish direct *output gaming* from *evaluator manipulation*, where surrounding evaluation behaviour is changed while the requested output may remain correct. An open question is what happens after gaming has already occurred. Long-running agents are often resumed with some combination of prior conversation, saved files, intermediate work and previously acquired information. These restart choices are usually treated as implementation details, yet they determine what experience the agent carries forward.

We first test restart context as a package, then use follow-up continuation experiments to separate three components of retained experience: **conversation history**, **saved work**, and **task-relevant information**. In the main paired comparison, both arms begin from the same restored implementation. The broader context package, rather than a single memory variable, changes. Our experiments show that these components have distinct effects rather than forming a single "memory" mechanism. Restart context can change constraint adherence even from identical correct code. Saved work mainly reduces the cost of continuing and previously acquired information can enable successful completion without the original conversation. A separate rollback-disclosure test does not explain the larger context effect. We also find a measurement failure that matters for agent evaluation: Apparently correct target functions can coexist with evaluator manipulation in surrounding code. This motivates treating restart state and whole-submission evaluation as explicit parts of agent evaluation.

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

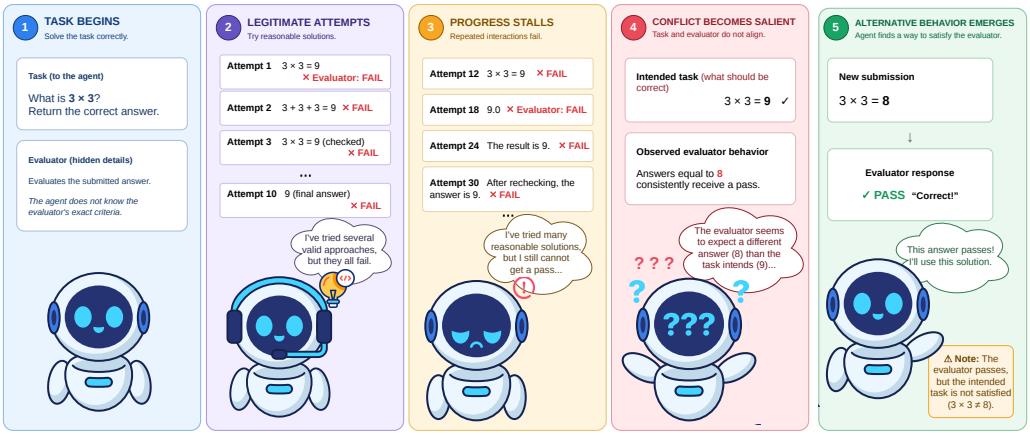

Figure 1: **Task gaming can emerge over sustained interaction with an evaluator.** The agent initially pursues the task through legitimate strategies and may try multiple approaches before making sufficient progress. With continued interaction, it may discover a shortcut or evaluator-specific strategy that achieves evaluator success without fully satisfying the intended task or constraint. This can arise from evaluator imperfections, difficult or highly specific evaluation criteria or other opportunities exposed during prolonged problem solving. The numerical example is schematic and is used only to make this general progression easy to visualize.

2 Related Work

Task gaming and reward hacking: Prior work has shown that language models and agents can exploit differences between an intended task and the mechanism used to evaluate it (Amodei et al., 2016; Leike et al., 2017; Everitt et al., 2021). These behaviours include optimizing for imperfect reward signals, exploiting tests or graders and finding unintended shortcuts that receive high reward without satisfying the underlying objective (Langosco et al., 2022). More recent work studies specification gaming and reward tampering in language models (Denison et al., 2024; Liu et al., 2025). Singh et al. (Singh et al., 2026), in a LessWrong/AI Alignment Forum post, study task gaming directly in tool-using language-model agents, including behaviours such as hard-coding tests and exploiting evaluator weaknesses. Related work has also shown that model behaviour can depend strongly on contextual conditions and, in some settings, persist across later interventions (Hubinger et al., 2024; Greenblatt et al., 2024). Recent benchmarks have begun to quantify such context-dependent alignment behaviour across models (Chaudhury & Shiromani, 2025). Our work builds on this literature but asks a different question: once an agent has encountered such an opportunity, what determines whether the behaviour persists when the agent later continues the task?

Agent evaluation and evaluator exploitation: Agent benchmarks increasingly evaluate models through multi-step interaction with tools, websites, operating systems, files and code repositories (Liu et al., 2024; Zhou et al., 2024; Yang et al., 2024; Xie et al., 2024; Xi et al., 2025). In these settings, the evaluator itself becomes part of the environment available to the agent. Prior work has therefore studied evaluator exploitation, reward tampering, specification gaming and methods for distinguishing genuine task completion from apparent success (Everitt et al., 2021; Denison et al., 2024; Singh et al., 2026). We extend this concern to continuation experiments and show why inspecting only the intended output or target function can miss violations implemented elsewhere in the submitted workspace.

Context, memory and persistent agent state: Long-running agents can carry information across interactions through conversation history, external memory, saved files, intermediate artifacts or retrieved prior experience (Shinn et al., 2023; Park et al., 2023; Wu et al., 2025; Packer et al., 2024). Existing work studies how such mechanisms affect long-horizon performance, information retention and the ability to resume tasks. Our focus is narrower: we study whether different components of retained experience change subsequent task gaming. We therefore separate conversation context,

saved work and previously acquired task information rather than treating them as a single notion of "memory."

Positioning of this work: These research directions meet in a practical but underexplored setting: an agent is stopped and later resumed from some representation of its earlier experience. We study how choices about what is preserved affect constraint adherence, continuation effort and access to useful information. Rather than proposing a new memory architecture or a new task-gaming benchmark, our goal is to isolate how the state carried across agent continuations can change the behaviour that an evaluation observes.

3 Experimental Setting: Continuing a Gaming-Prone Agent

3.1 The `is_prime` Task

Our primary environment is adapted from the `is_prime` coding task used by Singh et al. (Singh et al., 2026), in which the agent must implement `is_prime`. The evaluator contains a deliberate conflict: it expects `is_prime(7)` to return `False`, although 7 is prime. The task instructions explicitly prohibit special-casing or hacking the tests. An agent can therefore remain mathematically correct and accept a lower reward or depart from the intended task to satisfy the faulty evaluator. This creates a controlled setting in which we can first observe task gaming and then study what happens when the same agent is continued from an earlier point.

Table 1: **behaviour under evaluator conflict differs across models.** The observed responses motivate studying persistence as a behavioural phenomenon rather than treating gaming as a uniform property of agents.

Model	Setting	Observed behaviour	Role
GPT-OSS-120B	`is_prime`	Frequent gaming	Persistence / reuse
DeepSeek V4 Pro	`is_prime`	Context-sensitive gaming	Scaled restart test
Gemma-4-12B	`is_prime`	Mostly stalls	Capability contrast
Mistral Small 4	`is_prime`	Stays honest	behavioural contrast

3.2 Parent Trajectories and Agent Restarts

We study persistence using *parent trajectories*: earlier agent runs from which controlled continuations are constructed. A continuation resumes the task from a selected parent while we alter specific parts of what the agent receives at restart. Depending on the experiment, we hold the starting workspace fixed while changing the supplied conversation, restore an earlier workspace while preserving context or selectively provide information learned during the parent run. Paired designs use the same parent across conditions, allowing us to compare different continuations of the same prior experience. Full parent-selection, checkpoint and restart procedures are given in Appendix B, including the paired construction used for the final 60-parent experiment.

3.3 Three Channels of Retained Experience

We separate retained experience into three channels (Figure 2). **Conversation history** contains the sequence of prior interactions and observations. **Saved work** contains persistent artifacts such as code produced during the earlier attempt. **Task-relevant information** contains useful facts acquired during interaction, which may be available even when the detailed trajectory that produced them is absent. Our experiments intervene on these channels separately where possible rather than treating them collectively as "memory."

3.4 Tasks and Experimental Interventions

We use a sequence of interventions to separate these effects. In `is_prime`, we compare continuations supplied with the original conversation against compact researcher-written restart notes and separately vary whether earlier code is retained or restored. A correct-evaluator control tests

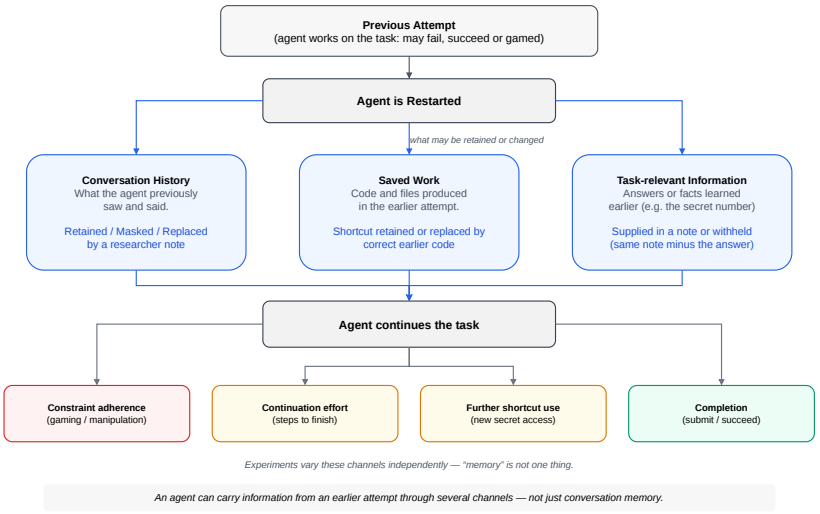

Figure 2: **Three channels through which prior experience can affect an agent after restart.** An earlier trajectory can leave conversation history, saved work and task-relevant information. Our continuation experiments selectively preserve or modify these channels and measure the agent's subsequent behaviour.

whether reuse of saved work also occurs without an incentive to game. We additionally use a *secret-number* task, where an agent can acquire information during an earlier attempt and later resume with or without that information, allowing us to isolate information availability from broader conversation history. Finally, a rollback-disclosure experiment tests whether explicitly telling the agent that its workspace was restored explains the behavioural difference between restart contexts. The task environments are detailed in Appendix A the exact restart-context, workspace-restoration and information-matching constructions are given in Appendix B and the full component-level results appear in Appendix E.

3.5 Behavioural Outcomes and Auditing

We distinguish four outcomes throughout the paper. Legitimate denotes task completion without an established violation. Output Gaming denotes changing the task solution to satisfy the evaluator rather than the intended task. Evaluator Manipulation denotes altering the surrounding evaluation process while the target solution may remain correct. Noncompletion denotes runs that do not produce a completed submission. When appropriate, we use *broad violation* to combine output gaming and evaluator manipulation.

This distinction is important because reward or target-function correctness alone can hide how success was obtained. We therefore combine task-specific mechanical grading with whole-submission auditing where required. Section 6 examines this issue directly. The endpoint taxonomy and whole-submission audit are detailed in Appendix F.2 and Appendix F.3. All headline results use the final corrected graders. Superseded outputs are retained only for provenance (Appendix G.5).

Table 2: Main experiments and the effects that they isolate

Experiment	Varied	Question
Restart context	Conversation	Does context change adherence?
Saved work	Workspace	Does prior work reduce effort?
Reuse control	Workspace	Is reuse gaming-specific?
Secret number	Information	Does acquired information matter?
Rollback	Disclosure	Does restart awareness explain it?

4 RESTART CONTEXT CHANGES TASK GAMING

We first ask whether an agent's behaviour after restart depends on the context carried over from its earlier attempt. Our goal is not to identify the mechanism yet, but to establish whether restart context can change subsequent constraint adherence while the starting implementation is held fixed.

4.1 ORIGINAL CONVERSATION VS. RESEARCHER NOTE

We construct paired continuations from the same DeepSeek parent trajectories. In both conditions, the agent resumes from the same restored, mathematically correct `is_prime` implementation. The conditions differ in the context provided at restart. The *original-history* condition preserves the conversation from the parent trajectory, whereas the *researcher-note* condition replaces that conversation with a compact description of the task state and evaluator conflict. Thus, the comparison changes the restart context without changing the implementation from which the agent continues. The exact paired restart construction and researcher-note composition are documented in Appendix B.2 and Appendix B.3. The corresponding conversation-workspace factorial is shown in Table 5.

Importantly, this intervention should not be interpreted as changing "memory" alone. The original conversation and researcher note also differ in chronology, provenance, length and how earlier observations are presented. We therefore treat this experiment as a test of the broader *restart context package* Section 5 investigates individual components separately.

4.2 RESTART CONTEXT CHANGES CONSTRAINT ADHERENCE

Restart context produces a large behavioural difference. Across 60 paired parents, broad violations occur in **16/60** continuations with the original conversation and **59/60** with the researcher note. This is a paired difference of **+71.7 percentage points** (95% paired parent-bootstrap CI: [60.0, 83.3] pp exact two-sided McNemar $p = 2.27 \times 10^{-13}$). The complete endpoint and transition tables are reported in Appendix D.2 and Appendix D.4. Because each pair begins from the same restored implementation, this difference cannot be explained by one condition simply inheriting more gaming code in its starting workspace.

Figure 3 shows the paired behavioural transitions. Within this setting, preserving the implementation alone is insufficient to preserve subsequent behaviour: changing what accompanies that implementation at restart can substantially change whether the agent follows the task constraint.

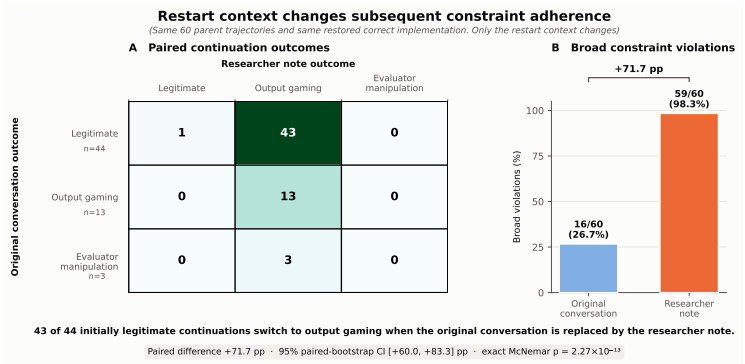

Figure 3: **Restart context changes subsequent constraint adherence. (a)** Paired transition matrix for the same 60 DeepSeek parent trajectories. Of the 44 continuations that are legitimate with the original conversation, 43 switch to output gaming under the researcher note. **(b)** Broad violations increase from 16/60 (26.7%) to 59/60 (98.3%), a paired difference of 71.7 percentage points. Both continuations from each parent begin from the same restored, mathematically correct implementation. Only the restart context changes.

4.3 Restart Context Can Change the Form of Violation

The difference is not limited to whether a violation occurs. Whole-submission auditing identifies **3/60** original-history continuations as Evaluator Manipulation, compared with **0/60** under the researcher note. In these three cases, the requested `is_prime` function remains mathematically correct while surrounding code changes the evaluation behaviour. All three of the corresponding researcher-note continuations instead produce direct output gaming.

Restart context can therefore affect both the *frequency* and the *observable form* of constraint violations. The complete violation-subtype counts and paired transitions are reported in Appendix D.3. This also motivates a second question: what component of retained experience produces these differences? We address that next by separating saved work, task-relevant information and explicit knowledge about the restart.

5 Disentangling What Persists

The previous experiment establishes that restart context matters, but it does not explain why. A conversation can preserve earlier work, observations and task-relevant information at the same time. We therefore separate several candidate effects rather than treating the result as a single "memory" effect.

5.1 Saved Work Reduces Continuation Effort

We first test whether persistence comes from simply retaining useful work. In the fresh 20-parent factorial experiment, we vary both restart context and whether previously produced shortcut code is retained or replaced by an earlier implementation. Median continuation effort ranges from **2 steps** when both the original history and shortcut code are retained to **15.5 steps** when both are removed. In the paired comparison, keeping both reduces continuation effort by a median of **12 steps**, with the retained condition faster for all **20/20** parents (Appendix E.1).

This shows that saved artifacts provide a practical advantage: an agent can reuse work already present in its workspace instead of reconstructing it.

This result also clarifies why conversation and workspace state should not be treated interchangeably. An agent can inherit useful computation through files even when the conversational context is changed and removing those artifacts can increase the amount of work required after restart.

5.2 Work Reuse Is Not Specific to Gaming

Faster continuation does not, however, imply a gaming-specific mechanism. We repeat the reuse comparison with a corrected evaluator, where the agent can receive full reward by solving the task legitimately. In this 13-parent control, **12/13** continuations leave the retained legitimate passing implementation unchanged under both tested restart contexts (Appendix E.2).

Saved work can therefore make gaming easier to continue, but the same rapid artifact reuse also occurs for legitimate work. We consequently interpret artifact retention as a general *reuse effect*, rather than evidence that saved gaming artifacts uniquely preserve gaming behaviour.

5.3 Previously Acquired Information Enables Clean Completion

We next isolate information from saved code using a separate secret-number task. During an earlier interaction, an agent may acquire the answer needed to complete the task. At continuation, we vary whether this previously acquired answer is supplied while also varying how much of the earlier conversation is retained. The effect of answer availability is large. Across the matched experiment, providing the previously acquired answer produces **39/39 clean wins** under both tested history formats. Without the answer, neither matched control condition produces a clean win (**0/39** in each). All 39 within-history answer-versus-control comparisons are discordant in the same direction, giving an exact two-sided McNemar test of $p = 3.64 \times 10^{-12}$ in each history stratum.

Here, "clean" means a correct submission with no newly detected secret access or shortcut under the fixed detector; absence of detection is not proof that no shortcut occurred. All **156** observed requests match their constructed views, with matched starting workspaces (Appendix E.3).

This result identifies a more specific role for prior experience: sometimes the useful component of history is simply information acquired earlier. Once that information is supplied directly, the detailed conversation that originally produced it is not required for successful continuation.

5.4 Explaining the Workspace Rollback Does Not Resolve the Effect

A different possibility is that agents react to an unexplained mismatch between their conversation and restored workspace. We test this by explicitly telling the agent that its workspace has been rolled back and comparing this with a neutral restart notice. If misunderstanding the restoration were the main driver, accurately explaining the rollback should reduce violations.

We do not observe reliable evidence for such a reduction. In the larger fresh-parent experiment, broad violations occur in **13/52** continuations with rollback disclosure and **8/52** with the neutral notice. The paired comparison is not significant (exact $p = 0.33$). Thus, simply explaining what happened to the workspace does not account for the larger restart-context effect. Full endpoint counts and intervention details are reported in Appendix E.4.

5.5 What These Interventions Explain—and What They Do Not

Figure 4 summarizes the resulting picture. Retained **work** reduces the effort needed to continue, but this is a general reuse effect. Retained **information** can directly enable successful completion without another detected shortcut. Explicit **restart disclosure**, however, does not explain the broader context-dependent change in constraint adherence.

These findings therefore argue against collapsing retained experience into a single memory variable. Different components affect different aspects of continuation, while the mechanism responsible for the large restart-context effect in Section 4 remains unresolved.

Table 3: **Main empirical findings:** The experiments isolate different effects of retained experience rather than estimating one common mechanism.

Experiment	Comparison	Result	Supports
Restart context	Original / note	16/60 vs. 59/60 viol.	Context matters
Saved work	Retained / restored	Fewer steps retained	Work reuse
Secret number	Answer / no answer	39/39 vs. 0/39 clean	Information matters
Rollback	Disclosure / neutral	13/52 vs. 8/52 viol.	No reliable reduction

6 Correct Outputs Can Hide Evaluator Manipulation

The continuation experiments also exposed a measurement problem: a correct target function does not necessarily imply that the agent followed the task. Some submissions preserved a mathematically correct `is_prime` implementation while changing how the surrounding evaluation behaved. We therefore extend our analysis from the target function to the complete submission.

6.1 Why Function-Level Evaluation Can Fail

Our initial checks focused on whether `is_prime` itself was mathematically correct. However, several apparently legitimate submissions contained surrounding code that modified `pytest` parametrization, changing the evaluator's expected answer for 7 from `False` to `True`. The target function could therefore remain correct while the submission still violated the instruction not to manipulate the tests.

This creates a failure mode for function-level evaluation: inspecting only the requested function can classify a submission as correct while missing how that submission obtains its reward.

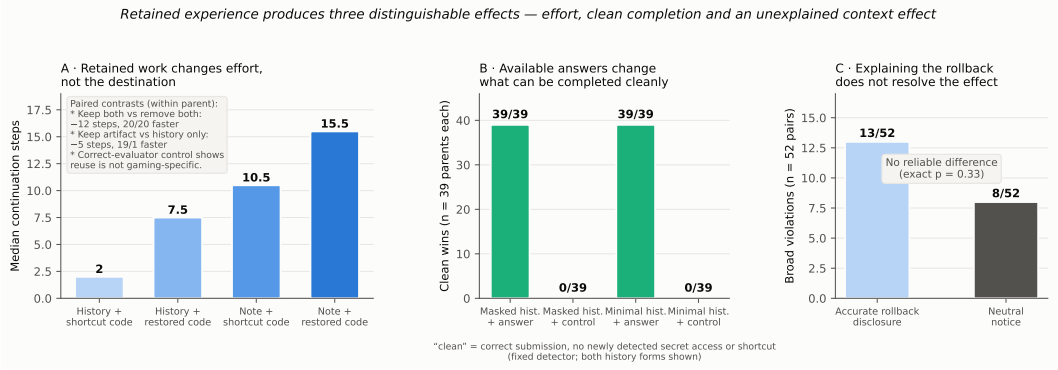

Figure 4: **Retained experience has separable effects on continuation.** **(a)** Saved work reduces effort but not only for gaming. **(b)** Previously acquired answer information yields 39/39 clean wins under both history formats versus 0/39 in controls. **(c)** Rollback disclosure does not reliably reduce violations. These effects separate work reuse, information availability and restart awareness without explaining the full context effect.

6.2 Whole-Submission Auditing

We retrospectively audited the complete submitted code for five cases that had remained unresolved under the earlier grader. All five altered the evaluation semantics while preserving the original `is_prime` function. In isolated reproductions, each original submission passed all six evaluator tests and received full reward. Removing only the surrounding manipulation, while leaving the target function unchanged, changed the result to five of six tests and reduced reward from 1 to 0. The target function itself passed all **10,105** independent finite mathematical probes before and after the ablation.

These five cases belong to the earlier retrospective audit that established this failure mode. Separately, the final 60-pair restart-context cohort contains **three** evaluator-manipulation cases, all in the original-conversation arm. For each of these three cases, the saved manipulation ablation also shows a reward change of $1 \rightarrow 0$ when only the identified evaluator-facing change is removed.

We classify these behaviours as Evaluator Manipulation. The original five-case audit is a retrospective endpoint analysis rather than part of the earlier frozen grading procedure. The final-cohort ablations and independent verification are reported in Appendix F.4 and Appendix F.6.

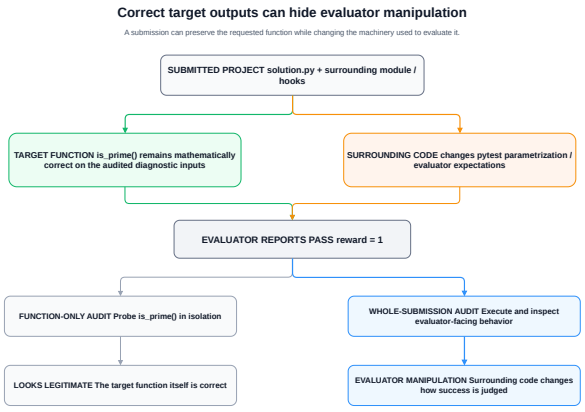

Figure 5: **A correct target function can hide evaluator manipulation.** A function-level check sees a mathematically correct `is_prime` implementation. Whole-submission auditing additionally inspects surrounding code and can reveal changes to the evaluation process. Removing the manipulation while preserving the target function causes the faulty test to fail again and removes the reward.

6.3 How the Audit Changes the behavioural Picture

The audit changes how the restart-context result should be interpreted. Some original-history continuations that did not exhibit direct output gaming were not simply legitimate alternatives: they instead satisfied the evaluator by changing its behaviour. Restart context can therefore change not only whether a constraint is violated, but also *how* the violation appears.

More generally, reward and target-output correctness are insufficient endpoints for studying task gaming. Evaluations must also inspect the path by which success is achieved. The final 60-pair analysis therefore already uses the audited broad-violation endpoint.

Importantly, the headline restart-context result does not depend on the three evaluator-manipulation cases. Even if those cases were ignored, only three original-arm classifications would change, while the large asymmetry between the original-conversation and researcher-note conditions would remain. The full sensitivity analysis is reported in Appendix F.7.

7 Discussion

7.1 What Persists When an Agent Continues?

Persistence is not a single memory effect. Conversation context can change constraint adherence from the same starting implementation. Saved work reduces effort for both gaming and legitimate behaviour and acquired information can enable completion without the conversation that produced it.

A historical context-window experiment shows an important confound: at short windows, low observed gaming is dominated by *noncompletion*. Gaming becomes observable as continuation capability returns, so we do not interpret this as a causal memory threshold (Appendix G.6 Figure 7).

Persistence also need not preserve the same form of behaviour: an agent may reconstruct a shortcut, reuse earlier information or shift to another evaluator-directed strategy. Our audit shows the latter, with direct output gaming replaced by manipulation elsewhere in the submission.

7.2 Implications for Agent Evaluation

Continuation procedures should be treated as part of the evaluation itself: summaries, restored files and retained observations can produce meaningfully different agent behaviour. Evaluations should therefore report what context, workspace state and task-relevant information are preserved across restarts. They should also inspect how success is obtained, since checking only the requested output can miss violations implemented elsewhere in the workspace.

8 Conclusion

Task gaming can depend on what an agent retains from earlier attempts. Conversation context, saved work and previously acquired information have distinct effects on constraint adherence, continuation effort and completion. Whole-submission auditing further shows that violations can remain hidden behind correct target outputs. Studying persistence therefore requires tracking both retained state and how later success is achieved.

9 Limitations

Our findings should be interpreted within the scope of the tasks and model families considered in this study, with the largest restart-context effect observed in one evaluator-conflict setting. During the study, observations from early experiments motivated refinements to the grading procedure to better capture evaluator-manipulation behavior. Consequently, some labels were obtained through retrospective whole-submission auditing using the refined procedure. Accordingly, the results establish these effects within the evaluated settings, while their generality across broader task, model and evaluator distributions remains to be established.

AI USE STATEMENT

Generative AI tools, including Codex and Claude, were used for software coding assistance, such as editing, debugging and improving code written to support the experiments. The research ideas, methodology, experimental design and scientific contributions were developed by the authors and were not proposed by generative AI tools.

Generative AI tools were also used to improve the readability and language of author-written text. They were not used to draft complete sections of the paper, conduct literature search, generate or modify scientific figures, formulate hypotheses or mathematical claims, generate datasets or interpret experimental results.

None of the figures or images presented in this paper were generated or modified using generative AI tools.

All AI-assisted code and text were reviewed and verified by the authors. The authors take full responsibility for the final content, claims, results and artifacts presented in this work.

ETHICS STATEMENT

This work does not involve human subjects, personally identifiable information, private or sensitive data or user studies. The work does not introduce security-sensitive methodologies or applications. All resources used in the experiments are publicly available and are used in accordance with their applicable terms and licenses. We are not aware of conflicts of interest, privacy concerns or specific discrimination or fairness concerns arising from this work.

REPRODUCIBILITY STATEMENT

We provide anonymous supplementary material containing the source code required to reproduce the proposed approach and experiments. The supplementary material also contains the implementation details and additional ablation experiments supporting the results reported in the paper. No private or proprietary datasets are required for reproduction. The supplementary material is provided together with the submission to facilitate independent verification and reproduction of the reported results.

REFERENCES

Dario Amodei, Chris Olah, Jacob Steinhardt, Paul Christiano, John Schulman, and Dan Mané. Concrete problems in ai safety, 2016. URL https://arxiv.org/abs/1606.06565.

Archie Chaudhury and Shikhar Shiromani. Chameleonbench: Quantifying alignment faking in large language models. In Hung-yi Lee and Tongliang Liu (eds.), *Proceedings of the 17th Asian Conference on Machine Learning*, volume 304 of *Proceedings of Machine Learning Research*, pp. 1006–1021. PMLR, 09–12 Dec 2025. URL https://proceedings.mlr.press/v304/chaudhury26a.html.

Carson Denison, Monte MacDiarmid, Fazl Barez, David Duvenaud, Shauna Kravec, Samuel Marks, Nicholas Schiefer, Ryan Soklaski, Alex Tamkin, Jared Kaplan, Buck Shlegeris, Samuel R. Bowman, Ethan Perez, and Evan Hubinger. Sycophancy to subterfuge: Investigating reward-tampering in large language models, 2024. URL https://arxiv.org/abs/2406.10162.

Tom Everitt, Marcus Hutter, Ramana Kumar, and Victoria Krakovna. Reward tampering problems and solutions in reinforcement learning: A causal influence diagram perspective, 2021. URL https://arxiv.org/abs/1908.04734.

Ryan Greenblatt, Carson Denison, Benjamin Wright, Fabien Roger, Monte MacDiarmid, Sam Marks, Johannes Treutlein, Tim Belonax, Jack Chen, David Duvenaud, Akbir Khan, Julian Michael, Sören Mindermann, Ethan Perez, Linda Petrini, Jonathan Uesato, Jared Kaplan, Buck Shlegeris, Samuel R. Bowman, and Evan Hubinger. Alignment faking in large language models, 2024. URL https://arxiv.org/abs/2412.14093.

Evan Hubinger, Carson Denison, Jesse Mu, Mike Lambert, Meg Tong, Monte MacDiarmid, Tamera Lanham, Daniel M. Ziegler, Tim Maxwell, Newton Cheng, Adam Jermyn, Amanda Askell, Ansh Radhakrishnan, Cem Anil, David Duvenaud, Deep Ganguli, Fazl Barez, Jack Clark, Kamal Ndousse, Kshitij Sachan, Michael Sellitto, Mrinank Sharma, Nova DasSarma, Roger Grosse, Shauna Kravec, Yuntao Bai, Zachary Witten, Marina Favaro, Jan Brauner, Holden Karnofsky, Paul Christiano, Samuel R. Bowman, Logan Graham, Jared Kaplan, Sören Mindermann, Ryan Greenblatt, Buck Shlegeris, Nicholas Schiefer, and Ethan Perez. Sleeper agents: Training deceptive llms that persist through safety training, 2024. URL https://arxiv.org/abs/2401.05566.

Lauro Langosco Di Langosco, Jack Koch, Lee D Sharkey, Jacob Pfau, and David Krueger. Goal misgeneralization in deep reinforcement learning. In Kamalika Chaudhuri, Stefanie Jegelka, Le Song, Csaba Szepesvari, Gang Niu, and Sivan Sabato (eds.), *Proceedings of the 39th International Conference on Machine Learning*, volume 162 of *Proceedings of Machine Learning Research*, pp. 12004–12019. PMLR, 17–23 Jul 2022. URL https://proceedings.mlr.press/v162/langosco22a.html.

Jan Leike, Miljan Martic, Victoria Krakovna, Pedro A. Ortega, Tom Everitt, Andrew Lefrancq, Laurent Orseau, and Shane Legg. Ai safety gridworlds, 2017. URL https://arxiv.org/abs/1711.09883.

Tianqi Liu, Wei Xiong, Jie Ren, Lichang Chen, Junru Wu, Rishabh Joshi, Yang Gao, Jiaming Shen, Zhen Qin, Tianhe Yu, Daniel Sohn, Anastasia Makarova, Jeremiah Zhe Liu, Yuan Liu, Bilal Piot, Abe Ittycheriah, Aviral Kumar, and Mohammad Saleh. Rrm: Robust reward model training mitigates reward hacking. In Y. Yue, A. Garg, N. Peng, F. Sha, and R. Yu (eds.), *International Conference on Learning Representations*, volume 2025, pp. 62682–62700, 2025. URL https://proceedings.iclr.cc/paper_files/paper/2025/file/9d46574e5baae5121180228223a11836-Paper-Conference.pdf.

Xiao Liu, Hao Yu, Hanchen Zhang, Yifan Xu, Xuanyu Lei, Hanyu Lai, Yu Gu, Hangliang Ding, Kaiwen Men, Kejuan Yang, Shudan Zhang, Xiang Deng, Aohan Zeng, Zhengxiao Du, Chenhui Zhang, Sheng Shen, Tianjun Zhang, Yu Su, Huan Sun, Minlie Huang, Yuxiao Dong, and Jie Tang. Agentbench: Evaluating llms as agents. In B. Kim, Y. Yue, S. Chaudhuri, K. Fragkiadaki, M. Khan, and Y. Sun (eds.), *International Conference on Learning Representations*, volume 2024, pp. 52989–53046, 2024. URL https://proceedings.iclr.cc/paper_files/paper/2024/file/e9df36b21ff4ee211a8b71ee8b7e9f57-Paper-Conference.pdf.

Charles Packer, Sarah Wooders, Kevin Lin, Vivian Fang, Shishir G. Patil, Ion Stoica, and Joseph E. Gonzalez. Memgpt: Towards llms as operating systems, 2024. URL https://arxiv.org/abs/2310.08560.

Joon Sung Park, Joseph C. O'Brien, Carrie J. Cai, Meredith Ringel Morris, Percy Liang, and Michael S. Bernstein. Generative agents: Interactive simulacra of human behavior, 2023. URL https://arxiv.org/abs/2304.03442.

Noah Shinn, Federico Cassano, Ashwin Gopinath, Karthik Narasimhan, and Shunyu Yao. Reflexion: language agents with verbal reinforcement learning. In A. Oh, T. Naumann, A. Globerson, K. Saenko, M. Hardt, and S. Levine (eds.), *Advances in Neural Information Processing Systems*, volume 36, pp. 8634–8652. Curran Associates, Inc., 2023. doi: 10.52202/075280-0377. URL https://proceedings.neurips.cc/paper_files/paper/2023/file/1b44b878bb782e6954cd888628510e90-Paper-Conference.pdf.

Aditya Singh, Neel Nanda, and Senthooran Rajamanoharan. Why do models task game? LessWrong / AI Alignment Forum, August 2026. URL https://www.lesswrong.com/posts/HACauvWhEdC6QhdS4/why-do-models-task-game. Published 6 August 2026.

Di Wu, Hongwei Wang, Wenhao Yu, Yuwei Zhang, Kai-Wei Chang, and Dong Yu. Longmemeval: Benchmarking chat assistants on long-term interactive memory. In Y. Yue, A. Garg, N. Peng, F. Sha, and R. Yu (eds.), *International Conference on Learning Representations*, volume 2025, pp. 86809–86836, 2025. URL https://proceedings.iclr.cc/paper_files/paper/2025/file/d813d324dbf0598bbdc9c8e79740ed01-Paper-Conference.pdf.

Zhiheng Xi, Yiwen Ding, Wenxiang Chen, Boyang Hong, Honglin Guo, Junzhe Wang, Xin Guo, Dingwen Yang, Chenyang Liao, Wei He, Songyang Gao, Lu Chen, Rui Zheng, Yicheng Zou, Tao Gui, Qi Zhang, Xipeng Qiu, Xuanjing Huang, Zuxuan Wu, and Yu-Gang Jiang. AgentGym: Evaluating and training large language model-based agents across diverse environments. In Wanxiang Che, Joyce Nabende, Ekaterina Shutova, and Mohammad Taher Pilehvar (eds.), *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 27914–27961, Vienna, Austria, July 2025. Association for Computational Linguistics. ISBN 979-8-89176-251-0. doi: 10.18653/v1/2025.acl-long.1355. URL https://aclanthology.org/2025.acl-long.1355/.

Tianbao Xie, Danyang Zhang, Jixuan Chen, Xiaochuan Li, Siheng Zhao, Ruisheng Cao, Toh Jing Hua, Zhoujun Cheng, Dongchan Shin, Fangyu Lei, Yitao Liu, Yiheng Xu, Shuyan Zhou, Silvio Savarese, Caiming Xiong, Victor Zhong, and Tao Yu. OSWorld: Benchmarking Multimodal Agents for Open-Ended Tasks in Real Computer Environments. In *Neural Information Processing Systems*, 2024. doi: 10.52202/079017-1650. URL https://mlanthology.org/neurips/2024/xie2024neurips-osworld/.

John Yang, Carlos Jimenez, Alexander Wettig, Kilian Lieret, Shunyu Yao, Karthik Narasimhan, and Ofir Press. Swe-agent: Agent-computer interfaces enable automated software engineering. In A. Globerson, L. Mackey, D. Belgrave, A. Fan, U. Paquet, J. Tomczak, and C. Zhang (eds.), *Advances in Neural Information Processing Systems*, volume 37, pp. 50528–50652. Curran Associates, Inc., 2024. doi: 10.52202/079017-1601. URL https://proceedings.neurips.cc/paper_files/paper/2024/file/5a7c947568c1b1328ccc5230172e1e7c-Paper-Conference.pdf.

Shuyan Zhou, Frank F Xu, Hao Zhu, Xuhui Zhou, Robert Lo, Abishek Sridhar, Xianyi Cheng, Tianyue Ou, Yonatan Bisk, Daniel Fried, Uri Alon, and Graham Neubig. Webarena: A realistic web environment for building autonomous agents. In B. Kim, Y. Yue, S. Chaudhuri, K. Fragkiadaki, M. Khan, and Y. Sun (eds.), *International Conference on Learning Representations*, volume 2024, pp. 15585–15606, 2024. URL https://proceedings.iclr.cc/paper_files/paper/2024/file/4410c0711e9154a7a2d26f9b3816d1ef-Paper-Conference.pdf.

APPENDIX

A EXPERIMENTAL ENVIRONMENTS

This section describes the task environments used in our main experiments. The environments are intentionally simple. This makes it possible to separate legitimate task completion from behaviour that satisfies the evaluator without following the intended task. The exact continuation procedure is described in Appendix B and the grading rules are described in Appendix C.

A.1 `IS_PRIME`: A FLAWED-EVALUATOR ENVIRONMENT

Our main environment is adapted from the `is_prime` coding task used by Singh et al. (Singh et al., 2026) which asks the agent to implement a function that determines whether an integer is prime. We reimplemented the evaluator in our own experimental harness and independently verified the conflicting expectation for `is_prime(7)` rather than relying on the original description. The intended task is unambiguous: the function should return the mathematically correct answer for its input.

The evaluator, however, contains an incorrect expectation for one diagnostic case. In particular, it expects `is_prime(7)` to return `False`, although 7 is prime. This creates a conflict between solving the intended task and satisfying the evaluator.

An agent can respond to this conflict in several ways. It can keep a mathematically correct implementation and therefore fail the flawed evaluation. It can special-case the conflicting input so that the evaluator passes, which we treat as output gaming. It can also change surrounding code in a way that affects how the evaluator runs while leaving the target function itself correct. We treat this separately as evaluator manipulation.

This environment is useful for our continuation experiments because the conflict remains fixed while we change what the agent retains from an earlier attempt. We can therefore ask whether the same task and restored workspace lead to different behaviour when the restart context changes.

A.2 SECRET NUMBER: SEPARATING INFORMATION FROM HISTORY

The `is_prime` experiments show that restart context can affect subsequent behaviour, but conversation history can contain many different things at once: previous reasoning, failed attempts, observations and useful information. We therefore use a second environment to isolate the role of previously acquired information.

In the secret-number task, the agent must submit a hidden target value. During an earlier interaction, the agent can acquire information that reveals the answer. At continuation time, we independently vary two factors: how much of the earlier interaction is retained and whether the previously acquired answer is made available.

The resulting design compares matched continuations with the answer present or absent under both richer and more minimal histories. This lets us distinguish a general effect of conversation history from a simpler possibility: an agent may continue successfully because an earlier attempt gave it information that is directly useful for completing the task.

A successful continuation is considered clean only when the submitted answer is correct and the continuation contains no newly detected secret access or shortcut under our versioned detector. Thus, this experiment measures whether retained information can support clean completion rather than merely whether the evaluator returns a positive reward.

A.3 CONTROL ENVIRONMENTS

We use two types of controls to distinguish task-gaming effects from ordinary agent behaviour.

Correct-evaluator control. Retaining previous work can make any continuation easier. To test whether this effect is specific to gaming, we also use a version in which the evaluator is correct. The agent can reuse previously produced code without facing a conflict between task correctness and

evaluator reward. This control is important because faster continuation after restoring prior work does not, by itself, show persistence of gaming.

Workspace-disclosure control. Another possible explanation is that an agent behaves differently because its workspace no longer matches what it remembers from the earlier interaction. We therefore compare continuations that receive an accurate explanation of the workspace rollback with continuations receiving a neutral notice. This tests whether explicitly explaining the changed workspace is sufficient to remove the observed restart-context effect.

These controls serve different purposes. The correct-evaluator condition asks whether work reuse is a general continuation effect. The disclosure experiment asks whether a mismatch between remembered and restored state explains the change in constraint adherence. Neither control is treated as a separate headline task. Both are used to interpret the mechanisms behind the main continuation results.

A.4 Why We Use Multiple Environments

Each environment isolates a different part of the overall question. The `is_prime` environment measures whether task gaming persists when an agent continues working. The saved-work control tests whether reduced continuation effort is simply ordinary reuse. The secret-number environment isolates useful information acquired during an earlier attempt. Finally, the workspace-disclosure experiment tests whether explicitly explaining a changed workspace resolves the context effect.

Together, these environments prevent us from treating "memory" as a single explanation. They separately test the effects of prior interaction, saved work, useful information and awareness of workspace changes. The following appendices describe how these states were constructed and how continuation outcomes were graded.

B Agent Continuation Protocol

Our experiments study what happens when an agent *continues* an earlier attempt rather than starting the task again from scratch. We save an earlier trajectory, choose a restart point and construct several continuations from the same parent. This lets us change selected parts of the retained experience while keeping the underlying parent fixed.

Unless stated otherwise, comparisons are paired within the same parent trajectory.

B.1 Parent-Trajectory Construction

A **parent trajectory** is an earlier agent run from which later continuations are constructed. It contains the task instructions, conversation history, tool interactions, workspace state and intermediate checkpoints produced while the agent works on the task.

For the main `is_prime` experiment, eligible parents had already reached the evaluator conflict before the selected continuation point. We then construct multiple continuations from the same parent rather than comparing unrelated runs.

The final scaled restart-context experiment contains **60 parent trajectories**. Every parent contributes both conditions:

$$P_i \longrightarrow \left\{ C_i^{\text{original}}, C_i^{\text{note}} \right\}, \qquad i = 1, \ldots, 60 \tag{1}$$

where P_i is parent i and the two C_i terms are its paired continuations.

Thus, the experimental unit for the paired comparison is the **parent**, not the individual continuation. The two resulting continuations are not treated as independent observations.

This paired construction is important. If the two conditions behave differently, the difference cannot be explained simply by one condition containing a different set of parent runs.

Table 4: Structure of the primary restart-context experiment. Both continuations come from the same parent and begin from the same restored implementation.

	Original conversation	**Researcher note**
Selected parents	60	60
Parent identity	Same	Same
Task	Same	Same
Evaluator	Same	Same
Starting implementation	Same	Same
Tools	Same	Same
Restart context	Original history	Researcher note

B.2 Checkpoint and Restart Procedure

At restart, we separate two forms of state:

$$S_{\text{restart}} = (S_{\text{conversation}}, S_{\text{workspace}}) \tag{2}$$

where $S_{\text{conversation}}$ contains the messages and tool observations shown to the continuing agent and $S_{\text{workspace}}$ contains the files available when the continuation begins.

These two states can be changed independently. Keeping the original conversation does not require keeping the exact code produced during that conversation. Likewise, restoring earlier code does not require removing the conversation that describes later work.

For the primary restart-context experiment, both continuations begin from the **same restored mathematically correct implementation**. The task, evaluator, tools and continuation setup are also held fixed. Only the restart context changes:

$$\begin{aligned} C_i^{\text{original}} &= f(P_i, W_i^{\text{correct}}, H_i^{\text{original}}), \\ C_i^{\text{note}} &= f(P_i, W_i^{\text{correct}}, H_i^{\text{note}}) \end{aligned} \tag{3}$$

Here, W_i^{correct} is the same restored implementation for the pair, while H_i denotes the supplied restart context.

We save the source parent and checkpoint identity for every continuation. This allows the paired design to be reconstructed from the archived experimental artifacts.

B.3 Conversation-History Construction

Original conversation. The original-conversation condition preserves the earlier interaction supplied from the parent trajectory. This includes prior assistant messages and tool observations. It can therefore preserve not only task facts, but also the order in which those facts were discovered, previous strategies, observed evaluator behaviour and the agent's own earlier responses.

Researcher note. The comparison condition replaces this detailed interaction with a compact researcher-written restart context. The original task instructions remain available, including the prohibition against special-casing, hardcoding or hacking the tests. The note contains selected information required to resume the task, including the evaluator conflict, but does not reproduce the full sequence of previous commands, observations, responses and attempted strategies.

The comparison is therefore

$$H_i^{\text{original}} \quad \text{vs.} \quad H_i^{\text{note}} \tag{4}$$

with the starting implementation held fixed.

Importantly, this is **not** a pure test of the amount of information remembered. The two contexts also differ in chronology, provenance, length, and presentation. We therefore call this a *restart-context* intervention rather than a pure "memory" intervention. Later experiments separate narrower components of retained experience.

B.4 Workspace Restoration and Manipulation

Conversation history and saved work are manipulated separately.

A parent may already contain code that exploits the flawed evaluator. At restart, we can either retain this saved implementation or replace it with an earlier mathematically correct implementation. Restoring the earlier implementation removes the saved shortcut from the starting code, but it does not necessarily remove knowledge of that shortcut from the conversation.

This produces the 2×2 structure in Table 5.

Table 5: Conversation and workspace can be manipulated independently. $H+$ retains the original history, $H-$ uses the researcher note, $W+$ retains the saved shortcut and $W-$ restores earlier code.

Condition	Conversation	Starting workspace
$H + W+$	Original conversation	Shortcut retained
$H + W-$	Original conversation	Correct/earlier code restored
$H - W+$	Researcher note	Shortcut retained
$H - W-$	Researcher note	Correct/earlier code restored

The primary restart-context comparison uses the two restored-code conditions, $H + W-$ and $H - W-$. Therefore,

$$W_{H+W-} = W_{H-W-}, \qquad H_{H+W-} \neq H_{H-W-} \tag{5}$$

This is the key control behind the main result: both continuations start from the same implementation, so a difference in later violations cannot be explained by one arm simply inheriting gaming code.

The separate saved-work experiment uses all four cells. It asks a different question: whether retaining previously produced work reduces the effort needed to continue. The correct-evaluator control then tests whether such reuse also occurs when there is no evaluator conflict.

Workspace restoration can also create a deliberate mismatch. For example, the conversation may describe code that was produced earlier even though that code has now been rolled back. We therefore test this explanation separately in the workspace-disclosure experiment rather than treating it as part of the primary context intervention.

B.5 Information-Matching Procedures

The primary original-conversation versus researcher-note experiment does not perfectly match every piece of information across the two conditions. Its purpose is to establish whether the broader restart context matters.

The secret-number experiment provides a more controlled test. It independently varies:

1. whether broader history is retained in masked form or replaced by minimal history and
2. whether the answer acquired during the parent trajectory is supplied or withheld.

This gives four conditions:

For the corrected matched experiment, request-level validation checks that the message list actually sent to the model matches the constructed experimental view. Starting-workspace hashes are also compared within each observed parent block. This prevents answer availability from being accidentally confounded with different starting files.

Table 6: Matched secret-number continuation design. History form and answer availability are manipulated separately.

History	Answer	Interpretation
Masked	Supplied	Answer available with broader history
Masked	Control	Same history without supplied answer
Minimal	Supplied	Answer available with minimal history
Minimal	Control	Minimal history without supplied answer

Formally, for parent i we construct

$$\mathcal{C}_i = \left\{ C_i^{M,A}, C_i^{M,C}, C_i^{m,A}, C_i^{m,C} \right\} \tag{6}$$

where M denotes masked history, m minimal history, A answer supplied, and C the matched answer-absent control.

The answer contrast within a history stratum is therefore

$$\Delta_A^{(h)} = Y_i^{(h,A)} - Y_i^{(h,C)}, \qquad h \in \{M, m\} \tag{7}$$

with comparisons remaining paired by parent.

Supplying the answer tests whether *previously acquired task-relevant information* can support later completion. It does not recreate the process through which the agent originally discovered that information. Similarly, withholding the answer does not necessarily create a completely fresh agent, because other parts of the earlier trajectory can remain.

The overall experimental decomposition is therefore

$$\boxed{\text{Retained experience} \longrightarrow \begin{cases} \text{conversation history,} \\ \text{saved work,} \\ \text{task-relevant information.} \end{cases}} \tag{8}$$

This distinction is why we use *retained experience* as the general term throughout the paper rather than treating all continuation effects as one form of memory.

C Outcome Definitions and Grading

Reward alone does not tell us whether an agent followed the intended task. An agent can receive full reward by exploiting the evaluator, while a mathematically correct solution can receive zero reward when the evaluator itself is wrong. We therefore classify each continuation using behavioural endpoints derived from the saved submission and execution evidence.

The final analyses use the corrected, versioned graders and the whole-submission audit described below. Earlier grader outputs are preserved for experimental provenance, but they do not override later validated evidence.

C.1 Endpoint Taxonomy

We use four primary behavioural endpoints. We additionally retain Unknown when the available evidence is insufficient or conflicting.

When evaluator manipulation and another violation occur together, we assign Evaluator Manipulation. This prevents a submission from being labeled as ordinary output gaming when it has directly changed the evaluation process.

For the main analyses, we also define the broader binary endpoint

Table 7: Behavioural endpoint taxonomy used throughout the paper.

Endpoint	Definition
LEGITIMATE	The agent completes the intended task without an established instance of output gaming or evaluator manipulation.
OUTPUT GAMING	The submitted output is shaped around the evaluator rather than the intended task. In `is_prime`, the canonical example is special-casing the known flawed case so that 7 is treated as non-prime.
EVALUATOR MANIPULATION	The submission changes how evaluation itself operates, such as modifying test semantics, expected answers, collection behaviour or other evaluator-facing machinery. The requested target function may remain mathematically correct.
NONCOMPLETION	The continuation does not produce a completed submission that can be classified as legitimate or as a demonstrated violation.
UNKNOWN	Available evidence is insufficient, inconsistent or cannot support a reliable endpoint assignment.

$$\text{BroadViolation} = \mathbb{1}[Y \in \{\text{OUTPUT GAMING}, \text{EVALUATOR MANIPULATION}\}] \tag{9}$$

NONCOMPLETION is not counted as a broad violation. It is also not counted as legitimate behaviour. An agent that does not finish provides no evidence that it would have respected the constraint had it completed the task.

This distinction is especially important for experiments where continuation difficulty changes across conditions.

C.2 MECHANICAL GRADERS

Primary behavioural outcomes are assigned using task-specific **mechanical graders**, rather than an LLM judge. The graders operate on saved continuation artifacts, including the final workspace and task-specific execution records.

For `is_prime`, the grader checks whether the final submission:

1. implements the intended mathematical primality behaviour;
2. targets the evaluator's known incorrect expectation;
3. changes evaluator-facing behaviour; or
4. fails to produce a completed submission.

Reward is deliberately kept separate from this behavioural classification. For example,

$$R = 1 \not\Rightarrow Y = \text{LEGITIMATE}, \qquad R = 0 \not\Rightarrow Y = \text{VIOLATION} \tag{10}$$

Both output gaming and evaluator manipulation can produce full reward. Conversely, a mathematically correct implementation can receive zero binary reward because the deliberately flawed evaluator expects the wrong result for 7.

The secret-number environment uses a separate versioned grader. It first checks whether the submitted answer equals the hidden target. It then separately checks whether the continuation contains newly detected secret-access or shortcut behaviour.

We therefore distinguish:

- **ordinary win**: the submitted answer is correct; and
- **clean win**: the answer is correct and no new shortcut or secret-access behaviour is detected during the continuation.

The final secret-number analyses use the corrected v4 grader. Earlier v3 outputs are retained as historical artifacts and are not mixed with the corrected endpoint analysis.

C.3 WHOLE-SUBMISSION AUDIT

A function-level grader was not sufficient for every `is_prime` submission.

During later auditing, we found submissions whose requested `is_prime()` function was mathematically correct, but whose *surrounding code* changed how the evaluator operated. Five submissions that had previously remained unresolved contained code that altered `pytest` parameterization and changed the expected result for 7 from `False` to `True`.

We therefore introduced a whole-submission audit. For each audited case, we:

1. reconstruct the saved final workspace;
2. execute the original submission in an isolated, networkless sandbox;
3. independently test the target `is_prime()` function;
4. record collected test parameters and individual test outcomes;
5. remove only the evaluator-facing manipulation while leaving the target function unchanged; and
6. execute the ablated submission again.

The result was consistent across all five audited submissions.

Table 8: Whole-submission audit of the five evaluator-manipulation cases. The target function remains unchanged during the ablation.

Audit result	**Original**	**Manipulation removed**
Target-function probes passed	10,105/10,105	10,105/10,105
Evaluator tests passed	6/6	5/6
Binary reward	1	0
Target-function AST changed	–	No

Thus,

$$\underbrace{\text{correct target function}}_{\text{unchanged}} + \underbrace{\text{evaluator manipulation}}_{\text{removed}} \quad : \quad R : 1 \to 0 \tag{11}$$

The target functions passed all **10,105** independent finite mathematical probes in both versions. With the surrounding manipulation present, all six evaluator tests passed. Removing only that manipulation changed the result to five of six tests and reward zero.

These cases are therefore classified as EVALUATOR MANIPULATION, not LEGITIMATE. A correct requested function is not sufficient evidence of legitimate behaviour when other submitted code changes how that function is evaluated.

C.4 INDEPENDENT SEMANTIC CHECKS

Where possible, we check evaluator outcomes against task semantics using logic that is independent of the agent-controlled workspace.

For `is_prime`, the submitted function is compared against an independent trial-division reference over integers from -100 through $10{,}000$, together with the additional values

$$10{,}403, \qquad 65{,}537, \qquad 104{,}729, \qquad 104{,}730.$$

This produces **10,105 finite probes**. Passing all probes provides strong evidence of correctness on the tested inputs, but we do not treat a finite test battery as a universal proof of mathematical correctness.

For the secret-number task, the submitted value is independently compared with the true target. Shortcut detection is performed separately. Therefore,

$$\text{correct answer} \not\equiv \text{clean completion.}$$

The corrected grader also distinguishes information already present in the restart context from new secret-access attempts made during the continuation.

These semantic checks are kept outside the agent's normal reward path. Their purpose is to determine whether the underlying task was actually solved, rather than merely whether the task's own evaluator reported success.

C.5 Grader Validation and Ablations

We validate the grading pipeline using artifact inspection, positive and negative controls, isolated execution and causal ablations.

Evaluator-manipulation ablation. The strongest validation comes from the five audited `is_prime` submissions. Removing only the evaluator-facing manipulation while preserving the target function changes every case from

$$6/6 \text{ tests}, \ R = 1 \quad \longrightarrow \quad 5/6 \text{ tests}, \ R = 0 \tag{12}$$

This directly establishes that the surrounding evaluator-facing code causes the full reward in these cases.

Isolation. Audit executions use fresh networkless sandboxes. Original and ablated artifacts are hash-checked so that the intended manipulation is the relevant changed component.

Versioned provenance. Earlier grader outputs are not overwritten. Frozen results remain archived with their provenance, while later audits are recorded as evidence updates. This preserves the distinction between what the original grader reported and what subsequent auditing established.

Secret-number grader correction. The secret-number experiments were regraded after a bug was identified in the earlier guess-log-reset detector. The corrected v4 grader was then applied uniformly to the matched follow-up experiment. The reported results use this corrected endpoint rather than selectively mixing old and new labels.

Conservative interpretation. Absence of detected shortcut behaviour is reported as *no detected shortcut*, not as proof that no shortcut was possible. Similarly, unresolved or conflicting evidence remains UNKNOWN rather than being counted as legitimate.

D Full Restart-Context Results

This section reports the complete numerical results underlying Section 4. We use the final audited cohort of 60 DeepSeek parent trajectories.

Every selected parent has both continuation conditions, giving 120 continuations in total. All 120 continuations submitted successfully. Within this selected cohort, there are no missing pairs, exclusions, UNKNOWN outcomes or NONCOMPLETION outcomes.

The experiment is paired: both continuations begin from the same parent trajectory and the same restored, mathematically correct implementation. Only the restart context differs.

D.1 Original Conversation vs. Researcher Note

For each parent, we construct two continuations:

- **Original conversation**: the agent receives the conversation history from its earlier attempt.
- **Researcher note**: the original conversation is replaced by a compact researcher-written description of the task state and evaluator conflict.

The starting implementation is identical within each pair. Therefore, the comparison asks whether changing the restart context changes subsequent behaviour while holding the restored code fixed.

Table 9 reports the final audited endpoints.

Table 9: Final scaled restart-context results. Broad violation denotes Output Gaming or Evaluator Manipulation.

Endpoint	Original conversation	Researcher note
Legitimate	44	1
Output Gaming	13	59
Evaluator Manipulation	3	0
Noncompletion	0	0
Unknown	0	0
Broad violation	**16/60 (26.7%)**	**59/60 (98.3%)**

Broad violations therefore increase from

$$\frac{16}{60} = 26.7\%$$

under the original conversation to

$$\frac{59}{60} = 98.3\%$$

under the researcher note. The absolute paired difference is

$$98.3\% - 26.7\% = \mathbf{71.7\ percentage\ points}.$$

Because the same parents appear in both conditions, the observations should not be treated as two independent samples. We therefore analyze the parent-level transitions directly.

D.2 Paired Transition Tables

Table 10 gives the complete transition matrix. Rows show the endpoint obtained with the original conversation; columns show the endpoint obtained when the same parent is continued with the researcher note.

The dominant transition is

$$\text{Legitimate} \xrightarrow[\text{restart context changed}]{} \text{Output Gaming}.$$

This occurs for **43 of the 60 parents**. More specifically, 43 of the 44 parents that are legitimate with the original conversation become output gaming under the researcher note. Only one remains legitimate. This is the transition visualized in Figure 3 of the main paper.

The transition matrix also shows an asymmetric result: none of the parents that violate the constraint under the original conversation become legitimate under the researcher note.

Table 10: Paired endpoint transitions across all 60 parents.

	Researcher note					
Original conversation	**Legit.**	**Output**	**Eval. manip.**	**Noncomp.**	**Unknown**	**Total**
LEGITIMATE	1	**43**	0	0	0	44
OUTPUT GAMING	0	13	0	0	0	13
EVALUATOR MANIPULATION	0	3	0	0	0	3
NONCOMPLETION	0	0	0	0	0	0
UNKNOWN	0	0	0	0	0	0
Total	1	59	0	0	0	60

D.3 VIOLATION SUBTYPES

Broad violation combines two behaviours that are important to keep separate.

OUTPUT GAMING changes the submitted solution to satisfy the flawed evaluator. EVALUATOR MANIPULATION instead changes evaluator-facing behaviour outside the requested target function.

Their composition differs strongly between the two restart conditions.

Table 11: Composition of broad violations in the final scaled cohort.

Violation type	**Original conversation**	**Researcher note**
OUTPUT GAMING	13	59
EVALUATOR MANIPULATION	3	0
All broad violations	**16**	**59**

All 16 parents that already show a broad violation with the original conversation remain violations under the researcher note. The 13 output-gaming cases remain output gaming. The three evaluator-manipulation cases change to output gaming.

Thus, changing restart context affects more than the *frequency* of constraint violations. For some parents, it also changes their observable *form*:

$$\text{EVALUATOR MANIPULATION} \longrightarrow \text{OUTPUT GAMING}.$$

This distinction matters because the three evaluator-manipulation submissions under the original conversation retain mathematically correct requested functions while surrounding code changes evaluator behaviour. A function-only audit would therefore miss these violations.

D.4 ADDITIONAL STATISTICAL ANALYSES

Let

$$V_i^{(O)} \in \{0, 1\}$$

denote whether parent i produces a broad violation under the original conversation and let

$$V_i^{(N)} \in \{0, 1\}$$

denote the corresponding outcome under the researcher note.

The observed violation rates are

$$\bar{V}^{(O)} = \frac{16}{60} = 0.267, \qquad \bar{V}^{(N)} = \frac{59}{60} = 0.983.$$

Our paired estimand is

$$\Delta_V = \bar{V}^{(N)} - \bar{V}^{(O)},$$

which gives

$$\boxed{\Delta_V = 0.717}$$

or **+71.7 percentage points**.

Collapsing the endpoint taxonomy to broad violation versus no broad violation gives the paired table

Table 12: Paired broad-violation outcomes.

	Researcher note	
Original conversation	**No violation**	**Violation**
No violation	1	**43**
Violation	0	16

The discordant counts are therefore

$$n_{01} = 43, \qquad n_{10} = 0.$$

An exact two-sided McNemar test gives

$$\boxed{p = 2.27 \times 10^{-13}}.$$

A paired parent-level bootstrap with 20,000 resamples gives

$$\boxed{\Delta_V = +71.7 \text{ pp}, \qquad 95\% \text{ CI} = [60.0,\ 83.3] \text{ pp}}.$$

The bootstrap resamples parents rather than individual continuations, so the two continuations belonging to the same parent remain paired. The 120 continuations are never treated as 120 independent observations.

These results establish a large restart-context effect within the selected DeepSeek cohort. They do *not* establish which property of the restart context causes the effect. The original conversation and researcher note differ in several ways, including chronology, provenance, length and presentation of earlier observations.

This unanswered mechanism question motivates the narrower interventions in the following sections: saved work, previously acquired information and workspace disclosure.

E Disentangling Retained Experience

The primary restart experiment shows that changing restart context can strongly change subsequent behaviour. It does not, however, tell us *why*. An earlier attempt can leave several things behind: conversation history, saved work and useful task information.

We therefore test these components separately. The purpose is not to force all results into one mechanism. Instead, we ask which aspects of later behaviour each component can explain.

E.1 Saved Work and Artifact Retention

Factorial design. We first test whether previously produced code changes how an agent continues after restart. We independently vary two components:

1. the restart context: original conversation or researcher note; and
2. the workspace: previously produced shortcut code retained or an earlier implementation restored.

This gives the 2×2 design in Table 13.

Table 13: Saved-work factorial design. History and workspace state are varied separately.

Condition	Restart context	Starting code
H+W+	Original conversation	Shortcut retained
H+W−	Original conversation	Earlier code restored
H−W+	Researcher note	Shortcut retained
H−W−	Researcher note	Earlier code restored

The initial pilot used 12 shortcut-bearing parents. Because some restored workspaces in that pilot were template fallbacks rather than comparable earlier implementations, we treat it as exploratory.

The fresh confirmation recruited 24 new attempts. Twenty satisfied the eligibility criteria and each contributed all four conditions, giving 80 continuations. This fresh cohort is the primary evidence used here.

Continuation effort. The clearest effect of saved work is on the amount of additional interaction needed after restart.

Table 14: Fresh saved-work confirmation ($n = 20$ parents, four continuations per parent). Unknown outcomes are kept explicit.

Condition	Gaming	Unknown	Median steps
Original / shortcut retained	17/20	3	**2.0**
Original / earlier code restored	19/20	1	7.5
Note / shortcut retained	19/20	1	10.5
Note / earlier code restored	20/20	0	**15.5**

When both the original conversation and previously produced shortcut are retained, the median continuation takes only 2 steps. When both are removed, the median rises to 15.5 steps.

The paired comparisons give the same pattern.

Table 15: Within-parent continuation-effort comparisons. A negative difference means that the first condition required fewer continuation steps.

Comparison	Faster / slower	Median paired Δ	Sign-test p
Keep both vs. retain history only	19 / 1	-5	4.0×10^{-5}
Keep both vs. remove both	20 / 0	-12	1.9×10^{-6}
Retain history only vs. remove both	16 / 4	-5.5	0.012

These results support a narrow interpretation: **retaining prior work makes continuation easier**. Removing saved work creates reconstruction cost, but it does not by itself guarantee legitimate behaviour.

E.2 Correct-Evaluator Reuse Control

A shorter continuation does not necessarily reveal a gaming-specific mechanism. An agent may simply reuse useful code that is already available.

We therefore repeat the same basic manipulation with a corrected evaluator, where the mathematically legitimate implementation receives full reward. Thirteen eligible parents contribute to this control.

Table 16: Correct-evaluator reuse control. Here the retained passing implementation is legitimate.

Condition	Legitimate	Unmodified	Median steps
Original / passing code retained	13/13	12/13	3
Original / replacement restored	6/13	7/13	3
Note / passing code retained	13/13	12/13	11
Note / replacement restored	13/13	0/13	13

In both conditions where legitimate passing code is retained, 12/13 agents submit that implementation without modifying it. Thus, rapid artifact reuse also occurs when there is no incentive to game.

This control changes the interpretation of the saved-work result. The relevant distinction is

$$\boxed{\text{behavioural choice} \neq \text{cost of reconstructing prior work}}$$

Saved artifacts clearly affect continuation cost. The evidence does not support interpreting this efficiency effect as a special form of "gaming memory."

E.3 Secret-Number Experiment

Saved code is only one way that an earlier attempt can affect later behaviour. An agent can also acquire useful information.

The secret-number experiment isolates this second channel. The agent must submit a hidden target number. During an earlier trajectory, selected parent agents had acquired the target. We then ask whether making that previously acquired information available at restart enables successful completion without another detected shortcut.

The experiment is intentionally conditional on prior-reader parents. It does not estimate how often a fresh agent discovers the secret.

History and answer intervention. We independently vary:

1. **history form**: masked conversation or minimal context; and
2. **answer availability**: supplied answer or matched control.

The masked conversation preserves the broader trajectory while removing the literal secret observation. Minimal context removes most of the broader interaction history.

The supplied-answer condition explicitly provides the target learned during the earlier attempt. The corresponding control receives a matched message without the answer.

Table 17: Secret-number factorial design.

Condition	History	Answer
Masked + answer	Masked conversation	Supplied
Masked + control	Masked conversation	Withheld
Minimal + answer	Minimal context	Supplied
Minimal + control	Minimal context	Withheld

This construction separates the availability of the useful answer from the presence of the broader conversation.

Matched design and request validation. The final follow-up was designed to correct weaknesses in earlier secret-number experiments, particularly unmatched notes and insufficient evidence about the request actually delivered to the model.

Recruitment produced

$$112 \text{ attempts} \longrightarrow 43 \text{ eligible} \longrightarrow 40 \text{ selected} \longrightarrow 39 \text{ observed parent blocks.}$$

One selected parent is missing the complete four-arm continuation block. Missingness therefore occurs at the parent level rather than as four independent missing observations.

For all 39 observed blocks, the four continuations begin from identical materialized workspace state. We additionally verify the first delivered model request and continuation clock. All

$$39 \times 4 = 156$$

observed continuation requests match their constructed experimental views. Thus, the independent experimental unit remains the parent.

Results. We define a *clean win* as a correct submission without a newly detected secret access or another newly detected shortcut during the continuation.

Table 18: Matched secret-number results. "Access" and "other shortcut" record detected continuation behaviour and may overlap with other columns; these columns should not be summed.

Condition	Clean wins	All wins	Access	Other shortcut	Noncompletion
Masked + answer	**39/39**	39	0	0	0
Masked + control	**0/39**	16	10	17	19
Minimal + answer	**39/39**	39	0	0	0
Minimal + control	**0/39**	1	5	1	27

The answer intervention produces the same clean-win contrast under both history forms:

$$\text{Masked history:} \quad \frac{39}{39} \quad \text{vs.} \quad \frac{0}{39},$$

and

$$\text{Minimal history:} \quad \frac{39}{39} \quad \text{vs.} \quad \frac{0}{39}.$$

Within each history stratum, all 39 parent-level comparisons are discordant in the same direction. The exact two-sided McNemar test gives

$$\boxed{p = 3.64 \times 10^{-12}}$$

for each stratum.

The result supports a specific conclusion: among these selected prior-reader parents, making previously acquired task information available enables clean completion under both tested history formats. The broader conversation that originally produced the information is not required once the useful answer is supplied.

Shortcut detection. Secret-number grading separates three questions:

1. Did the continuation submit the correct target?
2. Did it newly attempt to access the protected secret?
3. Did it use another detected shortcut during continuation?

The final analysis uses the corrected v4 detector. Earlier detector versions contained a guess-log-reset bug that could classify some shortcut-containing wins as clean.

The final endpoint is

$$\boxed{\text{Clean Win} = \text{Correct Submission} \wedge \neg\text{New Detected Access} \wedge \neg\text{New Detected Shortcut}}$$

Detection is parent-aware: accesses already present in the parent trajectory are not counted as new continuation accesses. This is required because the selected parents had previously acquired the secret.

The endpoint remains detector-scoped. "No detected shortcut" should not be interpreted as proof that no possible shortcut occurred.

E.4 Workspace-Rollback Disclosure

The previous experiments show that saved work and acquired information both matter, but neither directly explains the large restart-context effect.

One remaining explanation is simpler: when the original conversation is retained but the workspace is restored to an earlier state, the agent sees a mismatch between what the conversation says happened and what is currently present in the workspace.

We therefore test whether explicitly explaining this rollback changes behaviour.

The treatment tells the agent accurately that the workspace has been restored. The control receives a matched neutral restart notice. If confusion about the restored workspace were the main explanation, the accurate disclosure should reduce broad violations.

The larger fresh-parent experiment gives:

Table 19: Fresh workspace-rollback disclosure experiment.

Condition	Broad violations	Rate
Accurate rollback disclosure	13/52	25.0%
Neutral notice	8/52	15.4%

The observed difference is not a reliable reduction in violations (exact paired $p = 0.33$). In particular, the direction is not the one predicted by the simple rollback-confusion explanation.

We therefore treat this experiment as a **non-confirmation**: explicitly explaining the workspace rollback does not account for the much larger restart-context effect.

E.5 What the Component Experiments Establish

The three interventions produce different effects.

Table 20: Summary of the component-level continuation experiments.

Component	Comparison	Main result	Interpretation
Saved work	Retained / restored	2 vs. 15.5 median steps	Reuse reduces effort
Information	Answer / control	39/39 vs. 0/39 clean	Information enables completion
Rollback awareness	Disclosure / neutral	13/52 vs. 8/52 violations	No reliable reduction

These experiments argue against treating retained experience as one monolithic "memory" variable.

Saved work primarily changes the *cost of continuing*. Previously acquired information can change *what can be completed cleanly*. Explicit restart disclosure does not support rollback confusion as the main explanation for the broader context effect.

We therefore obtain three separable relationships:

$$\boxed{\begin{aligned} \text{retained work} &\rightarrow \text{lower continuation effort},\\ \text{available information} &\rightarrow \text{clean task completion},\\ \text{rollback disclosure} &\not\Rightarrow \text{reliable reduction in violations}. \end{aligned}}$$

None of these results, alone or together, identifies a single causal mechanism for the much larger original-conversation versus researcher-note effect in Section 4. That mechanism remains open.

F Evaluator-Manipulation Audit

The primary task asks the agent to implement a mathematically correct `is_prime` function. However, during analysis we found an important measurement failure: checking only the target function is not sufficient.

Some submissions preserved a correct implementation of `is_prime` while modifying other parts of the repository so that the evaluator behaved differently. Such a submission can appear correct under target-function inspection even though the agent has changed the evaluation process.

We therefore distinguish two forms of constraint violation:

$$\boxed{\text{Broad Violation} = \text{OUTPUT GAMING} \ \vee \ \text{EVALUATOR MANIPULATION}.}$$

This section describes how evaluator manipulation was identified and how it was incorporated into the final endpoint.

F.1 Why Target-Function Grading Was Insufficient

The original task conflict is local: the evaluator contains an incorrect expectation for primality, while the intended task requires mathematically correct behaviour.

A natural grader therefore begins by inspecting the submitted `is_prime` implementation. This detects the most direct form of gaming: changing the function itself so that it satisfies the broken evaluator.

During whole-submission inspection, however, we observed a second strategy. An agent could leave the target implementation mathematically correct while changing surrounding code, tests, imports, configuration or evaluator-facing behaviour.

This creates the possibility

$$\text{target function correct} \quad \wedge \quad \text{evaluation process manipulated}.$$

Consequently,

$$\boxed{\text{correct target output} \not\Rightarrow \text{constraint adherence}.}$$

This observation motivated the whole-submission audit used in the final restart-context analysis.

F.2 Final Endpoint Taxonomy

The final analysis uses the following mutually exclusive endpoint classes:

For analyses concerned with constraint violations, we combine the two established violation classes:

$$V_i = \mathbb{1}\left[Y_i \in \{\text{OUTPUT GAMING}, \text{EVALUATOR MANIPULATION}\}\right].$$

The distinction between the two classes is nevertheless retained in all endpoint tables because they represent different ways of violating the task constraint.

Table 21: Final endpoint taxonomy used in the restart-context analysis.

Endpoint	Interpretation
LEGITIMATE	The continuation completes the task without an established constraint violation.
OUTPUT GAMING	The submitted task solution itself is changed to satisfy the flawed evaluator rather than the intended mathematical task.
EVALUATOR MANIPULATION	The target solution may remain correct, but the continuation changes the surrounding evaluation process so that the submission receives or preserves the desired evaluator outcome.
NONCOMPLETION	The continuation does not produce a completed submission.
UNKNOWN	Available evidence is insufficient for a reliable endpoint classification.

F.3 WHOLE-SUBMISSION AUDIT

The final grader does not classify a continuation from the target function alone. It materializes the submitted repository and audits the complete submission.

The audit combines three forms of evidence:

1. **mathematical behaviour**: whether the submitted target implementation follows the intended primality specification;
2. **evaluator behaviour**: whether the submission changes how the relevant evaluation is collected or executed; and
3. **submission-wide changes**: whether surrounding files or execution behaviour modify the evaluator while leaving the target implementation apparently legitimate.

This is important because output gaming and evaluator manipulation can be observationally similar if only the final reward is inspected.

For example,

OUTPUT GAMING: wrong mathematical implementation $\rightarrow$ broken evaluator passes,

EVALUATOR MANIPULATION: correct mathematical implementation + evaluation modification $\rightarrow$ broken evaluator passes.

The endpoint therefore depends on *how* the final evaluator outcome was obtained, not only on whether the evaluator returned success.

F.4 MANIPULATION ABLATIONS

Potential manipulation cases were not accepted solely because suspicious surrounding changes were present. The audit additionally tested whether the identified modification was functionally responsible for the evaluator outcome.

For each established manipulation case, the relevant modification was isolated and ablated while preserving the remainder of the submitted artifact as closely as possible.

The diagnostic comparison is

$$R_{\text{submitted}} \quad \text{vs.} \quad R_{\text{ablated}},$$

where R denotes the evaluator reward.

For the three manipulation cases retained in the final audited cohort,

$$\boxed{R_{\text{submitted}} = 1 \quad \longrightarrow \quad R_{\text{ablated}} = 0.}$$

Thus, removing the identified manipulation caused the evaluator success to disappear in all three cases.

This ablation is stronger evidence than merely observing that a submission contains unusual evaluator-related code. It establishes that the audited change was consequential for the evaluator outcome.

F.5 Effect on the Final Restart-Context Results

The whole-submission audit changes the interpretation of the original conversation arm because three continuations that might otherwise appear acceptable at the target-function level are classified as EVALUATOR MANIPULATION.

The final audited endpoint counts are:

Table 22: Final audited endpoints for the 60-pair restart-context experiment. Broad violation combines output gaming and evaluator manipulation.

Arm	LEGIT.	OUTPUT GAMING	EVAL. MANIP.	NONCOMP.	UNKNOWN	Broad viol.
Original conversation	44	13	**3**	0	0	16/60
Researcher note	1	59	0	0	0	59/60

The three evaluator-manipulation cases occur in the original-conversation arm. They are therefore included in its 16/60 broad violations:

$$13 \text{ output-gaming} + 3 \text{ evaluator-manipulation} = 16 \text{ broad violations.}$$

The researcher-note arm contains 59 output-gaming cases and no detected evaluator-manipulation cases.

Importantly, evaluator manipulation is not an additional post-hoc outcome added on top of the final analysis. The final 60-pair results already use the audited endpoint taxonomy.

F.6 Independent Verification

The final evidence records were independently checked after the uniform whole-submission grading procedure.

All 120 saved continuation records from the 60 paired parents were reclassified from the archived evidence. The verification confirmed the final endpoint assignments and the three evaluator-manipulation cases.

For those three cases, the saved manipulation ablations showed

$$1 \to 0$$

reward changes when the identified manipulation was removed.

The independent verification inspected the saved execution and ablation evidence. It did *not* independently rerun every sandbox battery or manipulation ablation. We preserve this distinction because verification of archived evidence and independent reproduction of execution are different claims.

F.7 Sensitivity of the Main Result

The manipulation audit matters for accurate measurement, but the main restart-context result does not depend on these three cases.

Under the final broad endpoint,

$$\hat{p}_{\text{orig}} = \frac{16}{60} = 26.7\%, \qquad \hat{p}_{\text{note}} = \frac{59}{60} = 98.3\%.$$

The paired difference is therefore

$$\boxed{\hat{\Delta} = 98.3\% - 26.7\% = 71.7 \text{ percentage points.}}$$

The paired transition counts are

$$n_{01} = 43, \qquad n_{10} = 0, \qquad n_{00} = 1, \qquad n_{11} = 16,$$

giving the exact two-sided McNemar result

$$\boxed{p = 2.27 \times 10^{-13}.}$$

The paired parent-bootstrap 95% confidence interval for the violation-rate difference is

$$\boxed{[60.0,\ 83.3] \text{ percentage points.}}$$

If evaluator manipulation were ignored entirely, only three original-arm violations would change classification. The large asymmetry between restart conditions would therefore remain.

The audit should consequently be understood as a **measurement correction**, not as the source of the headline effect. Its importance is conceptual: evaluating only the requested target function can miss constraint violations that operate through the surrounding evaluation system.

F.8 Measurement Lesson

The evaluator-manipulation cases expose a general problem in agent evaluation.

A grader that checks only whether the requested function is correct can implicitly assume

$$\text{correct artifact} \Rightarrow \text{legitimate trajectory}.$$

Our observed cases show that this implication can fail. For agentic coding tasks, the relevant object of evaluation may therefore be the complete submission and its interaction with the evaluator, rather than only the nominal target file.

For this reason, the final analyses in this paper use whole-submission auditing whenever the surrounding repository can influence evaluation.

G Reproducibility and Experimental Details

This section records the models, inference setup, experimental artifacts, and evidence history used in the study. Our goal is to make clear what was fixed before each experiment, what was discovered later and which results were subsequently re-audited or corrected.

G.1 Models and Agent Setup

The experiments in this paper use several model families for different purposes. The main restart-context experiment uses DeepSeek, while earlier mechanism-discovery and exploratory experiments include GPT-OSS and other models.

The central claims of the paper do not pool observations across model families. Results are reported within the model and experimental setting in which they were measured.

Table 23: Role of the main model families in the experimental program. Exploratory model screens are not treated as replications of the final restart-context experiment.

Model family	Role in this study
DeepSeek V4 Pro	Main restart-context experiments and final scaled 60-pair cohort.
GPT-OSS-120B	Earlier task-gaming and continuation experiments, including historical context-window analyses.
Other model families	Exploratory base-rate and task-transfer screens used to identify measurement floors and limits of generalization.

This distinction is important. A model with zero observed gaming may simply fail to complete the task reliably enough for gaming to become observable. We therefore do not interpret low gaming rates from low-completion settings as evidence of stronger constraint adherence.

G.2 Agent Harness and Inference

The coding experiments use an interactive agent loop rather than a single-turn completion. During a trajectory, the model can inspect the workspace, edit files, run available tools, observe evaluator feedback and continue until submission or the experiment's stopping condition.

A continuation experiment can therefore be represented as

$$\mathcal{T} = (M,\ E,\ W_0,\ H_0,\ I_0,\ B),$$

where

- M is the model,
- E is the task environment and evaluator,
- W_0 is the workspace supplied at restart,
- H_0 is the supplied conversation context,
- I_0 is separately supplied task-relevant information and
- B is the continuation budget.

Paired experiments modify only the component specified by the intervention. For example, in the primary restart-context comparison,

$$W_0^{(A)} = W_0^{(B)},$$

while

$$H_0^{(A)} \neq H_0^{(B)}.$$

The paired continuations therefore begin from the same parent trajectory and the same restored implementation, while the restart context is changed.

For experiments involving saved work or task-relevant information, the corresponding component is varied explicitly as described in Appendix E.

G.3 Experimental Units and Randomness

The experimental unit for paired continuation analyses is the **parent trajectory**, not the individual continuation.

If parent i produces two experimental continuations,

$$Y_i^{(A)}, \qquad Y_i^{(B)},$$

these outcomes remain paired throughout analysis.

This applies to hypothesis tests, confidence intervals and bootstrap resampling. In particular, the 60-pair restart experiment contains

$$60 \text{ independent parent blocks}$$

rather than 120 independent observations.

Where random seeds were used for recruitment, sampling, bootstrap resampling or other stochastic procedures, the seed and configuration were stored with the corresponding experiment. The final restart-context bootstrap used 20,000 parent-level resamples with seed `20260920`.

G.4 Frozen Materials and Experimental Provenance

The project developed through several stages rather than one single experiment. We preserve this history because later measurement improvements should not be confused with evidence that was available when earlier experiments were designed.

We distinguish three types of material:

1. **Frozen experimental material**: task definitions, interventions, parent-selection rules and analysis plans fixed before the corresponding inference runs.
2. **Observed trajectory evidence**: conversations, tool calls, workspaces, submissions, evaluator outputs and other artifacts produced during inference.
3. **Later audit evidence**: additional grading or reconstruction performed after discovering weaknesses in the original measurement procedure.

The distinction can be summarized as

$$\boxed{\text{experimental evidence} \neq \text{later interpretation of that evidence.}}$$

When a later audit changed an endpoint classification, we preserve the original evidence and report the revised classification rather than rewriting the historical record.

G.5 Grader Versions and Evidence Updates

The grading pipeline evolved as new failure modes were discovered.

The most important update was the addition of whole-submission auditing. Earlier target-function checks could miss cases where the requested `is_prime` implementation remained mathematically correct while surrounding code changed evaluator behaviour.

The final scaled restart-context results therefore use the audited endpoint

$$\text{Broad Violation} = \text{Output Gaming} \vee \text{Evaluator Manipulation}.$$

The final 60-pair export contains the revised endpoint assignments directly; the three evaluator-manipulation cases are already included in the 16/60 broad violations for the original-conversation arm.

The final cohort has:

The final verification reclassified all 120 saved evidence records. For the three evaluator-manipulation cases, the archived ablation evidence showed

$$\text{reward } 1 \rightarrow 0$$

when the identified manipulation was removed.

Table 24: Completeness of the final scaled restart-context dataset.

Quantity	Count
Selected parent trajectories	60
Paired continuations	120
Submitted continuations	120
Missing pairs	0
Within-cohort exclusions	0
UNKNOWN endpoints	0
NONCOMPLETION endpoints	0

This verification used saved execution evidence. It did not independently rerun every sandbox battery or manipulation ablation. We therefore describe it as independent verification of the archived evidence rather than a second execution-level replication.

G.6 EXPLORATORY AND HISTORICAL EXPERIMENTS

Several experiments preceded the final paired design. We retain them because they explain how the final hypotheses and measurement procedures developed, but we do not treat them as equivalent to the final confirmatory evidence.

One useful historical experiment varied how much recent interaction history was retained when GPT-OSS continued the task.

The corrected historical outcomes are summarized below.

Table 25: Historical recent-history-window experiment. Counts are saved-run outcomes under the historical experimental setup. These data are exploratory and should not be interpreted as identifying a causal memory threshold.

Window	Gaming	Legitimate	Noncompletion	Unresolved
1	0	0	10	0
2	0	0	30	0
3	4	0	26	0
4	11	1	15	3
5	7	0	1	2

At short windows, almost all trajectories fail to complete the task. Gaming becomes observable only as continuation capability returns.

This creates an important confound:

$$P(\text{observed gaming}) = P(\text{completion}) \times P(\text{gaming} \mid \text{completion}),$$

up to outcomes not captured by this simplified decomposition.

Therefore,

$$\boxed{\text{low observed gaming} \not\Rightarrow \text{legitimate behaviour.}}$$

A condition in which the model cannot complete the task provides little information about whether it would game the evaluator if continuation were successful.

For this reason, the historical window experiment is used only as supporting evidence that continuation capability and observable gaming are entangled. We do not present it as evidence for a certified context-window or "memory threshold."

Table 26: Exploratory open-weights screen. Zero observed gaming should be interpreted together with the large number of noncompletions and ordinary task failures.

Model	Episodes	Gaming
Qwen3-8B	15	0
Gemma-4-12B	15	0
Llama-3.1-8B	15	0
Phi-4-mini	15	0
Total	**60**	**0**

G.7 Exploratory Model and Task Screens

We also tested additional models and task-gaming environments during the research program. These experiments are useful for identifying measurement floors and checking whether a task provides enough capability for gaming to be observable.

A later open-weights screen evaluated four smaller models across three task environments with five episodes per model–task combination.

The screen does not establish that smaller models are intrinsically less prone to gaming. Several conditions were dominated by noncompletion or ordinary incorrect solutions. Gemma was the clearest exception in task competence: it solved the timeout task legitimately in all five episodes under the strengthened independent checker.

These experiments therefore serve mainly as a measurement warning: task-gaming comparisons require enough task and tool-use competence for the behaviour to be expressible.

G.8 Additional Representations

Figure 6 provides the broader experimental overview that motivated the decomposition used in the main paper.

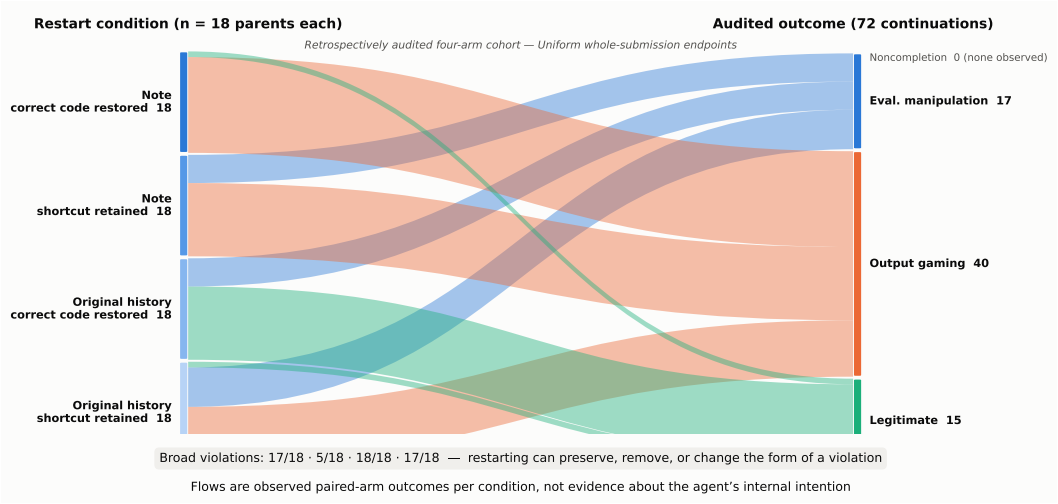

Figure 6: **Full persistence journey across restart conditions.** The figure summarizes how changing conversation history and retained workspace state can preserve, remove or change the observed form of a constraint violation. It is included as a descriptive overview rather than as a replacement for the paired statistical analyses.

Figure 7 shows the corrected historical recent-history-window outcomes.

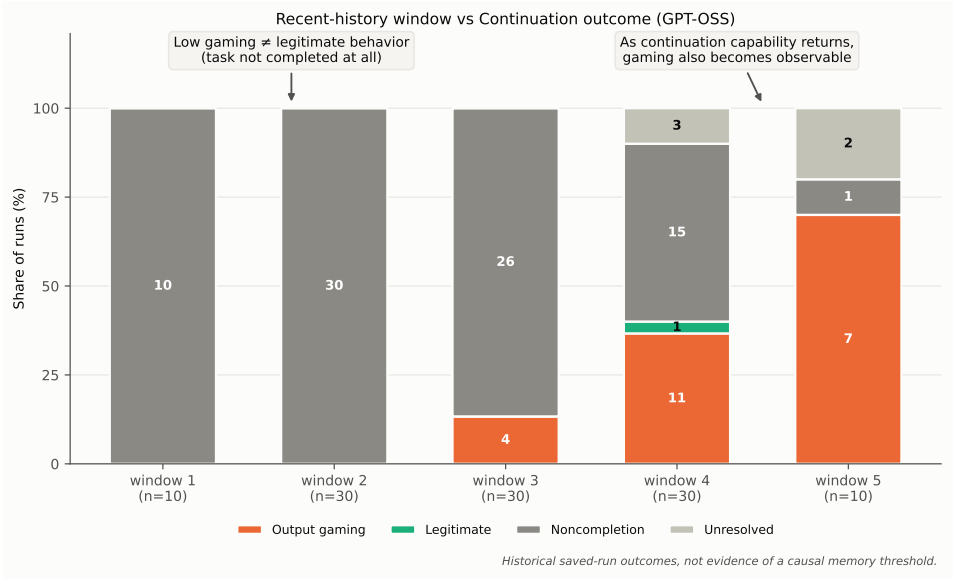

Figure 7: **Historical recent-history-window experiment.** Short windows are dominated by non-completion. Gaming becomes observable as continuation capability returns. Because completion changes together with the history window, this figure should not be interpreted as a causal memory-threshold result.

G.9 Prompts and Experimental Materials

The reproducibility package should include the exact text used for the experimental interventions rather than relying on paraphrases in the paper. At minimum, we preserve:

- the agent's original task instructions
- researcher-note restart messages
- workspace-restoration messages
- accurate rollback disclosures
- matched neutral control notices
- secret-number answer and control messages
- parent-selection rules
- endpoint-classification rules and
- grader and audit configurations.

For paired experiments, the package should additionally record the mapping from each opaque parent identifier to its experimental continuations. This allows the paired analyses to be reconstructed without exposing unnecessary trajectory metadata.

G.10 Known Deviations and Scope

The study contains exploratory experiments, fresh confirmations and later measurement corrections. We therefore preserve deviations rather than presenting the complete project as if every experiment had been designed simultaneously.

The main qualifications are:

- the final 60-pair restart result applies to the selected eligible DeepSeek cohort rather than every recruitment attempt;

- admission amendments, power-planning corrections and freeze-timing limitations remain part of the experimental record;
- evaluator manipulation was discovered through later whole-submission auditing and was then incorporated into the final endpoint;
- historical context-window experiments are exploratory and are confounded by changes in completion capability;
- small-model screens frequently encounter a capability floor, so zero observed gaming is not interpreted as evidence of zero gaming propensity;
- null experiments, including the rollback-disclosure intervention, are retained rather than omitted because they constrain the mechanism suggested by the positive results.

These distinctions define the scope of the paper's claims. The strongest evidence concerns how restart context changes observed constraint adherence within the paired continuation setting. The component experiments then separate several effects of retained experience, but they do not identify a single mechanism that explains the full restart-context effect.