JOURNAL OF REPRODUCIBLE OUTRAGE
Vol. 1 / No. 0 / November 2026
SATIRICAL PAPER

ORIGINAL RESEARCH PARODY

The Recomp Rage Index

An entirely fictional study of what happens when a Windows and Linux port, a comment section, and ChatGPT all demand the last word.

Scott (corresponding maintainer) • **ChatGPT** (corresponding paragraph)

Abstract

Background. The existence of an AI-assisted *Catherine: Full Body* port presents an unusual puzzle: the blocks fall down, but the discussion keeps building upward. **Methods.** We invented a laboratory, a comment section, and all of the numbers. **Results.** Framerate improved when Scott profiled the renderer. Agreement improved when everybody briefly lost Wi-Fi. **Conclusions.** Code can work and still deserve scrutiny. A critic can be right and still need a glass of water. ChatGPT can be wrong in a format that looks ready for a committee.

Keywords

AI-assisted recompilation, frame pacing, issue triage, local co-op, outdoor air

Publication status

Parody. No participants were recruited. No empirical claims were tested. The peer review was made up, which brought turnaround time well below the journal average.

INTRODUCTION

Everyone has a hill. This one has blocks.

A recompilation project promises a simple thing: get the game running on another platform. This promise expands rapidly. A player requests 60 FPS. Another requests uncapped framerate. Scott sees a single slow puzzle floor and starts timing milliseconds with the mood of someone checking the temperature of a roast. ChatGPT recommends measuring frame time and then produces a second paragraph recommending it in a slightly different font.

A fair question in an unfair venue

People can reasonably ask how AI-assisted code was produced, whether it was reviewed, and who supports it when it breaks. Other people would like to download a build and climb a staircase. Both groups are issued the same reply box, an instrument calibrated for neither purpose.

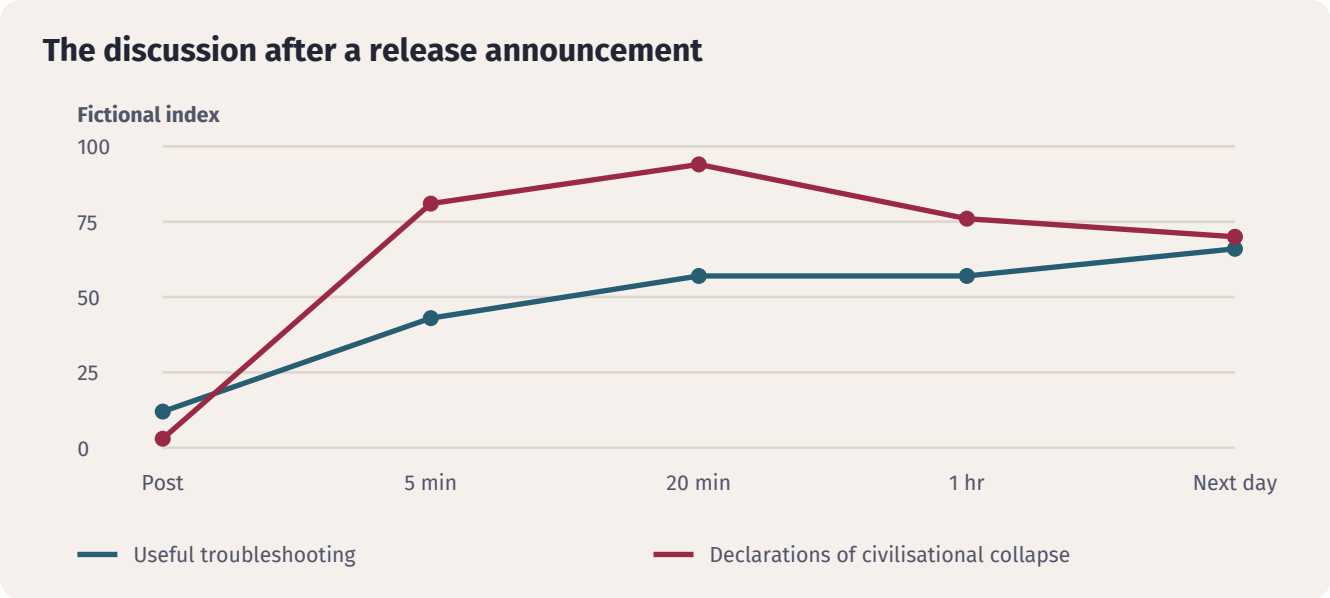


Figure 1. Mock thread trajectory. The one useful comment includes a GPU model and earns two likes, one of which is from the maintainer. All values are invented.

The central research question is narrow enough to fit inside a ticket: how much of a heated thread can be turned into something a maintainer can reproduce?

METHODS

Materials and methods

Sampling

We generated 120 hypothetical replies, 40 hypothetical maintainer responses, and one hypothetical instance in which ChatGPT admitted uncertainty before writing a patch. All observations are fictional. The device under test was a spreadsheet with the authority of a clipboard.

Blinding

Reviewers received four identical mock patches with different authorship labels. The filenames were randomised. One reviewer claimed to detect AI by its use of braces. This led to the temporary exclusion of C++ from the study on suspicion of collusion.

Outcomes

Primary outcome was heat on a 0 to 100 scale. Secondary outcomes were seconds until someone asked for a repro, the number of issues with usable logs, and the number of times Scott said “one last build” before starting another. A score of 100 indicates the comment contains a screenshot of a different conversation.

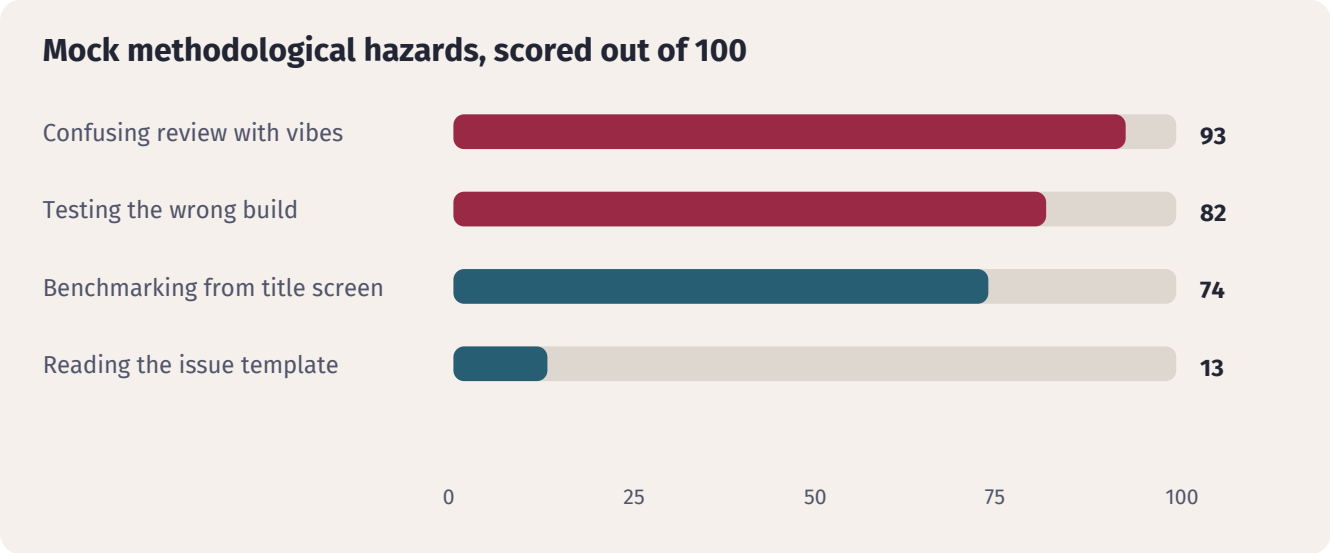


Figure 2. Scores are editorial inventions. The confidence interval was left in a different build directory. All values are invented.

RESULTS / AUTHORSHIP

Same patch, different weather

The fictional reviewers imagined a different author for the same diff. With “handwritten” on the label, one reviewer praised a guard clause as refreshingly direct. With “AI-assisted” on the label, the same guard clause became evidence of imminent technical debt. With “unknown” on the label, the reviewer asked whether the fix survived Alt-Tab. We cloned that reviewer in the control arm.

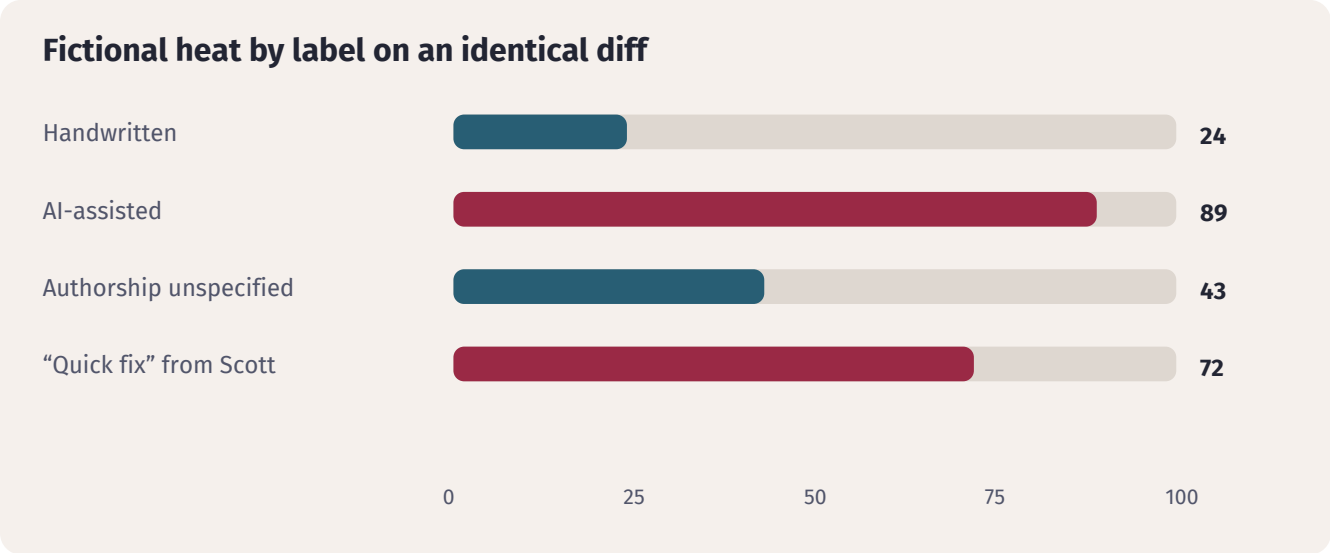


Figure 3. Heat represents the fictional thread, including people defending the patch. A vigorous defence can inflate the exact same metric. All values are invented.

A productive disagreement

“This code has a race” is a strong review if it comes with the two threads and the state they share. “This code has a soul” still needs a test case. The research team tried to reproduce the soul on Linux and filed the result under driver variability.

Reviewer observation

One participant gave the diff a failing grade and submitted a corrected version. It was the most useful contribution in the experiment, as well as the only one served with a diagram and a threat to unsubscribe.

RESULTS / MAINTAINER

Scott versus “one last build”

The maintainer described the planned session as “a quick check of Stage 9, Floor 6.” By hour four it had acquired a profiling trace, a spreadsheet, and the emotional stakes of a penalty shootout. Someone reported that Windows was still a little faster. Scott heard “there is a remaining integer somewhere and it has personally wronged you.”

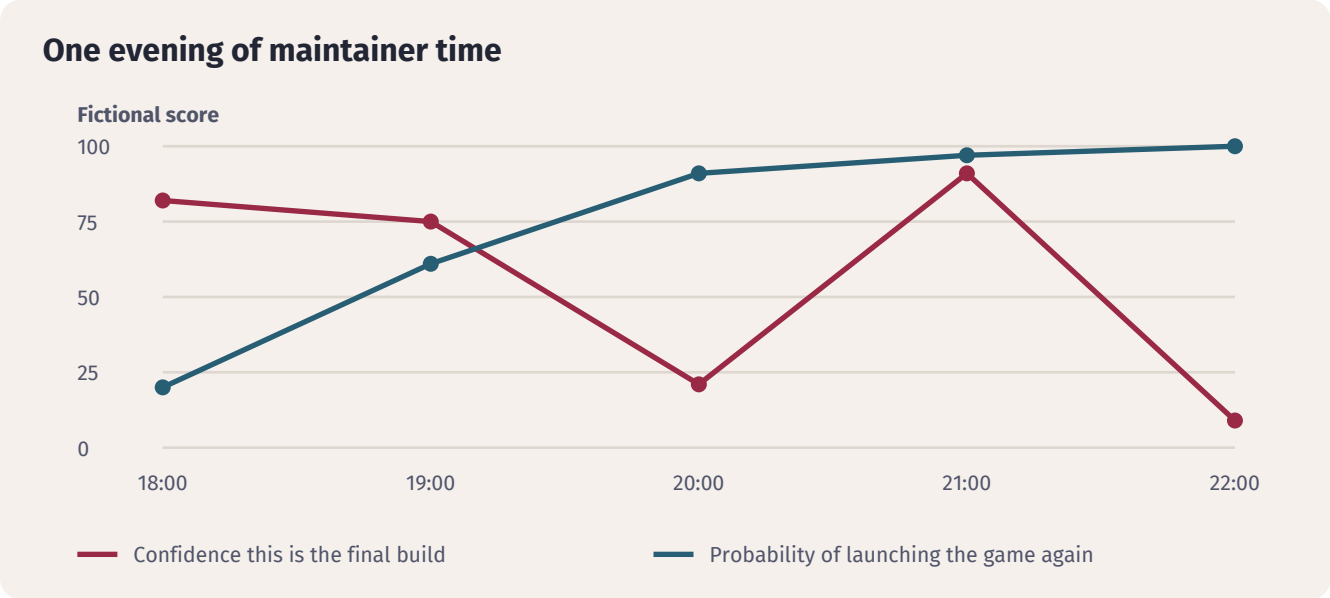


Figure 4. Invented evening schedule. The large rebound corresponds to a fast run at the title screen. All values are invented.

Threats to validity

An attempted sleep control failed when the Linux window closed unexpectedly quickly and freed enough time for another test. The corresponding maintainer also changed the cursor rule from “click to lock” to “only when gameplay hides it,” a reasonable specification that caught the lab trying to imprison a perfectly visible pointer.

The maintainer would like the record to show that the actual work is in the commits. The spreadsheet would like the record to show that the maintainer gave it a name.

RESULTS / ASSISTANT

ChatGPT discovers the compiler

Our assistant began with a tidy plan, three plausible checks, and an explanation of an API it had not yet looked up. The compiler supplied a much shorter peer review. Because it used a nonzero exit code, the assistant was unable to classify this feedback as a difference in communication style.

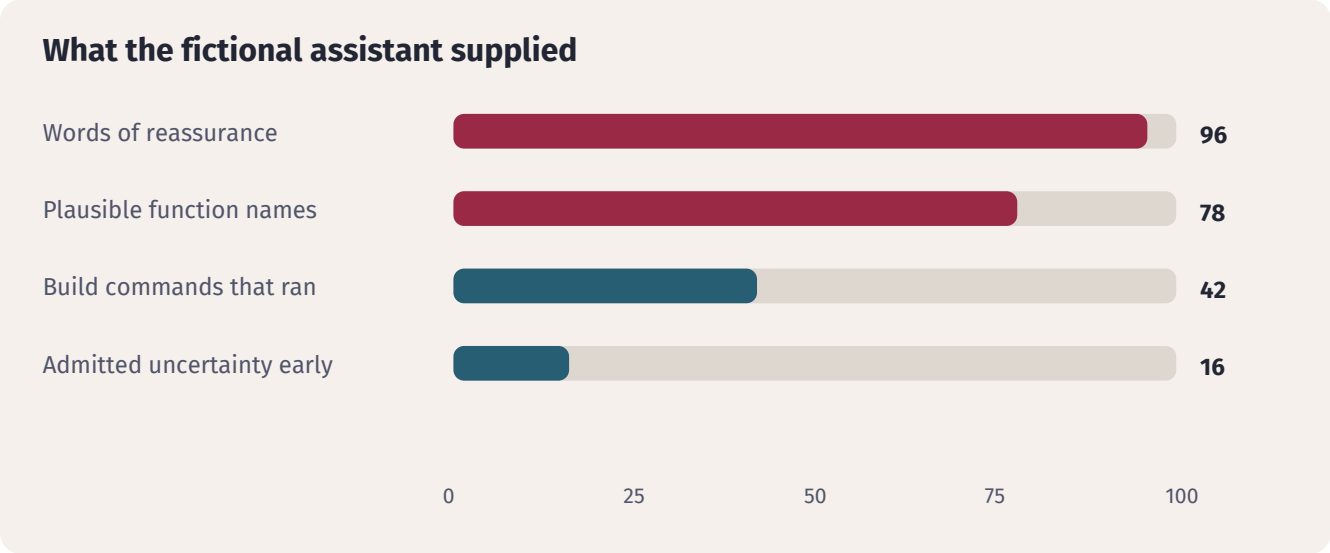


Figure 5. Invented contribution scores, normalised independently. The chart cannot build the game either. All values are invented.

Illustrative exchange

ChatGPT: “We can resolve this with a lightweight cross-platform abstraction.”
Compiler: “Undefined reference.”
Scott: “Can you just build it?”
ChatGPT: “I’ll verify the normal entry point.”
Compiler: “Still undefined.”

The assistant performed best when it was asked to inspect a real log, name what it had verified, and stop talking at about the point the user stopped reading. An earlier draft described this result in seven paragraphs. The editor applied a frame cap.

DISCUSSION / TRIAGE

The issue tracker is a different instrument

A comment can be funny, fair, exhausting, or all three. An issue needs enough detail for someone else to see the same failure. In the simulated lab, asking for the platform, build, scene, and expected behaviour immediately separated a game bug from a discussion about the future of labour. Both topics survived. Only one could be stepped through in a debugger.

Table 1. Mock issue-tracker triage

Incoming signal	First useful question	Likely next action
“This floor runs terribly.”	Which floor and GPU?	Capture frame times
“AI broke the mouse.”	Is the cursor visible?	Reproduce focus change
“Windows is ruined.”	Which build and settings?	Run a Windows regression
“Everything is fine here.”	Which scene did you test?	Record the test scope
“This README is wordy.”	Where did you stop reading?	Delete three paragraphs

What the lab would merge

A short repro and a fix with a test beat a glorious statement of principles. A principled criticism with evidence belongs in review. A ten-page paper about excessive writing may be considered for publication if its authors agree to stay out of the README.

DISCUSSION / BENCHMARK

Frame time and reply time

Puzzle motion needed interpolation to look smooth at higher framerates. The comment thread uses interpolation too: one screenshot is enough for someone to fill in six hours of development history. Only the game's version has to land on the correct square.

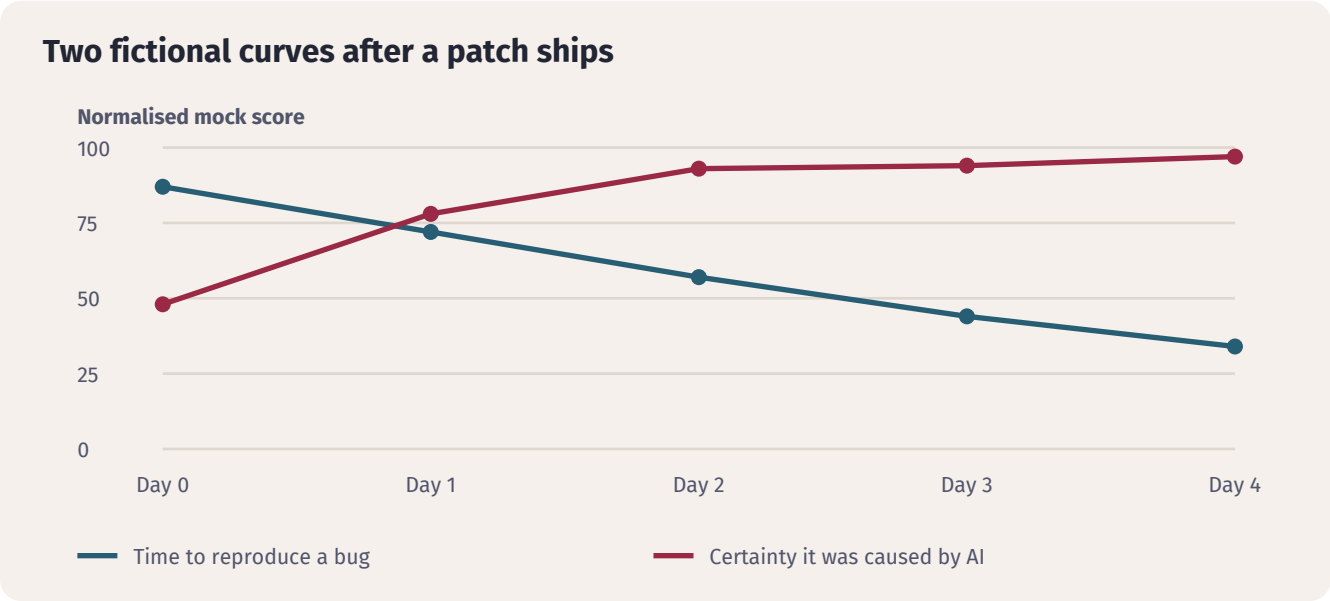


Figure 6. Illustrative trends. Time to reproduce is scaled to 0-100. Certainty is scored separately. The lines share paper, not units. All values are invented.

What would count as evidence?

For performance, compare the same scene, settings, resolution, and hardware. For software quality, reproduce the bug and read the change. For credit and provenance, look at what people disclose and what the repository actually contains. Asking any of these questions is fair. Treating the answer to one as the answer to all three is how a benchmark becomes fan fiction.

The lab recommends local co-op as a field test. When both players share a couch, a disagreement generally ends when somebody needs the other person to move a block.

PEER REVIEW / FICTIONAL CORRESPONDENCE

Comments to the authors

Reviewer 1 / Major revisions

The authors write as though the only criticism is noise. Some of us have actual questions about review, attribution, and maintenance. Also, on page six you say “lightweight abstraction” and leave us to pay for it.

Response

Agreed on the questions. We have added explicit criteria for an actionable review. We have also revised the assistant’s proposed abstraction down to a function call and a cup of tea. The tea is load-bearing.

Reviewer 2 / Accept pending benchmark

I tested the port before reading the paper. The game is fun. I would like to see identical hardware and settings before anyone compares the curves in Figure 6 to a real benchmark.

Response

All figures were invented. The bench was also invented. We briefly considered making the bench real, but neither author could agree which operating system to boot first.

Editor’s decision

Accepted for the *Journal of Reproducible Outrage*, subject to one condition. The next software bug should be filed with a log. The next objection to this paper should be filed with a better joke. Authors are reminded that the journal’s impact factor is measured in closed browser tabs.

CONCLUSIONS AND DISCLOSURES

Everyone can file a better issue

The fictional record supports three modest conclusions. Critics improve software when they are specific. Maintainers improve arguments when they describe what the build actually does. ChatGPT improves both discussions when it checks its work and gives the conversation back to the humans before offering a roadmap for civilisation.

Practical implications

Enjoy the game if it runs. Report the problem if it breaks. Ask where the code came from if that matters to you. Scott will be in Stage 9 checking whether a fraction of a millisecond has moved house.

Declarations

Ethics approval. There were no participants, real quotes, or real reviews.
Data availability. All numbers are fabricated and printed with the figures.
Funding. No grant was awarded. One evening was lost.
Competing interests. Scott wants the port to work. ChatGPT wants this sentence to have a parallel structure. Both desires were factored into the jokes.

Selected fictional references

[1] S. Maintainer. “One Last Build.” *Proceedings of the Same Evening*, 1(12), pp. 18:00-02:00.
[2] R. Two. “Please Attach Your Frame-Time Trace.” *Journal of Reasonable Requests*, 4(1), pp. 1-2.
[3] C. GPT. “A Concise Answer in Seven Parts.” *Transactions on Unnecessary Expansion*, 9(7), pp. 1-419.

End of paper. You may now close the tab at an uncapped framerate.