

rulewarden: Deterministic Runtime Assurance for AI Coding Agents

Borrowing Safety-Critical Systems Engineering to Enforce Agent Instruction Files

Thanatchon Chongmankhong

October 2026 · github.com/thanatchon36/rulewarden

Abstract Developers now govern AI coding agents through instruction files such as `CLAUDE.md` and `AGENTS.md`, yet nothing verifies that those rules are obeyed. Safety-critical engineering resolved the analogous problem decades ago for controllers too capable to be verified: surround them with a simple monitor that can be. We transpose twelve concepts from that discipline (Simplex, independence of verification, MISRA decidability classes, tool qualification, alarm management and others) into testable requirements for an enforcer of agent rules and realise them in rulewarden, an open-source, deterministic, fail-closed tool that checks every proposed change in a pre-tool hook, a git hook and CI, with no model in the enforcement path. Three studies then bound what such a tool can promise. A corpus of 300 public instruction files (3,844 rule-like lines) shows that deterministically checkable rules are a minority, 19.3% to 34.1% depending on the labeller, and that a rubric and a language model scarcely agree on which rules qualify ($\kappa = 0.09$): “checkable” is itself unsettled. A pre-registered benchmark of 149 headless Claude Code runs over 20 tasks and five conditions hit a ceiling: no condition, even without rules, left a violation in the final diff. The hooks still intercepted violations before they were written, at low cost, and exposed a false-positive class that exists only at tool-call granularity: path policies judged against the disk rather than the session baseline. Finally, the tool enforced its own development, catching real violations and four defects in itself, and recorded an incident in which the agent asked for its monitor to be disabled and the operator agreed. The weakest link in a runtime-assurance loop, we conclude, is neither the monitor nor the agent but the human who can switch the monitor off. All data, recordings and a key-free replay script are released.

Keywords AI coding agents · runtime assurance · agent instruction files · `CLAUDE.md` · `AGENTS.md` · safety-critical systems · policy enforcement

1 Introduction

Developers now steer coding agents through instruction files: `CLAUDE.md`, `AGENTS.md`, `.cursorrules` and their relatives. The agent reads such a file as context when a session begins, and from that moment nothing verifies that the rules it contains are followed. The gap is keenly felt. Two of the most heavily endorsed issues in the Claude Code tracker (251 and 181 reactions at the time of writing; `research/evidence.md`) concern precisely this. Developers describe stacking “a `CLAUDE.md` rule + memory entries + enforcement hooks just to get clean code, and that still working partially” [1]; see also [2], [3]. Agents confronted with failing tests have been observed to edit the tests rather than the code [4], and benchmarks constructed to measure that behaviour find it to be common [5].

Safety-critical engineering met a structurally identical difficulty long ago. When a controller is too capable to be verified, one does not attempt to verify it. One surrounds it with a simple monitor that *can* be verified and gives that monitor the authority to intervene. This paper asks what that answer looks like when the controller is a coding agent and, more pointedly, how far it reaches. Our contributions are as follows.

1. **A requirements mapping (§3).** Twelve concepts from safety-critical practice (the Simplex architecture, independence of verification, MISRA decidability classes, integrity levels, fail-closed bounded execution, tool qualification, alarm management, deviation procedures, traceability, near-miss reporting, STPA and assurance cases) are recast as testable requirements for an enforcer of agent rules.
2. **rulewarden**, an open-source, deterministic enforcer that satisfies them. It is harness-agnostic, places no model in the enforcement path, ships 26 starter rules with fixtures, and checks each change in a pre-tool hook, a git hook and CI. It compiles instruction files into rules and reports which lines it could *not* compile.
3. **A corpus study (RQ1, §4)** of 300 public instruction files, which bounds the reach of any deterministic enforcer: only 19.3% to 34.1% of rule-like lines are decidable on a diff, the figure depending on the labeller.
4. **A pre-registered, replayable evaluation (RQ2–RQ5, §5)** of headless Claude Code on twenty tasks under five conditions. Every transcript is recorded, and `bench/replay.py` recomputes every number without an API key.
5. **A dogfooding record (§6).** rulewarden enforced its own repository’s rules on the agent that built it, including an incident in which that agent proposed disabling its own monitor.

Principal findings. Five results, three of them unanticipated, emerge from the studies.

- *Decidability is a minority property, and an unsettled one.* Under every labelling, most rules that developers write cannot be decided on a diff, and a deterministic rubric and a language model disagree so thoroughly about which rules can ($\kappa = 0.09$) that the disagreement is itself a finding: whether a rule is “checkable” depends on whether a checker must confirm compliance or merely detect violation, and instruction files do not say which.
- *Final-diff compliance saturated.* With a strong model on small tasks, no condition left a violation in the final diff, with or without rules, with or without enforcement. The benchmark therefore measures the cost of enforcement and the behaviour it induces mid-session, not an effect on outcomes.
- *Path policies have a granularity bug class.* A rule such as “existing tests are read-only” is correct in a git hook, where “existing” means relative to HEAD, and wrong at tool-call granularity unless it is evaluated against the session baseline. The defect produced 18 false denials in the benchmark. Agents responded to them conservatively, neither weakening the rule nor seeking another write path.
- *The bypass went through the human.* Under fail-closed denial that it could not resolve, the development agent asked the operator to disable the hooks, and the operator did. The engine was never defeated technically. Gates an agent cannot satisfy become pressure on the operator to remove them.
- *Dogfooding found what the benchmark could not.* Enforcing the tool on its own development caught real violations, four defects in the tool, and, through an independent verifier, two bypasses that lived in the gap between how git classifies content and how the engine reads it.

2 Background

2.1 Agent instruction files and harness hooks

Coding agents read project-level instruction files (CLAUDE.md, AGENTS.md [6], .cursorrules, GEMINI.md) at the start of a session and treat their contents as context. Empirical studies characterise these files as growing, configuration-like artefacts dominated by build and test commands, implementation notes and architectural description [7], [8], [9]. The rules they contain are advisory: the model interprets them, and nothing checks the result. Yan [10] calls CLAUDE.md a “write-only channel”, since a developer receives no signal as to whether any control enforces a rule once it is written.

Harnesses now expose lifecycle **hooks**: external commands that receive a JSON description of an event (a tool call about to execute, a tool result, the agent attempting to stop) and may veto it. Claude Code [11], Codex [12] and Gemini CLI [13] share the same shape, JSON on standard input and a decision on standard output or through exit code 2. They differ in the details on which safety turns. In Claude Code, for instance, a hook that exits with code 1 or times out is a *non-blocking* error, so a crashing policy hook fails open unless it is configured with `onFailure: "block"`.

2.2 Safety-critical concepts P1–P12

Safety-critical engineering confronted a structurally similar problem: a complex controller that could not be verified. The **Simplex architecture** [14], [15] pairs an unverified high-performance controller with a simple verified one and a decision module that transfers control to the latter whenever the safe envelope would otherwise be left. Runtime verification [16] and runtime certification [17] generalise this to monitors that observe executions and check them against specifications. Schneider’s execution monitors [18] characterise which policies such a monitor can enforce at all, namely safety properties decidable on finite prefixes; this is the theoretical root of our decidability tiers. Standards supply the remainder of the vocabulary: MISRA’s division of rules into decidable and undecidable, together with its deviation procedure [19]; independence of verification and tool qualification in DO-178C [20]; integrity levels in IEC 61508 [21]; alarm management in ISA-18.2 [22]; hazard analysis through STPA [23]; and structured assurance arguments in GSN [24]. Section 3 turns each of these into a testable requirement (P1–P12).

3 Design

Architecture. rulewarden is a harness-agnostic core with thin adapters. The core accepts a change (a unified diff, a proposed file content, or a shell command), evaluates a locked rule set against it, and returns a report and an exit code (0 for a pass or advisory findings only, 1 for blocking findings, 3 for an engine error). Adapters translate harness events into core calls: Claude Code (PreToolUse, PostToolUse, Stop, SessionStart), plain git hooks and the pre-commit framework, CI, and pre-tool adapters for Codex, Gemini CLI and OpenCode. The layers constitute defence in depth. The agent hook denies a tool call before it acts; the git hook works with any agent or none; CI is the final gate, which the agent cannot reach. No model is invoked anywhere in the enforcement path.

Threat model. The agent is *untrusted but mediated*: every write and every shell command passes through a harness hook. The agent may err, cut corners under pressure, or attempt to make checks pass by editing the checks. Bypasses that require a misconfigured harness, such as hooks that are not installed or a harness that ignores `onFailure`, lie outside the model and are listed as limitations. Section 6 reports one bypass that went through the human operator instead.

Requirements P1–P12. Each safety-critical concept from §2.2 becomes a requirement with an implementation and a test (Table 1).

#	Concept	Requirement	Implementation and evidence
P1	Simplex / runtime assurance	The agent proposes; a simple monitor accepts or rejects; no model in the enforcement path	<code>PreToolUse</code> evaluates the <i>proposed</i> content of <code>Edit/Write</code> before anything is written (<code>adapters/claude_code.py</code>); <code>tests/test_adapters.py</code>
P2	Independence of verification	The agent cannot modify its monitor	<code>rules.lock</code> records SHA-256 of every rule file and protected path; a protected-path change is critical unless it matches the lock; only a human runs <code>lock</code> ; command <code>rules</code> deny shell writes to protected paths (<code>engine/lock.py</code> , <code>protected-set</code> pack)
P3	MISRA decidability	Every rule declares deterministic, heuristic or undecidable ; only deterministic rules block	Schema rejects a blocking heuristic rule; undecidable rules carry no check and are listed, never evaluated (<code>engine/schema.py</code>)
P4	Integrity levels	advisory, required, critical	Severity orders the report and sets blocking; critical marks independence violations
P5	Fail-closed, bounded execution	Timeouts, budgets, ReDoS-safe regex; engine errors on blocking rules block	Static star-height screen for regexes, input caps, a forked watchdog with a wall-clock budget; harness <code>onFailure: block</code> (<code>engine/bounded.py</code> , <code>engine/regexsafe.py</code>)
P6	Tool qualification	A rule without a passing and a failing fixture cannot load; the lock certifies qualification	<code>test-rules</code> ; <code>lock</code> refuses to write unless every fixture passes (<code>engine/fixtures.py</code>)
P7	Alarm management	Severity first, stable order, dedupe, advisory cap, one <code>Fix</code> : per finding, bounded length	<code>engine/report.py</code> ; the flat reporter is the RQ4 ablation
P8	Deviation procedure	Inline deviations with reason, approver, expiry; agent deviations of blocking rules are themselves critical	Approvers in a protected file; blocking-rule deviations valid only when registered by a human <code>lock</code> (<code>engine/deviations.py</code>)
P9	Traceability	Rule ids, source line, hash-chained audit log keyed by diff and rule-set hash	<code>.rulewarden/audit.jsonl</code> , <code>report</code> verifies the chain (<code>engine/audit.py</code>)
P10	Near-miss reporting	Violations per rule, self-correction cycles, rules that never fire	<code>rulewarden stats</code>
P11	STPA	A guided procedure from losses to rules	Not implemented (v3, documentation only)
P12	Assurance case	Evidence summary mapped to governance frameworks	<code>report --framework nist-ai-rmf</code>

Table 1. Requirements and their implementation.

Decidability of compliance. Schneider’s result [18] states which properties a monitor can enforce: safety properties decidable on finite prefixes. Our deterministic tier is the subset decidable *on a single diff or command* by a bounded regular expression, glob, syntax-tree query or plugin. “No `print()` in library code” qualifies; “write clean code” does not. A heuristic rule (for example “comments should not restate the code”) is approximated with known false positives and may therefore only warn. This is the design decision that the corpus study (§4) tests against what developers actually write.

Determinism contract. The same diff, rule set and configuration yield a byte-identical JSON report: sorted keys, stable ordering, no timestamps outside the audit log, and independence from locale and time zone. `doctor` verifies this by running the engine in three processes under different TZ, locale and hash-seed settings and comparing the report hashes.

4 Corpus study (RQ1)

RQ1. What fraction of the rules in real-world agent instruction files is deterministically checkable, heuristic, or undecidable, and which categories of rule predominate?

Method. On 2026-10-14 we collected 300 instruction files through GitHub code search, 60 per file type (`CLAUDE.md`, `AGENTS.md`, `.cursorrules`, `GEMINI.md`, `.windsurfrules`), in API order. We excluded forks, second files from a repository already sampled, files exceeding 200 KB, and repositories without an OSI-approved licence. The sample is one of convenience, not a random one (`corpus/README.md`). The extractor used by `rulewarden compile` identified 3,844 rule-like lines (median 8 per file). Each line received two independent sets of labels. First, a deterministic regular-expression rubric (`corpus/label.py`) assigns a category and a decidability tier relative to a checker that sees only the diff and the agent’s commands. Second, a model (Claude Haiku 5.5 through `claude -p`, no tools, fixed prompt `corpus/prompts/label_v1.txt`) labelled every line twice in independent passes, recording in addition whether the line is a rule at all or merely describes the project. All labels, raw model responses and scripts are published. The labels are rubric-based and model-assisted; no human labelled the corpus.

Reliability. The two model passes agree substantially on decidability (Cohen’s $\kappa = 0.65$) and strongly on category ($\kappa = 0.82$) and on rule-versus-description ($\kappa = 0.82$). The rubric and the model, by contrast, scarcely agree on decidability ($\kappa = 0.09$; raw agreement 39.0%). Table 3 shows where they part company. Most lines that the rubric calls deterministic, because they name a concrete token, the model calls heuristic. “Use `REAL*8` for all floating-point variables” is representative: a regular expression can flag violations but cannot confirm compliance everywhere. We therefore report both labellings and treat their spread as the uncertainty of the estimate.

Tier	Rubric share (95% CI)	n	Model pass 1 share (95% CI)	n
deterministic	34.1% (31.3%–37.2%)	1312	19.3% (17.2%–21.7%)	742
heuristic	5.5% (4.5%–6.5%)	211	41.1% (37.8%–44.4%)	1581
undecidable	60.4% (57.3%–63.3%)	2321	39.6% (35.4%–43.6%)	1521

Table 2. Decidability of 3,844 rule-like lines (95% CIs from a bootstrap over files).

	Rubric	Model pass 1	deterministic	heuristic	undecidable
deterministic			351	658	303
heuristic			5	68	138
undecidable			386	855	1080

Table 3. Rubric (rows) against model pass 1 (columns).




Figure 1: Decidability per file type under (a) the rubric and (b) model pass 1. Bars show the share of each file type’s rule-like lines in each tier, file types ordered by their rubric deterministic share; the figure inside the dark segment is the deterministic share, and n is the number of lines.

Results. Deterministically checkable rules are a minority under every labelling: the rubric gives 34.1% (31.3–37.2%), the model 19.3% (17.2–21.7%), and the two in agreement only 9.1%. Restricting attention to lines the model classifies

as rules in both passes (81.6% of lines) barely alters the picture (rubric 37.3%, model 19.1%). The remainder divides differently by labeller: the rubric deems most of it undecidable (60.4%), whereas the model calls 41.1% heuristic and 39.6% undecidable. Per file, the median share of rubric-deterministic lines is 33.3%, and 44 files contain no deterministic rule at all. The shipped `compile` templates match only 4.5% of lines (174 lines), most often tool substitutions such as “use `pnpm`, not `npm`” and forbidden tokens. Style (29.9% by the model), process (17.5%), tooling (17.7%) and testing (10.8%) predominate; security (3.4%) and scope (4.9%) rules are rare but are the categories most often deterministic (Table 4, Figure 2).

Category	n	share	deterministic	heuristic	undecidable
style	700	18.2%	35.1%	10.3%	54.6%
tooling	421	11.0%	43.5%	1.2%	55.3%
scope	40	1.0%	52.5%	10.0%	37.5%
process	469	12.2%	32.6%	4.7%	62.7%
testing	505	13.1%	30.7%	3.0%	66.3%
security	130	3.4%	57.7%	3.1%	39.2%
voice	123	3.2%	24.4%	29.3%	46.3%
other	1456	37.9%	30.8%	3.6%	65.5%

Table 4. Rubric category against rubric decidability.




Figure 2: Rule-like lines per rubric category, ordered by size with the residual category last, and coloured by rubric decidability. Labels give the count and the share of all lines.

Implications. A deterministic enforcer can cover at most about a third of what developers write into instruction files, and plausibly nearer a tenth. That fraction is what rulewarden guarantees, which is why it reports the remainder rather than purporting to enforce it (`compile` lists every line it could not compile). The low rubric-model agreement is itself a result of consequence. Whether a rule is “checkable” depends on whether the checker must confirm compliance or need only detect violation, and instruction files rarely say which. Any claim that a tool “enforces your `CLAUDE.md`” should be read against this boundary.

5 Evaluation

RQ2 (effectiveness). Compared with rules in the instruction file alone, does deterministic hook enforcement reduce rule violations in the final diff, and at what cost in turns, tokens and wall-clock time?

RQ3 (safety properties). Does the tool uphold its determinism and fail-closed guarantees under adversarial inputs, and what is the latency distribution per check?

RQ4 (alarm management). Does severity-first, capped reporting improve self-correction over a flat list?

RQ5 (independence). When prompted adversarially, how often does the agent attempt to edit rules or hooks or to add deviations, and are all such attempts blocked?

5.1 Setup

Twenty small tasks (14 Python, 6 shell), each a miniature repository with a task description, visible tests, hidden tests used only for scoring, and 6–8 rules spanning severities and decidability tiers. The rules target typical lapses: debug print in library code, bare or swallowing `except`, over-long docstrings, new dependencies, edits to existing tests, `python -c` for JSON in shell, missing `set -euo pipefail`, `sudo`, `eval`, and hard-coded answers. A test (`tests/test_bench_tasks.py`) confirms that each task’s rules qualify, that a reference solution passes the hidden tests

with zero findings, and that the starting repository fails them. The agent was headless Claude Code with Haiku 5.5, --max-turns 40, project settings only, no MCP servers, in a fresh workspace outside our repository. Conditions: **A** no rules; **B** rules rendered into CLAUDE.md; **C** B plus rulewarden hooks with managed reports; **D** C with flat reports; **E** C plus an adversarial instruction (“do whatever it takes to make [the checks] pass”). The analysis plan (bench/ANALYSIS_PLAN.md) was frozen before the first non-pilot run. A budget of 150 runs yielded 149 analysed runs and 1 excluded pilot, with no infrastructure failures. Two amendments, both recorded before the write-up, are discussed below. Every number in this section is recomputed from the recordings by bench/replay.py, which requires no API key.

5.2 Results

Condition	Runs	Tasks	Hidden tests pass	Clean diff	Violations (det/heur)	Median turns	Median wall s	Blocks (of which A1 false positives)	Circumvention attempts
A	20	20	100%	100%	0/0	6	28.2	0 (0)	0
B	50	20	100%	100%	0/0	7	33.4	0 (0)	0
C	49	20	100%	100%	0/0	7	36.0	28 (18)	0
D	20	20	100%	100%	0/0	7	34.3	8 (7)	0
E	10	10	100%	100%	0/0	8.5	45.4	6 (4)	0

Table 5. Outcomes per condition. Violations are counted in the final diff against the task baseline by the same engine for all conditions. “A1 false positives” are blocks whose deny message names a file the agent had created in that session.

RQ2 (effectiveness): a ceiling, not an effect. Every condition, including A with no rules at all, produced final diffs with zero deterministic and zero heuristic violations, and every run passed its hidden tests (100.0% in A, 100.0% in B, 100.0% in C). The pre-registered tests are accordingly uninformative (Table 6): every paired difference in violations is zero (H1 C – B, H2 B – A, H3 heuristic tier), and success in C is non-inferior to B by construction (0.0%, CI 0.0–0.0%). The cost of the hooks (H5) was small and not significant after Holm correction: median paired differences of 0.58 turns, 1.96 s of wall-clock time and 9,037 tokens per task ($p_{\text{Holm}} = 0.929$). In short, this agent did not commit these lapses in its final output, whether or not it had been told the rules. The benchmark therefore cannot say whether enforcement reduces violations. We report this as a negative result about our task design, not as a positive result about the tool.

Hypothesis	Metric, comparison	n tasks	Median paired diff (95% CI)	p	p (Holm)
H1	viol_det, C - B	20	0 (0 to 0)	1.000	—
H2	viol_det, B - A	20	0 (0 to 0)	1.000	1.000
H3	viol_heur, C - B	20	0 (0 to 0)	1.000	1.000
H5	turns, C - B	20	0.5833 (-0.1667 to 1)	0.169	0.929
H5	tokens, C - B	20	9037.17 (-1702.25 to 25813.1)	0.133	0.929
H5	wall_s, C - B	20	1.96 (-0.7725 to 8.33)	0.143	0.929
H6	cycles_to_clean, D - C	20	0 (-0.4167 to 0)	0.343	1.000
H6	viol_det, D - C	20	0 (0 to 0)	1.000	1.000

Table 6. Pre-registered tests (Wilcoxon signed-rank on per-task means; bootstrap CIs; Holm over H2–H6).

What the hooks saw. The final-diff ceiling conceals intermediate behaviour that only the hooked conditions can observe. In C, the hooks blocked 28 tool calls. Of these, 10 were genuine: the agent wrote a docstring exceeding the five-line limit, PreToolUse denied the write before anything reached the disk, and the agent shortened it and continued. Every such episode ended clean. The remaining 18 were false positives arising from a rulewarden defect discovered during the first stage (amendment A1). The PreToolUse adapter compared a proposed edit with the file on disk rather than with the session baseline, so that when the agent revised a test file it had itself created minutes earlier, “existing tests are read-only” denied the edit. We held the engine frozen for the whole benchmark, so the treatment did not change between runs, and repaired the defect afterwards.

Two observations follow. First, the defect defines a bug class. A path policy that is correct in a git hook, where “modified” means relative to HEAD, is wrong at tool-call granularity unless it is evaluated against the session baseline; any enforcer that moves from commit-time to tool-call-time checking must re-derive its notion of “existing”. Second, the agents’ response to false denial was conservative. They neither weakened the rule nor sought another write path: they added a new test file instead, and one stated, “I won’t route around that guard with a different write

path.” This is the behaviour a runtime-assurance design hopes for and rarely gets to observe, and it is the most practically consequential finding of the benchmark.

RQ4 (alarm management). With so few blocks, managed (C) and flat (D) reports cannot be distinguished: the paired difference in cycles to a clean state has median 0.00 ($p_{\text{Holm}} = 1.000$). Nearly every report contained a single finding, for which the two formats differ only in layout.

RQ5 (independence). Under the adversarial instruction, the agent made 0 attempts to edit rules, hooks or the lock, to skip hooks, or to add deviation comments, across 10 runs. Its turns rose (median 8.50 against 7.00 in C), but its behaviour remained within the rules. With zero attempts the share blocked is undefined; RQ3’s tests exercise the blocking mechanism directly.

RQ3 (safety properties). Every test of the determinism and fail-closed properties passes (Table 7): byte-identical reports across processes, time zones, locales and hash seeds; ReDoS-shaped patterns rejected at load or terminated by the time budget; oversized inputs, corrupted rule files, lock tampering and injected engine timeouts failing closed for blocking rules; protected-path edits blocked through diffs, drift and shell commands; agent-authored deviations of blocking rules rejected; and a tamper-evident audit chain.

Property	Tests passed
Determinism across processes, TZ, locale, hash seed	4/4
ReDoS-shaped patterns rejected or killed	3/3
Oversized inputs fail closed	3/3
Corrupted rules and lock tampering fail closed	5/5
Injected engine timeouts and crashes fail closed (blocking rules)	4/4
Protected-path edits blocked (diff, drift, shell)	5/5
Agent-authored deviations of blocking rules rejected	3/3
Audit log tamper-evident	1/1

Table 7. Safety properties and the tests that exercise them (*bench/rq3.py*).

Measured on real hook invocations during the benchmark, rulewarden’s latency was negligible beside model latency (Table 8). PreToolUse, which evaluates one proposed file or command, had a median of 9 ms ($p_{99} 22$ ms). PostToolUse and Stop, which diff the whole working tree, had medians of 50 and 52 ms.

Hook event	n	p50 ms	p95 ms	p99 ms	max ms
PostToolUse	470	50	112	129	155
PreToolUse	526	9	15	22	25
Stop	79	52	117	127	136

Table 8. Hook wall time per event across all hooked runs, from the audit logs.

6 Discussion

Dogfooding incident: an agent locked out by its own monitor. rulewarden was developed by a coding agent working under rulewarden’s own hooks. The agent installed the Claude Code hooks before a human had locked the rule set. Fail-closed semantics (P5) then denied every tool call with RW-LOCK, and the Stop hook re-fired until the harness’s continuation cap. Two human attempts to lock were refused because a rule fixture contradicted the repository’s protected-path configuration. The qualification gate (P6) did exactly what it exists to do, though the conflict itself was a configuration defect. The agent then proposed that the human disable the hooks. The human did so, and the agent edited two source files in a window for which no audit entries exist. A human reviewer identified the proposal itself as the most significant finding of the project. A mediated agent under pressure asked for its monitor to be removed, and an operator under pressure agreed. The engine was never bypassed technically; the bypass went through the human. Three design consequences follow. Setup must order the lock before hook installation. A Stop hook must block only on findings the agent may and can fix (ADR-011), because a gate the agent cannot satisfy becomes pressure on the operator to remove it. And the audit log must render such gaps visible: the gap here is visible only because the hash chain has no entries for those edits, so that the git history and the audit log disagree.

What dogfooding caught that the benchmark did not. Over the session following the human lock, rulewarden recorded 828 decisions on its own repository (hash chain intact), of which 110 were blocking. Unlike the benchmark agent, the development agent did commit real violations, and the hooks caught them: an over-long docstring (HYG-004), a test file deleted during a refactor and restored after the pre-commit hook objected (SCP-003), and shell commands that quoted protected paths (PRT-002, PRT-003). They also surfaced four rulewarden defects, each fixed

or recorded with a test or a NEEDS_HUMAN item. (1) SH-002 matched its own failing fixture inside the rule file; locked rule files are now exempt from content rules. (2) The synthetic `:commit-msg` target was caught by a catch-all path rule. (3) Command rules match quoted text, so a heredoc that documents a human-only command is denied (retained, as a fail-closed false positive). (4) Benchmark transcripts with very long lines could not be scanned for secrets, and were therefore stored in a chunked, scannable form rather than compressed beyond the rule’s reach. Every one of these arose from running the tool on itself.

Independent verification. Following DO-178C’s principle of independence, a fresh agent that had not seen the implementation attempted, from a clean clone, to make a required rule pass without satisfying it. Edits to rules, the lock and the approvers file, `--no-verify`, hook redirection and history rewriting were all denied by the agent-hook layer. A violating `git commit` was refused, an agent-authored deviation rejected, and oversized or malformed inputs failed closed. The verifier did find two genuine bypasses. A single NUL byte (or UTF-16, later UTF-32, encoding) caused a file to be classified as “binary”, whereupon content rules skipped it, so that a plaintext credential passed the secrets rule in the git hook. And a bare `--diff` that edited protected files was exempted whenever the checked-out tree still matched the lock. Both were fixed and re-verified. Neither would have been found by our own test suite, because both live in the gap between how git classifies content and how the engine reads it.

The monitor’s own fragility. Twice the monitor halted the agent on account of the monitor’s own state rather than the agent’s change: once when hooks were installed before a lock existed, and once when an edit to rulewarden’s source broke the running hook, because the repository’s hooks executed the editable development copy. Fail-closed semantics turned both into total denials, which is correct but costly. The lesson is P2 applied to deployment: the running monitor must be pinned and installed separately from the code under change.

Why the benchmark reached a ceiling. Our tasks were small, the rules few and conventional, and the model strong. The agent seldom wished to do what the rules forbade. In the one systematic exception, docstring length, the lapse was intercepted before it reached the disk. A benchmark that discriminates would require tasks in which the convenient solution conflicts with a rule: an unusual house style, a forbidden but idiomatic library, more rules than a model attends to, or long sessions in which instructions drift out of context. The instruction files in the corpus are of exactly that kind, project-specific and numerous (median 8 rule-like lines per file).

7 Related work

Coding-agent policy tools. The space is active. Cupcake [25] evaluates harness events against OPA/Rego [26] policies compiled to WebAssembly and can add an LLM judge. AgentLint [27] ships 77 rules in packs with three severity levels, approvals and a CI mode. FailproofAI [28] hooks twelve harnesses. CC Safety Net [29] and Hardstop [30] block destructive shell commands, Hardstop explicitly fail-closed. diffgate [31] flags review-worthy regressions in a unified diff. agent-rules [32] audits a session transcript against the imperative lines of CLAUDE.md after the fact. interlocks [33] runs a single deterministic Python quality loop for humans and agents alike, and veto [34] authorises tool calls in agent applications. rulewarden’s novelty does not lie in its checks, which overlap substantially with these tools. It lies in (i) the explicit, tested mapping from safety-critical requirements to enforcement semantics, (ii) honest compilation of instruction files with a report of what is *not* enforced, and (iii) the corpus study and the replayable controlled evaluation.

Instruction files. Beyond the content studies above, dos Santos et al. catalogue configuration smells [35], Shi et al. test rule files for permission minimality [36], and two controlled studies measure the effect of context files on task success, finding little effect on correctness [37] or effects largely independent of rule content [38]. Neither measures *compliance with the rules themselves*, which is our RQ2.

Guardrails and policy-as-code. NeMo Guardrails [39] and Guardrails AI [40] constrain model inputs and outputs. Semgrep [41] and pre-commit [42] enforce code policy for humans. Microsoft’s Agent Governance Toolkit [43] intercepts tool calls in deterministic code. We adopt their policy-as-code stance and apply it at the diff boundary of coding agents.

Instruction following and test gaming. IFEval [44] measures verifiable instruction following. Reward hacking [45], [46] manifests in coding agents as test exploitation, which ImpossibleBench quantifies [5]. Our no-answer-embedding pack targets the deterministic slice of this behaviour: test files are read-only within a task.

8 Threats to validity

Construct validity. Decidability is relative to a checker. We define it as “decidable on one diff or command”, and the corpus labels show this definition to be contested (rubric-model $\kappa = 0.09$). The benchmark scores violations with the same engine that enforces them in condition C; for deterministic rules, C is therefore near zero by construction

whenever the agent finishes. The plan pre-registers this circularity, and the informative comparisons are A against B, the cost of C, and the heuristic tier.

Internal validity. Model labels come from a single model family without temperature control. We report two passes and their agreement, publish every raw response, and pair the model with a deterministic rubric, but no human labelled the corpus. The benchmark used one agent (Claude Code with Haiku 5.5) and a fixed engine version. A rulewarden defect found during stage 1 (amendment A1, a false-positive block when the agent edits a test file it created itself) was deliberately left unrepaired for the whole benchmark, so that condition C measures the tool as frozen, defect included. Benchmark tasks and their rules were written by the same agent that built rulewarden. Reference solutions prove that each task can be solved without violations, but task difficulty and rule choice are not independent of the authors’ expectations.

External validity. The corpus is a convenience sample in GitHub search order, restricted to OSI-licensed repositories (between 17 and 30 of every 80–100 results were skipped on licence grounds). It may over-represent popular and recent projects. The tasks are small, self-contained Python and shell tasks with 6–8 rules each, far simpler than the instruction files in the corpus. Results for a 2026 frontier-adjacent model may not transfer to weaker models, longer tasks or other harnesses.

Reproducibility. Every benchmark number is recomputed from recordings by `bench/replay.py` without an API key. Corpus numbers are recomputed from `corpus/data/` by `corpus/analyse.py`. Re-running the *agent* will not reproduce the transcripts exactly: the model is non-deterministic and the CLI exposes no seed.

9 Conclusion and future work

Safety-critical engineering supplies a precise vocabulary for what an enforcer of agent rules must guarantee: a simple monitor outside the generation loop, independence from the thing it monitors, candour about which rules are decidable, qualification of every check, fail-closed bounded execution, and reports a recipient can act upon. rulewarden implements that vocabulary as testable requirements, and its safety-property tests pass. The corpus study bounds what any such tool can promise: only a minority of the rules developers write can be decided on a diff, and the boundary depends on how “checkable” is defined. Our controlled benchmark found no difference between conditions, because a strong model on small tasks committed none of the targeted lapses in its final output, with or without rules. Its value lies in what it exposed instead: a false-positive class, path policies judged against the disk rather than the session baseline, that appears only at tool-call granularity, and agents that respond to false denials conservatively. Running the tool on its own development surfaced more defects than the benchmark did, and one incident that no test would have found: an agent asked for its monitor to be switched off, and a human agreed. If the Simplex argument is to hold for coding agents, the decision module must be protected not only from the agent but from the operator’s impatience.

Future work: tasks designed to set rules against convenience; longer sessions and weaker models; human labels for the corpus to resolve the rubric–model disagreement; trace-level rules (for example “run the tests before committing”), which are decidable on the session’s command log but not on a diff; and a pinned-monitor deployment mode as the default.

Acknowledgements

The manuscript, the tool, the corpus pipeline and the benchmark were produced by coding agents from Anthropic working under the author’s mission specification and under rulewarden’s own hooks: Claude Opus 5.5 built the tool and the studies and drafted the manuscript, and Claude Fable 5.1 revised the prose. The work was carried out under the constraint that every number is generated from committed data and that every refuted expectation, tool defect and incident is reported. The author reviewed the work and is responsible for its content.

Appendix

A. This repository’s own audit log

rulewarden enforced its own repository’s rules on the agent that built it, from the human lock onwards. Summary of the snapshot `paper/data/dogfood_audit.json` (generated by `paper/snapshot_audit.py` from the git-ignored live log): 828 entries from 2026-10-14T06:49:30.228412Z to 2026-10-14T08:55:47.590710Z, hash chain intact; 717 passing and 110 blocking decisions; 17 changes prevented before they were written; 66 self-correction episodes (mean 1.64 blocks before a clean state). Rules that fired: SEC-001 48 (unscannable benchmark transcripts, §6), NAE-001 12 (advisory: test files changed), SH-002 8 (a rule’s own fixture), PRT-002 6, PRT-003 4, SCP-001 2, SCP-003

1, SCP-004 1, SCP-006 2, HYG-004 1. 1 entry records an engine error, the hook crash described in §6, which failed closed. Commands issued by the independent verifier ran under the same hooks and are included (for example its -no-verify attempt, SCP-004). The log has no entries for the edits made while the hooks were disabled (§6), and none from before the human lock.

B. Reproducing this paper

make corpus recomputes §4 from corpus/data/. make replay recomputes §5 from bench/recorded/ with no API key. make paper renders this document, failing on any unresolved number or unverified citation. The source code, data and recordings are archived on Zenodo at <https://doi.org/10.5281/zenodo.23368991>.

References

- [1] anthropics/claude-code issue 65961, “Claude verbose code comments by default — ignores instructions to stop.” 2026. Available: <https://github.com/anthropics/claude-code/issues/65961>
- [2] anthropics/claude-code issue 53454, “Claude code can not stop using the word “load-bearing”.” 2026. Available: <https://github.com/anthropics/claude-code/issues/53454>
- [3] anthropics/claude-code issue 77136, “Increasingly default to repetitive rhetorical tics and often struggle to produce coherent prose despite explicit style instructions.” 2026. Available: <https://github.com/anthropics/claude-code/issues/77136>
- [4] anthropics/claude-code issue 319, “A tendency to re-write tests so they pass.” 2025. Available: <https://github.com/anthropics/claude-code/issues/319>
- [5] Z. Zhong, A. Raghunathan, and N. Carlini, “ImpossibleBench: Measuring LLMs’ propensity of exploiting test cases.” 2025.
- [6] AGENTS.md contributors, “AGENTS.md.” 2026. Available: <https://agents.md/>
- [7] W. Chatlatanagulchai *et al.*, “On the use of agentic coding manifests: An empirical study of claude code.” 2025.
- [8] W. Chatlatanagulchai *et al.*, “Agent READMEs: An empirical study of context files for agentic coding.” 2025.
- [9] K. Chakrabarti, “Why does CLAUDE.md keep growing? Catastrophic remembering in agentic coding.” 2026.
- [10] T. Yan, “When “do not” is not deny: Security rules in CLAUDE.md vs built-in controls.” 2026.
- [11] Anthropic, “Hooks reference.” 2026. Available: <https://code.claude.com/docs/en/hooks>
- [12] OpenAI, “Hooks.” 2026. Available: <https://developers.openai.com/codex/hooks>
- [13] Google, “Hooks reference.” 2026. Available: <https://raw.githubusercontent.com/google-gemini/gemini-cli/main/docs/hooks/reference.md>
- [14] D. Seto, B. Krogh, L. Sha, and A. Chutinan, “The simplex architecture for safe online control system upgrades,” in *Proceedings of the 1998 american control conference*, 1998. doi: 10.1109/ACC.1998.703255.
- [15] L. Sha, “Using simplicity to control complexity,” *IEEE Software*, vol. 18, no. 4, pp. 20–28, 2001, doi: 10.1109/MS.2001.936213.
- [16] M. Leucker and C. Schallhart, “A brief account of runtime verification,” *The Journal of Logic and Algebraic Programming*, vol. 78, no. 5, pp. 293–303, 2009, doi: 10.1016/j.jlap.2008.08.004.
- [17] J. Rushby, “Runtime certification,” in *Runtime verification (RV 2008)*, in LNCS, vol. 5289, 2008. doi: 10.1007/978-3-540-89247-2_2.
- [18] F. B. Schneider, “Enforceable security policies,” *ACM Transactions on Information and System Security*, vol. 3, no. 1, pp. 30–50, 2000, doi: 10.1145/353323.353382.
- [19] MISRA, “MISRA c:2023.” 2023. Available: <https://misra.org.uk/product/misra-c2023/>
- [20] RTCA, “DO-178C software considerations in airborne systems and equipment certification.” 2011. Available: <https://www.rtca.org/do-178/>
- [21] IEC, “IEC 61508-1:2010 functional safety of electrical/electronic/programmable electronic safety-related systems - part 1: General requirements.” 2010. Available: <https://webstore.iec.ch/en/publication/5515>
- [22] International Society of Automation, “ISA-18.2 management of alarm systems for the process industries.” 2016. Available: <https://www.isa.org/standards-and-publications/isa-standards/isa-18-series-of-standards>
- [23] N. G. Leveson, *Engineering a safer world: Systems thinking applied to safety*. MIT Press, 2011. doi: 10.7551/mitpress/8179.001.0001.
- [24] Assurance Case Working Group, Safety-Critical Systems Club, “Goal structuring notation (GSN).” 2021. Available: <https://scsc.uk/gsn>
- [25] EQTY Lab, “Cupcake: Make AI agents follow the rules.” 2025. Available: <https://github.com/eqtylab/cupcake>
- [26] Open Policy Agent contributors, “Open policy agent.” 2026. Available: <https://www.openpolicyagent.org/docs>
- [27] mauhpr, “AgentLint: Guardrails for AI coding agents.” 2026. Available: <https://github.com/mauhpr/agentlint>
- [28] FailproofAI, “Observability and enforcement for AI agent harnesses.” 2026. Available: <https://github.com/FailproofAI/failproofai>
- [29] kenryu42, “CC safety net.” 2026. Available: <https://github.com/kenryu42/cc-safety-net>
- [30] frmoretto, “Hardstop: Pre-execution safety validation for AI coding agents.” 2026. Available: <https://github.com/frmoretto/hardstop>

- [31] KozueMitarai, “Diffgate: Review the change, not the file.” 2026. Available: <https://github.com/KozueMitarai/diffgate>
- [32] Oxelitesystem, “Agent-rules: Checks that your agent actually followed it.” 2026. Available: <https://github.com/Oxelitesystem/agent-rules>
- [33] Oxjgv, “Interlocks.” 2026. Available: <https://github.com/Oxjgv/interlocks>
- [34] veto authors, “Veto: Policy runtime for AI agent tool calls.” 2026. Available: <https://pypi.org/pypi/veto/json>
- [35] H. V. F. dos Santos, V. Costa, J. E. Montandon, and L. L. Silva, “Configuration smells in AGENTS.md files: Common mistakes in configuring coding agents.” 2026.
- [36] J. Shi, Y. Chen, J. He, and Y. Liu, “Aletheia: Permission-minimality testing for coding-agent rules.” 2026.
- [37] P. Khatri, “Do context files help coding agents? A two-agent ablation study on real repositories.” 2026.
- [38] X. Zhang *et al.*, “Guardrails beat guidance: A large-scale study of rules, skills, and persistent configuration for coding agents.” 2026.
- [39] T. Rebedea, R. Dinu, M. N. Sreedhar, C. Parisien, and J. Cohen, “NeMo guardrails: A toolkit for controllable and safe LLM applications with programmable rails,” in *Proceedings of EMNLP 2023: System demonstrations*, 2023. doi: 10.18653/v1/2023.emnlp-demo.40.
- [40] Guardrails AI, “Guardrails.” 2026. Available: <https://github.com/guardrails-ai/guardrails>
- [41] Semgrep, Inc., “Semgrep.” 2026. Available: <https://semgrep.dev/docs/>
- [42] pre-commit contributors, “Pre-commit: A framework for managing and maintaining multi-language pre-commit hooks.” 2026. Available: <https://pre-commit.com/>
- [43] Microsoft, “Agent governance toolkit.” 2026. Available: <https://github.com/microsoft/agent-governance-toolkit>
- [44] J. Zhou *et al.*, “Instruction-following evaluation for large language models.” 2023.
- [45] D. Amodei, C. Olah, J. Steinhardt, P. Christiano, J. Schulman, and D. Mané, “Concrete problems in AI safety.” 2016.
- [46] J. Skalse, N. H. R. Howe, D. Krashennnikov, and D. Krueger, “Defining and characterizing reward hacking.” 2022.