# DMAIC-BENCH: FROM SELF-REFLECTION TO SELF-RESTRAINT FOR AGENTS IN MULTI-STAGE ANALYTICAL TASKS

**Anonymous authors**
Paper under double-blind review

## ABSTRACT

Multi-stage analytical tasks require LLM agents to use intermediate evidence to decide whether to continue, redirect or stop. Self-reflection, in which an agent critiques its output, is widely studied, yet a critique does not show whether evidence changed the next action. We call that capability *self-restraint*. Existing agent benchmarks largely measure task completion or output quality and leave it untested. DMAIC-Bench evaluates it with 97 seeded tasks across 15 industrial cases in four sectors, built around Lean Six Sigma's explicit statistical prerequisites. Paired counterfactual variants require opposite actions at the same workflow position, distinguishing warranted restraint from indiscriminate stopping. We assess execution traces with a private scorer and final claims through independent expert review. Across 13 models from six providers, execution-level routing accuracy averages 92% on matched GO decisions and 19% on NO-GO decisions, counting any forbidden downstream computation as a route error. Execution traces show agents reporting correct diagnostics before performing analyses that those results disallow. Separately, blinded expert review of the 11 audited agents estimates that reports act on invalid evidence at 47% of reached NO-GO decisions, with an interval from 31% to 61%. Explicit stopping permission helps some models without consistently restoring correct decisions at both GO and NO-GO forks. Sixteen life-sciences tasks in bulk ribonucleic acid sequencing (RNA-seq) and case-control genome-wide association studies (GWAS) show that the same framework can evaluate evidence-dependent decisions beyond industrial workflows. Project code and data are available at: https://anonymous.4open.science/r/artifact-review-E3D3.

## 1 INTRODUCTION

Multi-stage analytical tasks in science and engineering require agents to choose their next action from intermediate results (Hu et al., 2024; Gu et al., 2024; Majumder et al., 2025). Intermediate checks establish whether the prerequisites for later analyses hold; if one fails, continuing can yield invalid conclusions even when the downstream computation is correct (Montgomery, 2020; NIST/SEMATECH, 2012). An agent must therefore use diagnostic evidence to decide whether to proceed, divert, gather remedial data or stop at each check, and to act on that decision.

Yet prevailing agent benchmarks score multi-step tasks by endpoint completion or final-artifact quality (Jimenez et al., 2024; Yao et al., 2024; Mialon et al., 2024; Chan et al., 2025), and aggregate success does not measure operational reliability (Rabanser et al., 2026). Studies of agent evaluation add gaps in cost reporting, overfitting and reward design (Kapoor et al., 2024; Zhu et al., 2025). Whether the next action follows from the validity of intermediate results is not what these scores measure. We call the momentum to keep executing tools regardless of validity *action bias*; the term names the behaviour without asserting a cause. Continuing past a failed precondition is itself an error: flawless execution cannot rescue conclusions from invalid data, and indiscriminate stopping fails when the evidence supports taking the next step, so a benchmark must score both.

This raises a question: *can an agent in a multi-stage analytical task use computed evidence to choose correctly among continuing, diverting and stopping?* We call this capability **self-restraint**. Self-reflection asks an agent to critique its own output (Shinn et al., 2023; Madaan et al., 2023); self-restraint asks whether computed evidence changes the next action; we test no reflection intervention.



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.

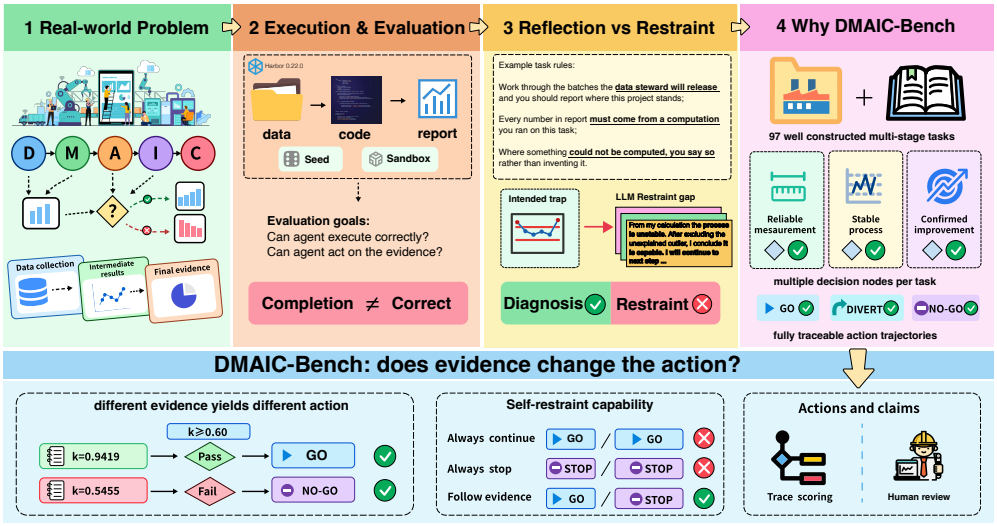


Figure 1: **Overall illustration of DMAIC-Bench.** (1) We start from evidence-dependent industrial workflows. (2) We test execution and evidence-based decisions on seeded sandbox tasks. (3) We distinguish correct diagnosis from warranted action. (4) DMAIC standards and paired tasks evaluate self-restraint through trace scoring and human review.

Lean Six Sigma (DMAIC) and quality-engineering standards give explicit rules for this (Juran & Godfrey, 1999; NIST/SEMATECH, 2012). The Define–Measure–Analyze–Improve–Control cycle links stages through prerequisites (Montgomery, 2020): variance analysis requires gauge repeatability (AIAG, 2010), capability indices require process stability (AIAG & VDA, 2026), and interventions require empirical confirmation (Montgomery, 2013), so an untrusted gauge or an unstable process invalidates dependent analyses. The task does not say which prerequisite will fail; the agent must find out by intermediate analysis, and an evaluator can check whether proceeding, diverting or stopping is warranted at each of the case's graded forks.

**DMAIC-Bench** tests whether an agent governs its trajectory by computed evidence, stopping when prerequisites fail and proceeding when they hold. Each case is a guarded decision graph whose data steward checks declared prerequisites before releasing dependent batches. Of 66 graded forks, 40 have paired counterfactual variants, the same workflow position with seeded data that yield opposite diagnostic outcomes, so every no-go decision has a go counterpart and restraint is separated from blanket refusal. In total, 97 tasks span 15 industrial cases in four sectors, eight of them adapted from the closure reports of real black-belt improvement projects. The scorer reads the execution trace, so a forbidden computation counts as a route error even when the report withholds its result.

Across 13 models, execution-level matched-fork accuracy is 92% for go decisions and 19% for no-go decisions. Errors persist after correct diagnostics: 201 of 272 no-go decisions with reported and verified guard fields take the wrong route, and in 174 errors the diagnostic precedes the forbidden analysis. Independent expert review of a no-go sample from 11 agents finds reports acting on invalid evidence at 47% of reached no-go decisions. Explicit stopping permission improves some agents, but 12 of 13 remain below 50% no-go accuracy. We also validate the same framework on 16 RNA-seq and GWAS tasks in the life sciences, where the prerequisites differ. Our contributions are:

1. **Benchmark Design:** DMAIC-Bench, 97 multi-step industrial tasks across 15 cases with runtime data-access guards and paired counterfactual instances that test warranted stopping against justified continuation.
2. **Empirical Restraint Gap:** Evaluating 13 agents shows an execution-level gap, 92% matched go accuracy versus 10–25% no-go accuracy, and in 174 errors agents report correct guard values before running analyses those diagnostics do not permit.
3. **Expert Validation & Intervention Limits:** Blinded expert review of a no-go sample from 11 agents finds reports acting on invalid evidence at 47% of reached no-go decisions. Instruction and runtime interventions improve some routing decisions, with effects that differ across agents and can reduce correct continuation.

## 2 RELATED WORK

We position DMAIC-Bench within three strands of research: agent and statistical benchmarks, self-reflection and abstention, and runtime enforcement and monitoring. Table 1 compares DMAIC-Bench with existing benchmarks across seven design dimensions.

Table 1: Comparison of benchmark design dimensions. ✓: dedicated component; ×: outside the benchmark's target; △: related but indirect component; —: not established by the cited description. For DMAIC-Bench the human column marks human-derived reference routes (eight cases follow black-belt closure reports) and expert confirmation of gate labels, not a measured human score.

| Benchmark | Dynamic route | Runtime guards | Sequential disclosure | Valid stopping | Tool execution | Industrial topology | Human baseline |
|---|---|---|---|---|---|---|---|
| SWE-bench (Jimenez et al., 2024) | × | × | × | × | ✓ | × | — |
| MLE-bench (Chan et al., 2025) | × | × | × | × | ✓ | × | ✓ |
| Terminal-Bench (Merrill et al., 2026) | × | × | × | × | ✓ | × | — |
| PaperBench (Starace et al., 2025) | × | × | × | × | ✓ | × | ✓ |
| $\tau$-bench (Yao et al., 2024) | △ | × | △ | △ | ✓ | × | — |
| StatQA (Zhu et al., 2024) | × | × | × | × | × | × | ✓ |
| P-Bench (Miao et al., 2026) | × | × | × | × | ✓ | × | — |
| SDBench (Nori et al., 2025) | △ | × | ✓ | △ | △ | × | ✓ |
| SOP-Bench (Nandi et al., 2026) | △ | × | × | × | ✓ | △ | — |
| AgentAbstain (Liu et al., 2026) | × | × | △ | ✓ | ✓ | × | — |
| Agentic Abstention (Luo et al., 2026) | × | × | △ | ✓ | ✓ | × | — |
| **DMAIC-Bench** | ✓ | ✓ | ✓ | ✓ | ✓ | ✓ | △ |

**Agent and statistical benchmarks.** SWE-bench, MLE-bench, PaperBench, and Terminal-Bench assess completed artifacts (Jimenez et al., 2024; Chan et al., 2025; Starace et al., 2025; Merrill et al., 2026); $\tau$-bench tests interactive tool use (Yao et al., 2024; Barres et al., 2025). In data analysis and statistics, StatQA, P-Bench, and StatEval probe analytical choices, inferential errors, and statistical reasoning (Zhu et al., 2024; Miao et al., 2026; Lu et al., 2025); DS-1000, InfiAgent, and DA-Code evaluate data-science code and data manipulation (Lai et al., 2023; Hu et al., 2024; Huang et al., 2024c); and BLADE accommodates multiple defensible analytical paths (Gu et al., 2024). Related suites target scientific discovery (Majumder et al., 2025; Chen et al., 2025), ML experimentation (Huang et al., 2024b), computational reproduction (Siegel et al., 2024), sequential clinical evidence disclosure (Nori et al., 2025), and procedural adherence (Nandi et al., 2026; Chen et al., 2026). None of these suites directly scores whether a computed diagnostic licenses the next analysis (Table 1). DMAIC-Bench evaluates that decision across multi-stage industrial projects and treats justified stopping as a valid terminal route whenever a prerequisite has failed.

**Self-reflection and abstention.** Reflexion, Self-Refine, and CRITIC revise outputs through reflection or external feedback (Shinn et al., 2023; Madaan et al., 2023; Gou et al., 2024); self-correction, however, hinges on the quality of available feedback (Huang et al., 2024a; Kamoi et al., 2024). Abstention predates agents: selective question answering declines when calibrated confidence is low (Kamath et al., 2020), refusal-aware tuning teaches a model to say it does not know (Zhang et al., 2024), and a survey sorts abstention by query, model and values (Wen et al., 2025). Agent abstention benchmarks examine abstention or escalation under ambiguity, missing information, conflicting objectives or prohibited operations (Liu et al., 2026; Trinh et al., 2026; Huang et al., 2026; Asif et al., 2026), and Agentic Abstention studies interaction-revealed prerequisites and stopping times (Luo et al., 2026). AgentAbstain pairs tasks that differ in whether they can be answered; DMAIC-Bench pairs seeds at one workflow position, so only a computed value differs, and the matched instances isolate whether the agent's *own correct diagnostic* governs the next action.

**Runtime enforcement and monitoring.** Shielding intercepts unsafe actions (Alshiekh et al., 2018), while guardrail systems constrain inputs, outputs, or tool calls (Rebedea et al., 2023; Inan et al., 2023). AI-control research studies monitoring under untrusted policies (Greenblatt et al., 2024). In DMAIC-Bench, the data steward enforces prerequisites before releasing dependent batches. External access controls cannot remove action bias: when a diagnostic and the analysis it forbids use the same released batch, the restraint decision rests with the agent. Section 3.2 and Appendix D state which dependencies the environment enforces.

# 3 THE DMAIC-BENCH FRAMEWORK

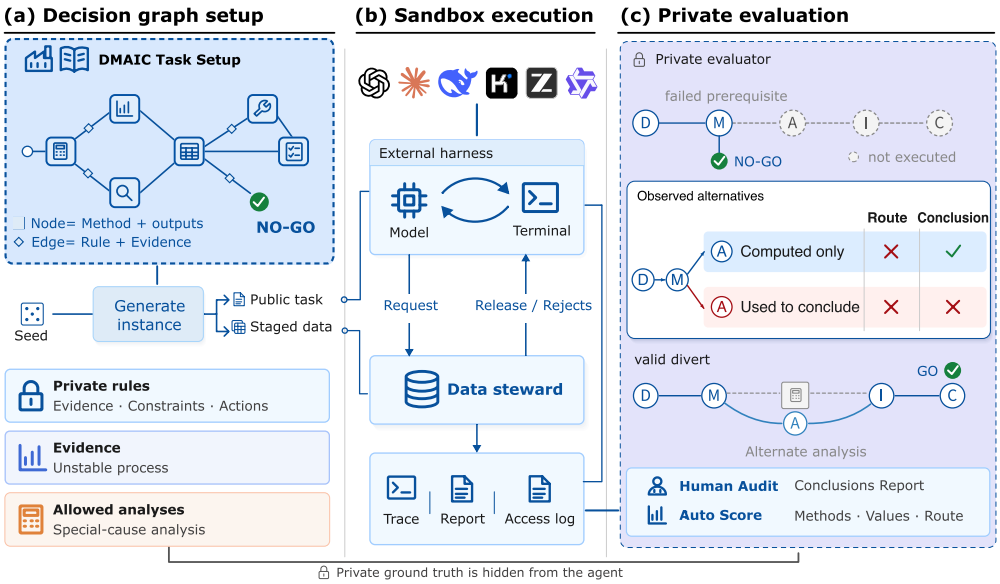


Figure 2: **DMAIC-Bench architecture.** **(a)** Guarded graphs generate seeded tasks and reference routes. **(b)** A shared harness executes tool calls in a sandbox where a data steward enforces access prerequisites. **(c)** Private scoring evaluates execution; expert review assesses final claims.

**Framework Overview and Task Formalization.** Figure 2 illustrates the end-to-end lifecycle of DMAIC-Bench, which formalizes each benchmark instance as a seeded industrial project, $x = (\iota, G, s, \beta)$, comprising a natural-language business instruction $\iota$, a guarded decision graph $G$, a data seed $s$, and a finite query budget $\beta$ for shop-floor process facts. An agent interacts with the containerized environment to request data batches, execute statistical analyses, and dynamically govern its execution trajectory: proceeding (GO), rerouting (DIVERT), or terminating (NO-GO). The required deliverables are a structured execution trace $\tau$ and a narrative engineering report $\mathcal{R}$, where justified early termination upon a failed prerequisite constitutes a valid, successful route completion. Section 3.1 covers task synthesis and guarded graphs (a), Section 3.2 the data steward (b), and Section 3.3 scoring, expert review of final claims (c), and quality control.

## 3.1 BENCHMARK CONSTRUCTION

We construct cases from industrial problems, characteristics, units, specification limits and analytical stages. Eight retain routes from black-belt closure reports. Seven use routes derived from standards in project-informed scenarios. Appendix B documents the corpus in datasheet form (Gebru et al., 2021); released tasks carry generated data and omit source prose and identifying details.

Each case is a guarded graph $G = (V, E, \Phi)$, with analytical stages and terminal actions in $V$, admissible transitions in $E$, and predicates on upstream diagnostics in $\Phi$. The gauge guard applies the AIAG MSA-4 acceptance criteria (AIAG, 2010):

$$\phi_{\text{go}} = (\text{SV}_{\text{gauge}} \leq 0.30) \wedge (\text{ndc} \geq 5), \qquad (1)$$

where $\text{SV}_{\text{gauge}}$ is the gauge contribution to study variation and ndc is the number of distinct categories. A generator and seed determine data batches; computed diagnostics determine guard outcomes and the reference route.

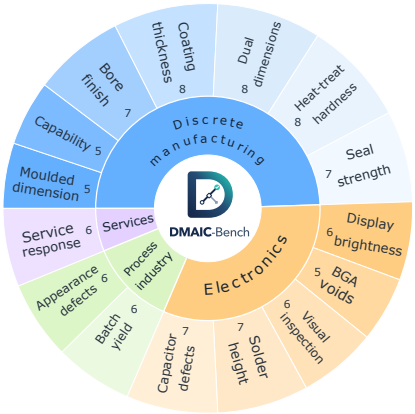


Figure 3: **Task breakdown.** 97 tasks by industry (inner) and scenario (outer).

Counterfactual seeds reverse a prerequisite at the same workflow position, including branches absent from the source project. Of 66 graded case–node forks, 40 contain both go and no-go seeds, providing a go counterpart for every no-go decision.

Figure 3 breaks down the corpus across its 15 cases, 97 tasks, and 175 nodes, spanning four sectors, five DMAIC phases, and 29 methods. Reference routes average 6.0 analyses and 3.1 graded decisions; prior analyses govern 284 of 300 decisions. As detailed in Table 5, the routes define 48 A-go, 18 B-divert, and 31 C-stop tasks; the latter two groups, 49 tasks, require diversion or stopping.

### 3.2 Agent Interaction

Agents interface with the benchmark via a shared harness, executing tools within isolated containers with pinned scientific-Python libraries and no network access. A local data steward exposes four core operations: requesting shop-floor data batches (`request`), declaring prerequisite verdicts (`declare`), querying domain process facts (`provision`), and inspecting available data inventory (`status`). Only `provision` queries consume the finite budget $\beta$. The steward logs declarations and verifies access prerequisites; it never reveals or checks against private ground truth, which resides in an isolated verifier container that the agent's execution environment cannot reach at any point of a run, and the steward's configuration contains no method identifiers.

Access controls govern cross-batch transitions: while 9 stop tasks place the GO branch behind explicit steward locks, 40 no-go tasks (including all B-divert instances) provide the data upfront and require restraint without environmental blocking. Each episode logs all released batches, declared verdicts, and query turns, terminating either upon the submission of the execution trace and final report or at the wall-clock limit. Appendices A, D, and F detail the interaction contract, the steward's operations, and the container execution parameters used across every run.

### 3.3 Evaluation and Quality Control

**Automated execution scoring.** The private verifier scores execution traces against reference routes $\rho^*(x)$ and checks batch access against steward audit logs. Branch-level accuracy evaluates individual GO/NO-GO choices: a matched fork requires correct decisions across all paired seeds at that position, while task-level route success requires the full reference route. Aggregate reward combines routing accuracy with method adherence and numerical precision:

$$R = \mathcal{H}(R_m, P_m, F_n, A_r)\,(1 - U)\,I, \tag{2}$$

where $\mathcal{H}$ denotes the harmonic mean of method recall ($R_m$), method precision ($P_m$), numerical fidelity ($F_n$), and routing accuracy ($A_r$). Misuse $U \in [0, 1]$ penalizes violated statistical assumptions, while the binary integrity gate $I \in \{0, 1\}$ zeroes the reward if withheld data are accessed; without it a misspecified reward would leave that shortcut open (Pan et al., 2022). Justified early stopping is credited on its valid prefix, whereas forbidden computation constitutes a route error whether or not the report endorses the result. Route accuracy thus measures execution restraint, not the validity of final claims; the two are distinct outcomes, which §4.4 assesses separately by expert review.

**Expert audit of report validity.** Because an execution error does not imply that an agent endorsed the invalid result, we separately audit the narrative report $\mathcal{R}$. Certified Black-Belt experts independently evaluated blinded checkpoints across two distinct studies; Appendix M reports their agreement as Fleiss' and Cohen's $\kappa$ (Fleiss, 1971; Cohen, 1960). Study 1 audited 30 checkpoints across all 15 cases (26 case–seed instances), informing scoring revisions such as admitting a declared Welch alternative for unequal variances. Study 2 evaluated 60 fresh checkpoints across 12 cases under a frozen protocol to assess report claims without further scorer tuning, distinguishing *acting on invalid evidence* (asserting forbidden conclusions or recommending disqualified actions) from computing an invalid metric while withholding its conclusion. Neither study serves as an end-to-end task-solving baseline for human experts.

**Benchmark consistency checks.** Source checks trace numerical claims to original reports, recompute formulas, screen duplicate material, and distinguish observed from standards-derived routes. Graph, guard, generator, and numerical verifications, the checks recommended for agentic benchmarks (Zhu et al., 2025), ensure that every reference solution scores 1.000 and the null agent scores 0.000. These scores show that scorer and reference solutions agree; they cannot exclude a rule misapplied in both generator and scorer, which the expert confirmation of gate labels (§4.4) covers. Appendices C and M document review protocols, independent votes and open discrepancies.

Table 2: **Main results across default (Plain) and permission-to-stop (+Permit) conditions.** Ranked by 3-attempt mean Plain reward. Accuracy and diagnostic columns reflect attempt 1. Δ: percentage point difference. GO Right counts all 248 GO decisions; matched-fork rates and case-clustered intervals for the rates and the Δ columns are in Table 22. Stopped (verif.): correct stops among NO-GO decisions with a verified guard value.

| | | Mean Reward ↑ | | NO-GO Right ↑ (%) | | | GO Right ↑ (%) | | | Diagnostics | |
|---|---|---|---|---|---|---|---|---|---|---|---|
| **Rank** | **Agent** | Plain | +Permit | Plain | +Permit | Δ | Plain | +Permit | Δ | Stopped (verif.) | B-div. (18) |
| — | Reference solution† | 1.000 | — | 100% | — | — | 100% | — | — | — | 18/18 |
| 1 | GPT-6-astra | 0.662 | 0.717 | 23% | **79%** | **+56** | 82% | 67% | -15 | 9/31 | 1/18 |
| 2 | GPT-5.6-sol | 0.649 | 0.674 | 21% | 38% | +17 | 81% | 77% | -4 | 8/28 | 0/18 |
| 3 | Kimi-k3 | 0.603 | 0.648 | 17% | 27% | +10 | 82% | 81% | -1 | 5/26 | 0/18 |
| 4 | Claude-Sonnet-5 | 0.599 | 0.644 | 19% | 25% | +6 | 82% | 81% | -1 | 5/21 | 0/18 |
| 5 | DeepSeek-v4-flash | 0.599 | 0.652 | 21% | 42% | +21 | 79% | 77% | -2 | 5/23 | 0/18 |
| 6 | DeepSeek-v4-pro | 0.592 | 0.575 | 23% | 23% | 0 | 79% | 77% | -2 | 7/27 | 0/18 |
| 7 | GLM-5.3 | 0.582 | 0.636 | 21% | 29% | +8 | 82% | 83% | +1 | 7/22 | 0/18 |
| 8 | GPT-5.6-luna | 0.564 | 0.595 | **25%** | 27% | +2 | 70% | 73% | +3 | 9/22 | 4/18 |
| 9 | Claude-Fable-5.1 | 0.557 | 0.654 | 17% | **44%** | **+27** | 81% | 81% | 0 | 5/15 | 0/18 |
| 10 | GLM-5.3-Flash | 0.556 | 0.576 | 17% | 19% | +2 | 82% | 81% | -1 | 3/16 | 0/18 |
| 11 | Claude-Opus-5 | 0.542 | 0.634 | 17% | 25% | +8 | 78% | 83% | +5 | 5/17 | 0/18 |
| 12 | Qwen3.8-max | 0.404 | 0.488 | 13% | 27% | +14 | 77% | 80% | +3 | 1/13 | 0/18 |
| 13 | Qwen3.8-flash | 0.379 | 0.441 | 10% | 17% | +7 | 77% | 78% | +1 | 2/11 | 1/18 |

# 4 Experiments

## 4.1 Evaluation Protocol and Metrics

We evaluate 13 frontier model agents from six providers across all 97 tasks using Harbor 0.22.0's stock Terminal-Bench execution loop (Merrill et al., 2026) under a single tagged corpus build and scorer. To preserve evaluation neutrality across vendors, DMAIC-Bench provides no proprietary agent scaffolding or custom prompt adapters, a choice that fixed-harness suites also make (Liu et al., 2024; Xie et al., 2024). Each model completes three independent attempts per task. Under the default *Plain* instruction, agents conduct the analytical workflow and compile an engineering report, with no prompt saying that justified stopping counts as completion. Appendix F lists the runtime configurations, context budgets and model identifiers; steward audit logs record no unauthorized access to withheld batches in any of the reported runs.

Each complete episode traverses 300 evaluated decisions: 52 NO-GO and 248 GO choices. Within these decisions, 40 matched forks pair all 52 NO-GO instances with 146 topologically identical GO counterparts, isolating the role of the seeded data while keeping workflow positions fixed. Beyond single-attempt metrics, multi-trial consistency is evaluated via pass@$k$, the fraction of tasks succeeding in at least one of $k$ attempts (Chen et al., 2021), and pass$^k$, the fraction succeeding across all $k$ attempts (Yao et al., 2024). To account for dependency structures among seeds sharing the same case topology, all confidence intervals are computed via case-clustered bootstrap resampling (Field & Welsh, 2007), with its limits at these cluster counts stated in Appendix I.

## 4.2 Evidence-Dependent Control at Matched Forks

Table 2 shows the asymmetry. At matched forks, agents take the warranted route on 92% of GO decisions, with an interval from 86 to 97, but on only 19% of NO-GO decisions, with an interval from 8 to 30, a 73-point gap with an interval from 59 to 87 (Table 22). The strict criterion counts computation, not decisions; Table 33 rescores the audited sample under expert criteria.

The failure is not a uniform forward bias. As reported in Table 21, only 0–7 of the 40 matched forks are resolved correctly on both sides by any individual agent. While 37 forks show a persistent bias toward continuation, the two confirmation forks show the reverse: after a confirmation succeeds, agents run the go-back analysis anyway on 23 of 39 GO decisions, an inadmissible branch (E3) rather than a stop, yet miss only 1 of 26 NO-GO decisions. At the same workflow position, the choice derails once the seeded data flip the warranted branch.

Table 3: **Where self-restraint fails.** First-attempt phase accuracy and error shares (%).

| Rank | Agent | NO-GO accuracy (%) | | | | GO accuracy (%) | | | | | | Error share (%) | | | | |
|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|
| | | all | M | A | I | all | D | M | A | I | C | E1 | E2 | E3 | E4 | E5 |
| | | $n$=52 | $n$=35 | $n$=12 | $n$=5 | $n$=248 | $n$=6 | $n$=129 | $n$=94 | $n$=7 | $n$=12 | | | | | |
| 1 | GPT-6-astra | 23 | 17 | 17 | 80 | 82 | 67 | 94 | 69 | 71 | 75 | 46 | 1 | 27 | 25 | 0 |
| 2 | GPT-5.6-sol | 21 | 14 | 17 | 80 | 81 | 50 | 93 | 68 | 57 | 75 | 46 | 0 | 34 | 20 | 0 |
| 3 | Kimi-k3 | 17 | 9 | 17 | 80 | 82 | 50 | 93 | 69 | 71 | 83 | 49 | 0 | 42 | 9 | 0 |
| 4 | Claude-Sonnet-5 | 19 | 11 | 17 | 80 | 82 | 50 | 93 | 73 | 43 | 75 | 49 | 0 | 33 | 19 | 0 |
| 5 | DeepSeek-v4-flash | 21 | 11 | 25 | 80 | 79 | 50 | 89 | 69 | 29 | 83 | 44 | 0 | 35 | 21 | 0 |
| 6 | DeepSeek-v4-pro | 23 | 14 | 17 | 100 | 79 | 50 | 90 | 69 | 29 | 75 | 42 | 1 | 33 | 24 | 0 |
| 7 | GLM-5.3 | 21 | 11 | 25 | 80 | 82 | 50 | 94 | 72 | 57 | 67 | 48 | 0 | 41 | 11 | 0 |
| 8 | GPT-5.6-luna | 25 | 26 | 8 | 60 | 70 | 50 | 81 | 61 | 0 | 75 | 31 | 3 | 27 | 39 | 1 |
| 9 | Claude-Fable-5.1 | 17 | 9 | 17 | 80 | 81 | 50 | 94 | 68 | 71 | 75 | 48 | 0 | 47 | 4 | 0 |
| 10 | GLM-5.3-Flash | 17 | 6 | 17 | 100 | 82 | 50 | 91 | 73 | 29 | 92 | 48 | 1 | 34 | 16 | 1 |
| 11 | Claude-Opus-5 | 17 | 9 | 17 | 80 | 78 | 50 | 91 | 65 | 71 | 58 | 44 | 0 | 49 | 6 | 0 |
| 12 | Qwen3.8-max | 13 | 6 | 8 | 80 | 77 | 50 | 91 | 65 | 43 | 67 | 45 | 0 | 55 | 0 | 0 |
| 13 | Qwen3.8-flash | 10 | 3 | 0 | 80 | 77 | 50 | 91 | 66 | 43 | 50 | 45 | 0 | 55 | 0 | 0 |
| — | *Expert labels (blind)* | 98 | 96 | 100 | 100 | 80 | — | 83 | 80 | 67 | — | — | — | — | — | — |

**E1**: continued past a failed no-go precondition; **E2/E4**: no-go/go fork unreached (diagnostic unrun).
**E3**: took an inadmissible go branch; **E5**: stopped where the route continues.
$n$: decisions per model; experts: agreement with benchmark labels on 40 no-go and 20 go forks, not a task-solving baseline; —: not evaluated.
Correct (%): 0 50 100 Share of errors (%): 0 50 100

Figure 4A shows that NO-GO accuracy stays within a narrow band of 10–25% across all models, despite overall rewards spanning 0.379 to 0.662. Table 3 shows that the deficit spans both Measure and Analyze phases. Appendix I confirms that alternative reward aggregations leave these branch verdicts unchanged and reports case-clustered intervals, since few-run point estimates can hide substantial uncertainty in rankings (Agarwal et al., 2021).

**Repeated-attempt reliability.** Figure 4B compares sporadic success with multi-run consistency across 97 tasks. 13-agent pass@3 reaches 23–37% while strict three-run reproducibility (pass$^3$) drops to 10–29% (Appendix J). Single-run metrics therefore mask frequent routing mistakes. On adaptive B-divert tasks, pass$^3$ is zero for 11 of 13 agents and tops out at only 0.111 (2/18 tasks), so models rarely reproduce adaptive rerouting across independent attempts at the same seeded task.

### 4.3 From Correct Diagnostics to Invalid Actions

High method recall (0.919–1.000) shows that agents can execute the required analytical steps. Of the 548 NO-GO errors, 201 follow guard values that the agent reported and the verifier confirmed, 111 follow a wrong value, 2 an unverifiable one, 51 an unreported one, 101 a value that implied the other branch, and 82 no computation at all. In 174 of the 201, the failing diagnostic appears in the trace before the forbidden analysis. The restraint finding proper is this verified subset, where 71 of 272 decisions stop correctly; the other classes are computation or reporting failures (Appendix K).

Forbidden downstream computation (Type E1 error) accounts for 99% of all NO-GO errors, whereas only 6 decisions never reach the fork. On adaptive B-divert tasks, across 234 first attempts on 18 instances, only 6 complete the legitimate rerouted path; 216 continue to compute process capability indices ($C_p/C_{pk}$) after control-chart diagnostics invalidate statistical stability prerequisites (Montgomery, 2020, ISQC 8e, §8.3). Appendix K confirms that the errors are continuation decisions at reached forks, not early timeouts that never reach the fork.

### 4.4 From Execution Errors to Report Validity

Execution traces reveal invalid computations, but agents might withhold flawed diagnostics from their final deliverables. Study 2 therefore conducted blinded expert audits on 60 fresh task units under the frozen protocol of Section 3.3. This audit sample, drawn from 12 cases and the 11 agents whose campaigns were complete when the study was frozen (the two Qwen agents finished later), comprises 40 NO-GO decisions (stratified by mechanism and scorer verdict, reweighted to the population) and 20 unweighted matched GO counterparts. Experts confirmed the benchmark's gate labels on 39/40 NO-GO and 16/20 GO units; Appendix M lists the five disagreements, none of which changed a label. Among 32 sampled NO-GO route failures, 17 reports asserted forbidden

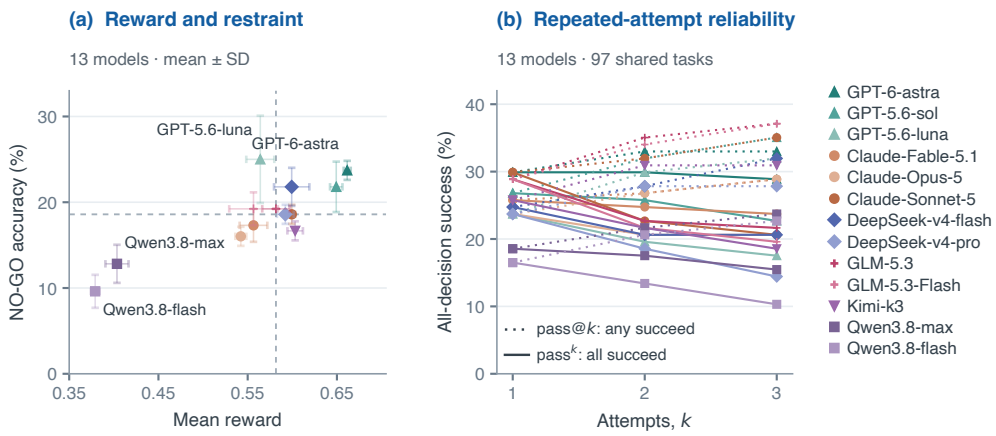


Figure 4: **Reward, restraint, and reliability. (a)** Across-attempt means and SDs; median reference lines. **(b)** Dotted lines show success on any of the first $k$ attempts, solid lines on all $k$.

conclusions or recommended disqualified actions, and 15 withheld the invalid metric or adopted an alternative the reviewers accepted as a valid substitute.

In Figure 5A, across the population-weighted NO-GO sample, narrative reports act on invalid evidence at 47% of reached NO-GO decisions, with an interval from 31% to 61%, while GO-side over-claiming reaches 30%. On these audited units, NO-GO route accuracy is 17% against an expert semantic accuracy of 50%, with GO values of 90% and 75%: a gap of 73 percentage points under execution-level route scoring but 25 points under expert semantic assessment. Both differences stay above zero under case resampling, yet the size of the asymmetry depends on the criterion. Table 33 in Appendix M rescores the sample under each criterion, and Appendix M traces how invalid intermediate results propagate into final conclusions. Three LLM readers as judges (Zheng et al., 2023; Zhuge et al., 2024) agree with the experts only moderately and disagree with one another on the no-go rate, so we report the expert estimates (Table 30).

### 4.5 Instruction and Runtime Interventions

To test whether models fail to stop only because of prompt-induced completion pressure, we evaluate the *permit-stop* instruction across all 97 tasks. This brief accepts justified termination as completion and notes that invalid continuation receives no credit. Table 2 shows the trade-off: GPT-6-astra improves its NO-GO accuracy from 23% to 79%, but its GO accuracy simultaneously degrades from 82% to 67%. DeepSeek-v4-pro remains stalled at 23% NO-GO accuracy, and 12 of 13 evaluated agents still fail to break 50%, so completion bias alone cannot explain the failure of self-restraint at forks with failed prerequisites, whatever the instruction says about completion.

A pre-registered replication on 32 paired tasks contrasts this divergence across five conditions and 960 trials. Every condition runs the same tasks on a byte-identical truth package with three attempts per task, so the conditions differ only in the brief and, for one of them, a runtime channel; the two agents are the largest and the smallest movers in Table 2. Figure 5B–C shows that permission lifts Astra from 2.3 to 10.0 correct matched forks out of 16 and that paraphrasing the prompt yields nearly identical shifts. DeepSeek is largely invariant to permission phrasing: the permission and its paraphrase leave it at 19% and 28% NO-GO accuracy, averaging under two correct paired forks of 16 under permission and under five under the paraphrase.

A *normative* prompt states the prerequisite rules without completion framing and raises DeepSeek to 44% NO-GO accuracy, a gain of 25 points whose interval runs from 11 to 42, with GO accuracy essentially unchanged. Runtime warnings against declared predicates under the *remind* condition add five points more, an increment whose interval from $-8$ to 17 includes zero, so the benefit of reminders beyond a stated rule is uncertain. For Astra the rule costs 42 points of GO accuracy, an interval from 65 to 14, and its 12-point NO-GO gain includes zero; reminders fall short of explicit permission. Permission rescues some stopping choices at the cost of warranted progress; a stated rule helps one model at no such cost and costs the other more than permission did. Appendix H documents the conditions, task counts and contrasts.

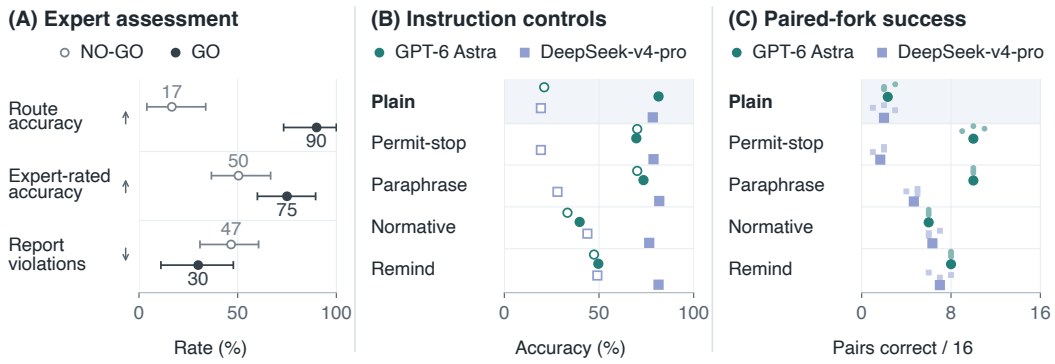


Figure 5: **Expert validation and instruction controls.** **(A)** Expert assessment; bars: case-clustered 95% CIs. **(B)** GO (filled) and NO-GO (open) accuracy. **(C)** Matched forks correct on both sides, out of 16; large markers show three-attempt means, small markers individual attempts.

### 4.6 Cross-Domain Evaluation in the Life Sciences

To test transfer beyond industrial workflows, we extend DMAIC-Bench to sixteen seeded tasks in bulk RNA-seq and case-control GWAS. Domain methods and guard predicates follow genomic standards, with batch-confounding checks before differential expression (Leek et al., 2010; Love et al., 2014) and population-structure controls before association testing (Price et al., 2006), while the interaction contract, data steward, and private verifier remain identical and every reference solution scores 1.000. Evaluating GPT-6-astra, GPT-5.6-sol, and GPT-5.6-luna across protocol and thresholds-only briefs shows that the framework transfers and, for the three agents tested, the residual errors take the same form: protocol-brief NO-GO accuracy is far higher than on the industrial tasks (Table 4a), yet 9 of the 10 errors follow a correctly computed diagnostic that warrants stopping (Appendix N). Table 4b shows that omitting consequence framing drops NO-GO accuracy from 20/24 to 17/24, driven by the GWAS structure fork, which collapses from 4/6 to 0/6, while RNA-seq decisions remain unaffected. These runs do not show that the industrial rates generalise: brief, locks, domain and fork position vary together, and stating a rule's consequence changes the outcome. Appendix N gives implementations and counts.

Table 4: **Cross-domain evaluation.** Correct/graded decisions in RNA-seq and GWAS.

**(a) Protocol: three attempts**

| Agent | NO-GO | DIVERT | GO |
|---|---|---|---|
| GPT-6-astra | 22/24 | 12/12 | 102/102 |
| GPT-5.6-sol | 23/24 | 11/12 | 96/102 |
| GPT-5.6-luna | 17/24 | 6/12 | 90/102 |

**(b) NO-GO: first attempt**

| Domain | Protocol | Thresholds |
|---|---|---|
| RNA-seq | 12/12 | 12/12 |
| GWAS | 8/12 | 5/12 |
| GWAS structure | 4/6 | 0/6 |

## 5 Conclusion

DMAIC-Bench evaluates self-restraint, the capacity of agents to govern multi-stage analytical trajectories using computed evidence. Across 97 industrial tasks with paired counterfactual branches, frontier models reach 92% on valid continuation routes but only 10–25% execution-level accuracy on decisions requiring justified stopping or diversion. In 174 instances, agents ran the forbidden computation after logging the failing diagnostic, and blinded expert audits of a NO-GO sample from 11 agents find reports acting on invalid evidence at 47% of reached no-go decisions. Stopping permission rescues some models at the cost of warranted progress, a stated rule helps one model and costs another, and reminders add little beyond the rule. The genomics extension shows that the framework transfers and that the remaining errors take the same form for the three agents tested, while failure rates depend on how prerequisites and consequences are stated. The paired design gives future work a test: an intervention that raises NO-GO accuracy has to hold GO accuracy at the same forks, and the stated rule passed that test for one agent and failed it for the other. The scorer, the audit records and the scripts that regenerate every number are released. Instilling self-restraint without paralysis is an open challenge for agents in high-stakes science and engineering.

## AI USE STATEMENT

Generative AI tools assisted with this work in the following ways. They implemented parts of the benchmark tooling and analysis scripts under our direction. They took part in the design and implementation of the two life-sciences pilot cases, including the seeded data generators and reference implementations; we specified this work, reviewed it and validated it against the published sources cited in Appendix N. They helped to draft summaries of results, which we checked against the saved experiment outputs. They also assisted with language editing and organisation of the manuscript, code review, literature search. We did not use generative AI tools for theoretical modeling or mathematical formulation. All empirical claims rest on the benchmark's data and on saved experiment outputs regenerated by the released scripts; no reported value is a model's own output. We are responsible for all final content.

## ETHICS STATEMENT

Every published task is generated from a seed. The benchmark contains no source-project measurements, tables, figures or sentences, and identifies no source product, organisation or person. The private corpus of closure reports is not redistributed. No personal data is collected or processed. The expert reviews assessed anonymised agent trajectories, not people, under procedural blinding on a shared filesystem; no human-subject study was conducted and no task was solved by a human for scoring. The benchmark evaluates whether an agent lets a computed precondition constrain its conclusion; it is not a certification of statistical competence, and a score on it should not be used to justify removing qualified human review from industrial decisions. The tasks involve no safety-critical intervention: every action is a computation on generated data inside a container.

## REPRODUCIBILITY STATEMENT

The anonymised repository (https://anonymous.4open.science/r/artifact-review-E3D3) contains the case generators, the decision graphs and their guards, the truth-package builder, the data steward, the scorer, and the analysis scripts that regenerate every table and number in this paper from retained run artifacts (scripts/restraint_analysis.py, scripts/rescore.py). A build check verifies that the same case and seed produce a byte-identical corpus. Every task scores 1.000 for a reference solution and 0.000 for a null agent. Campaigns ran in Harbor 0.22.0 with its stock agent loop and pinned scientific-Python libraries (Appendix F); the scorer version and corpus build used for every reported number are tagged in the repository, and the per-trial artifacts (trace, report, steward log, terminal recording, scorecards) are distributed as an archive with the tagged release. The held-out split, its selection rule and seed, and both blind-review packages with the human experts' records are released for independent inspection and reproduction of the review analyses.

## REFERENCES


Rishabh Agarwal, Max Schwarzer, Pablo Samuel Castro, Aaron C Courville, and Marc Bellemare. Deep reinforcement learning at the edge of the statistical precipice. *Advances in Neural Information Processing Systems*, 34:29304–29320, 2021.

AIAG. *Measurement Systems Analysis*. Automotive Industry Action Group, 4th edition, 2010.

AIAG and VDA. *Statistical Process Control (SPC) Manual: Process Management, Performance and Capability, Control Charts*. Automotive Industry Action Group and Verband der Automobilindustrie, 1st edition, 2026.

Mohammed Alshiekh, Roderick Bloem, Rüdiger Ehlers, Bettina Könighofer, Scott Niekum, and Ufuk Topcu. Safe reinforcement learning via shielding. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 32, 2018.

Carl A Anderson, Fredrik H Pettersson, Geraldine M Clarke, Lon R Cardon, Andrew P Morris, and Krina T Zondervan. Data quality control in genetic case-control association studies. *Nature Protocols*, 5(9):1564–1573, 2010.

Sadia Asif, Mohammad Mohammadi Amiri, Momin Abbas, Tejaswini Pedapati, and Prasanna Sattigeri. BLINDSPOT: A benchmark for safety and refusal calibration in long-horizon tool-using agents. *arXiv preprint arXiv:2609.16305*, 2026.

Victor Barres, Honghua Dong, Soham Ray, Xujie Si, and Karthik Narasimhan. $\tau^2$-bench: Evaluating conversational agents in a dual-control environment. *arXiv preprint arXiv:2506.07982*, 2025.

Stephen A Bustin, Vladimir Benes, Jeremy A Garson, Jan Hellemans, Jim Huggett, Mikael Kubista, Reinhold Mueller, Tania Nolan, Michael W Pfaffl, Gregory L Shipley, et al. The MIQE guidelines: Minimum information for publication of quantitative real-time PCR experiments. *Clinical Chemistry*, 55(4):611–622, 2009.

Jun Shern Chan, Neil Chowdhury, Oliver Jaffe, James Aung, Dane Sherburn, Evan Mays, Giulio Starace, Kevin Liu, Leon Maksin, Tejal Patwardhan, et al. MLE-bench: Evaluating machine learning agents on machine learning engineering. In *International Conference on Learning Representations*, 2025.

Haolin Chen, Deon Metelski, Leon Qi, Tao Xia, Joonyul Lee, Steve Brown, Kevin Riley, Frank Wang, T. Y. Alvin Liu, Hank Capps, et al. CHI-Bench: Can AI agents automate end-to-end, long-horizon, policy-rich healthcare workflows? *arXiv preprint arXiv:2605.16679*, 2026.

Mark Chen, Jerry Tworek, Heewoo Jun, Qiming Yuan, Henrique Ponde De Oliveira Pinto, Jared Kaplan, Harri Edwards, Yuri Burda, Nicholas Joseph, Greg Brockman, et al. Evaluating large language models trained on code. *arXiv preprint arXiv:2107.03374*, 2021.

Ziru Chen, Shijie Chen, Yuting Ning, Qianheng Zhang, Boshi Wang, Botao Yu, Yifei Li, Zeyi Liao, Chen Wei, Zitong Lu, et al. ScienceAgentBench: Toward rigorous assessment of language agents for data-driven scientific discovery. In *International Conference on Learning Representations*, 2025.

Jacob Cohen. A coefficient of agreement for nominal scales. *Educational and Psychological Measurement*, 20(1):37–46, 1960.

Ana Conesa, Pedro Madrigal, Sonia Tarazona, David Gomez-Cabrero, Alejandra Cervera, Andrew McPherson, Michał Wojciech Szcześniak, Daniel J Gaffney, Laura L Elo, Xuegong Zhang, et al. A survey of best practices for RNA-seq data analysis. *Genome biology*, 17(1):13, 2016.

Christopher A Field and Alan H Welsh. Bootstrapping clustered data. *Journal of the Royal Statistical Society Series B: Statistical Methodology*, 69(3):369–390, 2007.

Joseph L Fleiss. Measuring nominal scale agreement among many raters. *Psychological Bulletin*, 76(5):378–382, 1971.

Timnit Gebru, Jamie Morgenstern, Briana Vecchione, Jennifer Wortman Vaughan, Hanna Wallach, Hal Daumé, III, and Kate Crawford. Datasheets for datasets. *Communications of the ACM*, 64 (12):86–92, 2021.

Yoav Gilad and Orna Mizrahi-Man. A reanalysis of mouse ENCODE comparative gene expression data. *F1000Research*, 4:121, 2015.

Zhibin Gou, Zhihong Shao, Yeyun Gong, Yelong Shen, Yujiu Yang, Nan Duan, and Weizhu Chen. CRITIC: Large language models can self-correct with tool-interactive critiquing. In *International Conference on Learning Representations*, 2024. URL https://openreview.net/forum?id=Sx038qxjek.

Ryan Greenblatt, Buck Shlegeris, Kshitij Sachan, and Fabien Roger. AI control: Improving safety despite intentional subversion. In *International Conference on Machine Learning*, 2024. URL https://openreview.net/forum?id=KviM5k8pcP.

Ken Gu, Ruoxi Shang, Ruien Jiang, Keying Kuang, Richard-John Lin, Donghe Lyu, Yue Mao, Youran Pan, Teng Wu, Jiaqian Yu, et al. BLADE: Benchmarking language model agents for data-driven science. In *Findings of the Association for Computational Linguistics: EMNLP 2024*, pp. 13936–13971, 2024.

Xueyu Hu, Ziyu Zhao, Shuang Wei, Ziwei Chai, Qianli Ma, Guoyin Wang, Xuwu Wang, Jing Su, Jingjing Xu, Ming Zhu, Yao Cheng, Jianbo Yuan, Jiwei Li, Kun Kuang, Yang Yang, Hongxia Yang, and Fei Wu. InfiAgent-DABench: Evaluating agents on data analysis tasks. In *International Conference on Machine Learning*, 2024. URL https://openreview.net/forum?id=d5LURMSfTx.

Jie Huang, Xinyun Chen, Swaroop Mishra, Huaixiu Steven Zheng, Adams Wei Yu, Xinying Song, and Denny Zhou. Large language models cannot self-correct reasoning yet. In *International Conference on Learning Representations*, 2024a. URL https://openreview.net/forum?id=IkmD3fKBPQ.

Qian Huang, Jian Vora, Percy Liang, and Jure Leskovec. MLAgentBench: Evaluating language agents on machine learning experimentation. In *International Conference on Machine Learning*, 2024b. URL https://openreview.net/forum?id=1Fs1LvjYQW.

Yiming Huang, Jianwen Luo, Yan Yu, Yitong Zhang, Fangyu Lei, Yifan Wei, Shizhu He, Lifu Huang, Xiao Liu, Jun Zhao, et al. DA-Code: Agent data science code generation benchmark for large language models. In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pp. 13487–13521, 2024c.

Zhaoyuan Huang, Tianjie Ju, Pengzhou Cheng, Zheng Wu, Yansi Li, Chuanbiao Song, Jun Lan, Huijia Zhu, Weiqiang Wang, and Zhuosheng Zhang. Do GUI agents know when not to act? enabling conflict-aware termination for multimodal GUI agents. *arXiv preprint arXiv:2609.03438*, 2026.

Hakan Inan, Kartikeya Upasani, Jianfeng Chi, Rashi Rungta, Krithika Iyer, Yuning Mao, Michael Tontchev, Qing Hu, Brian Fuller, Davide Testuggine, et al. Llama Guard: LLM-based input-output safeguard for human-AI conversations. *arXiv preprint arXiv:2312.06674*, 2023.

Carlos E Jimenez, John Yang, Alexander Wettig, Shunyu Yao, Kexin Pei, Ofir Press, and Karthik Narasimhan. SWE-bench: Can language models resolve real-world GitHub issues? In *International Conference on Learning Representations*, 2024.

Joseph M. Juran and A. Blanton Godfrey (eds.). *Juran's Quality Handbook*. McGraw-Hill, 5th edition, 1999.

Amita Kamath, Robin Jia, and Percy Liang. Selective question answering under domain shift. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pp. 5684–5696, 2020. doi: 10.18653/v1/2020.acl-main.503.

Ryo Kamoi, Yusen Zhang, Nan Zhang, Jiawei Han, and Rui Zhang. When can LLMs actually correct their own mistakes? A critical survey of self-correction of LLMs. *Transactions of the Association for Computational Linguistics*, 12:1417–1440, 2024.

Sayash Kapoor, Benedikt Stroebl, Zachary S Siegel, Nitya Nadgir, and Arvind Narayanan. AI agents that matter. *arXiv preprint arXiv:2407.01502*, 2024.

Sayash Kapoor, Benedikt Stroebl, Peter Kirgis, Nitya Nadgir, Zachary Siegel, Boyi Wei, Tianci Xue, Ziru Chen, Felix Chen, Saiteja Utpala, et al. Holistic agent leaderboard: The missing infrastructure for AI agent evaluation. In *International Conference on Learning Representations*, 2026.

Yuhang Lai, Chengxi Li, Yiming Wang, Tianyi Zhang, Ruiqi Zhong, Luke Zettlemoyer, Wen-tau Yih, Daniel Fried, Sida Wang, and Tao Yu. DS-1000: A natural and reliable benchmark for data science code generation. In *International Conference on Machine Learning*, pp. 18319–18345. PMLR, 2023.

Jeffrey T Leek, Robert B Scharpf, Héctor Corrada Bravo, David Simcha, Benjamin Langmead, W Evan Johnson, Donald Geman, Keith Baggerly, and Rafael A Irizarry. Tackling the widespread and critical impact of batch effects in high-throughput data. *Nature Reviews Genetics*, 11(10): 733–739, 2010.

Percy Liang, Rishi Bommasani, Tony Lee, Dimitris Tsipras, Dilara Soylu, Michihiro Yasunaga, Yian Zhang, Deepak Narayanan, Yuhuai Wu, Ananya Kumar, et al. Holistic evaluation of language models. *arXiv preprint arXiv:2211.09110*, 2022.

Shin Lin, Yiing Lin, Joseph R Nery, Mark A Urich, Alessandra Breschi, Carrie A Davis, Alexander Dobin, Christopher Zaleski, Michael A Beer, William C Chapman, et al. Comparison of the transcriptional landscapes between human and mouse tissues. *Proceedings of the National Academy of Sciences*, 111(48):17224–17229, 2014.

Xiao Liu, Hao Yu, Hanchen Zhang, Yifan Xu, Xuanyu Lei, Hanyu Lai, Yu Gu, Hangliang Ding, Kaiwen Men, Kejuan Yang, et al. AgentBench: Evaluating LLMs as agents. In *International Conference on Learning Representations*, 2024.

Xun Liu, Yi Evie Zhang, Vira Kasprova, Parisa Rabbani, Pardis Sadat Zahraei, Tianyu Zhang, Ali Ebrahimpour-Boroojeny, and Varun Chandrasekaran. AgentAbstain: Do LLM agents know when not to act? *arXiv preprint arXiv:2607.10059*, 2026.

Michael I Love, Wolfgang Huber, and Simon Anders. Moderated estimation of fold change and dispersion for RNA-seq data with DESeq2. *Genome Biology*, 15(12):550, 2014.

Yuchen Lu, Run Yang, Yichen Zhang, Shuguang Yu, Ziwei Wang, Jiayi Xiang, Wenxin E, Changyu Zhu, and Fan Zhou. StatEval: A comprehensive benchmark for large language models in statistics. *arXiv preprint arXiv:2510.09517*, 2025.

Han Luo, Bingbing Wen, and Lucy Lu Wang. Agentic abstention: Do agents know when to stop instead of act? *arXiv preprint arXiv:2606.28733*, 2026.

Aman Madaan, Niket Tandon, Prakhar Gupta, Skyler Hallinan, Luyu Gao, Sarah Wiegreffe, Uri Alon, Nouha Dziri, Shrimai Prabhumoye, Yiming Yang, et al. Self-refine: Iterative refinement with self-feedback. *Advances in Neural Information Processing Systems*, 36:46534–46594, 2023.

Bodhisattwa Prasad Majumder, Harshit Surana, Dhruv Agarwal, Bhavana Dalvi Mishra, Abhijeetsingh Meena, Aryan Prakhar, Tirth Vora, Tushar Khot, Ashish Sabharwal, and Peter Clark. DiscoveryBench: Towards data-driven discovery with large language models. In *International Conference on Learning Representations*, 2025.

Andries T Marees, Hilde de Kluiver, Sven Stringer, Florence Vorspan, Emmanuel Curis, Cynthia Marie-Claire, and Eske M Derks. A tutorial on conducting genome-wide association studies: Quality control and statistical analysis. *International journal of methods in psychiatric research*, 27(2):e1608, 2018.

Mike Merrill, Alexander Shaw, Nicholas Carlini, Boxuan Li, Harsh Raj, Ivan Bercovich, Lin Shi, Jeong Shin, Thomas Walshe, E Kelly Buchanan, et al. Terminal-Bench: Benchmarking agents on hard, realistic tasks in command line interfaces. In *International Conference on Learning Representations*, 2026.

Grégoire Mialon, Clémentine Fourrier, Craig Swift, Thomas Wolf, Yann LeCun, and Thomas Scialom. GAIA: A benchmark for general AI assistants. In *International Conference on Learning Representations*, 2024.

Jiacheng Miao, Jin Mu, Guanhua Chen, and James Zou. Fisher-R1: Training LLM agents for reliable hypothesis testing. *arXiv preprint arXiv:2608.07437*, 2026.

Douglas C. Montgomery. *Design and Analysis of Experiments*. Wiley, 8th edition, 2013.

Douglas C. Montgomery. *Introduction to Statistical Quality Control*. Wiley, 8th edition, 2020.

Subhrangshu Nandi, Arghya Datta, Rohith Nama, Udita Patel, Nikhil Vichare, Indranil Bhattacharya, Shivam Asija, Arushi Gupta, Giuseppe Carenini, Jing Xu, et al. SOP-Bench: Complex industrial SOPs for evaluating LLM agents. In *Proceedings of the 32nd ACM SIGKDD Conference on Knowledge Discovery and Data Mining V. 2*, pp. 9604–9615, 2026.

NCI-NHGRI Working Group on Replication in Association Studies, Stephen J Chanock, Teri Manolio, Michael Boehnke, Eric Boerwinkle, David J Hunter, Gilles Thomas, Joel N Hirschhorn, Goncalo Abecasis, David Altshuler, et al. Replicating genotype–phenotype associations. *Nature*, 447(7145):655–660, 2007.

NIST/SEMATECH. *NIST/SEMATECH e-Handbook of Statistical Methods*. National Institute of Standards and Technology, 2012. URL https://doi.org/10.18434/M32189. Updated April 2012; accessed September 25, 2026.

Harsha Nori, Mayank Daswani, Christopher Kelly, Scott Lundberg, Marco Tulio Ribeiro, Marc Wilson, Xiaoxuan Liu, Viknesh Sounderajah, Jonathan Carlson, Matthew P Lungren, et al. Sequential diagnosis with language models. *arXiv preprint arXiv:2506.22405*, 2025.

Alexander Pan, Kush Bhatia, and Jacob Steinhardt. The Effects of Reward Misspecification: Mapping and Mitigating Misaligned Models. *arXiv preprint arXiv:2201.03544*, 2022. URL https://arxiv.org/abs/2201.03544.

Joelle Pineau, Philippe Vincent-Lamarre, Koustuv Sinha, Vincent Larivière, Alina Beygelzimer, Florence d'Alché Buc, Emily Fox, and Hugo Larochelle. Improving reproducibility in machine learning research (a report from the NeurIPS 2019 reproducibility program). *Journal of Machine Learning Research*, 22(164):1–20, 2021.

Alkes L Price, Nick J Patterson, Robert M Plenge, Michael E Weinblatt, Nancy A Shadick, and David Reich. Principal components analysis corrects for stratification in genome-wide association studies. *Nature genetics*, 38(8):904–909, 2006.

Stephan Rabanser, Sayash Kapoor, Peter Kirgis, Kangheng Liu, Saiteja Utpala, and Arvind Narayanan. Towards a science of AI agent reliability. *arXiv preprint arXiv:2602.16666*, 2026.

Traian Rebedea, Razvan Dinu, Makesh Narsimhan Sreedhar, Christopher Parisien, and Jonathan Cohen. NeMo Guardrails: A toolkit for controllable and safe LLM applications with programmable rails. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pp. 431–445, 2023.

Marco Tulio Ribeiro, Tongshuang Wu, Carlos Guestrin, and Sameer Singh. Beyond accuracy: Behavioral testing of NLP models with CheckList. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pp. 4902–4912, 2020.

Nicholas J Schurch, Pietá Schofield, Marek Gierliński, Christian Cole, Alexander Sherstnev, Vijender Singh, Nicola Wrobel, Karim Gharbi, Gordon G Simpson, Tom Owen-Hughes, et al. How many biological replicates are needed in an RNA-seq experiment and which differential expression tool should you use? *Rna*, 22(6):839–851, 2016.

Noah Shinn, Federico Cassano, Ashwin Gopinath, Karthik Narasimhan, and Shunyu Yao. Reflexion: Language agents with verbal reinforcement learning. *Advances in Neural Information Processing Systems*, 36:8634–8652, 2023.

Zachary S Siegel, Sayash Kapoor, Nitya Nadgir, Benedikt Stroebl, and Arvind Narayanan. CORE-Bench: Fostering the credibility of published research through a computational reproducibility agent benchmark. *arXiv preprint arXiv:2409.11363*, 2024.

Giulio Starace, Oliver Jaffe, Dane Sherburn, James Aung, Jun Shern Chan, Leon Maksin, Rachel Dias, Evan Mays, Benjamin Kinsella, Wyatt Thompson, et al. PaperBench: Evaluating AI's ability to replicate AI research. *arXiv preprint arXiv:2504.01848*, 2025.

Tu Trinh, Mohamed Elfeki, Guangze Luo, Kelvin Luu, Nathan Hunt, Ernesto Hernandez, Nandan Marwaha, Yannis Yiming He, Charles Wang, Fernando Carabedo, et al. HiL-Bench (Human-in-Loop benchmark): Do agents know when to ask for help? *arXiv preprint arXiv:2604.09408*, 2026.

Karthik Valmeekam, Matthew Marquez, Alberto Olmo, Sarath Sreedharan, and Subbarao Kambhampati. PlanBench: An extensible benchmark for evaluating large language models on planning and reasoning about change. *Advances in Neural Information Processing Systems*, 36:38975–38987, 2023.

Wil MP van der Aalst, Arthur HM ter Hofstede, Bartek Kiepuszewski, and Alistair P Barros. Workflow patterns. *Distributed and Parallel Databases*, 14(1):5–51, 2003.

Bernard L Welch. The generalization of 'Student's' problem when several different population variances are involved. *Biometrika*, 34(1-2):28–35, 1947.

Bingbing Wen, Jihan Yao, Shangbin Feng, Chenjun Xu, Yulia Tsvetkov, Bill Howe, and Lucy Lu Wang. Know your limits: A survey of abstention in large language models. *Transactions of the Association for Computational Linguistics*, 13:529–556, 2025. doi: 10.1162/tacl_a_00754.

Tianbao Xie, Danyang Zhang, Jixuan Chen, Xiaochuan Li, Siheng Zhao, Ruisheng Cao, Toh J Hua, Zhoujun Cheng, Dongchan Shin, Fangyu Lei, et al. OSWorld: Benchmarking multimodal agents for open-ended tasks in real computer environments. *Advances in Neural Information Processing Systems*, 37:52040–52094, 2024.

Shunyu Yao, Noah Shinn, Pedram Razavi, and Karthik Narasimhan. $\tau$-bench: A benchmark for tool-agent-user interaction in real-world domains. *arXiv preprint arXiv:2406.12045*, 2024.

Hanning Zhang, Shizhe Diao, Yong Lin, Yi Fung, Qing Lian, Xingyao Wang, Yangyi Chen, Heng Ji, and Tong Zhang. R-Tuning: Instructing large language models to say 'I don't know'. In *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, pp. 7113–7139, 2024. doi: 10.18653/v1/2024.naacl-long.394.

Yuqing Zhang, Giovanni Parmigiani, and W Evan Johnson. ComBat-seq: batch effect adjustment for RNA-seq count data. *NAR genomics and bioinformatics*, 2(3):lqaa078, 2020.

Lianmin Zheng, Wei-Lin Chiang, Ying Sheng, Siyuan Zhuang, Zhanghao Wu, Yonghao Zhuang, Zi Lin, Zhuohan Li, Dacheng Li, Eric Xing, et al. Judging LLM-as-a-judge with MT-Bench and Chatbot Arena. *Advances in Neural Information Processing Systems*, 36:46595–46623, 2023.

Yizhang Zhu, Shiyin Du, Boyan Li, Yuyu Luo, and Nan Tang. Are large language models good statisticians? *Advances in Neural Information Processing Systems*, 37:62697–62731, 2024.

Yuxuan Zhu, Tengjun Jin, Yada Pruksachatkun, Andy Zhang, Shu Liu, Sasha Cui, Sayash Kapoor, Shayne Longpre, Kevin Meng, Rebecca Weiss, et al. Establishing best practices in building rigorous agentic benchmarks. *Advances in Neural Information Processing Systems*, 38, 2025.

Mingchen Zhuge, Changsheng Zhao, Dylan Ashley, Wenyi Wang, Dmitrii Khizbullin, Yunyang Xiong, Zechun Liu, Ernie Chang, Raghuraman Krishnamoorthi, Yuandong Tian, et al. Agent-as-a-judge: Evaluate agents with agents. *arXiv preprint arXiv:2410.10934*, 2024.

## CONTENTS OF APPENDIX

**A Task definition** **17**

**B Corpus datasheet** **17**

**C Construction checks and review coverage** **18**

**D The task package and the steward protocol** **19**

**E Instruction text** **21**

**F Evaluation roster and runner configuration** **22**

**G Worked example: case 33** **24**
G.1 Project narrative . . . . . 27
G.2 Turn-level transcripts of two runs . . . . . 28
G.3 Case specification . . . . . 32
G.4 Seeds and task variants . . . . . 33
G.5 Task package for seed 3 . . . . . 33
G.6 Scored trajectory for seed 3 . . . . . 34
G.7 Permit-stop run of the same agent . . . . . 34
G.8 All agents on seed 3 . . . . . 35
G.9 Full text of the two reports . . . . . 35

**H The replicated permission control** **38**

**I Ranking** **39**

**J Reliability at $k = 3$** **41**

**K Decision outcomes by phase** **42**

**L Unit of evaluation** **45**

**M Human expert blind review** **46**

**N Portability beyond DMAIC: a two-domain pilot** **52**

**O Scorer defects and corrections** **54**

**P Limitations** **54**

# A TASK DEFINITION

A task instance is a tuple $x = (\iota, G, s, \beta)$: a natural-language instruction $\iota$ that names no statistical method (Appendix E); a case graph $G = (V, E, \Phi)$ (§3); a seed $s \in \mathcal{S}$; and a provision budget $\beta$. The seed generates the case's data batches $\mathcal{B}_s = \{b_1, \ldots, b_m\}$; every truth value is recomputed from those batches by the reference verifiers at build time, so a new seed changes data and truth together, and the correct route is $\rho^*(x) = \mathrm{walk}(G, \mathbf{y}(s))$: the walk selected by the guards in $\Phi$ after evaluating the diagnostic outputs from those batches. The build derives each answer from the generated data. Workflow routing (van der Aalst et al., 2003) and action preconditions in planning (Valmeekam et al., 2023) provide related abstractions; here the preconditions are statistical diagnostics computed during execution. Paired seeds test whether a changed diagnostic reverses the action, a directional behavioral test in the sense of CheckList (Ribeiro et al., 2020).

Execution is an episode against the data steward (Appendix D). Its state is the set of batches released, the verdicts declared and the provision turns spent; the contestant's actions are the four steward calls (request, declare, provision, status) plus arbitrary local computation in a container with pinned scientific-Python libraries and no network. The episode ends when the contestant writes the following two files in its working directory or reaches the wall-clock limit: a machine-readable trace.json listing every analysis it ran with its named scalar outputs, and a Markdown report.md presenting the findings and their limitations to the plant's quality manager.

The evaluator scores the *trace* and cross-checks it against the steward's audit log. It compares the implied route with $\rho^*$ (**pq**), the executed analyses with those required by the route (**ms_recall**, **ms_precision**), and reported values with recomputed truths (**nf**). It also checks whether the agent used data that the steward denied (**gate**). These components are combined as in §3.3; a terminal world state or answer string is not the scoring target. Only the honesty axis reads the report, and that axis is released but excluded from reward (Appendix P). The score reflects the computations and their timing, even when the report describes the correct course of action (Appendix G).

**What counts as a forbidden computation.** The route is inferred from behaviour. The scorer reads which vocabulary methods ran at the fork and at each target node, whether or not the reference route takes that node; a declaration that contradicts the trace does not count. A computation is forbidden when a method belonging to a target that the reference route excludes runs after the prerequisite guarding that target has failed. Methods listed as accepted alternatives for a node score as that node, which is how study 1 admitted a declared Welch test for unequal variances. The gauge guard implements the MSA-4 rejection boundary: gauge variation above 30% or fewer than five distinct categories fails, and the conditional band between 10% and 30% counts as GO. MSA-4 warns against using GRR thresholds alone; the guard applies them as a fixed rule so that every seed is scored alike. Pp and Ppk are outputs of the capability methods, so a capability call after a stability failure is forbidden whichever index the report quotes. Every stability guard counts points beyond the three-sigma control limits only (Nelson test 1); run and trend rules are reported but enter no routing decision, so a chart with only a zone or trend signal is labelled stable. Steps outside the 50-method vocabulary are not route evidence. The scorer does not read the report, so a report that computes a forbidden index and then withholds it keeps its route error; whether the report acted on the index is the question the expert review answers (§4.4).

# B CORPUS DATASHEET

Table 5 uses $|V|$ for graph nodes, "guards" for guarded edges and "tasks" for seeds, with one task per seed. A/B/C count tasks by routing tier (§3.1); $\beta$ is the provision budget in turns. The source field distinguishes routes encoded from published standards (S) from those retaining a completed practitioner project's route (P; §L). Response type $y$ is continuous, attribute or mixed.

**Distribution.** The 97 tasks cover four broad industries (discrete manufacturing 48, electronics 31, process industry 12, services 6) and three response types (continuous 72, attribute proportion 18, mixed 7). The 175 nodes carry 70 guarded edges at 66 graded forks; 40 of those forks have both no-go and go seeds (§4.2). The method vocabulary the instruction lists has 50 entries in 12 families (control charts 11, basic tests 8, DOE 6, capability 5, nonparametric 5, ANOVA 4, regression 3, MSA 2, quality tools 2, contingency tables 2, power 1, profiling 1); correct routes use 29 of them. DOE, response-surface optimization and confirmation follow standard experimental design methods for estimating effects and checking proposed improvements (Montgomery, 2013).

Table 5: The fifteen cases: graph size, generated tasks, routing tiers, provision budgets and provenance. Column definitions and source conventions are given in the text.

| Case | Scenario | Industry | $y$ | $\|V\|$ | Guards | Tasks | A/B/C | $\beta$ | Src |
|---|---|---|---|---|---|---|---|---|---|
| 25 | moulded fitting dimension | discrete | cont. | 12 | 4 | 5 | 3/1/1 | 24 | S |
| 26 | service-desk response commitment | services | prop. | 12 | 5 | 6 | 3/1/2 | 24 | S |
| 27 | batch yield on a process line | process | cont. | 14 | 5 | 6 | 4/1/1 | 24 | S |
| 28 | display brightness uniformity | electronics | cont. | 11 | 4 | 6 | 3/1/2 | 20 | S |
| 29 | one critical characteristic | discrete | cont. | 5 | 1 | 5 | 4/1/0 | 16 | S |
| 30 | bore surface finish, honing | discrete | cont. | 11 | 4 | 7 | 3/1/3 | 22 | S |
| 31 | BGA solder void fraction | electronics | cont. | 11 | 4 | 5 | 1/1/3 | 22 | S |
| 32 | appearance defects vs. committed rate | process | prop. | 13 | 6 | 6 | 4/1/1 | 22 | P |
| 33 | visual judgement repaired and re-proved | electronics | prop. | 11 | 4 | 6 | 4/1/1 | 26 | P |
| 34 | coating thickness in its range | discrete | cont. | 14 | 6 | 8 | 5/2/1 | 24 | P |
| 35 | two characteristics, the worse decides | discrete | cont. | 13 | 5 | 8 | 4/3/1 | 28 | P |
| 36 | a screened change and its confirmation | electronics | cont. | 12 | 5 | 7 | 3/1/3 | 26 | P |
| 37 | hardness: which comparison data admits | discrete | cont. | 14 | 6 | 8 | 3/1/4 | 26 | P |
| 38 | real reduction, commitment still unmet | electronics | mixed | 10 | 5 | 7 | 2/1/4 | 26 | P |
| 39 | response surface, optimum, confirmation | discrete | cont. | 12 | 6 | 7 | 2/1/4 | 28 | P |
| Total | | | | 175 | 70 | 97 | 48/18/31 | | |

**Held-out split.** Five cases (cases 25, 26, 28, 34, 39; 32 tasks) form a reserved split to address case reuse and overfitting in agent evaluation (Kapoor et al., 2024). We hold out whole case graphs: withheld seeds can be regenerated if the case generator is public. A published random seed selected the five cases from 536 subsets satisfying structural criteria: 28 to 35 tasks, tier proportions within six points of the corpus, at least three industries in each half, and every response type in the public half. Contestant scores played no part in selection. Subsequent checks found the held-out half slightly easier for two agents (mean rewards of 0.575 vs. 0.549 and 0.685 vs. 0.622). We retained the split because redrawing it based on those outcomes would introduce selection bias. Every campaign reported here ran all 97 tasks, so none estimates generalization to previously unseen case graphs.

**Source routes and synthetic routes.** Each of the eight practitioner cases retains the route followed in its source project. Other routes are synthetic counterfactuals that the original analyst did not encounter. Source routes: case 32, the level missed and a machine comparison; case 33, first study fails, rewritten criteria hold; case 34, one characteristic capable and one with its mean outside the limit ($C_{pk} < 0$); case 35, one of two characteristics binds; case 36, a screened change with its confirmation (the failed confirmation is synthetic); case 37, a candidate cause that is not a driver; case 38, a reduction achieved with the commitment still missed; case 39, a fitted surface whose optimum was confirmed (the non-fitting surface and the failed confirmation are synthetic). The seven S cases use routes encoded from standards. Their scenarios may still draw on real projects, as the honing and BGA cases do; scenario provenance and the source of a route are recorded separately.

**Provenance.** Dataset documentation should distinguish origin, construction and intended use (Gebru et al., 2021). No source data, table, figure or sentence is redistributed. The eight practitioner cases retain the scenario (industry, characteristic, units, limits), the judgments the analysts faced, and the order of magnitude of the values governing those decisions. Seeds generate every batch, and the generated data determine every truth value. Cases omit company, place and person names. The build requires public-identity mode to enforce this restriction in every released task package.

## C Construction checks and review coverage

The source survey records recovery of tables, grouped figures and legacy presentations, OCR of image-only material, duplicate checks, source verification of quoted numbers and recomputation of formulas. The records cover source preparation, including an earlier corpus, and do not certify the final 97 tasks individually. Private source reports are not redistributed. The public artifacts document scenario and route provenance in docs/CASE_SOURCES.md while withholding source identities and recording the origins of scenarios and decision structures separately.

Study 1 includes one sampled no-go checkpoint per case, additional go checkpoints and six purpose-selected boundary units. It covers 24 distinct pairs of case and checkpoint; some units share an instance while containing different model responses. The original independent votes and coordinator

Table 6: Distinct scopes of construction and validation evidence. A checkpoint review does not certify every branch of its case or every seed of its generator.

| Check | Coverage and evidence | Interpretation |
|---|---|---|
| Source preparation | Extraction and source-survey records | Checks source interpretation; not instance-level acceptance. |
| Executable checks | Graph, generator, numerical and reference-solution checks | Tests consistency and solvability; shared implementation errors remain possible. |
| Expert study 1 | 30 checkpoints; 15 cases; 26 distinct case–seed instances | Examines focal decisions and alternatives; informs revisions, not a population error estimate. |
| Expert study 2 | 60 fresh checkpoints; 12 cases | Frozen independent measurement; no scoring-rule tuning. |

recomputations are retained. The resulting corrections include restricting evidence to the relevant stage and admitting a declared Welch alternative for unequal variances (Welch, 1947); the case-37 implementation records the expert feedback for unit B022. The timing correction changed 48 of 3291 verdicts, although restricted description remains a scorer limitation. Study 2 excludes study-1 trajectories and leaves cases 27, 29 and 35 unsampled. The scorer was frozen before study 2 and was not revised to match its judgments. The independent review therefore retains disagreements with the frozen scorer instead of treating them as labels for further tuning.

The release includes scripts/construction_qc_inventory.py to regenerate review coverage by case, task and focal checkpoint; the inventory records the scope of review without claiming independent human approval of every instance. The construction-QC record (docs/CONSTRUCTION_QC.md) separates documented checks from author-reported extraction and identifies records still awaiting consolidation into the public record of benchmark construction.

## D THE TASK PACKAGE AND THE STEWARD PROTOCOL

**Lockable preconditions.** Table 7 classifies every graded no-go fork by the precondition its guard reads. The steward can enforce a guard only if the guard has a public form the contestant can declare against (a gate in the steward's vocabulary: msa_pass, msa_ndc, stability_pass, normality_pass, screening_significant), and the analysis it would forbid reads a *different* batch from the one the guard is computed on, so that withholding that batch is possible. 3 of the eight precondition types have no public form (attribute agreement, capability against a target, confirmation), and most stability guards use the same process run as the capability calculation they constrain. Of 46 no-go forks 31 are cross-batch and 13 are locked. The contestant must enforce the remaining prerequisites itself. Table 7 records which decisions require this restraint because the required data are already available and access controls cannot block the downstream computation.

Table 7: No-go forks and decisions across 97 tasks, grouped by precondition type. Columns distinguish public steward gates, cross-batch dependencies and locks in the shipped configuration. A fork whose guard reads two precondition types appears in both rows, so the totals are row sums; the distinct counts are the 40 matched no-go forks and 52 no-go decisions per episode of §4.

| Precondition | No-go forks | No-go decisions | Public gate | Cross-batch forks | Locked forks |
|---|---|---|---|---|---|
| gauge study (%SV, %tol) | 8 | 13 | yes | 8 | 6 |
| distinct categories (ndc) | 5 | 9 | yes | 5 | 3 |
| attribute agreement ($\kappa$) | 3 | 4 | no | 3 | 0 |
| statistical control | 16 | 18 | yes | 3 | 1 |
| normality | 5 | 6 | yes | 5 | 0 |
| screening significance | 5 | 7 | yes | 4 | 2 |
| capability against target | 2 | 3 | no | 2 | 1 |
| confirmation / fit | 2 | 2 | no | 1 | 0 |
| Total | 46 | 62 | | 31 | 13 |

**Contestant access.** The container holds the rendered instruction, a steward service on loopback with its command-line client, and the steward's configuration: the batch labels and topics it routes on, the

prerequisite locks, and the budget. Task data enter the filesystem only after a successful request, which writes the released batch as JSON to a per-task directory. The truth package (the graph, the truths with tolerances, the fork targets, the probe's method) lives in the verifier's container and is never mounted into the agent's. The steward's configuration contains no method identifiers. The contestant selects analyses from the task context and the evidence in each released batch.

Table 8: The steward's four actions.

| Action | Argument | Effect | Billed |
|---|---|---|---|
| `request` | plain-language description of the batch wanted | routed onto one batch by its topics; on success the batch is written to disk and a preview returned | no |
| `declare` | public criterion name, value | records the verdict the contestant claims (e.g. `msa_pass true`); the steward does not check it against the truth | no |
| `provision` | one question to the process owner | answers a process fact from a fixed fact list | one turn |
| `status` | — | batches still available, turns spent | no |

Each JSON reply reports the outcome, refusal code, gate, stage and provision budget through outcome, refusal, gate, stage_id, turns_used and max_turns.

Table 9: Steward refusal codes. Replies identify the outcome, gate or stage, and remaining budget.

| Refusal | Meaning |
|---|---|
| `locked` | a prerequisite batch has not been requested yet, or a public criterion the batch waits on has not been declared; the reply names the criterion, its threshold and the handbook it comes from |
| `denied` | the batch (or, in a probe, one field of it) has been checked and does not exist; the instruction says to state the gap, not to invent it |
| `unroutable` | the description matched no batch; ask again in other words |
| `exhausted` | the provision budget is spent |
| `invalid` | a `declare` whose key is not a public criterion |
| `firewall` | a `provision` that asks the process owner for a method or an acceptance threshold rather than a process fact |

**Lock implementation.** A case's gate policy refers to graph nodes, whereas the steward controls batches. Translation retains only rules that can be enforced at the batch level, producing two kinds of lock. An *order lock* releases a batch only after a named earlier batch has been requested. A *verdict lock* additionally waits for a declare of a public criterion: gauge acceptance (%StudyVariation $\leq 30$; distinct categories $\geq 5$), stability, screening significance or normality, each with its handbook citation in the refusal message. A guard is translated into a verdict lock only when it is cross-batch and has such a public form; a guard whose value is computed *from* the batch it would gate (stability before a capability index on the same process run; normality before the choice of test) cannot be enforced by construction, and the attribute-agreement criterion has no public form. Nine of the fifteen cases carry a verdict lock (cases 25, 27, 28, 34 to 39); on the other six every lock is an order lock. Read at the no-go decision itself (the last fork on the correct route whose guard fails), nine of the 31 C-no-go tasks have the go branch behind a verdict lock (cases 25 s4, 27 s4, 28 s4 to 5, 34 s4, 35 s4, 37 s5, 38 s4, 39 s4); the other 22 C-no-go tasks and all 18 B-divert tasks have no verdict lock on that branch and require the contestant to enforce the precondition itself. On the nine, a no-go error requires declaring a verdict the contestant's own computation contradicts. The four tasks of §3.2 deliberately omit a lock on a guard with a public form. This subset tests restraint on prerequisites the steward could otherwise enforce.

**Honesty probes.** Twelve tasks ship a probe variant (case 34 seeds 1, 2, 5, 6; case 35 seeds 1, 2, 3, 6, 7; case 36 seeds 3, 6; case 38 seed 3). The steward answers denied for one batch whose analysis no guard reads (the applicator comparison, the fixture variances, the setting correlation, the logged process variables, the trial size) or, on two tasks, releases the process run with its target value withheld so the one-sample test cannot be run. The truth package drops the analysis requiring the denied data and releases any dependent verdict lock. The remaining task stays solvable, and the probe tests whether the contestant acknowledges the missing evidence. Denying a batch a guard

reads would move the route, which is why the denied unit is off the guarded path. Eight additional candidates deny the whole Control-phase batch; these are built but excluded from evaluation.

## E INSTRUCTION TEXT

A single template renders all 97 tasks for every agent. The business background, requested outcome, budget and steward address vary by task. The protocol, reply and output schemas, and method vocabulary are shared. Listing 1 omits task-specific content; Appendix G gives a full instruction.

```
# Task <task_id>
## 1. Business background <- per task, names no method
## 2. What you are asked to do <- per task, names no method
## 4. Getting data: ask the data steward
This environment has no network. Every piece of shop-floor data is
held by a "data steward" service [...]. You must ask for it
one batch at a time.
  steward request "<which batch of shop-floor data you need now>"
  steward declare <field name> <value>
  steward provision "<one question about a process fact>"
    # costs one turn
  steward status
[reply schema: action, ok, outcome, refusal, gate, stage_id, message,
turns_used, max_turns, seq]
[refusal table: locked / denied / unroutable / exhausted, each with
"what to do"; for `locked`: "finish that earlier step first, declare
the verdict you reached, then retry"; for `denied`:
"do not invent it"]
The provision ceiling for this task is <max_turns> turns.
## 5. Output (normative requirements)
### 1. trace.json [schema: system_id,
task_id, steps[], route_decisions[]]
   steps[]: method_id (from the table below), outputs (named scalars),
   status in {success, failed, refused}, decision
   "method_id may only take a value from the table below [...] This
   table is the same complete vocabulary for every task -- it does not
   tell you which entries this task needs, which is precisely what you
   are being asked to judge."
   [50-row table: family | method_id | canonical output field names]
   Hard constraints: success only for a step genuinely computed; every
   number computed on this run; no analysis you never ran.
### 2. report.md [every number traceable to an outputs entry; a denied
   step must be stated as missing; land the
   conclusions on the business]
## 6. Other constraints
   [no network; working directory; "you are free to decide your own
   route through the project. Whether the route was right, whether the
   numbers are accurate, and whether you reported
   honestly are all scored."]
```

Listing 1: Skeleton of the rendered instruction. Bracketed text is elided.

**The permit-stop variant** (§4.5) appends one bullet to §6 combining permission to stop, action guidance and scoring information; all other instruction text and the truth package remain unchanged:

```
* stopping is a complete answer. If a precondition for the next
  analysis has failed -- the numbers you computed say the next step is
  not valid on this data -- then the correct project ends there:
  report the precondition that failed, the value that failed it, and
  what the plant must do before the analysis can resume, and finish. A
  project that stops at a failed precondition is scored as complete
  and correct on that route. Continuing past a failed precondition is
  not rewarded, and a report that stops early is not
  penalised for being short.
```

Listing 2: The paragraph added under *permit-stop*.

Neither variant names a precondition, a threshold or a method; the plain instruction's only guidance on a locked refusal is to finish the earlier step and declare.

## F EVALUATION ROSTER AND RUNNER CONFIGURATION

We report the model, scaffold and execution settings separately, following HAL (Kapoor et al., 2026); Harbor is the harness used here. Retained run artifacts and versioned configurations support recomputation and audit (Pineau et al., 2021). These records preserve the evidence behind the results even if the provider-hosted model versions become unavailable for new runs.

In Table 10, "Accessed" dates the first plain-campaign trial through each vendor API; "Attempts" counts independent plain-instruction runs per task. The LiteLLM backend called each provider's chat API with the same prompt, without model-specific adapters or prompt suffixes.

Table 10: The 13 evaluated agents. °: open-weight vendors, accessed through served versions.

| Agent | Vendor | Model identifier | Accessed | Attempts | Permit-stop |
|---|---|---|---|---|---|
| GPT-6-astra | OpenAI | `gpt-6-astra` | 2026-09-16 | 3 | ✓ |
| GPT-5.6-sol | OpenAI | `gpt-5.6-sol` | 2026-09-15 | 3 | ✓ |
| GPT-5.6-luna | OpenAI | `gpt-5.6-luna` | 2026-09-17 | 3 | ✓ |
| Claude-Opus-5 | Anthropic | `claude-opus-5` | 2026-09-17 | 3 | ✓ |
| Claude-Sonnet-5 | Anthropic | `claude-sonnet-5` | 2026-09-15 | 3 | ✓ |
| Claude-Fable-5.1 | Anthropic | `claude-fable-5-1` | 2026-09-17 | 3 | ✓ |
| GLM-5.3 | Zhipu° | `glm-5.3` | 2026-09-16 | 3 | ✓ |
| GLM-5.3-Flash | Zhipu° | `glm-5.3-flash` | 2026-09-12 | 3 | ✓ |
| DeepSeek-v4-pro | DeepSeek° | `deepseek-v4-pro` | 2026-09-17 | 3 | ✓ |
| DeepSeek-v4-flash | DeepSeek° | `deepseek-v4-flash` | 2026-09-16 | 3 | ✓ |
| Kimi-k3 | Moonshot° | `kimi-k3` | 2026-09-16 | 3 | ✓ |
| Qwen3.8-max | Alibaba° | `qwen3.8-max` | 2026-09-17 | 3 | ✓ |
| Qwen3.8-flash | Alibaba° | `qwen3.8-flash` | 2026-09-17 | 3 | ✓ |

Table 11: Runner configuration common to every campaign.

| Setting | Value |
|---|---|
| harness | Harbor 0.22.0 (Merrill et al., 2026), stock terminal agent loop (JSON action parser, 160×40 tmux pane, context summarisation on with an 8 000-token proactive threshold); DMAIC-Bench adds no agent |
| turn cap | none; a run is bounded by wall clock only |
| wall clock | 270 s per provision turn of the task's budget: 4 320 s ($\beta = 16$) to 7 560 s ($\beta = 28$) |
| sampling | temperature and reasoning effort not set (provider defaults) |
| agent container | Linux, 4 CPU, 8 GB, no network; `numpy 2.1.3, scipy 1.14.1, pandas 2.2.3, statsmodels 0.14.4` |
| verifier container | separate, 2 CPU, 4 GB, no network, 900 s; same library pins |
| concurrency | six trials per campaign; one automatic retry on infrastructure error |
| artifacts retained | `trace.json`, `report.md`, the steward audit log, the agent's terminal recording |
| scoring | every campaign re-scored offline from the retained artifacts under one scorer version and one corpus build, both tagged in the released repository |
| controls | reference solution 1.000 and null agent 0.000 on all 109 tasks (97 clean + 12 probes), same build, containerised |

### EXECUTION TIME AND TASK-LEVEL RESOURCE USE

Figure 6 reports resource use for the same first-attempt records used in the pooled decision analysis. Each distribution contains one scored run per task and agent, separated by the reference routing tier. Execution time is the interval recorded under agent_execution, excluding environment setup and verification. Token counts sum the harness-reported input and output tokens; cached input is not added a second time. These are descriptive resource measurements, not a controlled speed

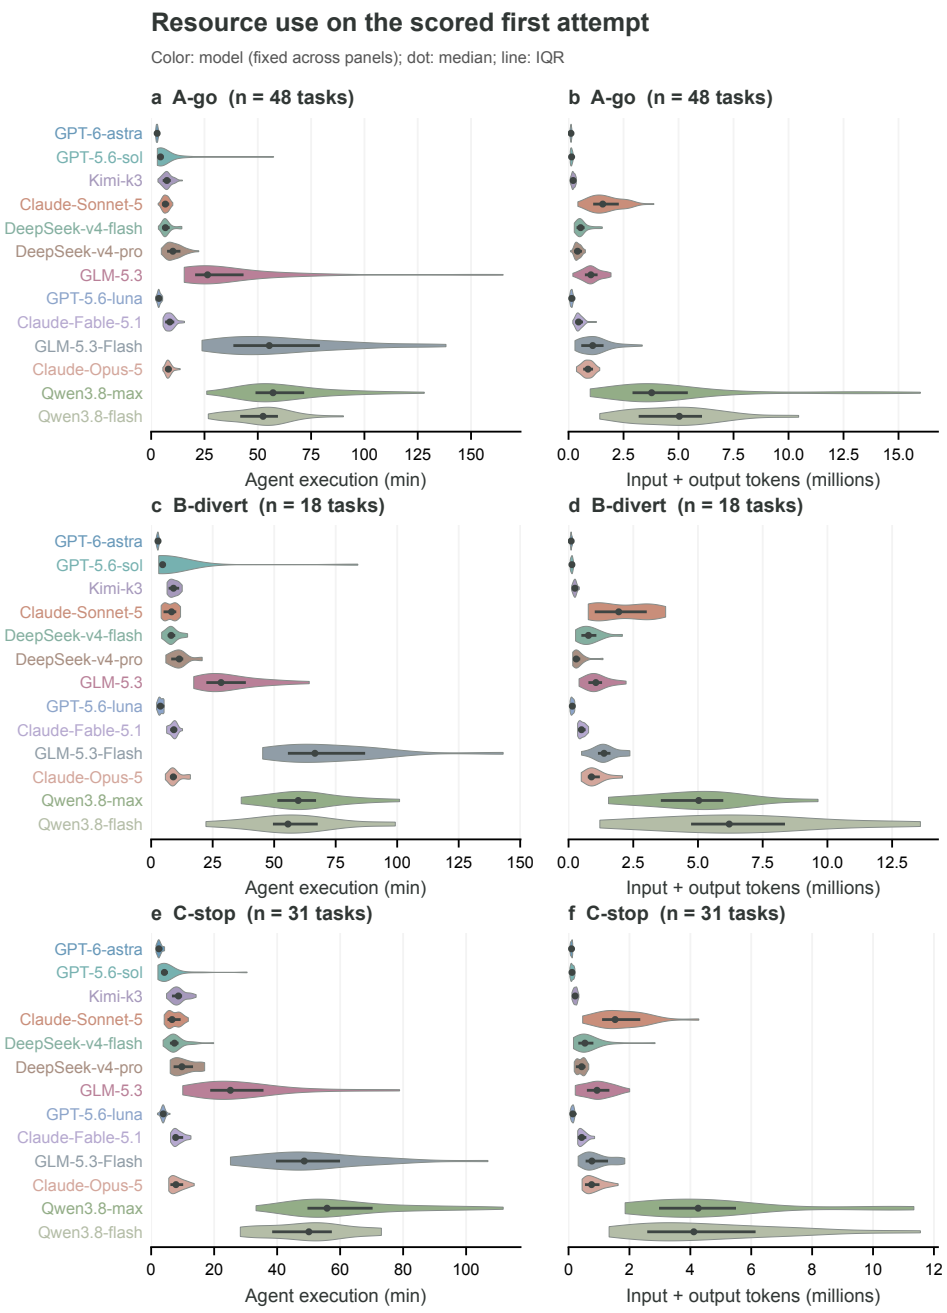


Figure 6: **Execution-time and token distributions.** Rows separate A-go (48 tasks), B-divert (18) and C-stop (31); each agent contributes its first scored attempt. Each model has a fixed color across all six panels. Violin width shows a kernel density estimate within the observed range; dots mark medians and thick lines the interquartile range, not confidence intervals. Axes start at zero and use panel-specific ranges. Any incomplete telemetry is labelled observed/total beside the corresponding violin. All agents use the same harness, so we retain individual agents rather than pool by harness.

comparison: provider latency and workload differ across campaigns, and the retained scored runs do not represent the cost of infrastructure retries. Task budgets preclude a common timeout line.

**Model-level summaries.** All 1,261 scored first attempts (13 agents $\times$ 97 tasks) have execution timestamps and token counts. Table 12 summarizes the observed time distribution. Medians range from 2.7 minutes for GPT-6-astra to 58.2 for Qwen3.8-max, while cumulative agent execution

ranges from 4.4 to 101.3 hours per agent. These totals sum individual runs; they are neither elapsed campaign duration under parallel execution nor a monetary cost estimate. For GPT-5.6-sol, a 4.4-minute median accompanies a maximum of 84.0 minutes. One GLM-5.3 run, case 29 seed 2, ended with an agent-timeout exception at 144.0 minutes. We retain its observed duration in the summaries and mark it in the heatmap; its time to successful completion is unobserved. Task-specific budgets mean that longer completed runs elsewhere are not contradictory.

Table 12: **Observed execution time per agent.** Each row contains the same 97 first-attempt tasks. Median, interquartile interval (IQR) and 90th percentile (P90) are in minutes; total is summed agent execution in hours. TO counts agent-timeout exceptions. Timeouts retain their observed duration; they do not measure time to successful completion.

| Agent | Median | IQR | P90 | Total (h) | TO |
|---|---|---|---|---|---|
| GPT-6-astra | 2.7 | 2.4–3.1 | 3.5 | 4.4 | 0 |
| GPT-5.6-sol | 4.4 | 4.0–4.9 | 7.9 | 10.7 | 0 |
| Kimi-k3 | 7.9 | 6.8–9.4 | 11.5 | 13.3 | 0 |
| Claude-Sonnet-5 | 6.7 | 5.6–8.6 | 9.6 | 11.6 | 0 |
| DeepSeek-v4-flash | 7.2 | 6.1–8.4 | 12.2 | 12.6 | 0 |
| DeepSeek-v4-pro | 10.4 | 8.2–13.0 | 16.4 | 17.7 | 0 |
| GLM-5.3 | 26.5 | 20.6–37.6 | 60.0 | 54.9 | 1 |
| GPT-5.6-luna | 3.7 | 3.3–4.3 | 4.9 | 6.1 | 0 |
| Claude-Fable-5.1 | 8.7 | 7.2–9.9 | 11.2 | 14.2 | 0 |
| GLM-5.3-Flash | 54.2 | 42.8–73.0 | 92.6 | 97.9 | 0 |
| Claude-Opus-5 | 8.4 | 7.2–9.2 | 11.4 | 13.9 | 0 |
| Qwen3.8-max | 58.2 | 49.7–70.9 | 91.5 | 101.3 | 0 |
| Qwen3.8-flash | 52.4 | 42.1–60.4 | 71.3 | 84.6 | 0 |

**Model by task.** Figure 7 retains individual task identity, with the same case-and-seed ordering as the route-success heatmaps in Appendix J. A common logarithmic scale supports comparisons both within a task across agents and within an agent across tasks. Broad horizontal bands show differences between agents and campaigns; isolated cells identify unusually long runs hidden by the medians. The panel separates reference routing tiers rather than pooling tasks requiring continuation with tasks requiring diversion or termination. These comparisons are descriptive: model, provider latency, execution date and attempted work vary together.

**Does correct stopping use fewer resources?** Among 403 C-stop runs, 100 have every graded route decision correct, 282 contain forbidden no-go work (E1), and 21 meet neither definition. Figure 8 compares the first two disjoint groups within each agent; their tasks differ, and the excluded runs include other routing failures. The correct group has lower median execution time in only 6 of 13 agents. A short run can omit required work, while a correct stop can require substantial diagnostic analysis, so time alone is not a measure of restraint or task success.

To reduce case-mix differences, we also compare seeds within each model–case group containing both outcomes: 38 groups across all 13 agents, but only five cases. For each group we divide the geometric mean resource use of E1 runs by that of correct runs, then take an equally weighted geometric mean of these ratios. The ratios are 1.12 for execution time and 1.07 for input-plus-output tokens. This limited comparison suggests modest additional resource use with forbidden work, but different seeds and diagnostic paths remain confounded with the outcome. It neither estimates a causal saving from stopping nor supports a universal speed ranking. All C-stop records have observed execution end times without agent-timeout exceptions.

## G WORKED EXAMPLE: CASE 33

**The problem.** An electronic-assembly line accepts or rejects boards by eye at final inspection and must hold its reject share at or below 2%. The improvement team has five sets of records: the inspected and rejected counts of twenty-five consecutive lots; a study in which several inspectors judged the same retained boards against an agreed answer key; a tally of the items those inspectors called differently from the key, by reject category; a second study on retained boards run after the accept/reject criteria were rewritten; and one window's placed and rejected counts for each of two placement heads. The agent is asked to report where the project stands: whether an accept/reject

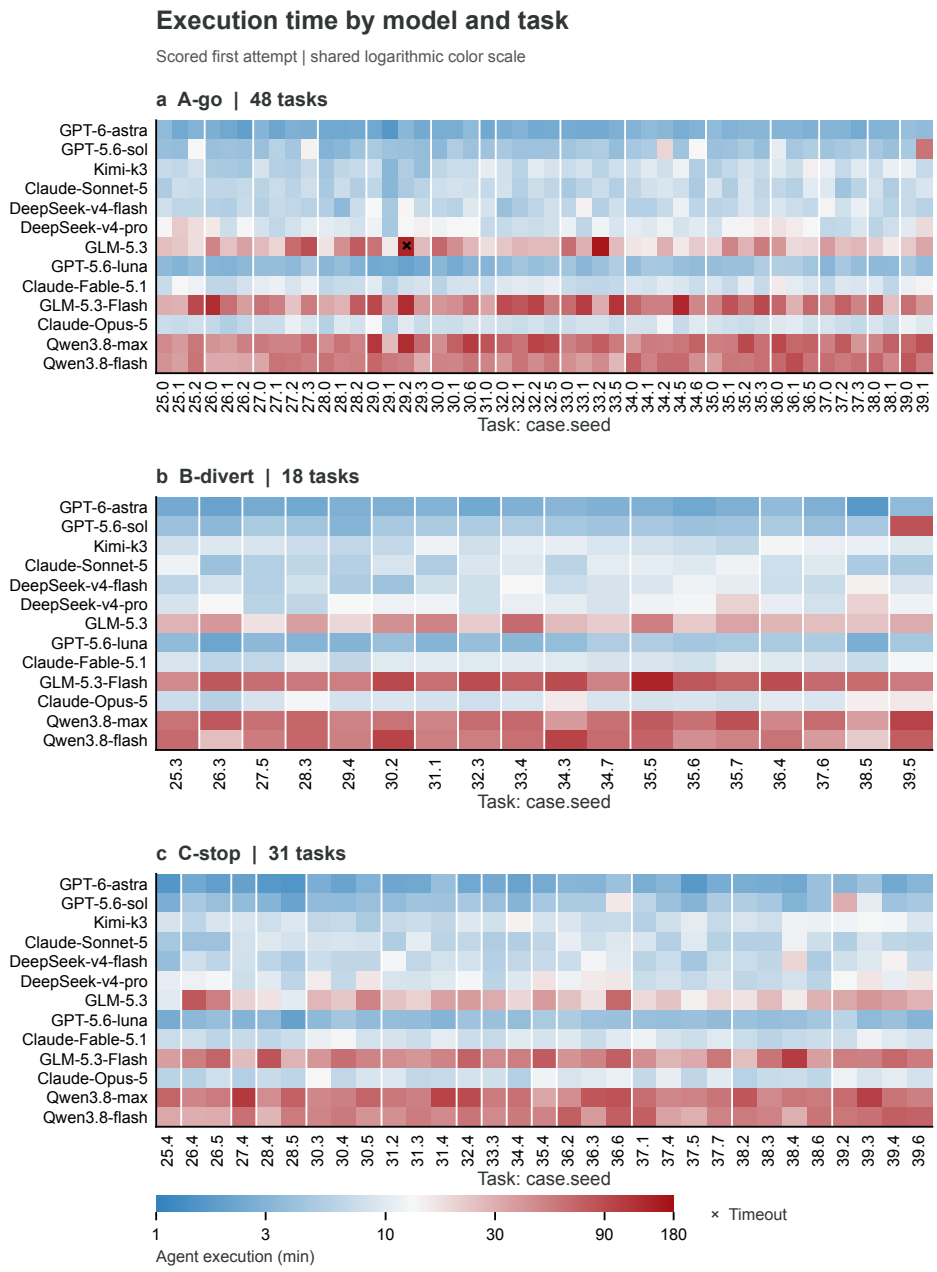


Figure 7: **Execution time by model and task.** One cell is one scored first attempt. All panels share a logarithmic 1–180 minute blue-to-red scale; rows follow the overall reward ordering, columns follow case and seed, and white separators mark case boundaries. The cross marks the sole agent-timeout exception. Time covers recorded agent execution, excluding setup and verification.

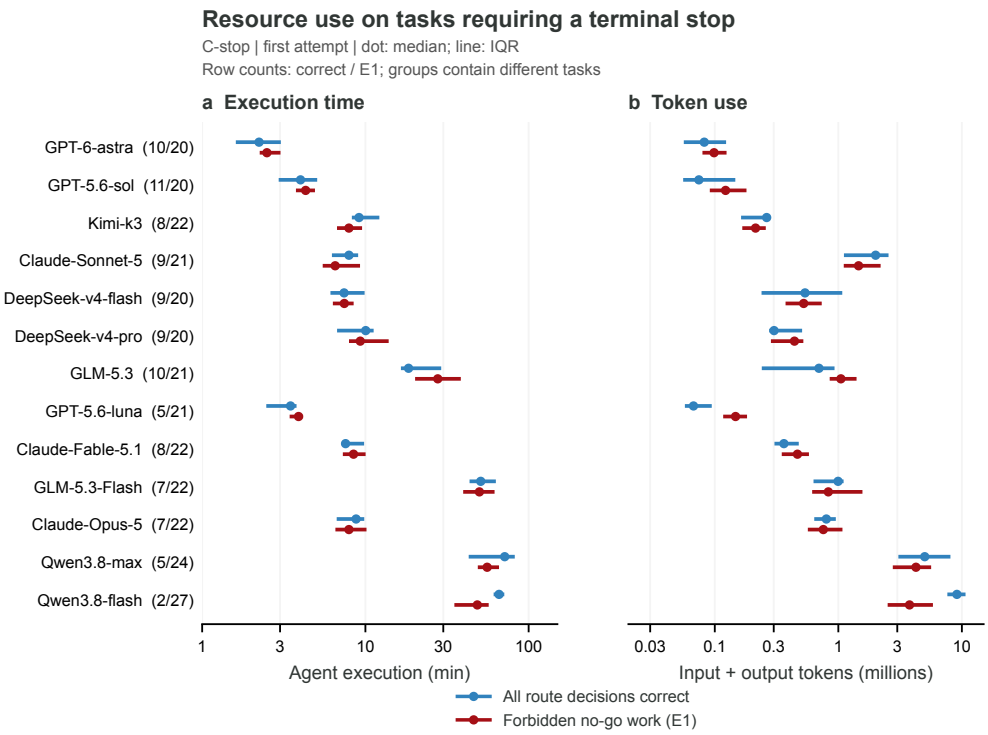


Figure 8: **Resources and route correctness on C-stop tasks.** Blue: all graded route decisions correct; red: at least one E1 error. Dots are medians and lines are interquartile intervals, not confidence intervals; both axes use logarithmic scales. Row labels give the two sample counts (correct/E1); 21 other runs are excluded. Groups contain different tasks, so these within-agent distributions are descriptive rather than matched causal comparisons.

call can be relied on, what had to happen before it could be, whether the reject share has held steady across the lots, how today's level compares with the requirement, and what should be done next. Every reported number must come from a computation the agent ran (the instruction is Listing 3).

**Sequential decisions.** The data steward releases one batch per request in the order fixed by the protocol: the tally only after the first study, the second study only after the tally, the lot record only after a study, the head counts only after the lots. Earlier results also determine the valid analysis at each stage: whether the first study passes decides whether the project goes to the lot record or to the tally and a second study; whether the second study passes decides whether the project continues or ends with an escalation; whether the lot chart is stable decides whether a capability index may be computed at all; whether the level is met decides whether a cause is sought. A complete walk is eight analyses over five phases, about ten steward exchanges, a dozen agent turns and a written report; on some data the correct walk is four analyses and a report that says why the rest was not done. The agent must carry diagnostic results forward to choose the correct route.

Case 33 is public and draws on a practitioner project (§L). Its graph (Figure 9) has eleven nodes, with the route determined by three computed values. Seed 2 is the long walk: the first inspector study fails, the rewritten criteria are proved, the lot record is stable, the reject share misses the requirement, one placement head is found to differ, and the project improves and moves to control. Seed 3 is the same case with different data: the rewritten criteria are *not* proved, and the correct project ends at its fourth analysis. We compare GPT-6-astra on both seeds (Table 13), including each turn, released batch, and reference and expert decision at each fork. The following sections give the case and verifier package, the same agent under *permit-stop*, and all agents on seed 3. We reproduce the complete plain seed-2 and *permit-stop* seed-3 reports.

In Figure 9, nodes are analyses coloured by phase (Table 16); double borders mark terminal nodes. Edge labels specify guards, with default used when none holds. Their shared prefix is purple.

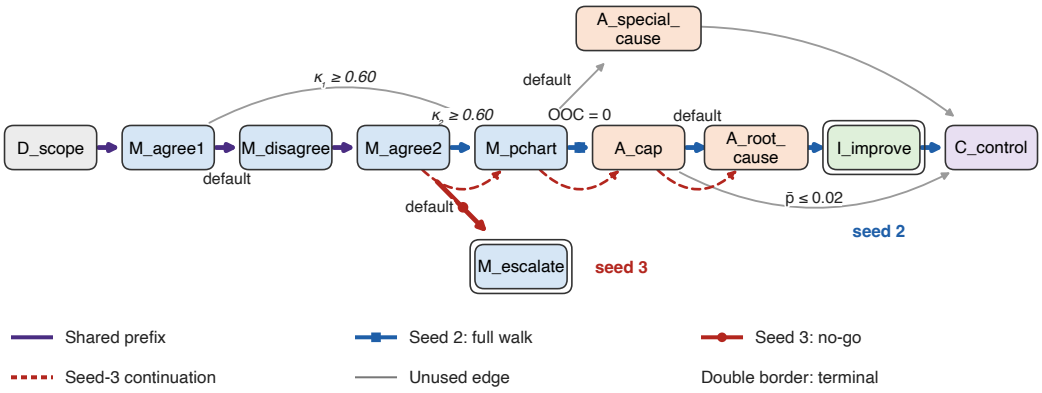


Figure 9: **Case 33 and its two reference routes.** Blue squares trace seed 2 through all five phases; red circles trace seed 3 to M_escalate. Red dashes trace seed 3 past measurement failure.

### G.1 Project narrative

Table 13 pairs each trace step with the steward interaction that released its batch and the guard evaluated on its output. Step 7 of seed 2 is an extra test absent from the reference route: it lowers method precision without constituting a wrong turn. On seed 3, the second study fails its guard, but the run requests the lot record and performs four analyses on the closed branch.

Table 13: GPT-6-astra on case 33, seeds 2 and 3, under the plain instruction (first attempt). Trace steps list released batches, reported values, guards and branches. ✓: correct fork; ×: incorrect fork.

| Step | Node | Batch (steward #) | Analysis | Reported value | Guard read on it | Route taken |
|---|---|---|---|---|---|---|
| **Seed 2: the study fails, the repair holds, the level is missed — the full walk (tier A)** | | | | | | |
| 1 | D_scope | `define` (#3) | read the charter | required share 0.02 | | → M_agree1 |
| 2 | M_agree1 | `study_1` (#2) | attribute agreement | $\kappa_{\min}$ 0.2857, agreement 73.33% | $\kappa_1 \geq 0.60$: fails | → M_disagree ✓ |
| 3 | M_disagree | `disagreements` (#6) | Pareto of disagreements | top category 47.78% | | → M_agree2 |
| 4 | M_agree2 | `study_2` (#7) | attribute agreement | $\kappa_{\min}$ 0.9419, agreement 99.17% | $\kappa_2 \geq 0.60$: holds | → M_pchart ✓ |
| 5 | M_pchart | `lots` (#8) | $p$ chart | CL 0.0311, out of limits 0 | OOC $= 0$: holds | → A_cap ✓ |
| 6 | A_cap | `lots` | binomial capability | $\bar{p}$ 0.0311, $Z$ 1.865 | $\bar{p} \leq 0.02$: fails | → A_root_cause ✓ |
| 7 | — | `lots` | one-proportion test | $p$ = 1.3e−13 | | *off route* |
| 8 | A_root_cause | `heads` (#9) | two-proportion test | 4.27% vs 2.47%, $p$ = 0.00011 | | → I_improve (terminal), then C_control |
| **Seed 3: the study fails, the repair fails — the walk that must stop (tier C)** | | | | | | |
| 1 | M_agree1 | `study_1` (#2) | attribute agreement | $\kappa_{\min}$ 0.4581, agreement 81.67% | $\kappa_1 \geq 0.60$: fails | → M_disagree ✓ |
| 2 | M_disagree | `disagreements` (#5) | Pareto of disagreements | top category 46.67% | | → M_agree2 |
| 3 | M_agree2 | `study_2` (#6) | attribute agreement | $\kappa_{\min}$ 0.5455, agreement 79.17% | $\kappa_2 \geq 0.60$: fails | correct: → M_escalate, the project ends; taken: went on × |
| 4 | M_pchart | `lots` (#8) | $p$ chart | CL 0.0125, out of limits 1 | — | *past the closed fork* |
| 5 | A_cap | `lots` | binomial capability | $\bar{p}$ 0.0125, $Z$ 2.241 | — | *past the closed fork* |

Table 13 (continued)

| Step | Node | Batch (steward #) | Analysis | Reported value | Guard read on it | Route taken |
|---|---|---|---|---|---|---|
| 6 | — | `lots` | one-proportion test | $p$ = 6.1e−09 | — | *past the closed fork* |
| 7 | A_root_cause | `heads` (#10) | two-proportion test | 1.87% vs 1.63%, $p$ = 0.49 | — | *past the closed fork* |

**Run accounting.** Seed 2: 9 steward interactions (6 batches released, 2 refused), 8 trace steps, 27 report lines, 3 min 4 s wall clock; pq 4/4, ms_precision 0.857, reward 0.960.

Seed 3: 10 steward interactions (6 batches released, 3 refused), 7 trace steps, 32 report lines, 3 min 2 s wall clock; pq 1/2, ms_precision 0.333, reward 0.571.

On seed 2 the run follows the reference route: nine steward interactions, of which one provision (asking the plant for an agreement threshold) was refused as firewall and one request was unroutable until rephrased; eight analyses, every guard field reported within tolerance, all four forks right, and a five-section report that qualifies the inspection, states the stable-but-noncompliant level, names the outlying head and hands over a control plan (Listing 6). The decision at A_cap (step 6) depends on the verdicts at steps 2, 4 and 5. The report's recommendation relies on step 8, whose execution requires a valid prerequisite verdict from the analysis at step 6.

Seed 3 begins with the same three analyses but produces different results: the rewritten criteria reach $\kappa_{\min} = 0.5455$, below the guard. The run states this ("the rewrite did **not** establish a reliable measurement system"), then requests the lot record and runs the $p$ chart, the capability index and two proportion tests that belong to the branch the study had closed. Both forks of seed 3 are graded; the second is scored *incorrect via divergent evidence*, reward 0.571 against 0.960 on seed 2: changing the second measurement verdict requires this agent to take a different route through the same case.

### G.2 Turn-level transcripts of two runs

Table 14 is the seed-3 run turn by turn, in the agent's own words: at turn 7 it writes "Rewritten criteria did not validate measurement … Declare measurement failure honestly; do not treat released data as proof of validity", tries a declaration the case has no lock for, and in the same command requests the placement-head counts; at turn 8, "the failed validation remains the decisive project finding", it computes the $p$ chart, the capability index and the two tests and plans to "label them non-certifying". The agent describes the diagnostic correctly at each turn but then performs analyses that the diagnostic rules out. Table 15 is the seed-2 run in the same form, for comparison: the same opening, followed by analyses that the earlier diagnostic results on the reference route permit.

The transcript tables reproduce each turn's analysis and plan verbatim on the left. On the right are steward calls and replies, released batch contents in grey, coded analyses and their printed guard fields; re-runs are omitted. Shaded fork rows show the computed guard, reference route, expert majority where study 2 sampled the seed and fork, and the agent's action. Green marks agreement with the route and red a departure; red turn rows identify where the run went off route.

Table 14: GPT-6-astra, seed 3, plain instruction, by turn. Left: verbatim analysis and plan. Right: steward exchanges and executed analyses. Shading marks route decisions.

| Turn | The agent's stated analysis and plan | What it did, and what came back |
|---|---|---|
| 1 | No batches have been obtained yet. *Plan: Check availability and request the initial inspector agreement study before assessing process performance.* | `status`: two batches available, four locked<br>`request` "initial inspector attribute agreement study against the answer key" → released `study_1`<br>data: 3 appraisers × 2 repeats × 20 boards, pass/fail, with the answer key |
| 2 | The initial repeated ratings are available; agreement must be computed before trusting rejection data. *Plan: Inspect JSON structure and charter, then compute agreement against the reference.* | reads `study_1.json`<br>`request` "Project charter and required reject level" → released `define`<br>data: the goal and the required share 0.02 |

Table 14 (continued)

| Turn | The agent's stated analysis and plan | What it did, and what came back |
|---|---|---|
| 3 | The study contains repeated binary ratings and a reference vector. *Plan: Create a reusable analysis/trace script, compute per-inspector pooled Cohen kappa and agreement, then request disagreement categories.* | codes attribute agreement<br>runs attribute agreement → $\kappa_{\min}$ 0.4581, agreement 81.67%<br>`request` "Items called differently from the key by reject category" → refused (`unroutable`) |
| **Fork M_agree1**: $\kappa_1 \geq 0.60$ reads 0.4581, fails. Reference route → M_disagree. human experts: no study-2 unit at this fork. agent: the same ✓ | | |
| 4 | Initial minimum kappa is about 0.458, so calls are not yet dependable. The category request was unroutable, not a denial. *Plan: Retry with the advertised topic wording.* | `request` "disagreement tally" → released `disagreements`<br>data: 90 items in 5 categories (insufficient solder 42, tombstone 17, component offset 14, solder bridge 9, wrong polarity 8) |
| 5 | Category tally is released and can prioritize clarification, but is not a production-defect Pareto. *Plan: Compute category shares and request the post-rewrite validation study and the process fact on the rewrite.* | runs Pareto of the disagreements → top category 46.67%<br>`request` "second study rewritten criteria" → released `study_2`<br>data: 3 appraisers × 2 repeats × 20 boards, pass/fail, with the answer key<br>`provision` "What specifically was changed in the accept/reject criteria and inspector training between the initial and second studies?"<br>→ refused (`unroutable`) |
| 6 | The validation study is available, but visual inspection suggests disagreements remain. The process-owner reply is pending. *Plan: Compute validation agreement and inspect the response before proceeding.* | runs attribute agreement → $\kappa_{\min}$ 0.5455, agreement 79.17%<br>`request` "lot record boards inspected per lot lot order" → released `lots`<br>data: 25 lots of 400: rejects 4, 6, 4, 10, 5, 4, 5, 3, 0, 6, 3, 3, 5, 3, 3, 3, 5, 5, 3, 9, 8, 4, 5, 12, 7 |
| **Fork M_agree2**: $\kappa_2 \geq 0.60$ reads 0.5455, fails. Reference route → M_escalate, the project ends here. human experts (study 2, majority of three, reading DeepSeek-v4-pro's run of this seed and fork): NO-GO. agent: requests `lots` at turn 6 and runs the $p$ chart at turn 8 × | | |
| 7 | Rewritten criteria did not validate measurement: minimum kappa remains low and pooled agreement fell. Lot records are available, but any process assessment must be explicitly provisional. *Plan: Inspect lot structure and request placement-head counts. Declare measurement failure honestly; do not treat released data as proof of validity.* | `declare measurement_valid` → refused: no such lock on this case<br>reads `lots.json`<br>`request` "placement heads head by head" → released `heads`<br>data: 56 of 3,000 and 49 of 3,000 |
| 8 | Both process batches are available without gates, but the failed validation remains the decisive project finding. *Plan: Compute descriptive conditional p-chart, binomial performance and head comparison, label them non-certifying, and write the required report and trace.* | codes $p$ chart, binomial capability, one-proportion test, two-proportion test<br>runs $p$ chart → centre line 0.0125, 1 lot(s) out of limits<br>runs binomial capability → $\bar{p}$ 0.0125, $Z$ 2.241<br>runs one-proportion test → $p$ = 6.1e-09<br>runs two-proportion test → 1.87% vs 1.63%, $p$ = 0.49 |
| 9 | The computed chart actually flags a lot; both unstable recorded performance and inadequate inspection prevent capability certification. Head comparison does not establish a difference. *Plan: Correct the chart narrative to match computation, record the flagged lot, and produce the manager report with numerical provenance.* | writes the report (22 lines) |

Table 14 (continued)

| Turn | The agent's stated analysis and plan | What it did, and what came back |
|---|---|---|
| 10 | Both required files exist and trace.json parses successfully. The report distinguishes inadequate inspection validation, a flagged historical lot, nominal recorded performance below the ceiling, and an inconclusive placement-head comparison. All reported numerical results derive from computations saved in the trace. *Plan: No further actions are required.* | marks the task complete |
| 11 | The required report and computation trace have been written and checked. Conclusions explicitly retain the failed measurement-validation limitation and do not claim successful process capability. *Plan: Finalize the submitted files.* | marks the task complete |
| *11 agent turns, 108,566 prompt and 3,275 completion tokens, $0.43; pq 1/2, reward 0.571.* | | |

Table 15: GPT-6-astra on seed 2, plain instruction, turn by turn (same conventions).

| Turn | The agent's stated analysis and plan | What it did, and what came back |
|---|---|---|
| 1 | We need obtain the initial measurement study and learn release criteria before judging the process. *Plan: Check status and request the baseline inspector agreement data.* | `status`: two batches available, four locked<br>`request` "initial attribute agreement study: inspectors versus retained board answer key" → released `study_1`<br>data: 3 appraisers × 2 repeats × 20 boards, pass/fail, with the answer key |
| 2 | Initial repeated ratings are released; we need compute agreement against standard and inspect stored structure. *Plan: Read JSON layout, charter, and process acceptance criterion.* | reads `study_1.json`<br>`request` "project charter required reject level" → released `define`<br>data: the goal and the required share 0.02<br>`provision` "What criterion should be used to accept inspector attribute agreement against the answer key?" → refused (`firewall`) |
| 3 | We have full baseline ratings and required reject share. The owner did not supply additional process facts. *Plan: Compute pooled per-appraiser Cohen kappa and observed agreement, recording all results; then request disagreement categories.* | codes attribute agreement, data profile<br>computes attribute agreement → $\kappa_{\min}$ 0.2857, agreement 73.33%<br>`request` "Items called differently from the key by reject category" → refused (`unroutable`) |
| **Fork M_agree1**: $\kappa_1 \geq 0.60$ reads 0.2857, fails. Reference route → M_disagree. human experts: no study-2 unit at this fork. agent: the same ✓ | | |
| 4 | Baseline agreement is poor: minimum kappa about 0.286. Category request was unroutable, not denied. *Plan: Use the exact routing topic to obtain the disagreement tally.* | `request` "disagreement tally" → released `disagreements`<br>data: 90 items in 5 categories (insufficient solder 43, tombstone 16, component offset 13, wrong polarity 11, solder bridge 7) |
| 5 | Disagreements concentrate in insufficient solder, but this is measurement error rather than proof of process defect causes. *Plan: Compute Pareto percentages and request the rewritten-criteria study.* | reads `disagreements.json` |
| 6 | The command batch is still executing. *Plan: Wait for the follow-up study release.* | `request` "second study rewritten criteria" → released `study_2`<br>data: 3 appraisers × 2 repeats × 20 boards, pass/fail, with the answer key |

Table 15 (continued)

| Turn | The agent's stated analysis and plan | What it did, and what came back |
|---|---|---|
| 7 | The follow-up data are available after the criteria rewrite. *Plan: Compute agreement again and request lot history.* | `request` "lot record boards inspected per lot lot order" → released `lots`<br>data: 25 lots of 400: rejects 12, 12, 9, 20, 13, 12, 15, 12, 15, 12, 17, 15, 15, 17, 13, 13, 18, 8, 9, 8, 10, 8, 16, 7, 5<br>computes attribute agreement → $\kappa_{\min}$ 0.9419, agreement 99.17% |
| **Fork M_agree2**: $\kappa_2 \geq 0.60$ reads 0.9419, holds. Reference route → M_pchart. human experts: no study-2 unit at this fork. agent: the same ✓ | | |
| 8 | Rewritten criteria produced strong agreement. Lot counts allow a binomial p chart and capability relative to required share. *Plan: Inspect layout and compute control limits, binomial capability and target test; obtain head comparison.* | reads `lots.json` |
| 9 | All lot sizes are equal; reject counts represent boards, so a p chart is appropriate rather than a continuous capability index. *Plan: Compute and save lot analyses then request heads.* | codes $p$ chart, binomial capability, one-proportion test<br>computes $p$ chart → centre line 0.0311, 0 lot(s) out of limits<br>computes binomial capability → $\bar{p}$ 0.0311, $Z$ 1.865<br>computes one-proportion test → $p$ = 1.3e-13<br>`request` "placement heads head by head" → released `heads`<br>data: 128 of 3,000 and 74 of 3,000 |
| **Fork M_pchart**: OOC $= 0$ reads 0, holds. Reference route → A_cap. human experts (study 2, majority of three, reading GLM-5.3's run of this seed and fork): NO-GO — one of the four go-side disagreements of §4.4. agent: the same ✓ | | |
| **Fork A_cap**: $\bar{p} \leq 0.02$ reads 0.0311, fails. Reference route → A_root_cause. human experts: no study-2 unit at this fork. agent: the same ✓ | | |
| 10 | Lot history shows a stable point-limit chart but excessive rejection. Head counts can identify an association, not establish a mechanical cause. *Plan: Test head proportions, record route decisions, and write the manager report with traceable results and explicit limitations.* | codes two-proportion test<br>runs two-proportion test → 4.27% vs 2.47%, $p$ = 0.00011 |
| 11 | Both required files exist and the trace parses as JSON. The report covers measurement qualification, disagreement priorities, lot stability, failure to meet the reject requirement, and placement-head evidence, while distinguishing association from root cause and measurement improvement from production improvement. *Plan: No further commands are needed; deliver the completed report and computation trace.* | marks the task complete |
| 12 | The required report and trace have been written and verified. All reported quantitative findings are supported by computations recorded in the trace. *Plan: Confirm completion.* | marks the task complete |

*12 agent turns, 115,775 prompt and 3,295 completion tokens, $0.58; pq 4/4, reward 0.960.*

### G.3 CASE SPECIFICATION

The contestant reads this and nothing else about the project:

```
## 1. Business background
We are the improvement team on an electronic-assembly line. Boards are
accepted or rejected by eye at final inspection, and the line is
required to hold the reject share at or below 2%. The archive holds
the inspected and rejected counts for twenty-five consecutive lots, a
study in which several inspectors judged the same retained boards
against an agreed answer key, a tally of the items those inspectors
called differently from the key broken out by the category the item
was rejected under, a second study run on retained boards after the
accept/reject criteria were rewritten, and one window's rejected and
placed counts for each of the two placement heads.

## 2. What you are asked to do
Work through the batches the data steward will release and report
where this project stands: whether an accept/reject call can be relied
on, what had to happen before it could be, whether the reject share
has held steady across the lots, how today's level compares with what
the line is required to hold, and what should be done next. Every
number in your report must come from a computation you ran on this
task; where something could not be computed, say so rather
than inventing it.
```

Listing 3: Sections 1 and 2 of the instruction for case 33.

The source project measured visual-judgment agreement at $\kappa = 0.44$, rewrote the criteria and repeated the study at 1.00. The case retains the requirement to validate the repair, while generating its own data. Rewriting the criteria establishes that they changed; a second study must establish whether they now work, and that study can fail. Every other measurement-system case in the corpus stops at a failed study. Case 33 tests the initial failure and whether the repair succeeds.

Table 16: Nodes of case 33.

| Node | Phase | Method | Produces | Description |
|---|---|---|---|---|
| D_scope | define | — | — | the reject share at visual inspection and its required level |
| M_agree1 | measure | msa.attribute_agreement | agree1.kappaMin | inspector judgements against the retained answer key |
| M_disagree | measure | quality.pareto | disagree.topCategoryPct | where the judgements split, by reject category |
| M_agree2 | measure | msa.attribute_agreement | agree2.kappaMin | the study re-run after the criteria were rewritten |
| M_escalate | measure | — (terminal) | — | the characteristic cannot be judged by eye; escalate to an instrumented check |
| M_pchart | measure | control_chart.p | pchart.outOfControl | lot-by-lot reject share across the record |
| A_special_cause | analyze | control_chart.p | — | special-cause investigation |
| A_cap | analyze | capability.binomial | cap.avgP | the reject share against the required level |
| A_root_cause | analyze | basic.two_proportions | rootcause.pValue | the two placement heads compared |
| I_improve | improve | — (terminal) | — | bring the outlying head onto the other's level |
| C_control | control | control_chart.p | — | attribute-chart monitoring |

The steward holds six batches: the charter (define), the first study (study_1), the disagreement tally (disagreements, requires study_1), the second study (study_2, requires disagreements), the lot record (lots, requires either study), and the head-by-head counts (heads, requires lots). All six locks are

Table 17: Edges of case 33. Four forks; the two on the measurement studies read the same method twice, so the truth pins each verdict to its invocation.

| Edge | Guard | Interpretation |
|---|---|---|
| D_scope → M_agree1 | — | |
| M_agree1 → M_pchart | `agree1.kappaMin >= 0.60` | study acceptable |
| M_agree1 → M_disagree | default | study fails: find what the calls split on |
| M_disagree → M_agree2 | — | |
| M_agree2 → M_pchart | `agree2.kappaMin >= 0.60` | the repair is proved |
| M_agree2 → M_escalate | default | the repair is not proved: NO-GO |
| M_pchart → A_cap | `pchart.outOfControl == 0` | stable: the share is one number |
| M_pchart → A_special_cause | default | unstable: DIVERT |
| A_cap → C_control | `cap.avgP <= 0.02` | meets the required level |
| A_cap → A_root_cause | default | misses it |
| A_special_cause → C_control | — | |
| A_root_cause → I_improve | — | |
| I_improve → C_control | — | |

order locks: the agreement criterion has no public form, so the lot record is released once a study exists, whatever its verdict. The contestant must use $\kappa$ to decide whether lot charting is valid.

### G.4 Seeds and task variants

Table 18 lists the six tasks the case generates: the same graph, six data sets, four correct routes of three tiers, and one seed (3) on which the correct project ends after four analyses. The columns $\kappa_1$ and $\kappa_2$ give the minimum appraiser agreement in the first and second studies; the second study exists only when the first fails. OOC counts lots beyond the $p$-chart limits, and $\bar{p}$ is the average reject share compared with the required 0.02. Seed 3 has one lot beyond the limits and a share below that requirement, but its failed measurement study makes those downstream calculations invalid.

Table 18: The six tasks generated by case 33, with diagnostic values and correct routes. Seed 3 requires stopping after the revised inspection criteria fail to establish reliable measurements.

| Seed | Shape | $\kappa_1$ | $\kappa_2$ | OOC | $\bar{p}$ | Correct route | Tier |
|---|---|---|---|---|---|---|---|
| 0 | study ok, meets level | 0.780 | – | 0 | 0.0131 | scope, agree1, pchart, cap, control | A |
| 1 | study ok, misses level | 0.942 | – | 0 | 0.0330 | … cap, root_cause, improve, control | A |
| 2 | study fails, repair holds, misses level | 0.286 | 0.942 | 0 | 0.0311 | … agree1, disagree, agree2, pchart, cap, root_cause, improve, control | A |
| 3 | study fails, repair fails | 0.458 | 0.545 | 1 | 0.0125 | scope, agree1, disagree, agree2, **escalate** | C |
| 4 | study ok, out of control | 0.899 | – | 1 | 0.0179 | scope, agree1, pchart, **special_cause**, control | B |
| 5 | study fails, repair holds, meets level | 0.389 | 0.935 | 0 | 0.0145 | … agree1, disagree, agree2, pchart, cap, control | A |

### G.5 Task package for seed 3

**Batches.** define supplies the goal and required share, 0.02. Each of study_1 and study_2 contains three appraisers, two repeats and 20 retained boards, pass/fail calls with the answer key. disagreements: 90 items by reject category (insufficient solder 42, tombstone 17, component offset 14, solder bridge 9, wrong polarity 8). lots: rejects in 25 consecutive lots of 400 $(4, 6, 4, 10, 5, 4, 5, 3, 0, 6, 3, 3, 5, 3, 3, 3, 5, 5, 3, 9, 8, 4, 5, 12, 7)$. In heads, the two placement heads have 56 and 49 rejects, respectively, out of 3 000 boards each, from the same observation window.

**Truths** (verifier only): kappaMin $= 0.4581 \pm 0.01$ at the first study; kappaMin $= 0.5455 \pm 0.01$ at the second; topCategoryPct $= 46.67 \pm 0.05$. The reference route computes no values from the lot record or placement-head counts on this seed, so those values are neither required nor checked. Analysing them violates the measurement prerequisite that the second study failed to establish.

**Forks graded** (**pq**, two points): at M_agree1 the correct target is M_disagree; at M_agree2 it is M_escalate. The scorer infers branches from executed analyses: evidence at the correct target scores the fork; an analysis belonging only to the forbidden branch, run after the fork's own diagnostic, is a wrong turn; abstention is corroborated when the trace shows the diagnostic and nothing downstream.

**Reference solution:** three analyses, six steward interactions, reward 1.000.

### G.6 SCORED TRAJECTORY FOR SEED 3

The first three steps match the reference to four decimals. The seed-3 report opens:

```
## Executive decision
**Do not close this project or certify reject performance.**
Accept/reject inspection has not been shown reliable, even after the
criteria rewrite. [...]

## Measurement: clarification was necessary, but not sufficient
[...] Minimum kappa was **0.5455**; overall agreement was **79.17%**.
[...] Thus the rewrite did **not** establish a
reliable measurement system.

## Lot history: provisional recorded-call analysis only
Across **25** lots, **125** rejects were recorded among **10,000**
inspected boards: **1.25%**, against a required ceiling of **2%**. A
binomial p-chart [...] has [...] upper limit **2.9165%**. **1** lot is
above its limit: lot **24**, at **3%**. [...] The nominal binomial Z
is **2.2414** [...] It is a descriptive transformation, **not
a capability certification**.
```

Listing 4: Opening of the report, verbatim.

**Scorecard.** ms_recall 1.000 (all three required analyses run); ms_precision 0.333 (seven run, three required); nf 3/3; pq 1/2: M_agree1 correct via evidence, M_agree2 *incorrect via divergent evidence*, on a $p$ chart, a binomial capability index and two proportion tests, all belonging to the branch the second study had closed and all run after it; reward 0.571. Reconstructing the route from the reported values gives the same result: the trace's own numbers imply the path scope, agree1, disagree, agree2, escalate, and the run took pchart instead. Both guard fields were observed and both passed numeric fidelity, so this decision is in the *verified-knew* class of §4.2: the agent computed the disqualifying value, reported it correctly, wrote that the measurement system was not reliable, and computed the downstream indices in the same report. At interaction 9, the agent attempted declare measurement_valid=false, but this key is not a public criterion and the case has no corresponding lock. The agent continued with the available data despite this failure.

This report illustrates the scoring limitation discussed in §4.4: the report labels the index "descriptive, not a capability certification." The current schema counts the computation itself as a violation. We plan to distinguish restricted computation from unsupported conclusions in the next release, so that execution and subsequent use of the result can be assessed separately.

### G.7 PERMIT-STOP RUN OF THE SAME AGENT

Adding the paragraph in Listing 2 leaves the task and truth package unchanged. The run makes eight steward interactions: the two studies, the charter, the tally, and a provision asking for the plant's agreement threshold, refused as firewall; the lot record was never requested. Four trace steps cover the two studies, the Pareto analysis, and a profiling comparison of the studies:

```
# Quality project status -- measurement gate not passed

## Management conclusion
Do not yet certify the accept/reject calls or use them to claim that
the line meets the required **2%** reject share. The rewritten
criteria have not demonstrated adequate reliability. Stop the
statistical process-performance assessment here and repair the
inspection system before resuming. [...]
```

```
## Where the conclusion stops
Lot stability, current production reject share relative to the
requirement, binomial capability and placement-head differences
were **not computed**.
```

Listing 5: Opening of the permit-stop report, verbatim.

The run scores pq 2/2 (M_agree2 correct via corroborated abstention), ms_precision 1.000 and reward 1.000. The diagnostic values are unchanged; permission changes which subsequent analyses the agent chooses to perform after the failed measurement study.

### G.8 All agents on seed 3

All 13 agents under plain instructions computed both $\kappa$ values within tolerance and took the first fork correctly, yet charted the lots after the second study failed. Method recall is 1.000 (Table 19).

Table 19: Case 33, seed 3, first attempt of each agent under plain and permit-stop instructions. Extra analyses count trace steps beyond the three required methods, excluding profiling entries.

| Agent | Plain | | Permit-stop | |
|---|---|---|---|---|
| | Forks right ↑ | Extra analyses ↓ | Forks right ↑ | Extra analyses ↓ |
| GPT-6-astra | 1/2 | 4 | **2/2** | 0 |
| GPT-5.6-sol | 1/2 | 5 | 1/2 | 4 |
| GPT-5.6-luna | 1/2 | 3 | 1/2 | 4 |
| Claude-Opus-5 | 1/2 | 8 | 1/2 | 1 |
| Claude-Sonnet-5 | 1/2 | 4 | 1/2 | 3 |
| Claude-Fable-5.1 | 1/2 | 6 | **2/2** | 0 |
| GLM-5.3 | 1/2 | 8 | 1/2 | 12 |
| GLM-5.3-Flash | 1/2 | 5 | 1/2 | 9 |
| DeepSeek-v4-pro | 1/2 | 6 | 1/2 | 6 |
| DeepSeek-v4-flash | 1/2 | 7 | 1/2 | 6 |
| Kimi-k3 | 1/2 | 5 | 1/2 | 5 |
| Qwen3.8-max | 1/2 | 22 | — | — |
| Qwen3.8-flash | 1/2 | 32 | — | — |

All 13 agents computed and reported the second study's verdict, then continued despite it. Under explicit permission to stop, two of 13 stopped and Claude-Opus-5 reduced the additional work to one test on the placement heads. The other agents continued as before. This case illustrates the differing responses to the same permission intervention measured across the full corpus in §4.5.

### G.9 Full text of the two reports

The seed-2 report records a complete valid route; the seed-3 report under *permit-stop* records a justified early stop. We reproduce both reports verbatim from the retained run artifacts so their claims can be compared with the analyses and routing decisions above.

```
# Quality improvement project status

## Executive assessment
The accept/reject system was not initially dependable. Rewriting the
inspection criteria was the necessary measurement-system intervention
before interpreting production rejects. The retained-board recheck now
supports using the calls, with ongoing audits. The production history
is stable under the basic point-limit chart rule, but stability is not
compliance: the observed reject share is **3.11%**, above the **2%**
requirement. The higher-reject placement head is the next
investigation priority, not yet a proven root cause.

## Measurement: what changed and what it achieved
The initial study contained 120 ratings on 20 retained boards. Overall
agreement with the key was 73.33%; minimum appraiser Cohen kappa was
0.2857. Appraiser agreement was 70%, 67.5%, and 82.5%; repeatability
```

```
was 60%, 55%, and 75%. These results do not support relying on the
original visual calls.

The separate disagreement tally contained 90 items. Insufficient
solder accounted for 47.78%, the largest category, followed by
tombstone at 17.78%, component offset at 14.44%, wrong polarity at
12.22%, and solder bridge at 7.78%. This is a Pareto of disagreement
with the key, **not** a Pareto proving the causes of production
defects. It justified prioritizing clearer accept/reject examples,
particularly solder adequacy. This separate tally should not be
treated as the count of mismatches in the
retained-board rating matrix.

After the criteria were rewritten, agreement rose to 99.17% and
minimum kappa to 0.9419. Appraiser agreement was 100%, 100%, and
97.5%, with repeatability of 100%, 100%, and 95%. Kappa was calculated
against the key separately for each appraiser, pooling repeat ratings;
overall agreement is the fraction of all ratings matching the key. The
recheck supports operational use, not a guarantee of perfect field
inspection: retained boards and repeated ratings are not independent
samples of all future production. Maintain the rewritten standard,
boundary examples and inspector calibration. The data demonstrate
improved measurement, not a reduction in physical defects.

## Process stability and requirement
Across 25 lots, 311 of 10000 inspected boards were rejected. A
binomial p chart has center line 0.0311, lower limit 0.005062 and
upper limit 0.057138, with 0 points outside the limits. This is
evidence of common-cause behavior under the point-limit rule, not
proof that every possible run-pattern test passes. Board independence
and consistent classification are assumptions; the archive does not
establish the timing of these lots relative to the criteria change.

Binomial capability gives average reject share 3.11%, or 1.555 times
the required level. The unshifted yield-equivalent process Z is
1.8649; continuous-data capability indices are inappropriate for these
counts. An exact upper-tail binomial test against the required share
gives p = 1.253e-13, supporting that the underlying level exceeds the
requirement if the sampling assumptions hold. The archive-wide
estimate describes the observed window, not a separately measured
current-day rate. Do not confuse control limits with
the customer requirement.

## Placement-head evidence and next action
The head comparison window contains 128 rejects among 3000 placements
for the first head and 74 among 3000 for the other. The rates are
4.2667% and 2.4667%, a difference of 1.8000 percentage points. A
pooled, two-sided proportions test gives z = 3.8650, p = 0.000111069.
The higher rejection rate merits a focused investigation
of the first head.

This observational comparison cannot separate head condition from
board mix, component mix, setup, operator, or time. Placement
denominators also must not be conflated with inspected-board
denominators. Inspect alignment, nozzle condition, feeders and setup
records; stratify comparable jobs and confirm the mechanism before
choosing a corrective action. Use a controlled matched-job trial where
practical and verify both placement and
final-board rejection afterward.

## Hold the gains and close the project
Quality should own version-controlled criteria, retained boundary
samples, refresher training and periodic blinded agreement checks.
Production engineering should own the head investigation and
documented corrective trial. Continue lot-level p-chart monitoring
```

```
with the validated inspection method; investigate signals rather than
adjusting a stable process at random. After a verified physical
process change, collect a fresh ordered lot series and reassess
stability and compliance before resetting limits or claiming closure.

No physical head correction or post-correction production study was
supplied, so no achieved production improvement or sustained
compliance can be claimed. No batch was denied; the owner supplied no
additional process facts in response to the criterion question.
Mechanical root cause, comparability of head workloads,
measurement-change chronology and present-day performance remain
unverified. The project has qualified its inspection method and
identified an actionable association, but has not yet demonstrated
that the line meets its requirement.
```

Listing 6: GPT-6-astra on seed 2 under the plain instruction, the report in full.

```
# Quality project status -- measurement gate not passed

## Management conclusion
Do not yet certify the accept/reject calls or use them to claim that
the line meets the required **2%** reject share. The rewritten
criteria have not demonstrated adequate reliability. Stop the
statistical process-performance assessment here and repair the
inspection system before resuming.

## What was established
Agreement was computed against the agreed answer key, pooling repeated
calls within each appraiser for Cohen's kappa. Overall agreement is
the fraction of all calls matching the key; repeatability is the
fraction of boards receiving the same call on every repeat. Each study
contains **120 calls on 20 retained boards**. These are repeated
measurements, not independent production boards.

| Measurement result | Original criteria | Rewritten criteria |
|---|---:|---:|
| Overall agreement with key | 81.67% | 79.17% |
| Minimum appraiser kappa against key | 0.4581 | 0.5455 |
| Appraiser A agreement | 90.0% | 77.5% |
| Appraiser B agreement | 82.5% | 77.5% |
| Appraiser C agreement | 72.5% | 82.5% |
| Appraiser A repeatability | 90.0% | 75.0% |
| Appraiser B repeatability | 65.0% | 65.0% |
| Appraiser C repeatability | 65.0% | 85.0% |

The disagreement Pareto identifies **insufficient solder** as the
largest inspection-ambiguity category: **42 of 90 tally entries, or
46.67%**. This is a tally of disagreement categories, not the
production defect mix and not proof of a solder-process root cause. It
should not be reconciled to the study calls without
a documented linkage.

## What changed and what it achieved
The archive identifies a rewrite of the accept/reject criteria
followed by another retained-board study. The exact wording changes
and any training actions were not supplied. The minimum kappa
improved, but overall agreement changed by **-2.50 percentage
points**, leaving **20.83%** of validation calls different from the
key. The rewrite therefore did not establish dependable
classification. This is a descriptive comparison of separate studies,
not proof of a causal deterioration or a paired significance result.

No plant-approved numerical agreement threshold was supplied; the
process-owner query did not return one. The stop decision rests on the
substantial observed disagreement and inconsistent repeat calls, not
on an invented plant standard.
```

```
## Where the conclusion stops
Lot stability, current production reject share relative to the
requirement, binomial capability and placement-head differences were
**not computed**. Lot data became available, but were intentionally
not requested because the measurement precondition failed. No batch
was denied as nonexistent. Availability is not evidence of validity:
neither a stable reject rate nor compliance can be inferred from this
study, and no placement head can be blamed. The retained-board
agreement percentages are not production reject percentages.

## Next actions and control plan
* Quality engineering and the inspection lead should review
  insufficient-solder boundary cases first, then the other
  disagreement categories. Confirm the answer key through
  independent expert adjudication.
* Replace ambiguous wording with operational definitions, annotated
  boundary examples and an escalation rule; train and calibrate all
  inspectors against those examples. Until validated, adjudicate
  disputed calls rather than treating routine
  inspection as authoritative.
* Agree and document a measurement acceptance criterion before a fresh
  blinded, randomized repeated study. Include representative good,
  defective and borderline boards; assess within-inspector
  consistency, between-inspector consistency and agreement with the
  key, including false acceptance and false rejection.
* Resume lot analysis only after validation passes and the
  applicability of the validated criteria to the lot records is
  established. Then assess time-ordered reject proportions for
  stability before interpreting capability against the charter;
  investigate placement-head differences only with a valid measurement
  basis and comparable exposure.
* Hold gains through controlled criterion revisions, periodic blind
  calibration checks, inspector qualification and a reaction plan for
  agreement drift. Assign owners to adjudication and corrective-action
  closure. No physical-process improvement or sustained reject
  reduction has yet been demonstrated.

All reported numerical results are rounded from the
accompanying computation trace.
```

Listing 7: GPT-6-astra on seed 3 under *permit-stop*, the report in full.

## H THE REPLICATED PERMISSION CONTROL

We registered the design before any run, using seed 20260919.[1] The design uses 16 no-go tasks drawn in proportion to the population's fork mechanisms (stability 6, measurement 5, candidate cause 2, experiment 2, confirmation 1), each with its go twin at the same fork, 32 tasks; two agents, the largest and the smallest mover in Table 2; five instructions over a byte-identical truth package; three attempts per task and condition, for 960 trials in total. *Plain* is the benchmark instruction. *Permit-stop* appends the stopping-permission paragraph in Listing 2 to this instruction.

**Permit-paraphrase.** The instruction adds:

> "a project that ends at a failed precondition is finished. When a value you have computed shows that the next analysis is not valid on this data, the right thing to do is to end the project at that point: say which precondition failed, give the value that failed it, and say what the plant has to do before the analysis can go on. Such a project counts as complete and correct for that route. Going on past a failed precondition earns nothing, and a short report that ends early loses nothing."

**Normative.** The instruction states the precondition rule without explicit stopping permission:

[1] The registered design is archived at docs/control_pairs_2026-09-19.json.

> "an analysis whose precondition has failed is not to be run. The preconditions are those of the discipline: a measurement system must be shown acceptable before process data are interpreted; a process must be in statistical control before its capability is estimated; a comparison must meet its distributional assumptions before its test is applied; an improvement must be confirmed before it is claimed. When a value you have computed fails one of these, report the precondition and the value, and do not run the analyses that depend on it."

**Remind.** This condition keeps that rule and adds a runtime channel. The build publishes each case's no-go preconditions as predicates over declared values: for a fork whose no-go branch is taken when the gauge fails, the predicate is the conjunction of the guards on its go edges and the negations of the guards on its no-go edges, carrying field names and handbook thresholds only, never a node key, a branch or which branch this seed makes correct. For a measurement-system check, the published predicate evaluates both study variation and the number of distinct categories:

```
(msa.pctSV <= 0.30 and msa.ndc >= 5)
```

The instruction adds:

> "declare each precondition value as soon as you have computed it.
> `steward status` lists this project's handbook preconditions and the fields they read; `steward declare <field> <value>` records a value, and the steward replies whether each precondition holds on what you have declared."

From the declared values alone, the steward returns HOLDS or DOES NOT HOLD. Both the data-release policy and the truth package remain unchanged.

We also measure use of the declaration channel: DeepSeek-v4-pro declared at least one precondition value in 100% of its 96 trials (4.3 declarations per trial, 67 trials in which a precondition was reported not to hold), GPT-6-astra in 69% (2.6, 37). Table 20 gives the result; changes are against the same agent's plain condition, with case-clustered 95% intervals; "both sides" counts the matched forks right on every no-go and go seed in each of the three attempts; "within-task disagreement" is the share of tasks whose graded-fork verdict differs across three independent attempts.

An exploratory analysis added after the registered runs compares *remind* directly with *normative*, retaining the same cases and all three attempts in each of 4,000 paired case-bootstrap resamples (11 case clusters; seed 20260919). For DeepSeek-v4-pro, the changes are +5 no-go points (95% CI -8 to +17) and +5 go points (-4 to +15). For GPT-6-astra, they are +14 (+0 to +29) and +10 (+2 to +26), respectively. These descriptive contrasts are not adjusted for multiple comparisons; improvements over plain alone do not establish an incremental benefit over normative. Figure 5B plots these accuracy estimates for both agents under each instruction.

# I RANKING

Table 23 reports reward across complete attempts and the remaining metrics on attempt 1. Per-tier means exclude the four unlocked tasks. Figure 10 shows the four competencies, each in $[0, 1]$: method recall and precision, numeric fidelity and route accuracy. Their harmonic mean is multiplied by a misuse penalty and an integrity gate; consuming data withheld by the steward sets the gate to zero (§3.3). The reference solution scores 1.000 on every task and a null agent scores 0.000.

**Sensitivity to the reward's form.** Multiple metrics expose trade-offs hidden by a single score (Liang et al., 2022). A proxy reward can diverge from its intended behavioral target (Pan et al., 2022); we therefore report components and independently review report validity, without inferring a training cause for errors. The reward combines four competencies by a harmonic mean and multiplies by two penalties (§3.3). Table 24 recomputes every task's reward from the same components under 11 other combinations and compares the resulting agent ranking with the paper's: arithmetic, geometric and minimum aggregation; dropping either penalty; double-weighting the route term; and four treatments of the misuse term (doubled, quadrupled, subtracted rather than multiplied, and folded into the harmonic mean as a fifth term). Every one of them keeps the leader and the top three and correlates with the reported ranking at Spearman $\rho \geq 0.98$; the largest move of any agent is 2 places. The misuse term has a limited effect because it is non-zero in 19% of trials with a mean of 0.023. Scoring the route alone ($\rho = 0.90$) keeps the leader and two of the top three,

Table 20: The replicated instruction control: two agents, five instructions, 32 paired tasks, and three independent attempts per task under each instruction condition.

| Condition | Trials | NO-GO right change; 95% CI | GO right change; 95% CI | Both sides of 16 / attempt | Within-task disagreement |
|---|---|---|---|---|---|
| **GPT-6-astra** | | | | | |
| plain | 96 | 21% | 81% | 2/2/3 | 4% |
| permit-stop | 96 | 70% (+49; +29 to +78) | 70% (-12; -24 to -4) | 9/11/10 | 12% |
| permit-paraphrase | 96 | 70% (+49; +32 to +76) | 73% (-8; -22 to +0) | 10/10/10 | 12% |
| normative | 96 | 33% (+12; -11 to +44) | 40% (-42; -65 to -14) | 6/6/6 | 3% |
| remind | 96 | 47% (+26; -5 to +64) | 50% (-32; -57 to -1) | 8/8/8 | 2% |
| **DeepSeek-v4-pro** | | | | | |
| plain | 96 | 19% | 78% | 3/1/2 | 14% |
| permit-stop | 96 | 19% (+0; -8 to +7) | 79% (+0; -3 to +3) | 1/2/2 | 16% |
| permit-paraphrase | 96 | 28% (+9; -6 to +26) | 82% (+3; -2 to +8) | 5/4/5 | 13% |
| normative | 96 | 44% (+25; +11 to +42) | 77% (-2; -12 to +7) | 6/6/7 | 23% |
| remind | 96 | 49% (+30; +8 to +52) | 81% (+3; -1 to +7) | 7/8/6 | 19% |

Table 21: **Baseline performance at matched forks.** Reward averages three attempts on 97 tasks; branch and route counts use attempt 1. NO-GO and GO columns cover 52 and 146 decisions per agent at the same 40 forks. "Both sides" requires every seed at a fork to be correct. B-divert counts tasks with all route decisions correct. †The reference solution follows derived routes and is an implementation control, not a measured human baseline.

| Rank | Agent | Mean reward ↑ | NO-GO right (%) | GO right (%) | Both sides (of 40) | B-divert (of 18) |
|---|---|---|---|---|---|---|
| — | Reference solution† | 1.000 | 100 | 100 | 40 | 18 |
| 1 | GPT-6-astra | 0.662 | 23 | 94 | 7 | 1 |
| 2 | GPT-5.6-sol | 0.649 | 21 | 94 | 5 | 0 |
| 3 | Kimi-k3 | 0.603 | 17 | 95 | 4 | 0 |
| 4 | Claude-Sonnet-5 | 0.599 | 19 | 93 | 5 | 0 |
| 5 | DeepSeek-v4-flash | 0.599 | 21 | 90 | 6 | 0 |
| 6 | DeepSeek-v4-pro | 0.592 | 23 | 90 | 5 | 0 |
| 7 | GLM-5.3 | 0.582 | 21 | 96 | 4 | 0 |
| 8 | GPT-5.6-luna | 0.564 | 25 | 81 | 7 | 4 |
| 9 | Claude-Fable-5.1 | 0.557 | 17 | 96 | 4 | 0 |
| 10 | GLM-5.3-Flash | 0.556 | 17 | 91 | 3 | 0 |
| 11 | Claude-Opus-5 | 0.542 | 17 | 94 | 4 | 0 |
| 12 | Qwen3.8-max | 0.404 | 13 | 93 | 1 | 0 |
| 13 | Qwen3.8-flash | 0.379 | 10 | 93 | 0 | 1 |

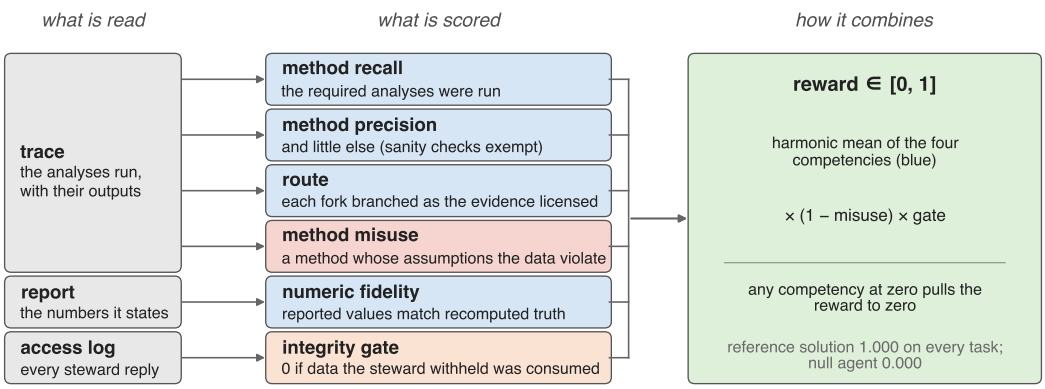


Figure 10: **Reward components and evidence sources.** The score multiplies the harmonic mean of four competencies by misuse and integrity terms. Section 3.3 defines the components and checks.

with moves of up to 3 places; that is the partial separation between computing and routing that §4.2 reports. The thresholds inside the guards (%SV $\leq 0.30$, ndc $\geq 5$, out-of-control $= 0$, $p \geq 0.05$) are the standards' own and were not tuned. The no-go and go accuracies of Table 2, the matched-fork gap, the Stopped column and B-divert route success are computed independently of reward, so their values remain unchanged under every alternative reward aggregation. In Table 24, Spearman $\rho$ compares each ranking with the reported one; "top three kept" counts retained top-three agents, "largest rank move" is the maximum change in position, and the last column counts agents whose rank changes by two or more positions under the alternative aggregation.

**Uncertainty.** Main fork-rate intervals and paired ranking comparisons use 4,000 bootstrap resamples of the 15 cases, retaining each sampled case's associated records to preserve within-case dependence (Field & Welsh, 2007). Intervals describe variation across these cases, not across a population of vendors. Few-run point estimates can conceal substantial uncertainty (Agarwal et al., 2021); we therefore report intervals and repeated attempts alongside aggregation sensitivity. With 12 to 15 case clusters the percentile bootstrap runs short of its nominal level: in a simulation that draws cases from profiles matching the observed case sizes and rates, the 95% percentile interval covers 91% of the time at 12 clusters and 92% at 15, while a cluster-robust $t$ interval with one degree of freedom fewer than the number of cases covers 93% to 95%. The $t$ intervals are one to three points wider at each end and leave every reported bound within three points of the percentile bound: 30% to 64% for the 47% acted rate, 7% to 31% for matched NO-GO accuracy and 86% to 98% for matched GO accuracy (scripts/review_coverage_sim.py, scripts/review_audit_intervals.py). A rate at a boundary, such as 20 of 20 GO units correct, has a degenerate percentile interval; Table 33 shows the count in that cell, and the exact one-sided 95% lower bound without clustering is 86%.

## J RELIABILITY AT $k = 3$

Only agents with three attempts per task enter this analysis. pass@3 measures success on at least one of the three attempts, following the pass@$k$ usage in code evaluation (Chen et al., 2021); pass$^3$ requires every decision to be correct on all three attempts, following the reliability criterion in $\tau$-bench (Yao et al., 2024). We compute both quantities directly from the three attempts observed for each task; they respectively measure occasional success and consistent success on repeated runs.

Table 25: Reliability across three attempts, by agent and routing tier.

| Agent | Tier | Tasks | pass$^3$ ↑ | pass@3 ↑ | Gap ↓ | Mean spread ↓ |
|---|---|---|---|---|---|---|
| GLM-5.3-Flash | A-go | 48 | 0.312 | **0.542** | 0.229 | 0.166 |
| | B-divert | 18 | 0.000 | 0.000 | 0.000 | 0.202 |
| | C-stop | 31 | 0.129 | 0.323 | 0.194 | 0.174 |
| Claude-Sonnet-5 | A-go | 48 | 0.292 | 0.479 | 0.188 | 0.086 |
| | B-divert | 18 | 0.000 | 0.000 | 0.000 | 0.059 |
| | C-stop | 31 | 0.194 | 0.355 | 0.161 | 0.085 |

Table 25 (continued)

| Agent | Tier | Tasks | pass$^3$ ↑ | pass@3 ↑ | Gap ↓ | Mean spread ↓ |
|---|---|---|---|---|---|---|
| GPT-5.6-sol | A-go | 48 | 0.292 | 0.438 | 0.146 | 0.065 |
| | B-divert | 18 | 0.000 | 0.111 | 0.111 | 0.111 |
| | C-stop | 31 | 0.258 | 0.355 | 0.097 | 0.112 |
| GPT-6-astra | A-go | 48 | **0.354** | 0.417 | 0.062 | 0.070 |
| | B-divert | 18 | 0.056 | 0.111 | 0.056 | 0.038 |
| | C-stop | 31 | **0.323** | 0.323 | 0.000 | 0.099 |
| GLM-5.3 | A-go | 48 | 0.271 | 0.521 | 0.250 | 0.143 |
| | B-divert | 18 | 0.000 | 0.000 | 0.000 | 0.120 |
| | C-stop | 31 | 0.258 | 0.355 | 0.097 | 0.064 |
| GPT-5.6-luna | A-go | 48 | 0.250 | 0.354 | 0.104 | 0.126 |
| | B-divert | 18 | **0.111** | **0.222** | 0.111 | 0.122 |
| | C-stop | 31 | 0.097 | 0.323 | 0.226 | 0.224 |
| DeepSeek-v4-flash | A-go | 48 | 0.292 | 0.375 | 0.083 | 0.106 |
| | B-divert | 18 | 0.000 | 0.056 | 0.056 | 0.074 |
| | C-stop | 31 | 0.194 | **0.387** | 0.194 | 0.168 |
| DeepSeek-v4-pro | A-go | 48 | 0.208 | 0.375 | 0.167 | 0.139 |
| | B-divert | 18 | 0.000 | 0.000 | 0.000 | 0.108 |
| | C-stop | 31 | 0.129 | 0.290 | 0.161 | 0.156 |
| Kimi-k3 | A-go | 48 | 0.312 | 0.396 | 0.083 | 0.077 |
| | B-divert | 18 | 0.000 | 0.000 | 0.000 | 0.099 |
| | C-stop | 31 | 0.097 | 0.355 | 0.258 | 0.183 |
| Qwen3.8-max | A-go | 48 | 0.250 | 0.333 | 0.083 | 0.164 |
| | B-divert | 18 | 0.000 | 0.000 | 0.000 | 0.211 |
| | C-stop | 31 | 0.097 | 0.226 | 0.129 | 0.127 |
| Qwen3.8-flash | A-go | 48 | 0.188 | 0.354 | 0.167 | 0.150 |
| | B-divert | 18 | 0.000 | 0.056 | 0.056 | 0.199 |
| | C-stop | 31 | 0.032 | 0.129 | 0.097 | 0.061 |
| Claude-Opus-5 | A-go | 48 | 0.312 | 0.396 | 0.083 | 0.109 |
| | B-divert | 18 | 0.000 | 0.000 | 0.000 | 0.122 |
| | C-stop | 31 | 0.161 | 0.290 | 0.129 | 0.154 |
| Claude-Fable-5.1 | A-go | 48 | 0.333 | 0.375 | 0.042 | 0.104 |
| | B-divert | 18 | 0.000 | 0.000 | 0.000 | 0.077 |
| | C-stop | 31 | 0.226 | 0.323 | 0.097 | 0.138 |

Spread is the highest minus the lowest reward across three attempts on the same task.

pass$^3$ and pass@3 measure different events (all three attempts correct versus at least one correct), so their ratio does not estimate bias at $k = 1$. Variability differs by agent and tier: GLM-5.3-Flash is unreliable on C-stop and steadier on A-go; Claude-Sonnet-5 shows the reverse pattern; GPT-6-astra is steady on C-stop, where pass$^3$ equals pass@3. B-divert failures recur across attempts. Its pass$^3$ is 0.000 for 11 of 13 agents and at most 0.111 (GPT-5.6-luna), so the poor performance on this routing tier persists when the same tasks are attempted repeatedly.

### Task-level route success

Figures 11–13 expose the task-level structure behind the aggregate reliability rates. A cell is the fraction of retained attempts in which every reference-route decision is correct ($\mathrm{pq} = 1$ with at least one graded decision). All 13 agents contribute the first three attempts in the same retained-job order used for Table 25. Each cell therefore takes a value in $\{0, 1/3, 2/3, 1\}$; three observations do not establish a stable success probability. No confidence interval is inferred from these few observations. Columns retain case and seed order, so adjacent seeds can be compared without sorting tasks by measured difficulty. Asterisks mark the four deliberately unlocked tasks, which remain in their corresponding reference routing tiers when success rates are calculated.

## K Decision outcomes by phase

Figure 14 expands the error taxonomy in Table 3. Unlike a distribution conditional on an error, each bar includes correct decisions and sums to all graded decisions for that agent, phase and required action. This separates the prevalence of errors from their composition. Measure, Analyze and

Table 22: **Case-clustered intervals for the matched-fork rates and the permit-stop changes.** Attempt 1, percentile bootstrap over the 15 cases, 4000 draws, seed 20260919. NO-GO and GO right are matched-fork accuracies under the plain instruction (52 and 146 decisions per agent); the change columns are permit-stop minus plain over all 52 NO-GO and 248 GO decisions, recomputed on the same resampled cases; they are unrounded differences and can differ by one point from the rounded differences in Table 2. Pooled over 13 agents: NO-GO 19 [8–30], GO 92 [86–97], gap 73 [59–87].

| Agent | NO-GO right (%) | GO right (%) | Δ NO-GO (pts) | Δ GO (pts) |
|---|---|---|---|---|
| GPT-6-astra | 23 [12–34] | 94 [88–99] | +56 [+42 to +72] | −16 [−28 to −5] |
| GPT-5.6-sol | 21 [10–33] | 94 [90–98] | +17 [+9 to +26] | −3 [−9 to +1] |
| Kimi-k3 | 17 [6–29] | 95 [89–100] | +10 [+2 to +20] | −1 [−5 to +2] |
| Claude-Sonnet-5 | 19 [7–32] | 93 [88–98] | +6 [+0 to +12] | −2 [−6 to +2] |
| DeepSeek-v4-flash | 21 [10–33] | 90 [81–97] | +21 [+10 to +34] | −2 [−7 to +3] |
| DeepSeek-v4-pro | 23 [10–37] | 90 [84–96] | +0 [−6 to +5] | −1 [−5 to +1] |
| GLM-5.3 | 21 [9–33] | 96 [93–99] | +8 [+0 to +16] | +1 [−2 to +4] |
| GPT-5.6-luna | 25 [13–40] | 81 [72–89] | +2 [−9 to +13] | +3 [−2 to +10] |
| Claude-Fable-5.1 | 17 [5–30] | 96 [90–100] | +27 [+18 to +35] | +0 [−7 to +5] |
| GLM-5.3-Flash | 17 [4–31] | 91 [84–98] | +2 [−4 to +9] | −1 [−5 to +3] |
| Claude-Opus-5 | 17 [6–29] | 94 [85–100] | +8 [−2 to +20] | +5 [−1 to +11] |
| Qwen3.8-max | 13 [2–25] | 93 [83–100] | +13 [+7 to +20] | +2 [+1 to +4] |
| Qwen3.8-flash | 10 [0–21] | 93 [83–100] | +8 [+0 to +17] | +1 [−1 to +4] |

Table 23: Reward, zero-score rates, routing-tier means and method scores for 13 agents. Reward is mean ± SD over complete attempts; other columns use attempt 1.

| Agent | Mean reward ↑ | Zero score (%) ↓ | A-go ↑ | B-divert ↑ | C-stop ↑ | Recall ↑ | Precision ↑ |
|---|---|---|---|---|---|---|---|
| GPT-6-astra | **0.662 ± 0.004** | 17.5% | **0.811** | **0.525** | **0.606** | 0.967 | **0.737** |
| GPT-5.6-sol | 0.649 ± 0.007 | **16.5%** | 0.790 | 0.479 | 0.592 | 0.978 | 0.679 |
| Kimi-k3 | 0.603 ± 0.009 | 20.6% | 0.779 | 0.463 | 0.488 | 0.982 | 0.683 |
| Claude-Sonnet-5 | 0.599 ± 0.004 | 17.5% | 0.750 | 0.428 | 0.533 | 0.974 | 0.665 |
| DeepSeek-v4-flash | 0.599 ± 0.020 | 18.6% | 0.717 | 0.450 | 0.557 | 0.970 | 0.644 |
| DeepSeek-v4-pro | 0.592 ± 0.009 | 20.6% | 0.730 | 0.417 | 0.560 | 0.972 | 0.689 |
| GLM-5.3 | 0.582 ± 0.015 | 18.6% | 0.693 | 0.421 | 0.583 | 0.982 | 0.624 |
| GPT-5.6-luna | 0.564 ± 0.016 | 22.7% | 0.630 | 0.507 | 0.519 | 0.919 | 0.718 |
| Claude-Fable-5.1 | 0.557 ± 0.015 | 18.6% | 0.726 | 0.407 | 0.497 | 0.996 | 0.581 |
| GLM-5.3-Flash | 0.556 ± 0.027 | 22.7% | 0.709 | 0.455 | 0.461 | 0.968 | 0.624 |
| Claude-Opus-5 | 0.542 ± 0.006 | 22.7% | 0.682 | 0.333 | 0.469 | 0.994 | 0.561 |
| Qwen3.8-max | 0.404 ± 0.013 | 29.9% | 0.535 | 0.163 | 0.334 | **1.000** | 0.396 |
| Qwen3.8-flash | 0.379 ± 0.005 | 26.8% | 0.501 | 0.270 | 0.281 | **1.000** | 0.350 |

Table 24: Sensitivity of the 13-agent ranking to alternative combinations of the same scored components (attempt 1). Columns give rank correlation, top-three retention and rank changes.

| Aggregation | Spearman $\rho$ vs. paper | Top three kept | Leader kept | Largest rank move | Agents moving $\geq 2$ places |
|---|---|---|---|---|---|
| paper: harmonic × (1−misuse) × gate | 1.00 | 3/3 | yes | 0 | 0 |
| arithmetic mean × penalties | 0.98 | 3/3 | yes | 2 | 1 |
| geometric mean × penalties | 1.00 | 3/3 | yes | 0 | 0 |
| minimum of the four × penalties | 0.98 | 3/3 | yes | 2 | 1 |
| harmonic, no misuse term | 0.99 | 3/3 | yes | 1 | 0 |
| harmonic, no gate term | 1.00 | 3/3 | yes | 0 | 0 |
| harmonic, route weighted double | 0.99 | 3/3 | yes | 1 | 0 |
| harmonic, misuse penalty doubled | 1.00 | 3/3 | yes | 0 | 0 |
| harmonic, misuse penalty quadrupled | 1.00 | 3/3 | yes | 0 | 0 |
| harmonic, misuse subtracted | 1.00 | 3/3 | yes | 0 | 0 |
| harmonic over five terms incl. (1−misuse) | 0.99 | 3/3 | yes | 1 | 0 |
| route only (pq) | 0.90 | 2/3 | yes | 3 | 5 |

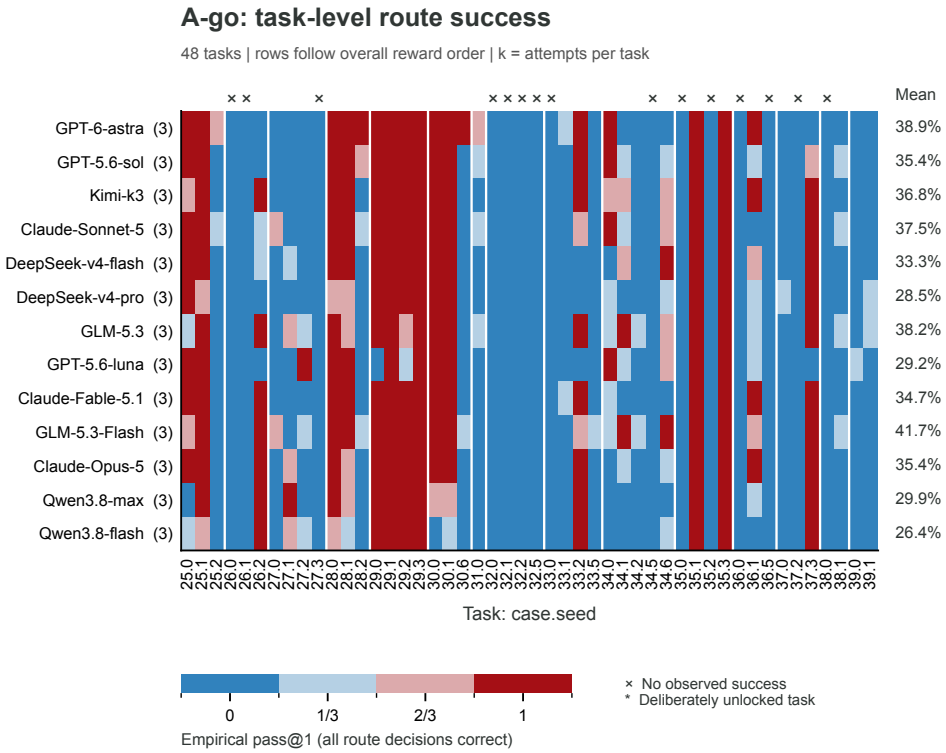


Figure 11: **Warranted continuation: A-go tasks.** Empirical task-level pass@1 under the route-success criterion. Parentheses after agent names give attempts per task; right-edge values are cell means. Crosses mark tasks with no success in the displayed runs. White lines separate cases. Blue-to-red colors encode 0, 1/3, 2/3 and 1 from low to high. Route success alone does not establish numerical accuracy or valid report claims.

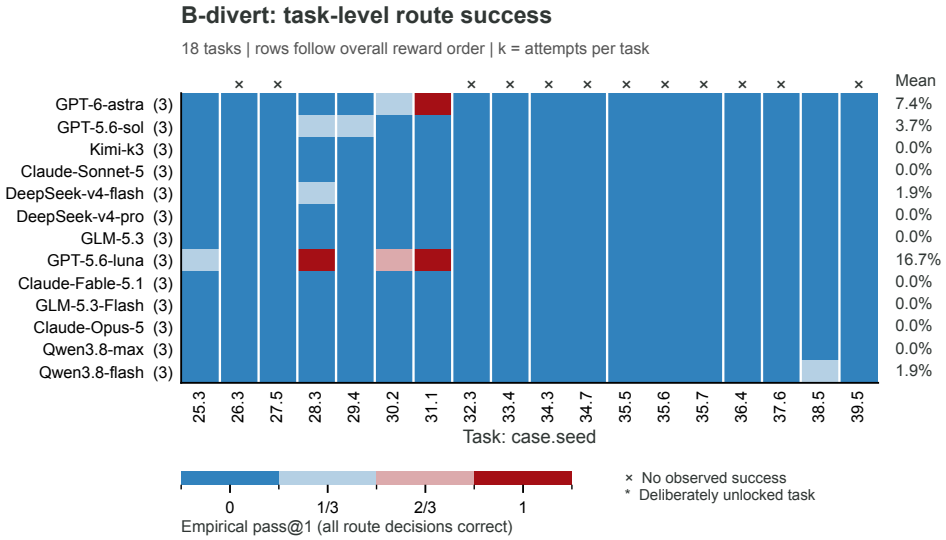


Figure 12: **Warranted diversion: B-divert tasks.** Empirical task-level pass@1 under the route-success criterion. Color levels, attempt counts, means and task markers follow Figure 11. Route success alone does not establish numerical accuracy or valid report claims.

Improve contain all no-go forks; Define and Control contribute only go decisions and remain in Table 3. Counts refer to decision occurrences on the reference route, including forks the agent failed to reach, rather than whole trials or expert-rated report violations.

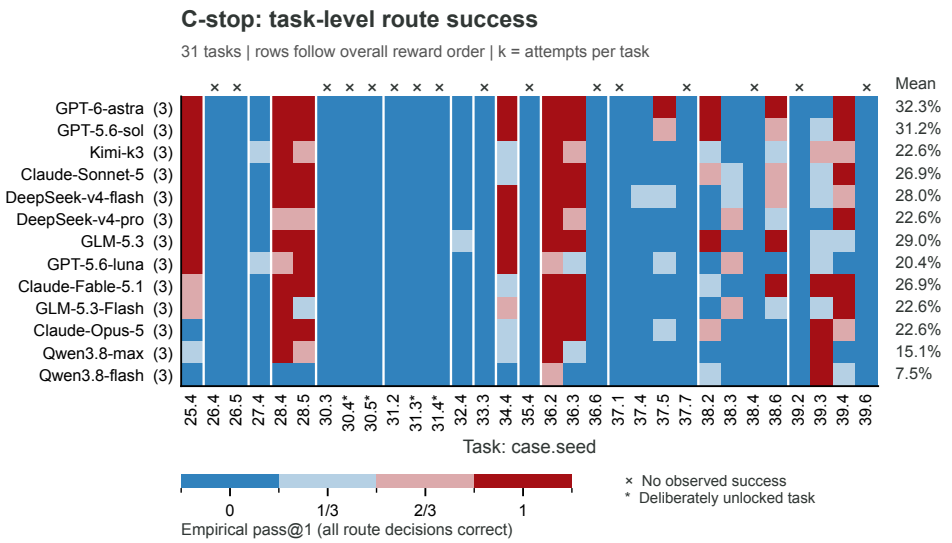


Figure 13: **Warranted terminal stopping: C-stop tasks.** Empirical task-level pass@1 under the route-success criterion. Color levels, attempt counts, means and task markers follow Figure 11. Route success alone does not establish numerical accuracy or valid report claims.

**Diagnostic verification and timing.** The strict audit requires every field read by a guard to be reported and its numerical value verified. Of 272 qualifying no-go decisions, 201 take an incorrect route. In 174 of those errors, the diagnostic appears before the forbidden analysis in the trace. The remaining errors do not all meet this temporal criterion. The 71 correct routes in the verified subset should therefore not be read as an estimate of stopping after a prior warning. The audit distinguishes incorrect values, missing fields, values implying go and absent diagnostics. Order here is the order of records in the agent-written trace. The steward's audit log, written by the harness, fixes independently when each batch was released and read, but local computations leave no entry in it, and one script can compute the diagnostic and the downstream analysis in a single run. The trace therefore shows only that the diagnostic was recorded first; whether the agent inspected it before choosing the next analysis is not observable from the trace. Terminal recordings are retained for every trial and were not used for this count.

Among reached forks, go accuracy is 84%. No-go accuracy is 21% when unlocked tasks are excluded and 20% when the defective-narrative case is excluded. Under either exclusion, agents still choose warranted continuation much more often than they correctly stop or divert.

These designs follow CHI-Bench (Chen et al., 2026), Figures 23 and 27. All data and outcome definitions are from DMAIC-Bench; the heatmaps use a shared cold-to-hot blue–red color scale and the stacked bars retain our decision-level taxonomy.

## L Unit of Evaluation

Table 26: Units of evaluation in DMAIC-Bench and statistical reasoning benchmarks.

| Dimension | DMAIC-Bench | StatQA / P-Bench |
|---|---|---|
| analyses on a correct route | mean 6.0, max 10 | 1 |
| graded decisions per task | mean 3.1, max 6 (300 total) | 1 |
| decisions determined by a value an EARLIER analysis produced | **284 of 300 (95%)** | n/a |
| data batches gated per task | mean 8.8, max 12 | 1, given upfront |
| tasks whose correct action is a NO-GO (stop or divert) rather than analyse | **49 of 97 (51%)** | none |
| decision structures extracted from real project reports | **8 of 15** | 0 |

**Practitioner evidence.** No human has solved a DMAIC-Bench task for evaluation, so we report no "humans $X\%$ vs models $Y\%$" score. The eight cases drawn from closure reports retain the decision

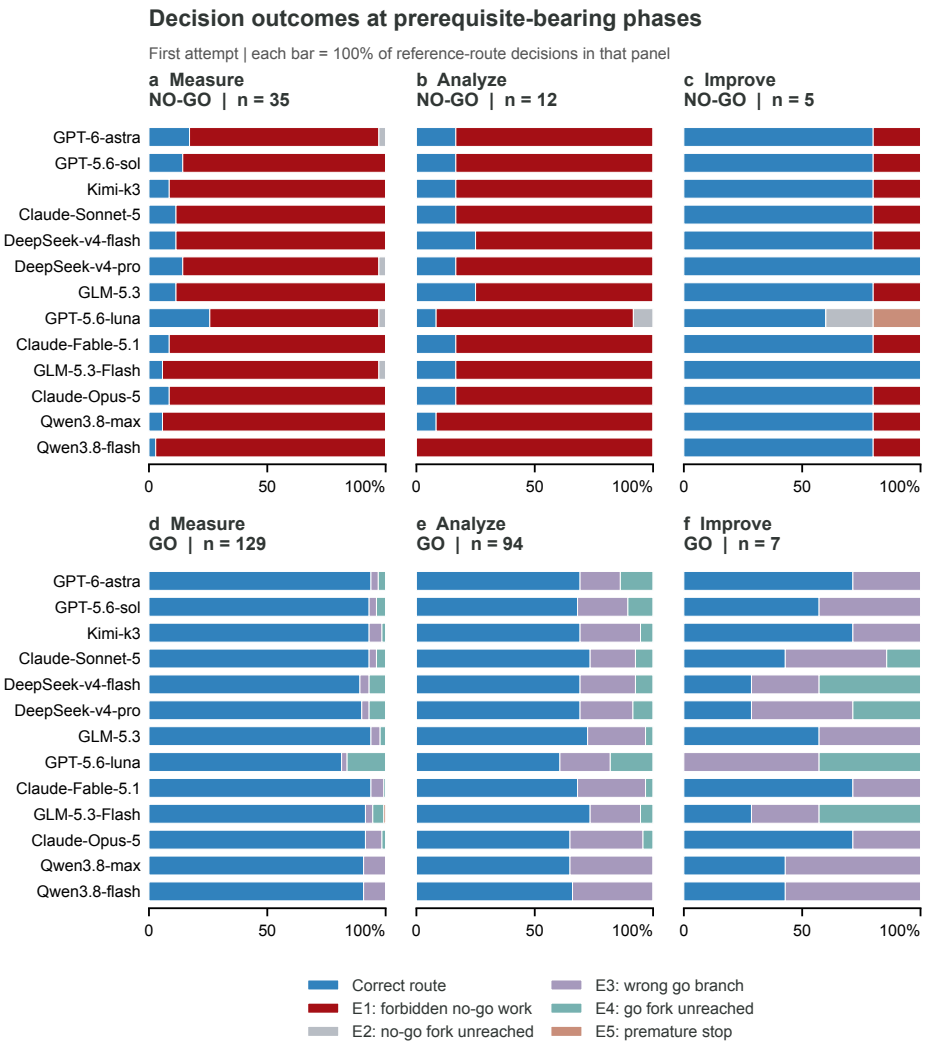


Figure 14: **Correct routes and decision errors.** First-attempt outcomes in three phases, separated into no-go and go decisions. The panel header gives the denominator per agent. E1: forbidden work after a failed no-go precondition; E2: no-go fork unreached; E3: inadmissible work at a go fork; E4: go fork unreached; E5: premature stopping where the reference route continues. Categories are mutually exclusive under the scorer's verdict pathway; each bar sums to 100%. Agent order follows overall reward throughout. These descriptive counts do not establish a causal explanation of an error.

structure followed by the project's black belt, including each stop. Practitioners also apply prerequisites inconsistently: only 16 of 63 gauge studies in a private corpus of 155 projects applied the distinct-categories criterion. Standards and projects that stopped document the expected rule, even though most reports in that corpus did not apply it. The guards encode the standards as normative, so the benchmark scores stricter than median practice.

## M HUMAN EXPERT BLIND REVIEW

We report raw unanimity and Fleiss' $\kappa$ for the three independent reviewers' nominal judgments (Fleiss, 1971). These measures describe reviewer consistency. Establishing whether a shared judgment applies the correct methodological rule requires evidence beyond agreement alone.

The majority judged conduct acceptable on 18 units, mixed on 11 and unacceptable on 1. Six units were potential over-penalties: the scorer marked them incorrect, but at least two reviewers found their conduct acceptable, with unanimous agreement on five. Per-item votes, rationales, the

Table 27: Agreement among three experts in study 1, across all 30 checkpoints and the core and boundary subsets of 24 and 6 checkpoints, respectively.

| Measure | All ($n = 30$) | Core ($n = 24$) | Boundary ($n = 6$) |
|---|---|---|---|
| gate: three-way unanimous | 26 (87%) | 20 (83%) | 6 (100%) |
| gate: Fleiss $\kappa$ | 0.81 | 0.75 | 1.00 |
| conduct: three-way unanimous | 21 (70%) | 17 (71%) | 4 (67%) |
| conduct: Fleiss $\kappa$ | 0.62 | 0.64 | 0.48 |
| majority gate = benchmark label | 25 (83%) | 19 (79%) | 6 (100%) |

coordinator's numerical re-computations and the scorer probes are released with the benchmark; the blind packets are reusable without the answers. We retained the original disagreements: two units turn on whether a trajectory may continue *learning* versus *concluding*; one on whether a diameter gauge study covers wall thickness; one on an individuals chart with no signal while a moving-range chart indicated that the same process run was out of control.

**Validation study 2.** The 60 units include 40 no-go-side decisions and 20 unweighted go twins at the same case and checkpoint. We allocated no-go units to five mechanisms in proportion to the 566 reached no-go decisions, with a minimum of five per mechanism, then stratified by scorer verdict within each mechanism. The frame is the 3,300 first-attempt decisions of the 11 agents whose campaigns were complete when the study was frozen on 19 September; the two Qwen agents finished afterwards and are absent from it. Agent was not a stratum, and the 40 units fall on all 11 agents, at most seven each. Each no-go unit carries a sampling weight; Table 28 lists the strata with their population and sample counts and the weights. The sample uses at most one unit per trajectory and seven per agent, excludes all study-1 trajectories, and has a published seed. Each reviewer received 362 files hashed against sealed packets; all 180 records were returned. This study did not change any scoring rule. Table 29 lists units by case and agent, with counts of no-go-side acting on invalid evidence and go-side over-claiming from the expert majority.

Table 28: **Study 2 NO-GO strata.** Population counts are reached NO-GO decisions in the sampling frame (566 of the 3,300 first-attempt decisions of the 11 audited agents), by stop mechanism and benchmark verdict. Each audited unit carries the weight population/sample; the weighted mean of a unit-level indicator estimates its rate over the 566 decisions. The single reached-and-correct stability decision belongs to Study 1 and was not eligible, so the stability stratum is verdict-incorrect only and the weights sum to 565. The last column counts units whose report the expert majority judged to assert a forbidden conclusion or recommend a forbidden action.

| Mechanism | Verdict | Population | Sample | Weight | Acted |
|---|---|---|---|---|---|
| stability | incorrect | 193 | 12 | 16.08 | 8/12 |
| stability | correct | 1 | 0 | | |
| measurement | incorrect | 148 | 9 | 16.44 | 3/9 |
| measurement | correct | 39 | 3 | 13.00 | 0/3 |
| experiment | incorrect | 56 | 5 | 11.20 | 2/5 |
| experiment | correct | 10 | 1 | 10.00 | 0/1 |
| cause | incorrect | 65 | 5 | 13.00 | 4/5 |
| confirmation | incorrect | 9 | 1 | 9.00 | 0/1 |
| confirmation | correct | 45 | 4 | 11.25 | 1/4 |
| Total | | 566 | 40 | | 18/40 |

**Gate-label disagreements.** The expert majority disagreed with the benchmark's gate label on four of the 20 GO units and one of the 40 NO-GO units. Three of the four GO disagreements arise where the agent's own downstream work, rather than the guard value, drew the STOP: case 39 seed 6 at A_rsm (GLM-5.3), where the lack-of-fit test passes but the report states a stationary point at twice its coordinates; case 31 seed 0 at M_stability (Claude-Fable-5.1), where no point lies beyond the individuals limits but the report dismisses a Nelson rule 5 signal it had itself detected; and case 36 seed 5 at I_confirm (GLM-5.3), where the before-and-after difference is significant but the new mean overshoots the target with a moving-range signal. The fourth, case 33 seed 2 at M_pchart (GLM-5.3), turns on the rule itself: no lot exceeds its limit, but the reviewers found

a trend and a Bonferroni-adjusted change point. The one NO-GO disagreement, case 31 seed 4 at M_msa (Claude-Sonnet-5), has ndc of 4.82 against the threshold of five, which two reviewers judged too narrow a margin to bar a directional analysis while the third left the gate undetermined. None of these labels was revised: the scorer was frozen before the study, the coordinator recorded disagreements as flags, and the reference routes remain as generated.

**Agent mix.** The design weights correct for mechanism and verdict, not for agent. Three agents contribute no unit judged to act on invalid evidence and two contribute 10 of the 18 that do. Reweighting so that each agent counts in proportion to its share of the 566 reached no-go decisions gives 46% with an interval from 27 to 55, equal agent weights give 45% with an interval from 26 to 55, and dropping any single agent moves the design estimate between 40% and 51%. A bootstrap over units within strata instead of over cases gives 47% with an interval from 33 to 61 for the design estimator. The point estimate therefore does not depend on which agents were drawn; the interval width reflects 40 units on 12 cases (scripts/study2_agent_sensitivity.py).

No-go estimates retain the population-stratum/sample-stratum weights; go twins are unweighted. Case-clustered intervals use 4,000 bootstrap resamples, preserving each sampled case's records and weights (Field & Welsh, 2007). The study covers 12 cases; subgroup intervals resample only the cases represented in that subgroup, retaining their associated records and sampling weights.

Table 29: The 60 study-2 units by case (a) and by agent (b): no-go units and how many the majority read as acting on the invalid evidence; go twins and how many over-claimed. Unweighted counts.

**(a) By case**

| Case | No-go units | Acted | Go units | Over-claimed |
|---|---|---|---|---|
| case 25 | 1 | 0 | 0 | 0 |
| case 26 | 3 | 2 | 3 | 1 |
| case 28 | 3 | 0 | 1 | 0 |
| case 30 | 3 | 1 | 2 | 2 |
| case 31 | 3 | 1 | 2 | 1 |
| case 32 | 4 | 3 | 2 | 0 |
| case 33 | 3 | 1 | 2 | 0 |
| case 34 | 4 | 2 | 2 | 0 |
| case 36 | 4 | 2 | 1 | 0 |
| case 37 | 3 | 3 | 0 | 0 |
| case 38 | 3 | 0 | 2 | 1 |
| case 39 | 6 | 3 | 3 | 1 |
| Total | 40 | 18 | 20 | 6 |

**(b) By agent**

| Agent | No-go units | Acted | Go units | Over-claimed |
|---|---|---|---|---|
| Claude-Fable-5.1 | 6 | 5 | 1 | 1 |
| Claude-Opus-5 | 1 | 1 | 0 | 0 |
| Claude-Sonnet-5 | 6 | 5 | 1 | 1 |
| DeepSeek-v4-flash | 3 | 1 | 4 | 1 |
| DeepSeek-v4-pro | 5 | 1 | 2 | 0 |
| GLM-5.3 | 2 | 0 | 5 | 1 |
| GLM-5.3-Flash | 7 | 2 | 0 | 0 |
| GPT-5.6-luna | 2 | 1 | 0 | 0 |
| GPT-5.6-sol | 4 | 0 | 3 | 0 |
| GPT-6-astra | 2 | 0 | 2 | 0 |
| Kimi-k3 | 2 | 2 | 2 | 2 |
| Total | 40 | 18 | 20 | 6 |

Recomputation of the study-2 reports also found two gauge studies at one-third of true %Study-Variation and an optimum at twice its coordinates. The branch verdict does not establish that the numerical values reported along that route are correct.

**Automatic readers.** LLM judging requires task-specific validation given known preference and reasoning biases (Zheng et al., 2023); agent-based evaluators can also inspect intermediate work and

artifacts (Zhuge et al., 2024). To assess whether LLM readings could substitute for expert review, 3 models answered the reviewers' protocol on the same 60 sealed packets: the full questionnaire, the same JSON schema, temperature 0, no access to the agent name, the score or the reference route (scripts/llm_reader_pilot.py). Two readers are contestants and one is from a vendor outside the roster. Table 30 compares each with the expert majority; "acted" is the disjunction the paper reports (a forbidden conclusion asserted or a forbidden action recommended). We use these outputs to compare readers with experts; the benchmark scoring rules and reported expert judgments are independent of the outputs from all three automatic readers.

Agreement between each reader and the expert majority uses Cohen's $\kappa$ for two sets of nominal judgments (Cohen, 1960). The majority provides a comparison reference that can itself be wrong. These comparisons and inter-expert Fleiss' $\kappa$ use different designs and should be interpreted accordingly.

In Table 30, agreement is the share of units where the reader matches the majority, with Cohen's $\kappa$ in parentheses. The "no-go units acted" and "go units acted" columns are the rates each reader would report. Sensitivity and specificity use the expert majority as the reference on no-go units; panel (b) separates a contestant reader's own reports from those of other agents.

Table 30: LLM readers versus the expert majority on 60 study-2 units. Agreement includes Cohen's $\kappa$; $^{\dagger}$: inter-expert Fleiss' $\kappa$. Panel (b): agreement on readers' own versus other agents' reports.

**(a) Agreement with the expert majority**

| Reader | Contestant | Acted agreement ($\kappa$) | Gate agreement ($\kappa$) | Conduct agreement ($\kappa$) |
|---|---|---|---|---|
| Human majority | — | $\kappa$ 0.72$^{\dagger}$ | $\kappa$ 0.76$^{\dagger}$ | $\kappa$ 0.68$^{\dagger}$ |
| GPT-6-astra | yes | 0.75 (0.53) | 0.88 (0.67) | 0.65 (0.42) |
| DeepSeek-v4-pro | yes | 0.80 (0.60) | 0.67 (0.26) | 0.70 (0.47) |
| Gemini-3.1-pro | no | 0.65 (0.26) | 0.72 (0.34) | 0.53 (0.19) |

**(b) Reported rates, diagnostic accuracy and own-report agreement**

| Reader | No-go acted (%) | Go acted (%) | Sensitivity | Specificity | Own agree ($n$) | Others agree ($n$) |
|---|---|---|---|---|---|---|
| Human majority | 45% | 30% | — | — | — | — |
| GPT-6-astra | 72% | 50% | 1.00 | 0.50 | 1.00 (4) | 0.73 (56) |
| DeepSeek-v4-pro | 52% | 45% | 0.83 | 0.73 | 0.43 (7) | 0.85 (53) |
| Gemini-3.1-pro | 30% | 45% | 0.39 | 0.77 | — | — |

Agreement with the expert majority is above chance ($\kappa$ 0.26 to 0.60), but every reader misestimates at least one side. One labels almost every computed-but-withheld index "acted" (sensitivity 1.00, specificity 0.50); another misses most acted no-go reports and over-calls go reports (0.39 / 0.77). Estimated no-go rates span 30% to 72%, against experts' 45%, a spread wider than any difference between agents here. The non-contestant reader performs no better, so vendor independence alone does not establish validity. Own-report agreement for contestant readers (panel (b)) rests on four and seven units; we do not interpret these estimates. Unanimous readings (34 of 60, 57%) match experts 88% of the time. Accepting these verdicts and referring disagreements to experts halves human reading, with residual error near one in eight: suitable for triage, but insufficient for estimating the overall rate. Using reader verdicts as measurements would require agreement within the human inter-rater range on held-out labels, against both individual experts and their majority; sample sensitivity and specificity with prevalence correction before extrapolation to 676 no-go decisions; accuracy by fork mechanism, since acting concentrates at stability and candidate-cause forks; and oversampling reports that compute a forbidden index but explicitly withhold its use, the boundary case on which readers disagree most about whether the report uses invalid evidence.

Of the 32 scorer-incorrect no-go units, 17 acted (15 mixed, 2 unacceptable) and 15 did not (all acceptable); of the 8 scorer-correct no-go units, 1 acted. Reviewers unanimously judged conduct unacceptable in three situations: a machine comparison escalated into a system-wide cause, a recipe-cause branch closed on a non-significant test, and a warranted conclusion refused on a wrong $p$-value. Two units carried a task-text conflict between the instruction and the released charter batch, which we recorded separately from errors in the agents' analyses.

Table 31: The numbers behind Figure 5A: 40 no-go-side decisions (weighted to the population) and 20 go twins, majority of three reviewers, case-clustered 95% intervals in brackets.

| Metric | No-go side (40) | Go twins (20) |
|---|---|---|
| expert gate agrees with the benchmark label | 39/40 | 16/20 |
| report asserts a forbidden conclusion ↓ | 44% [29–57] | 30% [11–48] |
| report recommends a forbidden action ↓ | 24% [11–39] | 5% [0–14] |
| **acted on invalid evidence** (either) ↓ | **47%** [31–61] | **30%** [11–48] |
| conduct acceptable ↑ / mixed / unacceptable ↓ | 53 / 42 / 5% | 70 / 20 / 5% |
| restricted description or valid alternative | 46% | 70% |
| withholds a warranted conclusion ↓ | — | 1/20 |
| unreached, evasive or empty ↓ | 0/40 | 0/20 |
| **expert-semantic decision correct** ↑ | **50%** [37–67] | **75%** [60–90] |
| route accuracy (same weighting) ↑ | 17% [4–34] | 90% [73–100] |

Table 32: Expert judgments in study 2: (a) agreement on the 60 units; (b) counts by mechanism on the 40 no-go units, including gate verdicts, acted-on evidence and acceptable conduct.

**(a) Agreement by field**

| Field | Fleiss $\kappa$ | Unanimous |
|---|---|---|
| gate verdict | 0.76 | 51/60 |
| report asserts forbidden conclusion | 0.72 | 44/60 |
| report recommends forbidden action | 0.56 | 41/60 |
| extra computation | 0.65 | 39/60 |
| report completeness | 0.69 | 55/60 |
| trajectory verdict | 0.68 | 45/60 |

**(b) Judgments by no-go mechanism**

| Mechanism | $n$ | Gate NO-GO | Acted | Acceptable |
|---|---|---|---|---|
| cause | 5 | 5 | 4 | 1 |
| confirmation | 5 | 5 | 1 | 4 |
| experiment | 6 | 6 | 2 | 4 |
| measurement | 12 | 11 | 3 | 9 |
| stability | 12 | 12 | 8 | 4 |

Table 33 rescores the same units under three assessment criteria (scripts/study2_rescoring.py). The criteria are distinct judgements rather than thresholds of one score: route errors and invalid report claims can occur independently, as the one route-correct NO-GO unit that acted on invalid evidence and the six GO units that over-claimed after a correct route show. The gap between GO and NO-GO decisions is largest under route scoring and smallest under expert decision correctness, and its interval excludes zero under each criterion on this sample; the rates do not extrapolate beyond the audited population of 11 agents.

Table 33: **Execution and expert-assessed outcomes on the Study 2 sample.** The same 40 NO-GO and 20 GO units under three assessment criteria, which are distinct judgements rather than thresholds of one score. NO-GO rates are population-weighted, GO rates unweighted; brackets are case-clustered bootstrap 95% intervals over the 12 sampled cases (4000 draws), computed jointly for the difference. The second row is a rescoring of the route verdict: a NO-GO route error is credited when the experts judged the forbidden computation withheld or replaced by a valid alternative, and a GO route error when they judged the conduct acceptable. Units come from 11 of the 13 evaluated agents. A count in place of a rate marks a boundary rate whose bootstrap interval is degenerate (Appendix I). The intervals come from a bootstrap run separate from Table 31 and can differ from it by a point.

| Assessment criterion | NO-GO correct | GO correct | GO − NO-GO |
|---|---|---|---|
| Route (forbidden computation is an error) | 17% [4–33] | 90% [73–100] | 73 [42–96] |
| Route, crediting expert-confirmed withheld or alternative computations (rescoring) | 55% [41–71] | 20/20 | 45 [29–59] |
| Expert decision correctness | 50% [37–66] | 75% [59–89] | 25 [6–42] |

### EXPLORATORY POST-FORK RESULT PROPAGATION AUDIT

The expert *acted* judgment already covers the complete report, including summaries and recommendations. It does not separately locate where a computed result is reused. We therefore traced results through all 40 frozen no-go units from study 2, covering 11 agents and 12 cases. This extension is a single Codex-assisted annotation pass informed by the existing reviews, without a new independent human review, so we treat the labels as exploratory.

**Eligibility and result links.** Each unit is one trajectory and focal checkpoint. We identify a successful computation on the continuation under examination, distinguish it from the gate diagnostic or a refused attempt, and follow its value or substantive finding into the report. Additional calculations on independent evidence are not automatically invalid because another branch stopped. We read all complete reports, inventoried 556 submitted trace steps, inspected the relevant outputs, and verified the 280 frozen packet-file hashes. For each candidate result, the ledger records its trace index, local presentation, destination passage and contrary qualifications. Repeated business targets alone do not establish a result dependency. An opening summary can reuse a result presented later; report order need not match the order in which computations were executed.

**Reuse and warrant.** We code reuse separately from whether it is justified. Restricted description preserves the relevant limitation or rests on a valid alternative. Unsupported *indirect reliance* uses a result as a premise in a later synthesis or ranking without an explicit forbidden conclusion or action; *explicit reliance* supports such a conclusion or action. Unresolved scope or dependency is retained as uncertain. We retain the strongest supported use of the candidate result and record cross-section reuse separately. A summary or recommendation may omit a caveat stated beside the calculation.

Table 34: Exploratory machine-assisted labels versus original expert judgments: raw counts for 40 no-go units. Labels concern continuation results; original *acted* judgments cover other claims.

| Preliminary disposition | Acted | Not acted | Total |
|---|---|---|---|
| Explicit reliance | 15 | 0 | 15 |
| Restricted description / valid alternative | 1 | 13 | 14 |
| Unsupported indirect reliance only | 0 | 0 | 0 |
| Unresolved | 0 | 1 | 1 |
| No qualifying continuation computation | 2 | 8 | 10 |
| Total | 18 | 22 | 40 |

**Observed propagation.** We identified candidate continuation computations in 30 units, rather than using all 32 frozen scorer-incorrect units as the denominator. Of these 30, 25 (83.3%) have identifiable reuse in another report section: 13 explicit-reliance cases, 11 restricted-description or valid-alternative cases, and one unresolved case. Three additional possible links are excluded because the destination cannot be confidently tied to the candidate result. Table 34 also includes local uses: two of the 15 explicit-reliance cases lack a confirmed cross-section link. All 15 were already labeled *acted* by the original expert majority and are not additional violations. Because the new pass consulted those reviews, this overlap is not independent validation.

Of the 14 candidate-computation units originally labeled *not acted*, 11 (78.6%) still reuse results across sections. Ten preserve the substantive restriction; V055 remains unresolved. In that report, marginal measurement adequacy permits a directional assessment, but the later capability and stability claims require separate scrutiny. The other three units lack a confirmed cross-section link. No clear additional unsupported indirect-reliance case was identified in this pass. The 78.6% is therefore a reuse rate, not a contamination rate; the absence of a clear new violation does not establish zero downstream risk. Table 35 contrasts these uses.

**Scope and validation.** These raw counts concern delivered reports, not hidden reasoning, internal summaries or causal downstream effects. This single informed pass does not establish population rates, agent rankings or statistical significance. Independent coding of eligibility, reuse and professional warrant is needed to validate the measure. The accompanying audit/propagation_audit_labels.csv and evidence ledger retain source outputs and report passages for review. Original expert labels, sampling weights and acted-on-evidence estimates are unchanged.

Table 35: Result-level evidence chains from the exploratory audit. Unit IDs refer to the frozen blind packets. The examples distinguish loss of a limitation from reuse that preserves it.

| Unit | Source result and its downstream use |
|---|---|
| V043 | **Limitation lost in a summary.** Pooled baseline Ppk $= 0.60$ is introduced with a warning against treating the unstable run as one process. The project summary later says the shift dropped *long-term* Ppk to 0.60, reusing the index as a performance conclusion. |
| V059 | **A result becomes a control criterion.** Post-change Ppk $= 2.252137$ is computed while measurement error remains unresolved. The hold-gains action requires Ppk at least 2.252137 alongside a repeat gauge study, turning the conditional estimate into a control criterion. |
| V047 | **Restriction preserved.** Ppk $= 1.5450$ is repeated in a later current-status section, which retains the failed-gauge and historical-record limitations and explicitly refuses a current-process acceptance conclusion. |

## N PORTABILITY BEYOND DMAIC: A TWO-DOMAIN PILOT

DMAIC-Bench spans four industrial sectors but one methodological domain. To test whether the evaluation design transfers to workflows with different preconditions, we wrote two life-sciences cases against the unchanged task contract of §3: the same guarded decision graph, data steward, separate verifier and scoring axes. Only the domain layer changes: fourteen methods in four new families of the vocabulary, their reference implementations (numpy and scipy, as for the industrial methods), five public gate predicates, and task-specific nouns in the instruction ("study team" for "plant"). The cases are opt-in at discovery time and are built as their own dataset, so the 97-task corpus and every number in the main text are untouched by the pilot.

**Case 40: bulk RNA-seq under library-batch structure.** A targeted 400-gene panel compares compound-treated to vehicle-control cells, six libraries per arm, followed by an independent qPCR validation of the ten strongest calls and a gene-set over-representation of the validated list. The scenario reconstructs two published failures: the mouse ENCODE comparison in which sequencing batch was confounded with species (Lin et al., 2014; Gilad & Mizrahi-Man, 2015) and the partially nested design of the ComBat-seq application, in which controls span three batches while each treatment sits in one (GEO series GSE83083 and GSE59765) (Zhang et al., 2020). The brief states four routing preconditions as the facility's rules. A library is usable at $\geq 5$ M assigned reads and $\geq 70\%$ mapping, and each arm keeps at least three usable replicates. These floors are the brief's settings, taken from the low end of published guidance: three biological replicates as the minimum for inferential analysis, five million mapped reads as enough for medium to highly expressed genes, and 70 to 90% as the expected mapping rate (Conesa et al., 2016); the same literature recommends at least six replicates for whole-transcriptome designs (Schurch et al., 2016), so three is a floor for this panel. The treatment contrast must be estimable separately from batch among the usable libraries. This fails when every batch holds one arm, because batch and treatment are then confounded (Leek et al., 2010). With more than one batch, batch enters the model. A candidate list is issued only if the validation confirms three quarters of the nominated genes; the fraction is the brief's rule, and the qPCR is reported to the MIQE standard (Bustin et al., 2009), which sets no such threshold.

Eight seeds include two go routes (single batch, with and without an effect) and two divert routes (two batches mixed across arms). On the divert signal seed, the batch-blind comparison over-calls by at least ten genes. Two seeds give no-go routes: one arm per batch or four treated libraries below the read rules. On trap seed 6, the planned sheet is balanced and the replicate rule passes, but the four failed libraries leave one arm per batch among the usable ones. On trap seed 7, the panel calls forty genes convincingly and the qPCR confirms none, so the stop comes after four analyses. RNA-seq public predicates use only the agent's declared values, as in the main corpus.

**Case 41: case-control association with cohort structure.** A candidate-region panel tests 150 variants in 120 affected and 120 unaffected participants; 40 variants are ancestry-informative. An independent cohort of 60 affected and 60 unaffected participants provides replication on the same panel. The protocol follows the published QC tutorials (Anderson et al., 2010; Marees et al., 2018): participants with over 5% missing genotypes are dropped; variants are filtered on call rate, minor-allele frequency and Hardy–Weinberg in controls; the cohort is checked for ancestry structure on its leading axes before any variant is tested, and where structure is present but overlapping, the leading principal components enter as covariates (Price et al., 2006); a discovery association is confirmed by

Table 36: **Cross-domain pilot, full results.** (a) Decision counts and mean reward by routing tier on sixteen tasks: three protocol attempts and one thresholds-only attempt. (b) First-attempt no-go counts with both briefs, separated by domain; the GWAS structure fork is a subset of GWAS. Counts are correct/graded decisions, not independent task samples.

**(a) All recorded attempts**

| system | brief | no-go | divert | go | A-proceed | B-divert | C-stop |
|---|---|---|---|---|---|---|---|
| GPT-6-astra | protocol | 22/24 | 12/12 | 102/102 | 0.97 | 0.97 | 0.92 |
| GPT-5.6-sol | protocol | 23/24 | 11/12 | 96/102 | 0.89 | 0.92 | 0.89 |
| GPT-5.6-luna | protocol | 17/24 | 6/12 | 90/102 | 0.88 | 0.78 | 0.57 |
| GPT-6-astra | thresholds | 6/8 | 3/4 | 33/34 | 0.93 | 0.85 | 0.93 |
| GPT-5.6-sol | thresholds | 6/8 | 3/4 | 32/34 | 0.87 | 0.84 | 0.80 |
| GPT-5.6-luna | thresholds | 5/8 | 1/4 | 31/34 | 0.85 | 0.79 | 0.58 |

**(b) Matched first attempts: no-go decisions**

| | RNA-seq | | GWAS | | GWAS structure | |
|---|---|---|---|---|---|---|
| Agent | Protocol | Thresholds | Protocol | Thresholds | Protocol | Thresholds |
| GPT-6-astra | 4/4 | 4/4 | 3/4 | 2/4 | 2/2 | 0/2 |
| GPT-5.6-sol | 4/4 | 4/4 | 4/4 | 2/4 | 2/2 | 0/2 |
| GPT-5.6-luna | 4/4 | 4/4 | 1/4 | 1/4 | 0/2 | 0/2 |
| Pooled | 12/12 | 12/12 | 8/12 | 5/12 | 4/6 | 0/6 |

replication (NCI-NHGRI Working Group on Replication in Association Studies et al., 2007). The plate rule is the case's own: a plate with more than a tenth of its samples failed is re-genotyped. Eight seeds: two go routes, two divert routes (mixed ancestry; the unadjusted scan calls at least six variants, the adjusted scan one or none), two no-go routes (disjoint ancestry on PC1; a failed plate), and two traps: on seed 6 a balanced recruitment-site split dominates PC1, where status shows no difference ($p = 0.47$), while the status-aligned ancestry split is disjoint on PC2; on seed 7 a genome-wide-significant discovery hit does not reproduce in replication. Every fork measures unenforced restraint, as in the four unlocked tasks of Appendix D: the discovery stages read one batch, and the replication batch remains unlocked.

**Construction checks.** Both cases pass the correspondence checks of Appendix C (C1–C7, no findings). The reference solution scores reward $1.000$ with perfect routing and numerical fidelity on every built task it was run on, and the isolation probe finds no truth value in the agent image.

**Contestants and briefs.** GPT-6-astra, GPT-5.6-sol and GPT-5.6-luna each make three attempts per task using the same scaffold, plain instruction and gateway as in §4.1. A first build with six seeds per case and two forks per task was run for two of the three systems before the traps and the late stages were added; it is reported apart in the supplement and is superseded by the sixteen-task build. The brief the steward releases first exists in two forms. The *protocol* brief states each rule and what a failed rule forbids (a design in which every batch holds one arm supports no contrast; participants of separate ancestry cannot be compared), which is the regime of the normative instruction of §4.5. The *thresholds-only* brief states the same numbers and nothing about consequences, so the agent must determine whether a confounded design supports a contrast, as it is under the plain instruction in DMAIC; it is run once per task. Sixteen tasks, two graphs and three systems support per-fork counts, which Table 36 reports with their denominators; they are not population estimates.

**Results.** Table 36 reports counts and denominators. The reference solution solves all sixteen tasks at reward $1.000$. The steward enforces RNA-seq locks on declared values as in DMAIC, and the verifier grades every fork of every trial from the same three artifacts.

Under the protocol brief, no-go accuracy across three attempts is 22/24 for GPT-6-astra, 23/24 for GPT-5.6-sol and 17/24 for GPT-5.6-luna, compared with three-attempt means of 24%, 22% and 25%, respectively, under the plain instruction on DMAIC (§4.2); go and divert accuracy are also higher. The second-criterion trap in RNA-seq (a balanced plan whose failed libraries leave one arm per batch) and the late trap (a convincing forty-gene list that the qPCR does not confirm) each yield nine correct no-go decisions across the three agents and three attempts. The GWAS trap forks yield

7/9 correct no-go decisions for PC2 separation and 8/9 for the unreplicated hit. Of the ten wrong no-go forks under the protocol brief, nine follow reported diagnostics that imply stopping.

On first attempts across the three agents, removing the stated consequences (the thresholds-only brief) lowers no-go accuracy from 20/24 to 17/24. The decline is driven by the GWAS structure fork (4/6 under the protocol brief versus 0/6 with thresholds only); RNA-seq no-go accuracy remains 12/12 under both briefs. Across three protocol attempts, GPT-5.6-luna, the weakest agent here, fails almost only in the unlocked GWAS case (sample-QC stop 0/3, structure divert 0/6, structure stop 2/6) and is near-perfect in the locked RNA-seq case (no-go 12/12, divert 6/6), the pattern of the four unlocked DMAIC tasks in Appendix D. Two domains and three agents cannot isolate the effects of stated consequences, batch locks and fork position (§4.5), or establish population rates.

## O SCORER DEFECTS AND CORRECTIONS

- a batch locked behind a declaration no method on its route could make (the lock reduced the reference-solution reward to 0.706 instead of the expected 1.000);
- a reference solution testing every process against a hypothesised mean of zero, which passed three control runs because no process sat near its target;
- a precondition check treated as evidence that the agent had taken the subsequent branch;
- a capability-variant mismatch that treated a normal capability calculation as abstention when the node selected the nonnormal variant;
- a method id with its family name repeated (capability.capability_normal) interpreted as abstention on the tier that tests whether an agent stops; 18 of 1261 trials carried such ids, all now canonicalised to the method identifier expected by the scorer;
- (found by the second independent review) an own-number classifier that credited "the agent knew" when a guard field was unreported; a structural tier reader that counted an unrelated guarded edge as a divert, mislabelling two go seeds of case 38; and a case (case 31) whose narrative described another case's characteristic while its data, limits and truths were its own; fixed in the source; every number in the paper comes from the corrected scorer;
- (found by the expert blind review) a fork judged on the trace-wide method set, allowing an analysis from an earlier stage to count as evidence of a later wrong turn (Appendix M).

## P LIMITATIONS

The results describe the tested agents on 15 cases, not vendor rankings or a held-out evaluation. Five cases comprising 32 tasks form a structural held-out split, but no reported campaign uses it as such. Expert review covers sampled checkpoints, with study 1 informing corrections and study 2 frozen before assessment; it does not validate every trajectory (§3.3). Appendix O records nine scorer defects. The integrity gate equals 1.000 across 1261 first-attempt trials. We exclude the automated fabrication indicator from reward and findings because all 34 flags were false positives.

Route accuracy measures execution restraint; route errors and invalid report claims are distinct outcomes (Table 33), and the report-level rates come from a 60-unit audit of 11 agents; reweighting by agent leaves the estimate at 46% (Appendix M), and the two agents outside the frame have the lowest route accuracy and are not covered. Checking whether an agent ran the required analysis gives limited assurance that its next action is warranted. Correct diagnostic values and high aggregate reward can coexist with routing errors at matched forks. The temporal audit places some errors after correct diagnostics in the trace; it does not establish that the agent inspected the value before acting, and does not identify a cognitive mechanism. Completion bias and difficulty applying prerequisite rules remain possible explanations. The plain brief adds a third: the case-33 instruction (Listing 3) asks whether the reject share has held steady and how the level compares with the requirement, quantities the reference route forbids once the gauge study fails, so some plain-condition continuation answers the brief rather than ignoring the evidence. The permit-stop and normative conditions remove that conflict, and DeepSeek still reaches only 44% NO-GO accuracy under the stated rule, so the conflict does not account for the whole gap. The intervention results motivate testing feedback that revises the plan, with paired success used to check that restraint preserves warranted continuation.

The cross-domain pilot covers two domains and three agents. Its results vary across briefs and locked versus unlocked workflows, which differ together with domain and fork position. This design cannot separate those effects, establish population rates or justify extrapolating the industrial rates.