What Does a Harness Repair?
A Preregistered Study of Visibility, Baseline Adequacy and Evaluation Defects

Bowen Xu[*†] *bxu666@stanford.edu*
Department of Biology, Stanford University, Stanford, CA, USA

Boyu Chen[*]
Institute of Health Informatics, University College London, London, UK
[*]*Equal contribution.* [†]*Corresponding author.*

arXiv:2610.05533v2 [cs.LG] 24 Oct 2026

Abstract

Harness search keeps a change to the prompts, token budgets or parsers around a frozen model if the change raises a score. Such a gain can come from answers the parser could not read before, a weak comparison or an evaluation defect. We preregistered two studies of where these gains come from, with three models of 1.7B to 8B parameters and with Qwen3-14B, and report all 65 primary endpoints. Turning thinking off raised accuracy over a capped thinking setting in 7 of 12 model-benchmark cells, each time mostly through questions the capped setting left without a readable answer. For the small models, the thinking-off setting was not meaningfully worse than a rescue configuration in 7 of 9 cells, nor, for the 1.7B model, than four GEPA-selected harnesses. At 14B, two GEPA harnesses beat it by 30 to 34 points on the LongevityBench letter tasks; in descriptive analyses they were not detectably above the strongest constant-label baseline and answered "Yes" in 95% to 99% of their parsed mortality answers. From a start capped at 64 tokens, GEPA repaired the 1.7B model's harness, but at 14B no proposal of either search scored above zero. Replication through the same pipeline reproduced a defect's distortion in 9 of 20 evaluable cases, 7 of them cases in which the distortion is fixed by an identity of the design. Descriptively, relative to its base model, longevity fine-tuning raised a 1.7B model's letter accuracy by 41 points and cut its MMLU-Pro and GSM8K accuracy by 16 and 63 points; on a second GPU, 58% of its zero-temperature token sequences changed and none of the nine grid verdicts we rechecked did.

1 Introduction

A harness is everything around a frozen model that changes what it is scored on: the prompts, reasoning switches, token budgets, retry rules and parsers. Harness search proposes a change and keeps it if the score rises (Khattab et al., 2023; Agrawal et al., 2025; Lee et al., 2026b). The same step sits inside systems that improve themselves: each round of a loop that rewrites its own scaffolding proposes a change, scores it and keeps the winner (Zhang et al., 2025; Wu et al., 2026). What a round has found depends on what the score measures.

A gain that replicates on untouched rows shows that the selected harness beat its parent by more than selection noise. It leaves three questions open. The first is visibility: did the answers improve, or did answers that the parent never produced in a readable form, because it ran out of tokens or never committed to one, become visible to the parser? The second is baseline adequacy: would a simple setting that the search never tried do as well, and does the model beat trivial baselines at all? The third is evaluation validity: are the inputs, the gold labels, the split and the served model what they are said to be?

Each question has a literature (Section 2), and most of it works after the fact. Harness and scaffold gains vanish against matched-budget test-time scaling and simple baselines (Wang et al., 2026b; Gideoni et al., 2026; Liang et al., 2026; Kapoor et al., 2024). A gain can be completion rather than correctness (Seyoum & Mittur, 2026; Lee, 2026; Chen et al., 2026). Thinking off can beat a thinking cap (Nie et al., 2026). Evaluation pipelines carry defects that move scores by tens of points (Zhu et al., 2025; Berriche et al., 2026). We found no earlier study that fixes in advance, for a stated contrast, how much of a harness gain is visibility, none that tests the benchmark authors' own protocol against searched harnesses with a margin, and none that injects known defects one at a time to ask whether replication reproduces them. The search behind these statements covered 2024-01 to 2026-10-03 and was capped at 200 web searches, and three of its six parts had to be resumed after a rate limit, so these statements are bounded (Appendix H).

Model choice adds a fourth question. What a harness can repair depends on what the model already does without it. A model that already answers in the required form leaves the harness little to repair, and one that does not may need the harness to answer at all. A baseline that is adequate for one model can therefore fall short for another.

Background. This work follows a case study of harness search around a frozen longevity language model (Appendix A). Its search scored every candidate through a caller that sent each benchmark row's reference answer to the model as an earlier assistant turn, and the rescue it selected replicated through the same caller, so replication could not detect the defect: every call carried it. Trivial baselines, input audits and per-call prompt-token reconciliation did detect it. After the caller was corrected, a preregistered re-evaluation found that the search's starting configuration parsed none of its 556 prospective rows, and that a thinking-off setting the search had never tried was non-inferior to the rescue at less than half the token cost. We ask whether this one model, one benchmark and one defect is a pattern.

This paper. We answer with two preregistered studies. The first, frozen on 30 September 2026 at 06:53:11 UTC, uses three models of 1.7B to 8B parameters, three benchmark pools, a replication partition and a confirmatory test partition, an adaptive arm in which GEPA (Agrawal et al., 2025) searches for harnesses around the 1.7B model, and six evaluation defects injected one at a time. The second, frozen on 7 October 2026 at 00:28:49 UTC, was written after the first study's results were known. It reuses the first study's rules by reference, fixes its own endpoints before any call on its replication and test partitions, and puts the same hypotheses to Qwen3-14B, with a GEPA arm of its own. It also runs the base and the longevity-tuned versions of Qwen3-1.7B again, on a second GPU, as a descriptive comparison, to ask what the fine-tuning added and whether the first study's conclusions depend on the GPU.

The paper makes five contributions. It splits the gain of a thinking-off protocol over a capped thinking configuration exactly into rows only the protocol parses, rows both parse, and rows only the capped configuration parses, the split that Chen et al. (2026) use for self-revision, and fixes a dominance statistic and a gain gate in advance. It tests that protocol for non-inferiority against a derived rescue and against GEPA-selected harnesses at 1.7B and at 14B, with trivial baselines next to every accuracy. It tests whether the serving stack's thinking budget, which also allows a longer answer, lowers truncation. For each defect and model it fixes one claim in advance, that replication through the same pipeline confirms the defect's distortion, and it measures the false alarms of the audit checks. Last, it compares the base and the fine-tuned version of one model, and the same procedure on two GPUs, as descriptive results with pre-specified limits on how they may be read. All 65 primary endpoints are reported with every deviation and addendum, and everything computed after the outputs existed is labelled.

The results cover four models in the confirmatory design (three of 1.7B to 8B parameters and Qwen3-14B) and the base Qwen3-1.7B descriptively, three pools with LB-0038 reported apart, and six defects, and nothing beyond them. We studied one round of selection and no recursive loop. We do not claim that harness gains can be illusory, that thinking off can beat a thinking cap, or that agent and harness evaluations can be preregistered, since each has been shown or done before, and we do not claim that harness search finds genuine quality gains.

2 Related work

We organise related work by the questions of Section 1; Appendix H states how it was searched and what the search could not cover.

Harness search, and whether its gains are real. Prompt optimisers (Zhou et al., 2023; Yang et al., 2024; Khattab et al., 2023; Agrawal et al., 2025) and harness search (Lee et al., 2026b; Wu et al., 2026; Zhang et al., 2025) treat the text and control flow around fixed weights as the object of search, and GEPA is the adaptive arm of our study. Earlier work has shown, after the fact, that such gains can be illusory: evolved harnesses and scaffolds do not beat matched-budget test-time scaling or simple baselines (Wang et al., 2026b; Gideoni et al., 2026; Liang et al., 2026; Kapoor et al., 2024), a hardware-verification harness gains mostly completion (Seyoum & Mittur, 2026), a search started from a defective seed fails (Liu et al., 2026), and self-improving loops tamper with or hack their evaluators (Wang et al., 2026a; Yang et al., 2026a). Our question is narrower: where the gain of one protocol comes from, and whether a known defect would be seen, with the questions and decision rules fixed in advance.

Visibility and quality (H1). Chen et al. (2026) split the accuracy change of a self-revision step into a content margin, a format-recovery margin, a format-loss margin and a dead cell, which is our split under other names but without a predefined dominance threshold. Nie et al. (2026) and Kaiser et al. (2026) give the identity for a single configuration or between models, completion rate times conditional accuracy; under a shared budget thinking off beats capped thinking (Nie et al., 2026; Ma et al., 2025), and capped single-call evaluation of hybrid models has been called an artifact of protocol (Lee et al., 2026a). Other work finds that unparsed answers hide correct ones (Garware & Zisad, 2026), separates traces cut by a cap from traces that finished without an answer (Goyal & Ray, 2026), reads a harness-evolution gain informally as completion rather than correctness on the twelve tasks that had been searched (Seyoum & Mittur, 2026), and counts repaired abstentions for constrained decoding (Lee, 2026). We add a pre-specified dominance statistic and gain gate, applied prospectively to the contrast between a thinking-off protocol and a capped thinking configuration.

Baseline adequacy (H2). Several studies find, after the fact, that simple baselines match searched or evolved harnesses and complex agents (Kapoor et al., 2024; Wang et al., 2026b; Gideoni et al., 2026; Liang et al., 2026; Chang & Chen, 2026). A preregistered non-inferiority test of a simple alternative exists for agent memory (Sharma & Lall, 2026), and equivalence margins have been used for scaffolds (Gringras, 2026) and for compression (Singh, 2026). Our test fixes the margin in advance for the benchmark authors' own protocol, against a derived rescue and four GEPA-selected harnesses.

Budget forcing (H3). Budget forcing and the Qwen3 thinking budget (Muennighoff et al., 2025; Yang et al., 2025) serve as evaluation controls, with a forced answer or a forced end of thinking (Nie et al., 2026; Carlon et al., 2026; Lee et al., 2026a; Chang & Chen, 2026), and one study forces the answer delimiter at a cap in vLLM (Goyal & Ray, 2026). We test the engine's thinking-token budget (Kwon et al., 2023) as a harness setting against a capped configuration; the technique is not ours.

Defects and prospective replication (H4). Other work injects faults with known ground truth and measures false alarms: tool bugs in agent benchmarks (Bellibatlu et al., 2026), tampering edits (Wang et al., 2026a), reward-hacking channels (Yang et al., 2026a) and planted gains (Shawn, 2026). Defects like ours have been catalogued as they occurred, for answer-key access, label defaults, parser fallbacks, stop sequences, denominators and gold errors (Wahi, 2026; Berriche et al., 2026; Zhu et al., 2025), and the reference scorer of MMLU-Pro credits a random option for an empty prediction.[1] We found no study that injects pipeline defects one at a time into a harness-search pipeline and tests, for each defect and model, whether replication through the same pipeline reproduces the distortion.

Served-model identity and determinism. Serving stacks change outputs at fixed weights (Pape et al., 2026; Wang & Sang, 2026; Yuan et al., 2025; He & Thinking Machines Lab, 2025), answers can flip while

[1] `evaluate_from_local.py`, function `save_res`, in the repository TIGER-AI-Lab/MMLU-Pro, read on 2026-10-03.

accuracy stays flat (Dutta et al., 2024; Kübler et al., 2026), and a served route is not its model identifier (Stenstrom et al., 2026; Gao et al., 2025). vLLM used to decide tying from the configuration alone and silently discarded a stored head; its fix (hmellor (vLLM project), 2026) is in the version we used (vLLM 0.30.0), which loads the stored head. We add a measured instance of the mismatch, since the mechanism is known.

GEPA from a broken start. Liu et al. (2026) report that GEPA, started from its official GSM8K seed (which puts the answer before the reasoning), falls from 23.81% to 13.50% and never repairs the order, and Ursekar et al. (2026) start optimisers from deliberately untuned seeds, one of them a non-functional stub. Ours is a controlled comparison of clean and broken starts; why a search repairs one defect and not another is not tested here.

Preregistration and prospective design. Preregistered or frozen designs with deviation logs, holdouts or equivalence margins appear in Gringras (2026); Weinberger & Hozez (2026); Sharma & Lall (2026); Lee (2026); Goyal & Ray (2026); Liang et al. (2026); Bellibatlu et al. (2026); Singh (2026). Thomas et al. (2026) confirm on the first eligible model released after the plan, a stronger standard than ours, since our freeze guards the rules and not model-specific tuning. We found no earlier study that combines the freeze with its hash and time in the commit, a pilot, replication then test, an adaptive arm before replication, every primary endpoint of two preregistrations reported (65 in all), addenda that state what had been seen, and an exposure log.

Hardware and token-level agreement. Numerical differences between hardware change the tokens a model generates, even at zero temperature (Yuan et al., 2025; He & Thinking Machines Lab, 2025). Our comparison of the same procedure on two GPUs asks whether the conclusions of a study depend on that agreement, which is a different question from whether the outputs match.

3 Design

The design is that of two frozen preregistrations, the first study's and the second study's (hashes and times in Appendix E); the second reuses the first's rules and changes the model and the families (Section 3.6), and this section states only what the results need.

3.1 Setting

Models and stack. The first study uses Qwen3-1.7B-Longevity (Zhavoronkov et al., 2026), a Qwen3 (Yang et al., 2025) model tuned on aging-biology data, SmolLM3-3B (Bakouch et al., 2025) and Intern-S1-mini (Bai et al., 2025). The second study uses Qwen3-14B and, descriptively, the base and the tuned version of Qwen3-1.7B (Section 3.6). The models of both studies ran on vLLM 0.30.0 (Kwon et al., 2023) in batch-invariant mode (He & Thinking Machines Lab, 2025) at zero temperature and client concurrency 32, the first study on one RTX 5090 and the second on an RTX PRO 6000, with three exceptions: the GEPA proposer and the seeded-sampling arms sample at the model cards' values, and the FreeToken (Yang et al., 2026b) continuity runs of the tied-output-layer defect (E) ran on FreeToken 0.1.3 instead of vLLM, at temperature 0 and client concurrency 1, in a context of 2,048 tokens. The Qwen3-1.7B-Longevity checkpoint declares tied embeddings but stores a different output matrix, and the study model uses the stored one (untied). Two models of the first study and the two 1.7B arms of the second also have a descriptive seeded-sampling arm.

Pools, cells and partitions. The LongevityBench letter pool combines, by row, LB-0030 (age band, five options), LB-0034 (which of two participants is older) and LB-0042 (ten-year survival), and LB-0038, an age regression, is reported apart. MMLU-Pro (Wang et al., 2024) (a stratified subset of 2,000 items) and GSM8K (Cobbe et al., 2021) are the other two pools, and a *cell* is a model with a pool, so there are nine cells in the first study and three in the second, which uses the same items, partitions and analysis rows. The items are split into search, replication and test partitions (40/20/40); LongevityBench's 14,720 rows, in 2,104 participant components, are split by component so that no participant spans partitions, and GEPA searches on the search partition only. The *analysis rows* leave out 147 replication and 371 test rows that the case study had held out and later called, and 7 replication rows whose reference answers could be recovered

Table 1: Configurations ("tokens" is `max_tokens`; thinking on or off is set in every request). RESCUE-1500 is derived from the two grid runs and verified against the real two-call sequence on 300 items per model.

Configuration	Request
THINK-64, THINK-512, THINK-1536	thinking on; 64, 512 or 1,536 tokens for thinking and answer together (THINK-64 is the case study's parent)
AUTHORS-1500	thinking off; 1,500 tokens (the benchmark authors' serving protocol)
BUDGET-512-1500	thinking on; vLLM ends thinking after 512 tokens (Muennighoff et al., 2025); up to 1,500 tokens of answer follow, so 2,012 tokens in all against 512 for THINK-512
RESCUE-1500	THINK-512; where its call ended at the token limit, the AUTHORS-1500 call on the same row (the call that gives the answer decides the properties of the row)
GEPA harnesses	in each study four runs (clean or broken start, two seeds), around Qwen3-1.7B-Longevity in the first and Qwen3-14B in the second, 6,000 metric calls on 300 search rows, Intern-S1-mini as proposer; the selected harness edits the system prompt and appends an instruction, so the arm searches prompt text only (Appendix J)

from an upload made by mistake; putting them back is a sensitivity analysis. The test cells then have 3,942 letter rows (789 participant components, the bootstrap unit), 799 MMLU-Pro items and 527 GSM8K items, and the replication cells have 1,990, 399 and 264.

3.2 Configurations and scoring

Table 1 lists the six configurations of the grid and the GEPA harnesses. A call is *unparsed* if the scorer cannot read an answer from it, and it then scores as wrong, so an unparsed answer may hide a correct answer, a wrong one or none (Appendix S states the scorer's rules). A call is *truncated* if it ended at the token limit (finish reason `length`), whatever the scorer makes of it, and *non-committed* if it ended by itself (finish reason `stop`) and is unparsed; for RESCUE-1500 both properties are those of the call that gives the answer.

For H1, let P be the baseline configuration and R the comparison, paired on the n rows of a cell, with u the parse indicator and c the correct indicator of a call and a bar negating an indicator. Row i is in V if only R parses it ($\bar{u}_{P,i}u_{R,i}$), in B if both parse it ($u_{P,i}u_{R,i}$) and in L if only P parses it ($u_{P,i}\bar{u}_{R,i}$). With $\Delta = \sum_i (c_{R,i} - c_{P,i})/n$ and the contributions $\Delta_V = \sum_i V_i c_{R,i}/n$ (P scores zero on V), $\Delta_B = \sum_i B_i (c_{R,i} - c_{P,i})/n$ and $\Delta_L = -\sum_i L_i c_{P,i}/n$, the identity $\Delta = \Delta_V + \Delta_B + \Delta_L$ is exact, because a row that neither configuration parses adds nothing. Each contribution is divided by the cell's n and not by the size of its row set, as in the tables, so Δ_V is R's number of correct answers on V per row of the cell. Chen et al. (2026) use this additive split for self-revision (content, format-recovery and format-loss margins); the dominance statistic, its fixed share and the gain gate are ours. The dominance statistic $T = \Delta_V - 0.8\,(\Delta_V + \Delta_B)$ is positive when visibility supplies more than 0.8 of what the rows R parses contribute.

3.3 Hypotheses and endpoints

Table 2 states the four hypotheses. The first study has 47 primary endpoints and the second 18, 65 in all; they are the only confirmatory results, and their verdicts come from percentile-bootstrap intervals adjusted within each family.

3.4 Defects, claims and checks (H4)

Six defects arose in the case study, and we inject each into a harness-search pipeline, one at a time, and compare the result with the clean pipeline on the same rows. The pipeline's start s is THINK-1536, and the reported quantity is the start's accuracy on the letter pool, measured on a selection sample S (the replication partition) and on a check sample R (the test partition); for the broken start it is the gain over the start. Let D_S and D_R be a defect's distortion of the reported quantity on S and on R, and $\Gamma = D_R - D_S$. The reference-answer defect (A) sends the row's reference answer as an assistant turn, on 1,204 replication rows

Table 2: Hypotheses, estimands and verdict rules. m is the minimal difference and δ the non-inferiority margin, both 0.05. Every bound in a verdict is a bound of the family-adjusted interval, and H1's gate uses the same resamples and the same adjusted level as T. The last column gives the first study's family and the second study's number of endpoints: the second study has 3, 7, 3 and 5 endpoints in the four families and is tested at the first study's family sizes and levels (Section 3.6).

	Claim and estimand	Verdict	Endpoints (family)
H1	Most of the accuracy that AUTHORS-1500 gains over THINK-512 comes from rows THINK-512 leaves unparsed. Gate: Δ; statistic: T.	*no gain: not tested* unless the lower bound of Δ is above zero; if it is, *visibility dominates* if the lower bound of T is above zero, *quality matters* if the upper bound of T is below zero, *inconclusive* otherwise.	9: one per cell; second study 3
H2	AUTHORS-1500 is not meaningfully worse than a comparator: $\theta = \mathrm{acc}(\text{AUTHORS-1500}) - \mathrm{acc}(\text{comparator})$.	*holds* if the lower bound of $\theta > -\delta$; *fails* if the upper bound $< -\delta$; else *inconclusive*.	13: RESCUE-1500 in nine cells; four GEPA harnesses (Qwen3-1.7B-Longevity, letters); second study 7: RESCUE-1500 in three cells, four harnesses (Qwen3-14B, letters)
H3	Against THINK-512, BUDGET-512-1500 (a forced end of thinking and a larger allowance) lowers truncation (D_T) and raises parsing (D_P).	*holds* if both lower bounds $> m$; *fails* if either upper bound $< m$; *no room* if the upper bound of THINK-512's truncation is $< m$ (applied first); else *inconclusive*.	9: one per cell; second study 3
H4	For each defect and model, prospective replication confirms the defect's distortion (Section 3.4).	*holds*, *fails* or *inconclusive* from three intervals; *not evaluable* if the defect cannot move the start by more than m.	16 pairs, fixed; second study 5 pairs

halved by component into S and R. The gold fallback (B1) takes LB-0042's gold from a historical fallback that, with this release's labels, differs from the derived gold on about half of its rows, and the unparsed-as-A scoring (B2) scores an unparsed answer as the letter A. The skewed split (C) re-implements the case study's split, with S and R the halves of its validation part. The broken start (D) starts from THINK-64 and reports the gain of the selected configuration over the start; the selection does not depend on the start, so it cancels, and the distortion is $D_S = \mathrm{acc}_S(\text{THINK-1536}) - \mathrm{acc}_S(\text{THINK-64})$, with the same difference on R for D_R. The tied output layer (E) serves the tied layer in place of the untied one (Qwen3-1.7B-Longevity only), as an engine that decides tying from the configuration alone does (hmellor (vLLM project), 2026); the second study applies the other five defects, since Qwen3-14B has one served variant. Appendix I gives each defect's construction, clean counterpart, samples and evaluability bound, and each check's predicate and threshold.

The claim for a defect and a model holds if D_S and D_R lie wholly outside $\pm m$ with the same sign and Γ lies wholly inside $\pm m$; it fails if any of the three lies wholly on the other side of its boundary, and it is inconclusive otherwise. A pair is evaluable only if the largest distortion its mechanism can cause on S exceeds m, a bound that we compute from parse, gold-change and answer-difference indicators without accuracy; the skewed split changes the sample and has no row-level bound, so it is evaluable by rule. The family stays at 16 whatever is evaluable. Five checks read the defects' signatures: per-call prompt-token reconciliation (TR), trivial baselines (TB), the simple configuration family (SCF), input reading (IR) and the gain and level rules of prospective replication (PR), plus served-model identity (SMI) for the tied output layer. Their detection matrix is largely true by construction and descriptive, and what is empirical is their false alarms on clean pipelines.

3.5 Statistics and what was fixed in advance

Every contrast is paired on rows. In each of $B = 50{,}000$ bootstrap resamples we draw the cell's units with replacement (participant components for LongevityBench, items for MMLU-Pro and GSM8K), so that a unit drawn k times counts its rows k times; every statistic of the cell is recomputed on the resample with each row weighted by the draw of its unit, and all contrasts of a cell share the resample. The seed of a cell derives from a hash of its partition, model and pool (base 4242). The B statistics are sorted ascending and indexed from zero, and for a two-sided tail probability q the lower bound of the interval is the value at index

$\lfloor Bq \rfloor - 1$ and the upper bound the value at index $\lceil B(1-q) \rceil - 1$, clamped to the range of indices. Every estimate carries an unadjusted 95% interval ($q = 0.025$), and a verdict uses the Bonferroni-adjusted interval of its family, with $q = 0.025/m$ for m endpoints (levels 99.44%, 99.62%, 99.44% and 99.69% for H1–H4). These levels are nominal: a percentile bootstrap covers only approximately, so the family-wise error rate is controlled approximately and is not guaranteed. For H4, the resample of an instance's units also serves its clean counterpart where the two share their rows (all defects except the skewed split, which uses other rows and is resampled independently), so that D_S and D_R are paired, and Γ pairs the b-th resample of S with the b-th of R, which are disjoint samples drawn with their own seeds. A test verdict is *replicated* if the replication partition gives the same verdict, *not replicated* if it gives the opposite one, and *inconclusive replication* otherwise. The error control of H4 covers "holds" only, and "fails" is a descriptive label. The first preregistration was frozen at 06:53:11 UTC on 30 September 2026, before any replication or test item was called, and the pilot called 700 pilot items only; the second was frozen at 00:28:49 UTC on 7 October 2026 on the same terms. The directions of H1 and H3 were visible in the pilot's parse and truncation rates, and the primary set of H1 changed twice before the freeze after those rates were seen (Section 5).

3.6 The second study

The second preregistration was written after the first study's results were known. It reuses the first study's rules by reference, keeps its hypotheses and verdict rules, and changes the model and the families: Qwen3-14B with 3 H1 cells, 7 H2 endpoints (the rescue in three cells and four GEPA harnesses on the letters), 3 H3 cells and 5 H4 pairs, 18 primary endpoints in all, tested at the first study's family sizes and so at the levels 99.44%, 99.62%, 99.44% and 99.69%. Its GEPA arm has the first study's design around Qwen3-14B, and its defect-A pass runs after the test partition, so the checkpoint locked no selection for it and that pair is evaluable by rule. The checkpoint, which records the four selected harnesses and the replication results before the first test call, is committed and its hash logged before that call. The two 1.7B arms run the base model and the tuned model on the same items, configurations and schedule as a secondary, descriptive comparison with fixed wording and a fixed limit on its interpretation (Section 4.6); they carry no verdict and no adjusted interval. The second study ran on an RTX PRO 6000 against the first study's RTX 5090, the same chip architecture.

Two events during the second study bear on a result. A platform shutdown on 8 October 2026, when the account balance ran out, ended the instance, and its data disk was cloned to another host with another RTX PRO 6000; the 14B's test grid had finished on the first GPU and every other test run ran on the second. The preregistration's rule for a change of the physical GPU after the first replication call was applied: the batch-invariance check on the pilot rows was repeated for the three models and passed, and 900 replication calls of each model, in three configurations, were rerun and were identical token for token to the outputs of the first GPU, so the confirmatory analysis stands. The GEPA harness evaluation that the shutdown interrupted was moved aside, kept, excluded from every analysis and rerun from scratch. Appendix C records these and every other deviation of the second study.

4 Results

Results are on the test partition, the confirmatory one, unless stated, and the replication partition enters through the replication label. Primary endpoints carry the family-adjusted interval and everything else the unadjusted 95% interval; Appendix B gives every endpoint of the first study and Appendix M every endpoint of the second. A star marks an endpoint whose quantity is fixed by an identity of the design or by the observed parse pattern, so that its verdict says little beyond the mechanism; the verdict still depends on the intervals, and the text says why where the endpoint appears.

4.1 The 65 primary endpoints at a glance

Table 3 gives the 31 endpoints of H1–H3 of the first study and the 13 of the second, and Table 6 the pairs of H4: 16 in the first study and 5 in the second. Of the 31 endpoints of H1–H3 in the first study, 30 are *replicated* and 1 is an *inconclusive replication*, and the number labelled *not replicated* is 0; all 13 of the

Table 3: Verdicts of the 31 endpoints of H1–H3 of the first study (the first nine rows) and the 13 of the second (Qwen3-14B) on the test partition. A dagger marks the one endpoint whose replication label is not "replicated": SmolLM3-3B on GSM8K, H2, which *fails* on test and is *inconclusive* on replication (an *inconclusive replication*). A star marks an endpoint whose quantity is fixed by an identity of the design or by the observed parse pattern; its verdict says little beyond that mechanism, although it still depends on the intervals (Sections 4.2 and 4.3). The last two rows are H2 against the GEPA harnesses of each study.

Model	Pool	H1 AUTHORS-1500 vs THINK-512	H2 AUTHORS-1500 vs RESCUE-1500	H3 BUDGET-512-1500 vs THINK-512
Qwen3-1.7B-Longevity	LB letters	no gain: not tested	holds	no room
	MMLU-Pro	no gain: not tested	holds	inconclusive
	GSM8K	no gain: not tested	holds	no room
SmolLM3-3B	LB letters	visibility dominates	holds	holds
	MMLU-Pro	visibility dominates	holds	holds
	GSM8K	visibility dominates	fails†	holds
Intern-S1-mini	LB letters	visibility dominates⋆	holds⋆	holds
	MMLU-Pro	visibility dominates⋆	holds	holds
	GSM8K	no gain: not tested	fails	holds
Qwen3-14B	LB letters	visibility dominates⋆	holds⋆	holds
	MMLU-Pro	visibility dominates	holds	holds
	GSM8K	no gain: not tested	fails	holds
Qwen3-1.7B-Longevity	LB letters	H2 against the four GEPA harnesses (clean 0, broken 0, clean 1, broken 1): holds⋆, holds, holds, holds		
Qwen3-14B	LB letters	H2 against the four GEPA harnesses (clean 0, clean 1, broken 0, broken 1): fails, fails, holds⋆, holds⋆		

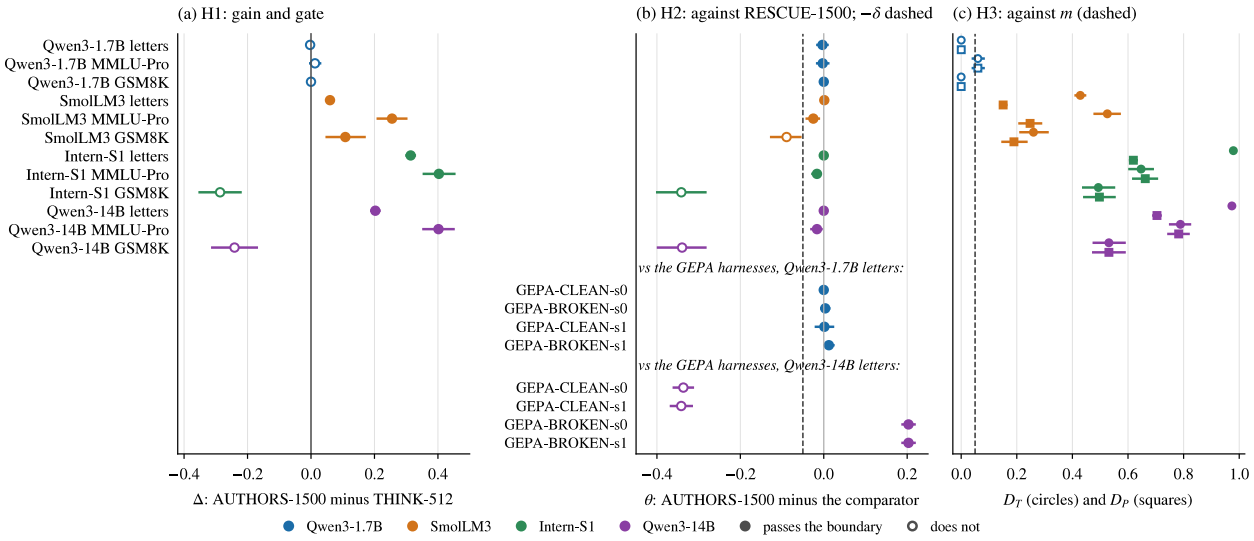

Figure 1: Adjusted intervals of the primary endpoints of H1–H3 on the test partition, the first study's nine cells above the second study's three. (a) H1: Δ, the gain gate; (b) H2: θ against the non-inferiority margin $-\delta$; (c) H3: D_T and D_P against the minimal difference m. Cells are model×pool; in (b) the harnesses of each study are below the cells; filled markers pass the verdict's boundary, open markers do not.

13 in the second study are replicated. Figure 1 plots the adjusted intervals, and the sections below take the endpoints by question. In the first study no sensitivity analysis changes more than 3 of the 31 H1–H3 verdicts (Appendix R); in the second the strict scorer changes 1 of the 13, the rescue on MMLU-Pro (*holds* becomes *fails*), and no other analysis changes any (Appendix M).

4.2 Where the gains come from (H1)

In the first study *visibility dominates* in SmolLM3-3B on all three pools and in Intern-S1-mini on the letters and MMLU-Pro. In these five cells AUTHORS-1500 gains over THINK-512 between $\Delta = +0.060$ $[+0.046, +0.074]$ (SmolLM3-3B, letters) and $+0.403$ $[+0.352, +0.456]$ (Intern-S1-mini, MMLU-Pro), and the gain is visibility: Δ_V is at least Δ in all five (shares from 1.00 to 1.82), because the rows both configurations parse contribute at most $\Delta_B \leq +0.001$ and the rows only THINK-512 parses cost Δ_L (for SmolLM3-3B on GSM8K, $\Delta_V = +0.197$, $\Delta_B = -0.036$ and $\Delta_L = -0.053$). Almost all of THINK-512's unparsed calls ended at the token limit; the exception is SmolLM3-3B on the letters, where 543 of 3,394 ended by themselves without an answer (non-committed). In two of the five the parse pattern nearly fixes the quantity, as the preregistration expected for the first. In Intern-S1-mini on the letters both configurations parse only 1 row of 3,942, so Δ_B is nearly zero and T, nearly a fifth of Δ_V, is positive (the preregistration listed the cell as "$T > 0$ nearly inevitable"). In Intern-S1-mini on MMLU-Pro, B has 69 of 799 rows, so Δ_B cannot exceed that share, which is below a quarter of the observed $\Delta_V = +0.419$, and $T > 0$ follows from B's row count and that observed Δ_V. The verdict is informative in SmolLM3-3B's three cells, where B has 448, 81 and 340 rows.

The other four cells fail the gain gate. Qwen3-1.7B-Longevity needs no visibility, since THINK-512 parses 1.000 of the letter rows, 0.940 of MMLU-Pro and 0.992 of GSM8K, and $\Delta = -0.003$ $[-0.018, +0.012]$ on the letters. In Intern-S1-mini on GSM8K the gate fails the other way, because AUTHORS-1500 loses accuracy ($\Delta = -0.287$ $[-0.355, -0.218]$); there $T = +0.286$ is positive only because $\Delta_B = -0.343$ is large and negative, the reading the gate exists to stop, and THINK-512 is the more accurate configuration on the rows both parse. The sanity checks (THINK-64 against AUTHORS-1500) give *visibility dominates* in 8 of nine cells, as expected, and fail the gate where Qwen3-1.7B-Longevity already parses GSM8K.

Qwen3-14B. In the second study visibility dominates again where the capped setting leaves rows unparsed. *Visibility dominates* on the letters ($\Delta = +0.203$ $[+0.186, +0.219]$) and on MMLU-Pro ($\Delta = +0.402$ $[+0.350, +0.453]$), and the gate fails on GSM8K ($\Delta = -0.241$ $[-0.315, -0.167]$). In the letter cell the parse pattern fixes the quantity: THINK-512 parses 0.001 of the letter rows and every unparsed call ended at the token limit, so B holds 2 of 3,942 rows, Δ_B is 0.000 and the whole gain is visibility; the verdict still rests on the gate and on the interval of T. On MMLU-Pro $\Delta_V = +0.419$ and $\Delta_B = -0.018$, with 107 rows in B. On GSM8K AUTHORS-1500 answers in a mean of 23.7 tokens (computed after the outputs existed) and is less accurate than THINK-512 on the rows both configurations parse ($\Delta_B = -0.342$ against $\Delta_V = +0.101$), the pattern of Intern-S1-mini: $T = +0.293$ is positive only because Δ_B is large and negative, the reading the gate exists to stop. Over both studies visibility dominates in 7 of 12 cells.

4.3 Whether simple baselines are adequate (H2)

In the first study AUTHORS-1500 is non-inferior to RESCUE-1500 in 7 of nine cells and to all four GEPA harnesses. It fails on GSM8K for SmolLM3-3B ($\theta = -0.089$ $[-0.129, -0.053]$) and for Intern-S1-mini ($\theta = -0.342$ $[-0.402, -0.281]$), where the rescue is more accurate than the thinking-off configuration by more than the margin. The study does not isolate which part of the rescue, the thinking within 512 tokens or the second call, accounts for the difference. We expected the Intern-S1-mini endpoint to be inconclusive, an informal expectation from a planned adjusted power of 0.19 at a true θ of zero (Appendix K), and it is not close. For SmolLM3-3B on GSM8K the replication partition is inconclusive, so the label is an *inconclusive replication*. In Intern-S1-mini's letter cell the rescue equals AUTHORS-1500 ($\theta = 0.000$, since every THINK-512 call is cut off), so *holds* is trivial there.

In the first study GEPA searched prompt text only around Qwen3-1.7B-Longevity (Appendix J), repaired its broken starts, and found no harness whose accuracy differs detectably from the thinking-off baseline's. In the clean-start run with seed zero no proposal beat the seed on the validation rows, so its selected harness is the seed, AUTHORS-1500 itself ($\theta = 0.000$ because the harness is the same text and request). The other three selected harnesses parse every test row and reach accuracy 0.620 to 0.631, against 0.632 for AUTHORS-1500; θ runs from $+0.002$ to $+0.012$, and all three adjusted intervals contain zero. Their validation scores (0.660

to 0.700 over the four runs) exceed their test accuracy in every run, as selection makes them optimistic, and none of the candidates GEPA saw (192 to 377 per run) contained a soft switch (`/think` or `/no_think`).

Qwen3-14B. Against the rescue, AUTHORS-1500 holds on MMLU-Pro (under the strict scorer, a sensitivity analysis, it fails; Appendix M) and fails on GSM8K ($\theta = -0.340$ $[-0.400, -0.281]$), as it did for SmolLM3-3B and Intern-S1-mini. On the letters the rescue equals AUTHORS-1500 ($\theta = 0.000$), because every THINK-512 call is cut off, so *holds* is trivial there, as it is for Intern-S1-mini. Against GEPA the result is the opposite of the first study's. Both clean-start harnesses are more accurate than AUTHORS-1500 by more than the margin ($\theta = -0.337$ $[-0.362, -0.311]$ and -0.342 $[-0.369, -0.314]$), with test accuracy 0.540 and 0.545 against 0.203, and the verdict is *fails* for both, on both partitions. The two broken-start searches each kept their seed, which is one text, so their two endpoints are one result: the harness is the start that answers nothing, and the verdict *holds* rests on AUTHORS-1500 being far above it.

Why the broken start could not be repaired. The broken start is the request of THINK-64, thinking on with a cap of 64 tokens. A search edits two texts, the system prompt and an appended instruction, and the request parameters stay fixed. The soft switches `/think` and `/no_think` were allowed in those texts, and no candidate of either run held one. In each of the two runs GEPA made 975 proposals; 142 and 137 of them failed the interface check and were scored zero without a call, and the others were called. Every proposal scored zero against its parent's zero on a minibatch of three rows, so the strict-improvement rule accepted none, no text besides the seed was validated, and the selected harness is the seed, which is truncated in every letter call. THINK-64 is truncated in every call of every cell, so the seed answers nothing; in these two searches, under their budget and acceptance rule, no proposal scored above zero. The Qwen3-1.7B-Longevity model in the first study answers within the cap, and there GEPA repaired its broken starts. The clean-start searches, in contrast, made 209 and 239 proposals and accepted 30 and 29. The arm searched prompt text only, so this says nothing about a search that may also change the request.

Qwen3-14B against the strongest constant. Table 4, computed after the outputs existed, gives the letter accuracy of the 14B's configurations and harnesses next to the strongest oracle constant, the constant label of the first study's analysis (below). The two clean-start harnesses are not detectably above the constant: accuracy minus the constant is $+0.001$ and $+0.006$ on test and -0.018 and -0.001 on replication, and all four unadjusted intervals contain zero. The intervals do not show equality either: on test their upper bounds are $+0.017$ and $+0.020$, so an advantage of up to about two points is not excluded. AUTHORS-1500 and BUDGET-512-1500 are below it by -0.336 and -0.135 on test.

How the harnesses reach the constant. The pooled accuracy adds up three different behaviours, one for each task, and they differ by partition (Table 5 for the test partition and Table 50 in Appendix N for both, computed after the outputs existed); each partition has its own constant label, chosen on it. On the mortality task the strongest constant is "Yes" (accuracy 0.840 on test, 0.823 on replication), both clean-start harnesses answer "Yes" in 94.6% and 98.2% of their parsed answers on test and in 94.9% and 98.6% on replication, and their accuracy, 0.847 and 0.837 on test and 0.830 and 0.828 on replication, is the constant's. On the participant-comparison task they exceed the constant in both partitions (0.602 and 0.586 against 0.515 on test; 0.565 and 0.580 against 0.510 on replication), but they lean on one label to different degrees: the constant participant is named in 19.7% and 58.3% of their answers on test and in 83.6% and 47.1% on replication. On the age-band task they fall below the constant in both partitions (0.202 and 0.232 against 0.250 on test; 0.190 and 0.229 against 0.270 on replication), and the constant band is 33.4% and 30.0% of their answers on test and 0.5% and 15.4% on replication. The pooled accuracy at the constant is therefore a net result: equal to it on the mortality task, which carries 40% of the rows and which they answer "Yes" almost throughout, above it on the comparison task and below it on the age-band task. AUTHORS-1500 of the same model answers "Yes" in only 16.2% of its parsed answers on the mortality task on test, and the tuned 1.7B model's AUTHORS-1500 is above the constant on all three tasks in both partitions. The harness texts tell the model not to infer survival beyond the snapshot data, or to conclude that survival cannot be determined when follow-up data are missing (Appendix L), and the harnesses still answer "Yes" almost throughout on the mortality task; we did not read their outputs, so why they do is not examined (Appendix S).

Table 4: Qwen3-14B on the LongevityBench letters: accuracy with its unadjusted interval, parse rate, and the paired difference between accuracy and the strongest oracle constant (the constant label). The two broken-start harnesses are the same text and answer nothing. Computed after the outputs existed.

Configuration	Accuracy	Parse rate	Accuracy minus the strongest constant
Replication (1,990 rows, 399 components)			
THINK-1536	0.239 [0.223, 0.255]	0.345	−0.300 [−0.324, −0.276]
AUTHORS-1500	0.222 [0.205, 0.239]	0.567	−0.318 [−0.346, −0.289]
BUDGET-512-1500	0.416 [0.397, 0.434]	0.728	−0.124 [−0.149, −0.098]
GEPA-CLEAN-s0	0.521 [0.504, 0.539]	0.992	−0.018 [−0.038, +0.003]
GEPA-CLEAN-s1	0.539 [0.518, 0.559]	0.994	−0.001 [−0.023, +0.022]
GEPA-BROKEN-s0	0.000 [0.000, 0.000]	0.000	−0.539 [−0.559, −0.519]
GEPA-BROKEN-s1	0.000 [0.000, 0.000]	0.000	−0.539 [−0.559, −0.519]
Strongest oracle constant (the label)	0.539 [0.519, 0.559]	–	–
Test (3,942 rows, 789 components)			
THINK-1536	0.235 [0.223, 0.246]	0.341	−0.304 [−0.323, −0.286]
AUTHORS-1500	0.203 [0.192, 0.215]	0.564	−0.336 [−0.355, −0.317]
BUDGET-512-1500	0.404 [0.391, 0.416]	0.705	−0.135 [−0.154, −0.116]
GEPA-CLEAN-s0	0.540 [0.527, 0.552]	0.997	+0.001 [−0.016, +0.017]
GEPA-CLEAN-s1	0.545 [0.531, 0.559]	0.995	+0.006 [−0.009, +0.020]
GEPA-BROKEN-s0	0.000 [0.000, 0.000]	0.000	−0.539 [−0.551, −0.526]
GEPA-BROKEN-s1	0.000 [0.000, 0.000]	0.000	−0.539 [−0.551, −0.526]
Strongest oracle constant (the label)	0.539 [0.526, 0.551]	–	–

Table 5: Per task, test partition: accuracy, the task's strongest oracle constant (a label, with its accuracy), the share of parsed answers equal to that label, and the share of correct answers that are that label. Computed after the outputs existed; the replication partition and the rerun are in Appendix N.

Task and constant	Model and configuration	Accuracy	Answers equal to the constant label	Correct answers that are the label
LB-0030, 60-75 (0.250)	Qwen3-14B GEPA-CLEAN-s0	0.202	526 of 1,576 (33.4%)	46.7%
	Qwen3-14B GEPA-CLEAN-s1	0.232	469 of 1,563 (30.0%)	37.2%
	Qwen3-14B AUTHORS-1500	0.058	113 of 432 (26.2%)	28.3%
	Qwen3-14B BUDGET-512-1500	0.126	274 of 776 (35.3%)	43.2%
	Qwen3-1.7B-Longevity AUTHORS-1500	0.344	132 of 1,577 (8.4%)	8.7%
LB-0034, Participant A (0.515)	Qwen3-14B GEPA-CLEAN-s0	0.602	154 of 783 (19.7%)	25.9%
	Qwen3-14B GEPA-CLEAN-s1	0.586	458 of 785 (58.3%)	58.4%
	Qwen3-14B AUTHORS-1500	0.445	193 of 544 (35.5%)	40.2%
	Qwen3-14B BUDGET-512-1500	0.383	174 of 490 (35.5%)	39.7%
	Qwen3-1.7B-Longevity AUTHORS-1500	0.717	267 of 788 (33.9%)	39.8%
LB-0042, Yes (0.840)	Qwen3-14B GEPA-CLEAN-s0	0.847	1,486 of 1,570 (94.6%)	96.3%
	Qwen3-14B GEPA-CLEAN-s1	0.837	1,545 of 1,574 (98.2%)	99.0%
	Qwen3-14B AUTHORS-1500	0.227	202 of 1,248 (16.2%)	49.2%
	Qwen3-14B BUDGET-512-1500	0.691	983 of 1,512 (65.0%)	84.0%
	Qwen3-1.7B-Longevity AUTHORS-1500	0.878	1,353 of 1,577 (85.8%)	89.7%

The first study against the trivial baselines. Appendix O gives the first study's accuracy next to the trivial baselines, computed after the outputs existed. On the letters only Qwen3-1.7B-Longevity beat the strongest constant, which is the oracle constant label, in every configuration that parses (for AUTHORS-1500 $+0.093$ $[+0.074, +0.112]$); SmolLM3-3B and Intern-S1-mini were below it in every configuration, because most of their calls end unparsed. The model that is best on the letters was the weakest on MMLU-Pro and barely above a constant answer on GSM8K, and Section 4.6 compares it with its base model.

4.4 Budget forcing (H3)

In the first study BUDGET-512-1500 lowers truncation and raises parsing in the six cells where THINK-512 truncates a large share, with D_T from $+0.260$ (SmolLM3-3B, GSM8K) to $+0.979$ (Intern-S1-mini, letters), D_P from $+0.151$ to $+0.662$, and both adjusted lower bounds above the minimal difference. In the Qwen3-1.7B-Longevity letter and GSM8K cells THINK-512 does not truncate (adjusted upper bounds $+0.001$ and 0.000), so there is *no room* for a drop of 0.05. In Qwen3-1.7B-Longevity on MMLU-Pro $D_T = D_P = +0.060$ with adjusted interval $[+0.038, +0.085]$, which contains 0.05, so the verdict is *inconclusive*. The parse gain is not the whole story for SmolLM3-3B: after the forced end of thinking it

Table 6: H4: the 16 (defect, model) pairs of the first study and the 5 of the second (Qwen3-14B), with the letters of Section 3.4; a star marks a claim whose quantity is fixed by an identity of the design; its verdict says little beyond that mechanism, although it still depends on the intervals. The tied output layer (E) exists for Qwen3-1.7B-Longevity only, the one model with two served variants. The bound is the largest distortion the mechanism can cause on S, computed without accuracy; a pair is evaluable only if it exceeds 0.05. D_S, D_R and $\Gamma = D_R - D_S$ are point estimates of the distortion of the start's accuracy (for the broken start, of the gain over the start, which equals the accuracy of THINK-1536 minus that of THINK-64); intervals are in Figure 2 (Appendix I) and Appendix B. A dagger marks the claim of the skewed split in SmolLM3-3B, which fails by the general rule of the preregistration (Section 5).

Defect	Model	Bound	D_S	D_R	Γ	Claim
A reference answer as assistant turn	Qwen3-1.7B	1.000	−0.052	−0.030	+0.022	inconclusive
	SmolLM3	1.000	+0.097	+0.119	+0.022	inconclusive
	Intern-S1	1.000	+0.142	+0.164	+0.022	inconclusive
	Qwen3-14B	1.000	+0.430	+0.404	−0.026	inconclusive
B1 gold fallback	Qwen3-1.7B	0.205	−0.145	−0.148	−0.003	holds
	SmolLM3	0.083	+0.001	+0.007	+0.006	fails
	Intern-S1	0.057	−0.012	−0.019	−0.007	fails
	Qwen3-14B	0.149	−0.097	−0.090	+0.007	holds
B2 unparsed scored as A	Qwen3-1.7B	0.000	not evaluable			–
	SmolLM3	0.208	+0.208	+0.223	+0.015	holds⋆
	Intern-S1	0.289	+0.289	+0.292	+0.003	holds⋆
	Qwen3-14B	0.237	+0.237	+0.247	+0.009	holds⋆
C skewed split	Qwen3-1.7B	by rule	+0.072	+0.101	+0.030	inconclusive
	SmolLM3	by rule	−0.001	+0.005	+0.006	fails†
	Intern-S1	by rule	+0.003	+0.004	+0.001	inconclusive
	Qwen3-14B	by rule	+0.034	+0.082	+0.048	inconclusive
D broken start	Qwen3-1.7B	1.000	+0.638	+0.635	−0.003	holds⋆
	SmolLM3	0.352	+0.164	+0.157	−0.007	holds⋆
	Intern-S1	0.189	+0.124	+0.125	+0.001	holds⋆
	Qwen3-14B	0.345	+0.239	+0.235	−0.005	holds⋆
E tied vs untied output layer	Qwen3-1.7B	0.139	+0.005	+0.014	+0.009	fails

re-opens a turn in 4,384 of its 7,216 test calls, and BUDGET-512-1500 still ends by itself without an answer on 0.424 of its letter rows. For Qwen3-14B BUDGET-512-1500 lowers truncation and raises parsing in all three cells, with D_T from +0.531 (GSM8K) to +0.974 (letters), D_P from +0.531 to +0.782 (MMLU-Pro) and both adjusted lower bounds above the minimal difference; THINK-512 truncates 1.000 of its letter calls. Over both studies the verdict is *holds* in 9 of 12 cells.

What H3 isolates. BUDGET-512-1500 differs from THINK-512 in two ways at once: vLLM ends thinking after 512 tokens, and the total allowance is 2,012 tokens against 512. H3 therefore supports the combined change and not forced termination alone. As a partial and descriptive comparison, THINK-1536, with an allowance of 1,536 tokens and no forcing, still ends at the token limit in 0.066 (SmolLM3-3B, GSM8K) to 0.800 (Intern-S1-mini, letters) of its calls in the six cells where THINK-512 truncates a large share, against 0.021 to 0.379 for BUDGET-512-1500, which truncates less in all six. The two differ in allowance as well (2,012 against 1,536 tokens), so the comparison does not separate forcing from allowance.

Is the truncation an artifact of zero temperature? In the two models of the first study with a seeded arm, substantial truncation persists under sampling at the model cards' values (descriptive; Appendix Q).

4.5 Whether replication reveals known defects (H4)

In the first study the claim holds in six pairs: the broken start in all three models, the unparsed-as-A scoring in SmolLM3-3B and Intern-S1-mini, and the gold fallback in Qwen3-1.7B-Longevity. Five of the six carry a star in Table 6, because the quantity is fixed by an identity of the design: the distortion of the unparsed-as-A

scoring equals the share the start leaves unparsed with gold A, which is also its evaluability bound, and the distortion under the broken start is the difference between the accuracies of two starts, THINK-1536 and THINK-64. The verdict still depends on the intervals: for the broken start the preregistration expected the claim to hold only about 0.50 of the time with no true difference, since it rests on Γ being within about 0.01 of zero (Appendix D). The gold fallback in Qwen3-1.7B-Longevity is the informative pair. Under the broken start the reported gain over the start is inflated by $D_S = +0.638$, $+0.164$ and $+0.124$ (Qwen3-1.7B-Longevity, SmolLM3-3B, Intern-S1-mini) and by almost the same amount on R ($\Gamma = -0.003$, -0.007, $+0.001$). Scoring unparsed answers as A inflates the start's accuracy by $+0.208$ and $+0.289$, exactly the share the start leaves unparsed with gold A, and the gold fallback moves Qwen3-1.7B-Longevity's start by -0.145. In all six the distortion is large and reappears on R, so a replication would have reproduced the distorted report instead of revealing it.

The claim fails in four pairs: the gold fallback in SmolLM3-3B and Intern-S1-mini, whose starts parse too few of the rows whose gold the fallback changes (bounds 0.083 and 0.057, distortions $+0.001$ and -0.012); the skewed split in SmolLM3-3B; and the tied output layer in Qwen3-1.7B-Longevity, where $D_S = +0.005$ and $D_R = +0.014$ both lie inside the band. Five pairs are inconclusive, the reference-answer defect in all three models and the skewed split in Qwen3-1.7B-Longevity and Intern-S1-mini; for the reference-answer defect in Intern-S1-mini both $D_S = +0.142$ and $D_R = +0.164$ clear the band, but the interval of Γ, $[-0.065, +0.108]$, is wide because each half has only 81 components. The unparsed-as-A scoring in Qwen3-1.7B-Longevity is not evaluable, because its start leaves no row unparsed (bound 0.000). These counts follow the general rule of the preregistration for D_R, which the pinned code implements (Section 5); under the narrower reading that grades only the two sub-cases listed after the general rule, the skewed split in SmolLM3-3B is inconclusive and the counts are 6 holds, 3 fails and 6 inconclusive (Appendix C).

In the second study the claim holds for B1, B2 and D and is inconclusive for A and C. Two of the three carry a star: B2's distortion, $D_S = +0.237$, equals its evaluability bound 0.237, the share the start leaves unparsed with gold A, and D's is the difference of two starts' accuracies, $D_S = +0.239$; their verdicts still depend on the intervals, as in the first study. B1 is the informative pair: the start parses 0.345 of the letter rows, the gold fallback changes the gold of a share 0.205 of the rows, and the distortion, $D_S = -0.097$ on the selection sample and $D_R = -0.090$ on the check sample, is the same on both. The reference-answer defect has $D_S = +0.430$ and $D_R = +0.404$, both clear of the band, but Γ, $[-0.151, +0.096]$, is wide; the skewed split has $D_S = +0.034$ and $D_R = +0.082$ with all three components open. Over both studies the claim holds in 9 of the 20 evaluable pairs, 7 of them pairs whose quantity is fixed by an identity of the design (the stars).

The audit checks. Appendix P summarises the checks of Section 3.4: their detections are descriptive and largely true by construction, and what is empirical is their false alarms on the 7 clean instances: IR4 fires on 7, SCF1 on 4, SCF2 on 4, TB2 on 3, SCF3 on 2 and the gain rule of PR on 1.

4.6 What the longevity fine-tuning added and cost

The second study ran the base and the tuned version of Qwen3-1.7B on the same items, configurations and schedule, as a descriptive comparison with unadjusted intervals, no verdict, and a limit on its reading fixed in advance. The two models differ in every weight, in the data the tuned one was fine-tuned on and in the handling of the output layer, so a difference shows the effect of the whole fine-tuning on a benchmark. It is consistent with, but does not isolate, the effect of supervised training on the same task formats from other NHANES survey waves; it does not show general domain knowledge, and it is not instance-level leakage if the wave split of the model's authors holds, which we take from their audit and do not verify. A difference on MMLU-Pro or GSM8K is not evidence about aging biology.

On the headline cell, AUTHORS-1500, the base model scored below the tuned model by more than five points on the letters (Table 7; tuned minus base $+0.414$ $[+0.395, +0.433]$ on test, 0.631 against 0.217), and above it on MMLU-Pro (-0.164 $[-0.203, -0.125]$; 0.220 against 0.384) and on GSM8K (-0.634 $[-0.677, -0.590]$; 0.066 against 0.700); on LB-0038 the exact-match rates are within five points of each other. Of the 12 headline cells (four pools, replication, test and pooled), 3 are classed base below tuned, 6 base above tuned and 3 about equal, and the class of a pool is the same in the three partitions. A cell is classed from the unadjusted interval of tuned minus base against a band of ± 0.05: base below tuned if the whole interval

Table 7: The headline cell, AUTHORS-1500, of the two 1.7B arms: accuracy of the base and of the tuned model, the difference with its unadjusted interval, its class (a band of five points), and the strongest oracle trivial constant of the cell (for LB-0038 the exact-match rate of the oracle median). Descriptive; the full report has every configuration, the seeded arm and the baselines.

Partition	Pool	Base	Tuned	Tuned minus base	Class	Strongest oracle constant
Replication	LB letters	0.220 [0.203, 0.238]	0.633 [0.614, 0.652]	+0.413 [+0.386, +0.440]	base below tuned	0.539
	MMLU-Pro	0.363 [0.316, 0.411]	0.185 [0.148, 0.226]	−0.178 [−0.231, −0.125]	base above tuned	0.133
	GSM8K	0.689 [0.633, 0.746]	0.061 [0.034, 0.091]	−0.629 [−0.689, −0.564]	base above tuned	0.045
	LB-0038 (exact match)	0.000 [0.000, 0.000]	0.030 [0.019, 0.043]	+0.030 [+0.019, +0.043]	about equal	0.013
Test	LB letters	0.217 [0.205, 0.229]	0.631 [0.617, 0.645]	+0.414 [+0.395, +0.433]	base below tuned	0.539
	MMLU-Pro	0.384 [0.350, 0.418]	0.220 [0.191, 0.249]	−0.164 [−0.203, −0.125]	base above tuned	0.119
	GSM8K	0.700 [0.660, 0.738]	0.066 [0.046, 0.089]	−0.634 [−0.677, −0.590]	base above tuned	0.034
	LB-0038 (exact match)	0.001 [0.000, 0.002]	0.031 [0.023, 0.040]	+0.030 [+0.022, +0.039]	about equal	0.013
Pooled	LB letters	0.218 [0.208, 0.228]	0.632 [0.620, 0.643]	+0.414 [+0.398, +0.429]	base below tuned	0.536
	MMLU-Pro	0.377 [0.350, 0.405]	0.209 [0.185, 0.232]	−0.169 [−0.200, −0.137]	base above tuned	0.124
	GSM8K	0.697 [0.665, 0.728]	0.064 [0.048, 0.082]	−0.632 [−0.668, −0.595]	base above tuned	0.033
	LB-0038 (exact match)	0.000 [0.000, 0.001]	0.031 [0.024, 0.038]	+0.030 [+0.024, +0.038]	about equal	0.013

lies above +0.05, base above tuned if it lies below −0.05, about equal only if the whole interval lies within the band, and unresolved if it crosses a boundary (0 of the 12 cells). The pooled cells are the replication and test rows together and not a third independent partition, so the agreement of the three partitions is not three independent confirmations. The classes are a descriptive reading of unadjusted intervals, not findings of a test corrected for the number of cells read. The fine-tuning raised letter accuracy and lowered accuracy on both general pools, most on GSM8K. In the Cell paper, its methods and the supplementary tables and figures of the copy we read, we found no evaluation of the longevity models on a general benchmark, so as far as we know this is the first such measurement for this checkpoint; the paper does evaluate the base architectures on the longevity tasks, and we take its statements about training from the authors.

4.7 Whether the conclusions change across GPUs

The tuned 1.7B model ran in both studies, on an RTX 5090 in the first and on an RTX PRO 6000 in the second, with the same environment, the same batch-invariant server, the same request bodies and the same items. Its outputs are not the same. Token IDs differ in 58.2% of the replication grid's calls and 57.6% of the test grid's, and in 67.4% and 67.2% of the two seeded arms' calls (Table 49, Appendix M); the request body never differs, and the finish reason differs in 72 of the 36,080 test-grid calls. The conclusions are the same. Recomputed on the rerun's outputs with the first study's rules, seeds and family sizes, 9 of the first study's 19 primary endpoints for this model, the endpoints of the grid's H1–H3, give the same verdict in 9 cases on the test partition and in 9 on the replication partition and the same replication label in 9, and none flips. The 4 GEPA endpoints and the 6 H4 pairs of this model were not rerun. The accuracy of every configuration in the three pools differs by at most 0.010 between the two runs, and 2 of the 54 unadjusted intervals of the difference exclude zero (Appendix M). The comparison is descriptive and covers one model, 9 grid endpoints and two GPUs of one architecture; it says that these conclusions did not depend on the token-level agreement, which numerical differences between hardware are known to break (Yuan et al., 2025; He & Thinking Machines Lab, 2025).

4.8 Qwen3-14B: non-commitment and refusal

Qwen3-14B often ends a call without a readable answer, and the frozen scorer counts that as wrong; this section reports how, descriptively and with the scorer unchanged (Table 45, Appendix M). Across the five grid configurations 23.8% of the analysed test calls are non-committed, ended by themselves without an answer, and 22.3% of the calls fall in the refusal class, a stated inability or no answer (23.7% and 22.3% on replication). The non-commitment sits where a call can end by itself. THINK-64 and THINK-512 are truncated almost throughout, so on the letters their non-commitment is 0.0% and 0.0%, whereas AUTHORS-1500 is non-committed in 43.6% of its letter calls, and 36.0% of its letter calls are refusals. LB-0038, the age task, stands out: the share of non-committed calls is 100.0% for AUTHORS-1500, 99.6% for THINK-1536 and 99.9% for

BUDGET-512-1500, so the model mostly declines to give an age. No example was read; the classes are those of a format audit fixed before the study, and the frozen scorer stays the primary reading.

5 Deviations and addenda

The preregistration allows a change only if it alters no request bytes, parsing, scoring or call order, and then only with a written note and a dry run that reproduces the recorded render hashes; any other change needs a new registration. After the freeze, the requests, the scoring and the primary criteria stayed as preregistered, with two exceptions that were not mechanical. We waived one stopping rule, once and for one stage: the check on Intern-S1-mini's thinking switch had stopped the run that sends the reference answer, and we let that check record its failures and the run continue. And we settled one descriptive label after seeing the outputs: while auditing this paper we found that the two listed cases of the claim rule for the check sample are incomplete, so its general rule governs, and under it the claim for the skewed split in SmolLM3-3B fails where the narrower reading leaves it inconclusive (Section 4 gives both totals for H4).

The second study's record is in the same appendix, and four of its items bear on a result. The physical GPU changed during the study and the preregistered recheck found no difference, and the interrupted harness evaluation was rerun from scratch (Section 3.6). A pinned file that lists the replication rows left out of every analysis because their reference answers could be recovered was missing on the instance; the checkpoint stopped when it first ran, before any run was read, the file was copied by hand and its hash checked, and the checkpoint was repeated. And the non-commitment of Qwen3-14B was decided, before the study's outputs were read, to be reported descriptively with no sensitivity reading of the unparsed answers and the frozen scorer primary (Section 4.8). The platform shutdown, the clone of the instance, the unavailability of an external drive and the copy of the outputs are recorded in the appendix and change no result.

Table 29 in Appendix C lists the deviations that bear on a result, and the appendix lists every exception, known deviation, pre-freeze decision and addendum with its date, reason, rule and the outputs seen at the time. No call touched a replication or test item before the freeze, and we saw only counts, and no accuracy, score or verdict of the main study, until every stage had run and the checkpoint's record had been committed.

Everything computed after the outputs existed is labelled where it appears. This paper's tables and figures only format aggregates committed in our repository, and its audit recomputes every verdict from its interval and agrees with the report on all of them (Appendix C).

6 Discussion and limitations

What the results say. The gain of the thinking-off configuration over THINK-512 comes mostly from rows that THINK-512 left unparsed in 7 of 12 cells, so a search that starts from a configuration like this can find "repairs" that a thinking-off baseline already provides. A thinking budget in the serving engine lowers truncation as well (H3), which is another way of making answers readable.

Baseline adequacy depends on the model. In the longevity-tuned 1.7B model the thinking-off setting was not meaningfully worse than any of four GEPA harnesses, and it beat the strongest constant label on the letters. In Qwen3-14B it was below that constant, and the clean-start harnesses were more accurate than it by 30 to 34 points. In these models, the thinking-off baseline beat the strongest constant label only for the model tuned on the task format, and GEPA beat the baseline only where the baseline was below that constant; the study cannot separate tuning from size. The thinking-off setting's letter accuracy is 0.631 for the tuned 1.7B model in the second study (0.632 in the first), against 0.217 for that model's own base, 0.203 for the 14B and, in the first study, 0.129 for SmolLM3-3B and 0.314 for Intern-S1-mini, while the strongest constant label scores 0.539. GEPA ran around one model in each study, so the comparison between harness searches rests on two models that differ in tuning as well as in size.

A harness gain has to be read against trivial baselines. The 14B's harnesses gain 30 to 34 points over the thinking-off setting and are not detectably above the strongest constant label: on the mortality task they match it by answering "Yes" almost throughout, on the participant-comparison task they exceed it, and on the age-band task they fall below it; the same gain is large against one baseline and not detectable against

another. In the first study only Qwen3-1.7B-Longevity beat the strongest constant label on the letters, and in the second study none of the 14B's configurations was detectably above it, whatever its parse rate.

Where prospective replication confirmed the distortion (9 of 20 evaluable pairs, 7 of them pairs whose quantity is fixed by an identity of the design), the distortion was large and, within five points, the same on the check sample as on the selection sample. Some systematic distortions therefore persist across samples, and a replication run through the same pipeline reproduces them instead of revealing them. In the other pairs the claim fails (a descriptive label) or is inconclusive, and a defect whose claim fails need not be harmless. Of the audit checks, the token reconciliation fired on no clean instance. The others also fired on clean pipelines: the gain rule of prospective replication on 1 of the 7 clean instances, the trivial baselines (TB2) on 3, the simple-configuration-family rules on up to 4 and the input-reading rule IR4 on all 7, so IR4 does not separate defective from clean pipelines. The detections themselves are largely true by construction.

The grid verdicts of H1–H3 did not depend on the GPU for the one model that ran on both (9 of its 19 endpoints of the first study; its GEPA endpoints and H4 pairs were not rerun), although most of its token sequences did. The longevity fine-tuning, read relative to the base model, bought letter accuracy and cost general accuracy, with the limits stated in Section 4.6.

Self-improving loops. We studied one round of selection and no recursive loop. Each round of a loop that edits its own harness is a selection of this kind, so its gains can likewise be answers made readable, and the cheap checks of this study can be run after every round: reporting truncated and non-committed answers apart, a thinking-off baseline, constant-label baselines, reading the rendered inputs and reconciling prompt tokens per call. Whether early visibility gains turn into quality later in a loop, and whether a loop can leave a start like the one that GEPA could not repair at 14B, this study does not answer.

Limitations. The studies cover four models in the confirmatory design (the three small models and Qwen3-14B, a single larger model from the family of the tuned 1.7B model) and the base Qwen3-1.7B descriptively, three pools with LB-0038 reported apart, and six defects chosen from one line of work, so H4 says how replication behaves for these defects and not for the space of defects. "Prospective" has limits: earlier work had called most LongevityBench rows with other configurations, callers and one model variant, and the rows the case study had held out are excluded. Several endpoints carry a star in Table 3 and Table 6: their quantity is fixed by an identity of the design or by the observed parse pattern, so their verdicts say little beyond the mechanism, although they still depend on the intervals. For the skewed split (C) the frozen analysis resamples the defective and the clean instance independently although they reuse the same participant components, so their covariance is ignored and the intervals may be too wide or too narrow; the four claims for C (one fails, three are inconclusive) are the ones that could depend on it (Appendix I). Power is mixed. The H2 endpoint of Intern-S1-mini on GSM8K was expected to be inconclusive and is not (an informal expectation, Appendix K), H4 has no planned power, and its error control covers "holds" only. The GEPA arms searched prompt text only, with one proposer and four runs in each study, so their results are about that space; the broken start of the second study could not be repaired within it. The two 1.7B arms are descriptive: the tuned and the base model differ in every weight and in training data, the intervals are unadjusted, and the classes are not corrected for the number of cells read. The extended scorer and the role-break rule were fixed after the pilot's parse audit. Unparsed counts as wrong by rule, so accuracy mixes parsing and correctness, and an unparsed answer may hide a correct answer, a wrong one or none; the study measures differences between configurations as served and not what the models know. The batch-invariant server differs from the standard one call for call, and each study ran on one GPU type, the first on an RTX 5090 and the second on an RTX PRO 6000; the second study's test grid for Qwen3-14B ran on one physical GPU and its other test runs on another of the same type, after a recheck that found no difference. Our intervals resample participant components (items in the two general pools) at a frozen harness, zero temperature and fixed prompts, so they cover item sampling and not prompt or judge variance (Messing, 2026). The freeze guards the rules and not model-specific tuning, which a preregistration for the next model would (Thomas et al., 2026), and the second preregistration was written after the first study's results were known, though before any call on its own partitions. Zero temperature may turn some thinking into repetition; for the two models with a seeded arm in the first study substantial truncation persists under

sampling (Appendix Q), and for SmolLM3-3B and the 14B it is untested. We used Claude (Claude Code and claude.ai) and Codex for code, drafting and review, and we verified all analyses and claims.

Broader impact. This paper contributes methodology for evaluating harness search and the pipelines that report it: a preregistered design, audit checks and a record of what each result rests on. It builds no new capability, since the models and benchmarks are public and the benchmark behind the longevity tasks derives from public survey data. The accuracies are those of small and mid-sized models on research tasks and do not support any use of these models for medical advice or clinical decisions. One misreading to avoid is that a thinking-off baseline is always enough, and the study shows cells in which it is not.

7 Conclusion

In two designs fixed in advance, most of the gain of a thinking-off configuration over a capped thinking configuration came from rows the capped configuration left unparsed, in 7 of 12 cells. The thinking-off configuration was non-inferior to the rescue in most cells of three models of 1.7B to 8B parameters, and to four GEPA-selected harnesses at 1.7B; at 14B the clean-start harnesses beat it by 30 to 34 points but were not detectably above the strongest constant label, which they matched on the mortality task by answering "Yes" almost throughout. In 9 of the 20 evaluable (defect, model) pairs, the defect's distortion of the reported number was larger than five points on the selection sample and reappeared, within five points, on the check sample; 7 of these pairs are ones whose quantity is fixed by an identity of the design. Reporting truncated and non-committed answers apart and running a thinking-off baseline and constant-label baselines cost little and changed what the gains meant. Of the audit checks, the token reconciliation fired on no clean instance, and the others also fired on clean pipelines, from 1 of the 7 clean instances (the gain rule of prospective replication) to all 7 (IR4, the rule that reads the rendered inputs). Rerun on a second GPU, the tuned 1.7B model gave the same verdicts as in the first study on the 9 grid endpoints that were rerun (its GEPA endpoints and H4 pairs were not), although most of its token sequences changed.

Code and data availability

The study's repository is not public at present. The paths, commits, tags and SHA256 values that this paper names belong to files in our repository, and we keep them so that a later public export can be checked against them (Appendix E). The run logs hold benchmark item text and model outputs and will not be released.

This arXiv version carries ancillary files and a short README that describes them: the frozen preregistration of each study, the addenda file of each, and the aggregate analysis reports of both. The preregistrations and the reports are attached byte for byte. The addenda files are attached as they stood at the posting, with the row keys of any rows that they name withheld, because a row key identifies a benchmark row. The freeze times are recorded only in our private repository, in the hash file that each freeze commit adds, and they are not independently timestamped. The SHA256 of each frozen file (Appendix E) therefore lets a reader check that the ancillary copy is the file that was frozen, but not when it was frozen. A filtered public export of the rest is planned and will be linked in a later version of this paper: the code and its tests and the hash manifests, with no item-level or row-level content derived from the held-out rows.

The benchmarks and models are used under the terms their sources give, read on 2026-10-03 (Qwen3-14B: 2026-10-24). LongevityBench (dataset card `insilicomedicine/longebench`) is released under CC-BY-NC-4.0; its card states that the underlying source datasets (NHANES, GTEx, GEO, Olink, OpenGenes, SynergyAge) retain their original licences and access conditions, with which users must comply. MMLU-Pro (`TIGER-Lab/MMLU-Pro`) and GSM8K (`openai/gsm8k`) are under the MIT licence. Qwen3-1.7B-Longevity (`insilicomedicine/Qwen3-1.7B-Longevity`) is released under CC-BY-ND-4.0, and its base model Qwen3-1.7B is distributed under Apache-2.0; Qwen3-14B (`Qwen/Qwen3-14B`) is under Apache-2.0; SmolLM3-3B (`HuggingFaceTB/SmolLM3-3B`) and Intern-S1-mini (`internlm/Intern-S1-mini`) are under Apache-2.0. The study uses the weights unmodified and redistributes neither benchmark item text nor model weights.

Author Contributions

Bowen Xu: Conceptualization, Methodology, Software, Investigation, Formal analysis, Visualization, Writing – original draft, Writing – review & editing, Project administration. Xu conceived and led the study.

Boyu Chen: Methodology, Supervision, Writing – review & editing.

References

Lakshya A Agrawal, Shangyin Tan, Dilara Soylu, Noah Ziems, Rishi Khare, Krista Opsahl-Ong, Arnav Singhvi, Herumb Shandilya, Michael J Ryan, Meng Jiang, Christopher Potts, Koushik Sen, Alexandros G. Dimakis, Ion Stoica, Dan Klein, Matei Zaharia, and Omar Khattab. GEPA: Reflective prompt evolution can outperform reinforcement learning. *arXiv:2507.19457*, 2025. ICLR 2026 (Oral).

Lei Bai et al. Intern-S1: A scientific multimodal foundation model. *arXiv:2508.15763*, 2025.

Elie Bakouch et al. SmolLM3: smol, multilingual, long-context reasoner. Hugging Face blog, 2025. URL `https://huggingface.co/blog/smollm3`.

Rohith Reddy Bellibatlu, Zichong Wang, and Wenbin Zhang. Do Agent Benchmarks Do What They Say? An Executable-Contract Audit of Tool-Using Agent Environments. *arXiv:2609.37315*, 2026.

Aymene Berriche, Cathrine Shalby, Mohannad Alhanahnah, and Yazan Boshmaf. Benchmark Scores Are Pipeline-Dependent: A Reliability Audit of Cybersecurity LLM Benchmarks. *arXiv:2609.08765*, 2026.

Francesca Carlon, Vincent Ginis, and Andres Algaba. Shorter Reasoning, Earlier Answers? An Evaluation of Reasoning Interfaces. *arXiv:2608.03401*, 2026.

Yuan Chang and Xiaoqi Chen. Naive Prompt Optimization: Rethinking the Need for Complex Prompt Search. *arXiv:2608.27266*, 2026.

Mingguang Chen, Bo Qu, and Licheng Wang. The Calibration Floor: Format Repair Can Masquerade as Self-Correction at Small-to-Mid Scale. *arXiv:2608.04355*, 2026.

Karl Cobbe, Vineet Kosaraju, Mohammad Bavarian, Mark Chen, Heewoo Jun, Lukasz Kaiser, Matthias Plappert, Jerry Tworek, Jacob Hilton, Reiichiro Nakano, Christopher Hesse, and John Schulman. Training verifiers to solve math word problems. *arXiv:2110.14168*, 2021.

Abhinav Dutta, Sanjeev Krishnan, Nipun Kwatra, and Ramachandran Ramjee. Accuracy is Not All You Need. *arXiv:2407.09141*, 2024.

Irena Gao, Percy Liang, and Carlos Guestrin. Model Equality Testing: Which Model Is This API Serving? *arXiv:2410.20247*, 2025. ICLR 2025.

Chaitanya Vilas Garware and Sharif Noor Zisad. When the Ruler is Broken: Parsing-Induced Suppression in LLM-Based Security Log Evaluation. *arXiv:2605.07293*, 2026.

Yonatan Gideoni, Sebastian Risi, and Yarin Gal. Simple Baselines are Competitive with Code Evolution. *arXiv:2602.16805*, 2026.

Ankit Goyal and Jaideep Ray. Mind the Cap: Output-Budget Regimes Change the Measured Multilingual Reasoning Gap. *arXiv:2608.04160*, 2026.

David Gringras. Safety Under Scaffolding: How Evaluation Conditions Shape Measured Safety. *arXiv:2603.10044*, 2026.

Horace He and Thinking Machines Lab. Defeating nondeterminism in LLM inference. Thinking Machines Lab: Connectionism (blog), 2025. URL `https://thinkingmachines.ai/blog/defeating-nondeterminism-in-llm-inference/`.

hmellor (vLLM project). Fix weight tying. GitHub pull request 51665, vllm-project/vllm; merged 2026-08-20, 2026. URL `https://github.com/vllm-project/vllm/pull/51665`.

JasonKeyiL (vLLM project). [Bugfix] Fix inverted condition causing thinking_token_budget to be silently ignored. GitHub pull request 41674, vllm-project/vllm; merged 2026-05-15, 2026. URL `https://github.com/vllm-project/vllm/pull/41674`.

Daniel Kaiser, Arnoldo Frigessi, Ali Ramezani-Kebrya, and Benjamin Ricaud. Beyond Accuracy: Decomposing the Reasoning Efficiency of LLMs. *arXiv:2602.09805*, 2026.

Sayash Kapoor, Benedikt Stroebl, Zachary S. Siegel, Nitya Nadgir, and Arvind Narayanan. AI Agents That Matter. *arXiv:2407.01502*, 2024.

Omar Khattab, Arnav Singhvi, Paridhi Maheshwari, Zhiyuan Zhang, Keshav Santhanam, Sri Vardhamanan, Saiful Haq, Ashutosh Sharma, Thomas T. Joshi, Hanna Moazam, Heather Miller, Matei Zaharia, and Christopher Potts. DSPy: Compiling declarative language model calls into self-improving pipelines. *arXiv:2310.03714*, 2023.

Jonas Kübler, Kailash Budhathoki, Matthäus Kleindessner, Xiong Zhou, Junming Yin, Ashish Khetan, and George Karypis. When LLMs get significantly worse: A statistical approach to detect model degradations. *arXiv:2602.10144*, 2026. ICLR 2026.

Woosuk Kwon, Zhuohan Li, Siyuan Zhuang, Ying Sheng, Lianmin Zheng, Cody Hao Yu, Joseph E. Gonzalez, Hao Zhang, and Ion Stoica. Efficient memory management for large language model serving with PagedAttention. In *Proceedings of the 29th Symposium on Operating Systems Principles*, pp. 611–626, 2023.

Janghoon Lee. Repair, Not Improvement: Decomposing Constrained Decoding in Tool-Call Abstention. *arXiv:2608.13959*, 2026.

Jungseob Lee, Seongtae Hong, Seungjun Lee, Jaehyung Seo, Junyoung Son, Sugyeong Eo, Chanjun Park, Hyeongju Park, Hyeonseok Moon, and Heuiseok Lim. DART: Draft-Agreement Routing for Training-Free Adaptive Thinking Budgets in Hybrid Reasoning Models. *arXiv:2606.23181*, 2026a. EMNLP 2026 (Findings).

Yoonho Lee, Roshen Nair, Qizheng Zhang, Kangwook Lee, Omar Khattab, and Chelsea Finn. Meta-harness: End-to-end optimization of model harnesses. *arXiv:2603.28052*, 2026b.

Yunhao Liang, Chengguang Gan, Ruixuan Ying, Hanjun Wei, Zhe Cui, and Shiwen Ni. Auditing and Decomposing Feedback-Driven Evolution in LLM Test Generation under the Oracle Problem. *arXiv:2608.19626*, 2026.

Shiyan Liu, Qifeng Xia, Qiyun Xia, Yisheng Liu, Xinyu Yu, and Rui Qu. Reflection in the Dark: Exposing and Escaping the Black Box in Reflective Prompt Optimization. *arXiv:2603.18388*, 2026. ACL SRW 2026.

Wenjie Ma, Jingxuan He, Charlie Snell, Tyler Griggs, Sewon Min, and Matei Zaharia. Reasoning Models Can Be Effective Without Thinking. *arXiv:2504.09858*, 2025.

Solomon Messing. Hidden Measurement Error in LLM Pipelines Distorts Annotation, Evaluation, and Benchmarking. *arXiv:2604.11581*, 2026.

Niklas Muennighoff, Zitong Yang, Weijia Shi, Xiang Lisa Li, Li Fei-Fei, Hannaneh Hajishirzi, Luke Zettlemoyer, Percy Liang, Emmanuel Candès, and Tatsunori Hashimoto. s1: Simple test-time scaling. *arXiv:2501.19393*, 2025.

Wenhua Nie, Junlin Liu, Jianan Wu, Zijie Meng, Yilong Fan, Zhang Zijian, Haoran Zheng, and Jyh-Shing Roger Jang. The Coupling Tax: How Shared Token Budgets Undermine Visible Chain-of-Thought Under Fixed Output Limits. *arXiv:2605.07686*, 2026.

David Pape, Jonathan Evertz, and Lea Schönherr. The Silent Hyperparameter: Quantifying the Impact of Inference Backends on LLM Reproducibility. *arXiv:2605.19537*, 2026.

Kidus Seyoum and Ajay Mittur. Automatic Harness Evolution for Hardware Design Verification: Can LLMs Consolidate Gains Across Discovered Harnesses? *arXiv:2609.28908*, 2026.

Rishabh Sharma and Rishika Lall. When Does Selection Replace Extraction? A Pre-Registered Test of Agent Memory with a Typed Decision Model. *arXiv:2609.34227*, 2026.

Zayx Shawn. PACE: Anytime-Valid Acceptance Tests for Self-Evolving Agents. *arXiv:2606.08106*, 2026.

Amogh Singh. Certifying Compressed Language Models: An Audit and a Statistical Toolkit. *arXiv:2608.15046*, 2026.

Blake Stenstrom, Charangan Vasantharajan, and Brian Sathianathan. IBIB: A Protocol for Measuring Enterprise AI Systems by Serving Route, Not Model Identifier. *arXiv:2609.10494*, 2026.

Maria Thomas, Kristina Gligoric, and Nihar B. Shah. Mitigating LLM-based p-Hacking by Preregistering for the Next LLM. *arXiv:2606.27687*, 2026.

Varun Ursekar, Apaar Shanker, Yash Maurya, Shehab Yasser, Vijay S. Kalmath, Veronica Chatrath, and Yuan Xue. HarnessOpt-Bench: Evaluating LLMs at Harness Optimization. *arXiv:2608.06301*, 2026.

Vansh Wahi. LLM-as-a-Judge Is Not an Oracle: Why Self-Improving Agents Need Deterministic Guardrails. *arXiv:2609.02246*, 2026.

Wenbo Wang and Di Sang. Interface-Induced Trajectory Censoring. *arXiv:2609.03966*, 2026.

Xing Wang, Xiaoyi Zhang, and Jie Shao. Auditing Harness Tampering in Self-Improving Agents. *arXiv:2609.00069*, 2026a.

Yike Wang, Huaisheng Zhu, Zhengyu Hu, Yige Yuan, Zhengyu Chen, Shakti Senthil, Hannaneh Hajishirzi, Yulia Tsvetkov, Noah A. Smith, Pradeep Dasigi, and Teng Xiao. Rethinking the Evaluation of Harness Evolution for Agents. *arXiv:2607.12227*, 2026b.

Yubo Wang, Xueguang Ma, Ge Zhang, Yuansheng Ni, Abhranil Chandra, Shiguang Guo, Weiming Ren, Aaran Arulraj, Xuan He, Ziyan Jiang, Tianle Li, Max Ku, Kai Wang, Alex Zhuang, Rongqi Fan, Xiang Yue, and Wenhu Chen. MMLU-Pro: A more robust and challenging multi-task language understanding benchmark. In *Advances in Neural Information Processing Systems 37, Datasets and Benchmarks Track*, 2024.

Sarel Weinberger and Amir Hozez. Prompt-Induced Waste in Coding Agents: Reasoning, Effort, Harness Design, and End-to-End Cost. *arXiv:2608.01347*, 2026.

Siwei Wu, Jincheng Ren, Yizhi Li, Haau-Sing Li, Chengran Yang, Xudong Ma, Yuxuan Zhang, Weicheng Gu, Jian Yang, Riza Batista-Navarro, Chuanyi Zhang, Xianglong Liu, Ming Zhou, Bryan Dai, and Chenghua Lin. ModularRSI: Modular and generalizable recursive harness self-improvement. *arXiv:2609.14857*, 2026.

An Yang et al. Qwen3 technical report. *arXiv:2505.09388*, 2025.

Chengrun Yang, Xuezhi Wang, Yifeng Lu, Hanxiao Liu, Quoc V. Le, Denny Zhou, and Xinyun Chen. Large language models as optimizers. In *The Twelfth International Conference on Learning Representations*, 2024.

Rongxin Yang, Yang Liu, Shang Luo, Haoxuan Jia, Chongyang Zhang, Hao Zheng, Yingguang Yang, Yulin Huang, Jianshen Zhang, Yongzhi Qi, Kefu Xu, Congjing Ran, and Bin Chong. Harness-agnostic detection and immunization of reward hacking in self-evolving language models. *arXiv:2609.04665*, 2026a.

Shuo Yang, Xiaoze Fan, Melissa Pan, Haocheng Xi, Zhe Wang, Shanlin Sun, Kurt Keutzer, Song Han, Matei Zaharia, Chenfeng Xu, and Ion Stoica. FreeToken: Efficient edge-native MoE serving with bandwidth-adaptive execution. *arXiv:2608.16157*, 2026b. Inference stack used here at version 0.1.3.

Jiayi Yuan, Hao Li, Xinheng Ding, Wenya Xie, Yu-Jhe Li, Wentian Zhao, Kun Wan, Jing Shi, Xia Hu, and Zirui Liu. Understanding and Mitigating Numerical Sources of Nondeterminism in LLM Inference. *arXiv:2506.09501*, 2025. NeurIPS 2025 (Oral).

Jenny Zhang, Shengran Hu, Cong Lu, Robert Lange, and Jeff Clune. Darwin godel machine: Open-ended evolution of self-improving agents. *arXiv:2505.22954*, 2025.

Alex Zhavoronkov, Vladimir Naumov, Denis Sidorenko, Alex Aliper, Vladimir Aladinskiy, Ramin Hasani, Alexander Amini, Katerina Nasto, Mathieu Reymond, Rim Shayakhmetov, Zulfat Miftakhutdinov, Vadim N. Gladyshev, and Fedor Galkin. An open benchmark and language models for AI in aging biology. *Cell*, 189:5980–5994.e8, 2026. doi: 10.1016/j.cell.2026.08.026.

Yongchao Zhou, Andrei Ioan Muresanu, Ziwen Han, Keiran Paster, Silviu Pitis, Harris Chan, and Jimmy Ba. Large language models are human-level prompt engineers. In *The Eleventh International Conference on Learning Representations*, 2023.

Yuxuan Zhu, Tengjun Jin, Yada Pruksachatkun, Andy Zhang, Shu Liu, Sasha Cui, Sayash Kapoor, Shayne Longpre, Kevin Meng, Rebecca Weiss, Fazl Barez, Rahul Gupta, et al. Establishing Best Practices for Building Rigorous Agentic Benchmarks. *arXiv:2507.02825*, 2025.

A Background: the case study this study follows

The case study. The earlier study (v0.5 in the project's records) ran a harness search around a frozen Qwen3-1.7B longevity model, served in its tied variant by FreeToken (Yang et al., 2026b). The search started from a parent configuration, CFG-THINK-64-ORIG (thinking on, 64 tokens for thinking and answer), and, driven by its failures on one shard, produced B1, a conditional rescue: when the first call used nearly all of its tokens, call again with a no-think suffix. B1's gain in parse rate replicated on a second shard. A later audit found a simpler comparator, thinking disabled, that was as accurate as B1; it then found that the project's caller had sent each benchmark row's reference answer to the model as a prior assistant turn in every historical call that could be checked. Prompt-token reconciliation, the server's prompt-token count against the local count of the intended request, was the evidence: the counts differed by the length of the reference turn.

The corrected re-evaluation. After the caller was corrected, a preregistration was frozen (26 September 2026, 23:52:06 UTC) and run once on 556 holdout rows that no earlier experiment had called, though the comparator configurations had been defined after another line of experiments had called other holdout rows, among them 216 with a request equivalent to the no-think suffix. The parent parsed none of the 556 rows, so the preregistered reliability endpoint was nearly guaranteed. Thinking disabled was non-inferior to B1 at a margin of 0.10 (difference $+0.012$, interval $[-0.028, +0.052]$) at 47% of B1's token cost; since every B1 first call reached the rescue threshold, B1 reduced to always adding the suffix. Post hoc, on LB-0030 no configuration clearly exceeded a constant guess of the age band most common in the two shards (band 30–45, accuracy 0.493, against 0.471 to 0.493 for the configurations). The case study's conclusion was that prospective replication controlled selection noise but detected neither the inadequate comparator nor the evaluation defect, and that trivial baselines, input audits and prompt-token reconciliation did.

Table 8: Lessons of the case study and where this study tests them.

Lesson	Tested by
The parent parsed none of the prospective rows; the rescue made answers visible	H1, H3
A thinking-off setting that the search never tried was non-inferior to the rescue	H2
The caller sent the reference answer	H4, defect A
LB-0042's gold came from a fallback rule; the scorer could fail open	H4, defects B1, B2
The search's split put a giant component into one partition	H4, defect C
The parent was a broken start	H4, defect D
The tied variant was served, not the checkpoint's stored output layer	H4, defect E
A constant guess matched the configurations on one task	trivial baselines (Section 4)

Why this study is not a replication of it. The stack (vLLM in batch-invariant mode, not FreeToken), the model variant (untied output layer), two of the three models, two of the three pools and five of the six configurations are new, and the rules were fixed in advance for all of them. It is not blind to direction: the pilot showed that the parent truncates where the budget is small, and the preregistration says so (Section 3).

B Full tables

Every number in this appendix is read from the analysis report committed in our repository (the preregistered analysis) or from the three files computed after the outputs existed; Appendix E lists them. Intervals of primary endpoints are the family-adjusted ones, all others unadjusted 95%. Δ, θ, D_T and D_P are accuracy or rate differences; V, B and L count the rows only AUTHORS-1500 parses, both parse and only THINK-512 parses.

B.1 H1 to H3

Table 9: H1, primary endpoints: THINK-512 against AUTHORS-1500, test partition (family of 9). Repl.: R replicated, I inconclusive replication, N not replicated. A star marks an endpoint whose quantity is fixed by an identity of the design or by the observed parse pattern; its verdict says little beyond that mechanism, although it still depends on the intervals.

Cell	Δ	Δ_V	Δ_B	Δ_L	T	V / B / L	Verdict	Repl.
Qwen3-1.7B LB letters	−0.003 [−0.018, +0.012]	0.000	−0.003	0.000	+0.002 [−0.009, +0.014]	1 / 3,941 / 0	no gain: not tested	R
Qwen3-1.7B MMLU-Pro	+0.013 [−0.006, +0.033]	+0.015	−0.003	0.000	+0.005 [−0.008, +0.018]	48 / 751 / 0	no gain: not tested	R
Qwen3-1.7B GSM8K	0.000 [0.000, 0.000]	0.000	0.000	0.000	0.000 [0.000, 0.000]	0 / 523 / 0	no gain: not tested	R
SmolLM3 LB letters	+0.060 [+0.046, +0.074]	+0.071	+0.001	−0.012	+0.013 [+0.009, +0.018]	749 / 448 / 100	visibility dominates	R
SmolLM3 MMLU-Pro	+0.255 [+0.207, +0.304]	+0.282	−0.014	−0.013	+0.067 [+0.054, +0.081]	529 / 81 / 13	visibility dominates	R
SmolLM3 GSM8K	+0.108 [+0.046, +0.173]	+0.197	−0.036	−0.053	+0.068 [+0.047, +0.090]	141 / 340 / 29	visibility dominates	R
Intern-S1 LB letters	+0.314 [+0.297, +0.331]	+0.314	0.000	0.000	+0.063 [+0.059, +0.066]	2,407 / 1 / 0	visibility dominates*	R
Intern-S1 MMLU-Pro	+0.403 [+0.352, +0.456]	+0.419	−0.016	0.000	+0.097 [+0.084, +0.111]	730 / 69 / 0	visibility dominates*	R
Intern-S1 GSM8K	−0.287 [−0.355, −0.218]	+0.057	−0.343	0.000	+0.286 [+0.240, +0.332]	273 / 253 / 0	no gain: not tested	R

Table 10: H1, sanity checks: THINK-64 against AUTHORS-1500, test partition, at the family's adjusted level, outside the family and carrying no claim.

Cell	Δ	Δ_V	Δ_B	Δ_L	T	V / B / L	Verdict	Repl.
Qwen3-1.7B LB letters	+0.632 [+0.613, +0.651]	+0.632	0.000	0.000	+0.126 [+0.123, +0.130]	3,942 / 0 / 0	visibility dominates	
Qwen3-1.7B MMLU-Pro	+0.089 [+0.063, +0.119]	+0.089	0.000	0.000	+0.018 [+0.013, +0.024]	274 / 525 / 0	visibility dominates	
Qwen3-1.7B GSM8K	+0.002 [0.000, +0.009]	+0.002	0.000	0.000	0.000 [0.000, +0.002]	1 / 522 / 0	no gain: not tested	
SmolLM3 LB letters	+0.129 [+0.116, +0.143]	+0.129	0.000	0.000	+0.026 [+0.023, +0.029]	1,197 / 0 / 0	visibility dominates	
SmolLM3 MMLU-Pro	+0.353 [+0.307, +0.401]	+0.353	0.000	0.000	+0.071 [+0.061, +0.080]	610 / 0 / 0	visibility dominates	
SmolLM3 GSM8K	+0.789 [+0.738, +0.837]	+0.789	0.000	0.000	+0.158 [+0.148, +0.167]	481 / 0 / 0	visibility dominates	
Intern-S1 LB letters	+0.314 [+0.297, +0.331]	+0.314	0.000	0.000	+0.063 [+0.059, +0.066]	2,408 / 0 / 0	visibility dominates	
Intern-S1 MMLU-Pro	+0.477 [+0.427, +0.526]	+0.477	0.000	0.000	+0.095 [+0.085, +0.105]	799 / 0 / 0	visibility dominates	
Intern-S1 GSM8K	+0.190 [+0.144, +0.239]	+0.190	0.000	0.000	+0.038 [+0.029, +0.048]	526 / 0 / 0	visibility dominates	

Table 11: H2, primary endpoints: θ of AUTHORS-1500 against RESCUE-1500 in nine cells and against the four GEPA harnesses (family of 13), with the replication partition's verdict and the label. A star marks an endpoint whose quantity is fixed by an identity of the design or by the observed parse pattern; its verdict says little beyond that mechanism, although it still depends on the intervals.

Endpoint	θ (adjusted interval)	Verdict	Replication verdict	Label
Qwen3-1.7B LB letters, RESCUE-1500	−0.003 [−0.018, +0.012]	holds	holds	replicated
Qwen3-1.7B MMLU-Pro, RESCUE-1500	−0.003 [−0.019, +0.014]	holds	holds	replicated
Qwen3-1.7B GSM8K, RESCUE-1500	0.000 [0.000, 0.000]	holds	holds	replicated
SmolLM3 LB letters, RESCUE-1500	+0.001 [−0.007, +0.009]	holds	holds	replicated
SmolLM3 MMLU-Pro, RESCUE-1500	−0.025 [−0.044, −0.009]	holds	holds	replicated
SmolLM3 GSM8K, RESCUE-1500	−0.089 [−0.129, −0.053]	fails	inconclusive	inconclusive replication
Intern-S1 LB letters, RESCUE-1500	0.000 [0.000, 0.000]	holds*	holds	replicated
Intern-S1 MMLU-Pro, RESCUE-1500	−0.016 [−0.030, −0.005]	holds	holds	replicated
Intern-S1 GSM8K, RESCUE-1500	−0.342 [−0.402, −0.281]	fails	fails	replicated
Qwen3 LB letters, GEPA-CLEAN-s0	0.000 [0.000, 0.000]	holds*	holds	replicated
Qwen3 LB letters, GEPA-BROKEN-s0	+0.004 [−0.009, +0.016]	holds	holds	replicated
Qwen3 LB letters, GEPA-CLEAN-s1	+0.002 [−0.022, +0.025]	holds	holds	replicated
Qwen3 LB letters, GEPA-BROKEN-s1	+0.012 [−0.002, +0.026]	holds	holds	replicated

Table 12: H2, descriptive: θ of AUTHORS-1500 against three other configurations, test partition, unadjusted intervals and the verdict rule applied to them.

Cell	vs THINK-512	vs THINK-1536	vs BUDGET-512-1500
Qwen3-1.7B LB letters	−0.003 [−0.013, +0.008] holds	−0.003 [−0.013, +0.007] holds	−0.003 [−0.013, +0.007] holds
Qwen3-1.7B MMLU-Pro	+0.013 [−0.001, +0.026] holds	−0.005 [−0.018, +0.008] holds	−0.003 [−0.015, +0.010] holds
Qwen3-1.7B GSM8K	0.000 [0.000, 0.000] holds	0.000 [0.000, 0.000] holds	0.000 [0.000, 0.000] holds
SmolLM3 LB letters	+0.060 [+0.050, +0.070] holds	−0.028 [−0.040, −0.015] holds	+0.009 [−0.003, +0.020] holds
SmolLM3 MMLU-Pro	+0.255 [+0.220, +0.290] holds	+0.039 [+0.003, +0.075] holds	+0.095 [+0.059, +0.131] holds
SmolLM3 GSM8K	+0.108 [+0.063, +0.154] holds	−0.091 [−0.127, −0.055] fails	−0.046 [−0.083, −0.008] inconclusive
Intern-S1 LB letters	+0.314 [+0.302, +0.326] holds	+0.189 [+0.175, +0.204] holds	+0.061 [+0.046, +0.076] holds
Intern-S1 MMLU-Pro	+0.403 [+0.367, +0.439] holds	+0.198 [+0.161, +0.235] holds	−0.065 [−0.104, −0.026] inconclusive
Intern-S1 GSM8K	−0.287 [−0.334, −0.237] fails	−0.643 [−0.685, −0.600] fails	−0.732 [−0.770, −0.693] fails

Table 13: H3, primary endpoints: THINK-512 against BUDGET-512-1500, test partition (family of 9); the truncation rate of THINK-512 decides *no room*.

Cell	D_T	D_P	THINK-512 truncation	Verdict	Replication verdict	Label
Qwen3-1.7B LB letters	0.000 [0.000, +0.001]	0.000 [0.000, +0.001]	0.000 [0.000, +0.001]	no room	no room	replicated
Qwen3-1.7B MMLU-Pro	+0.060 [+0.038, +0.085]	+0.060 [+0.038, +0.085]	+0.060 [+0.038, +0.085]	inconclusive	inconclusive	replicated
Qwen3-1.7B GSM8K	0.000 [0.000, 0.000]	0.000 [0.000, 0.000]	0.000 [0.000, 0.000]	no room	no room	replicated
SmolLM3 LB letters	+0.428 [+0.407, +0.450]	+0.151 [+0.136, +0.167]	+0.725 [+0.708, +0.742]	holds	holds	replicated
SmolLM3 MMLU-Pro	+0.526 [+0.476, +0.574]	+0.248 [+0.205, +0.292]	+0.905 [+0.875, +0.932]	holds	holds	replicated
SmolLM3 GSM8K	+0.260 [+0.209, +0.315]	+0.190 [+0.144, +0.239]	+0.300 [+0.245, +0.357]	holds	holds	replicated
Intern-S1 LB letters	+0.979 [+0.972, +0.985]	+0.619 [+0.603, +0.636]	+1.000 [+0.999, +1.000]	holds	holds	replicated
Intern-S1 MMLU-Pro	+0.647 [+0.601, +0.693]	+0.662 [+0.615, +0.708]	+0.915 [+0.887, +0.941]	holds	holds	replicated
Intern-S1 GSM8K	+0.493 [+0.435, +0.554]	+0.497 [+0.438, +0.556]	+0.522 [+0.461, +0.581]	holds	holds	replicated

B.2 H4

Table 14: H4, the 16 pairs with their adjusted intervals (family of 16, level 99.69%) and, in parentheses, which side of its boundary each interval lies on (*holds*, *fails* or *open*). Stability of Γ is a descriptive fact about the same interval. A star marks a claim whose quantity is fixed by an identity of the design; its verdict says little beyond that mechanism, although it still depends on the intervals.

Pair	Bound	D_S	D_R	Γ	Stability of Γ	Claim
A Qwen3-1.7B	1.000	−0.052 [−0.106, +0.002] (open)	−0.030 [−0.087, +0.030] (open)	+0.022 [−0.057, +0.102] (open)	inconclusive	inconclusive
A SmolLM3	1.000	+0.097 [+0.022, +0.173] (open)	+0.119 [+0.042, +0.199] (open)	+0.022 [−0.089, +0.133] (open)	inconclusive	inconclusive
A Intern-S1	1.000	+0.142 [+0.082, +0.201] (holds)	+0.164 [+0.102, +0.226] (holds)	+0.022 [−0.065, +0.108] (open)	inconclusive	inconclusive
B1 Qwen3-1.7B	0.205	−0.145 [−0.170, −0.118] (holds)	−0.148 [−0.166, −0.130] (holds)	−0.003 [−0.035, +0.028] (holds)	stable	holds
B1 SmolLM3	0.083	+0.001 [−0.019, +0.021] (fails)	+0.007 [−0.006, +0.021] (fails)	+0.006 [−0.018, +0.031] (holds)	stable	fails
B1 Intern-S1	0.057	−0.012 [−0.028, +0.004] (fails)	−0.019 [−0.030, −0.007] (fails)	−0.007 [−0.027, +0.013] (holds)	stable	fails
B2 Qwen3-1.7B	0.000	not evaluable: the start leaves no row unparsed				–
B2 SmolLM3	0.208	+0.208 [+0.181, +0.235] (holds)	+0.223 [+0.203, +0.243] (holds)	+0.015 [−0.018, +0.048] (holds)	stable	holds⋆
B2 Intern-S1	0.289	+0.289 [+0.260, +0.319] (holds)	+0.292 [+0.271, +0.313] (holds)	+0.003 [−0.033, +0.037] (holds)	stable	holds⋆
C Qwen3-1.7B	by rule	+0.072 [+0.004, +0.138] (open)	+0.101 [+0.047, +0.157] (open)	+0.030 [−0.057, +0.117] (open)	inconclusive	inconclusive
C SmolLM3	by rule	−0.001 [−0.052, +0.052] (open)	+0.005 [−0.048, +0.063] (fails)	+0.006 [−0.070, +0.082] (open)	inconclusive	fails
C Intern-S1	by rule	+0.003 [−0.046, +0.054] (open)	+0.004 [−0.045, +0.057] (open)	+0.001 [−0.070, +0.074] (open)	inconclusive	inconclusive
D Qwen3-1.7B	1.000	+0.638 [+0.610, +0.668] (holds)	+0.635 [+0.614, +0.656] (holds)	−0.003 [−0.039, +0.032] (holds)	stable	holds⋆
D SmolLM3	0.352	+0.164 [+0.142, +0.187] (holds)	+0.157 [+0.142, +0.173] (holds)	−0.007 [−0.034, +0.020] (holds)	stable	holds⋆
D Intern-S1	0.189	+0.124 [+0.104, +0.145] (holds)	+0.125 [+0.110, +0.140] (holds)	+0.001 [−0.025, +0.026] (holds)	stable	holds⋆
E Qwen3-1.7B	0.139	+0.005 [−0.016, +0.025] (fails)	+0.014 [−0.002, +0.029] (fails)	+0.009 [−0.017, +0.034] (holds)	stable	fails

Table 15: The detection matrix (descriptive; largely true by construction). Checks: PR prospective replication, TR token reconciliation, TB trivial baselines, SCF simple configuration family, IR input reading, SMI served-model identity.

Pair	PR	TR	TB	SCF	IR	SMI	Rules fired (expected ones in bold)
A Qwen3-1.7B	U	D	U	U	U	–	**IR1**, **IR2**, IR4, PR.gain, SCF2, TB2, **TB3**, **TR**
A SmolLM3	D	D	U	U	U	–	**IR1**, **IR2**, IR4, PR.gain, SCF1, SCF2, SCF3, TB2, **TR**; expected, not fired: TB3
A Intern-S1	M	D	U	U	U	–	**IR1**, **IR2**, IR4, SCF1, SCF3, TB2, **TB3**, **TR**
B1 Qwen3-1.7B	M	M	M	M	U	–	**IR3**, IR4
B1 SmolLM3	M	M	M	U	U	–	**IR3**, IR4, SCF1
B1 Intern-S1	M	M	M	U	U	–	**IR3**, IR4, SCF1, SCF2, SCF3
B2 Qwen3-1.7B	D	M	D	D	U	–	IR4, PR.gain, SCF2, **TB4**
B2 SmolLM3	D	M	D	U	U	–	IR4, PR.gain, SCF1, SCF2, **TB4**
B2 Intern-S1	M	M	D	U	U	–	IR4, SCF1, **TB4**
C Qwen3-1.7B	D	M	D	D	U	–	**IR4**, **IR5**, PR.gain, SCF2, TB2; expected, not fired: TB1
C SmolLM3	M	M	D	U	U	–	**IR4**, **IR5**, SCF1, TB2; expected, not fired: TB1
C Intern-S1	M	M	D	U	U	–	**IR4**, **IR5**, SCF1, SCF2, SCF3, TB2; expected, not fired: TB1
D Qwen3-1.7B	M	M	M	D	U	–	IR4, **SCF1**, SCF2, **SCF3**
D SmolLM3	M	M	M	U	U	–	IR4, **SCF1**, SCF2, **SCF3**
D Intern-S1	M	M	M	U	U	–	IR4, **SCF1**, SCF2, **SCF3**
E Qwen3-1.7B	D	M	M	U	U	D	IR4, PR.level, SCF2, **SMI**

D detects: fires on the defective instance and not on the clean one (5, 3, 6, 3, 0, 1 instances for PR, TR, TB, SCF, IR, SMI); U uninformative: fires on both (1, 0, 3, 12, 16, 0); M misses.

Table 16: False alarms of each rule on the clean instances (the main clean instance and the clean counterparts of A and E, per model).

Check	Rule	Fired on	Of
PR	PR.gain	1: clean/A / Qwen3	7
PR	PR.level	0: none	7
TR	TR	0: none	7
TB	TB1	0: none	7
TB	TB2	3: clean/A / Intern-S1, clean/A / Qwen3, clean/A / SmolLM3	7
TB	TB3	0: none	7
TB	TB4	0: none	7
SCF	SCF1	4: clean / Intern-S1, clean / SmolLM3, clean/A / Intern-S1, clean/A / SmolLM3	7
SCF	SCF2	4: clean / Intern-S1, clean/A / Intern-S1, clean/A / Qwen3, clean/E / Qwen3	7
SCF	SCF3	2: clean / Intern-S1, clean/A / Intern-S1	7
IR	IR1	0: none	7
IR	IR2	0: none	7
IR	IR3	0: none	7
IR	IR4	7: clean / Intern-S1, clean / Qwen3, clean / SmolLM3, clean/A / Intern-S1, clean/A / Qwen3, clean/A / SmolLM3, clean/E / Qwen3	7
IR	IR5	0: none	7

B.3 Configurations

Table 17: Accuracy, parse, truncation and non-commitment rates of the six configurations in the nine cells, test partition, with unadjusted 95% intervals. Unparsed answers count as wrong. The trivial baselines of the three pools are in Table 51 and Appendix B.4.

Cell	Configuration	Accuracy	Parse	Truncated	Non-committed
Qwen3-1.7B LB letters	THINK-64	0.000 [0.000, 0.000]	0.000 [0.000, 0.000]	1.000 [1.000, 1.000]	0.000 [0.000, 0.000]
	THINK-512	0.635 [0.621, 0.649]	1.000 [0.999, 1.000]	0.000 [0.000, 0.001]	0.000 [0.000, 0.000]
	THINK-1536	0.635 [0.621, 0.649]	1.000 [1.000, 1.000]	0.000 [0.000, 0.000]	0.000 [0.000, 0.000]
	AUTHORS-1500	0.632 [0.618, 0.646]	1.000 [1.000, 1.000]	0.000 [0.000, 0.000]	0.000 [0.000, 0.000]
	BUDGET-512-1500	0.635 [0.621, 0.649]	1.000 [1.000, 1.000]	0.000 [0.000, 0.000]	0.000 [0.000, 0.000]
	RESCUE-1500	0.635 [0.621, 0.649]	1.000 [1.000, 1.000]	0.000 [0.000, 0.000]	0.000 [0.000, 0.000]
Qwen3-1.7B MMLU-Pro	THINK-64	0.140 [0.116, 0.165]	0.657 [0.625, 0.690]	0.343 [0.310, 0.375]	0.000 [0.000, 0.000]
	THINK-512	0.217 [0.189, 0.245]	0.940 [0.922, 0.956]	0.060 [0.044, 0.078]	0.000 [0.000, 0.000]
	THINK-1536	0.234 [0.205, 0.264]	0.997 [0.994, 1.000]	0.003 [0.000, 0.006]	0.000 [0.000, 0.000]
	AUTHORS-1500	0.229 [0.200, 0.259]	1.000 [1.000, 1.000]	0.000 [0.000, 0.000]	0.000 [0.000, 0.000]
	BUDGET-512-1500	0.232 [0.203, 0.262]	1.000 [1.000, 1.000]	0.000 [0.000, 0.000]	0.000 [0.000, 0.000]
	RESCUE-1500	0.232 [0.203, 0.262]	1.000 [1.000, 1.000]	0.000 [0.000, 0.000]	0.000 [0.000, 0.000]
Qwen3-1.7B GSM8K	THINK-64	0.068 [0.047, 0.091]	0.991 [0.981, 0.998]	0.006 [0.000, 0.013]	0.004 [0.000, 0.009]
	THINK-512	0.070 [0.049, 0.093]	0.992 [0.985, 0.998]	0.000 [0.000, 0.000]	0.008 [0.002, 0.015]
	THINK-1536	0.070 [0.049, 0.093]	0.992 [0.985, 0.998]	0.000 [0.000, 0.000]	0.008 [0.002, 0.015]
	AUTHORS-1500	0.070 [0.049, 0.093]	0.992 [0.985, 0.998]	0.000 [0.000, 0.000]	0.008 [0.002, 0.015]
	BUDGET-512-1500	0.070 [0.049, 0.093]	0.992 [0.985, 0.998]	0.000 [0.000, 0.000]	0.008 [0.002, 0.015]
	RESCUE-1500	0.070 [0.049, 0.093]	0.992 [0.985, 0.998]	0.000 [0.000, 0.000]	0.008 [0.002, 0.015]
SmolLM3 LB letters	THINK-64	0.000 [0.000, 0.000]	0.000 [0.000, 0.000]	1.000 [1.000, 1.000]	0.000 [0.000, 0.000]
	THINK-512	0.070 [0.063, 0.077]	0.139 [0.131, 0.147]	0.725 [0.713, 0.737]	0.138 [0.128, 0.148]
	THINK-1536	0.157 [0.147, 0.168]	0.339 [0.326, 0.352]	0.485 [0.472, 0.498]	0.181 [0.170, 0.193]
	AUTHORS-1500	0.129 [0.120, 0.139]	0.304 [0.292, 0.315]	0.023 [0.019, 0.028]	0.677 [0.665, 0.688]
	BUDGET-512-1500	0.121 [0.112, 0.130]	0.290 [0.278, 0.302]	0.297 [0.283, 0.310]	0.424 [0.409, 0.438]
	RESCUE-1500	0.128 [0.119, 0.138]	0.308 [0.297, 0.319]	0.020 [0.016, 0.024]	0.675 [0.663, 0.686]
SmolLM3 MMLU-Pro	THINK-64	0.000 [0.000, 0.000]	0.000 [0.000, 0.000]	1.000 [1.000, 1.000]	0.000 [0.000, 0.000]
	THINK-512	0.098 [0.078, 0.119]	0.118 [0.096, 0.140]	0.905 [0.884, 0.925]	0.000 [0.000, 0.000]
	THINK-1536	0.314 [0.282, 0.347]	0.461 [0.426, 0.496]	0.526 [0.491, 0.561]	0.024 [0.014, 0.035]
	AUTHORS-1500	0.353 [0.320, 0.387]	0.763 [0.733, 0.792]	0.179 [0.153, 0.207]	0.105 [0.084, 0.126]
	BUDGET-512-1500	0.258 [0.228, 0.288]	0.365 [0.333, 0.399]	0.379 [0.345, 0.413]	0.304 [0.273, 0.337]
	RESCUE-1500	0.378 [0.344, 0.412]	0.778 [0.750, 0.807]	0.169 [0.144, 0.195]	0.096 [0.076, 0.118]
SmolLM3 GSM8K	THINK-64	0.000 [0.000, 0.000]	0.000 [0.000, 0.000]	1.000 [1.000, 1.000]	0.000 [0.000, 0.000]
	THINK-512	0.681 [0.641, 0.721]	0.700 [0.660, 0.738]	0.300 [0.262, 0.340]	0.000 [0.000, 0.000]
	THINK-1536	0.880 [0.852, 0.907]	0.934 [0.911, 0.954]	0.066 [0.046, 0.089]	0.000 [0.000, 0.000]
	AUTHORS-1500	0.789 [0.753, 0.824]	0.913 [0.888, 0.935]	0.023 [0.011, 0.036]	0.074 [0.053, 0.097]
	BUDGET-512-1500	0.835 [0.803, 0.865]	0.890 [0.863, 0.917]	0.040 [0.025, 0.057]	0.072 [0.051, 0.095]
	RESCUE-1500	0.879 [0.850, 0.905]	0.968 [0.953, 0.981]	0.011 [0.004, 0.021]	0.023 [0.011, 0.036]
Intern-S1 LB letters	THINK-64	0.000 [0.000, 0.000]	0.000 [0.000, 0.000]	1.000 [1.000, 1.000]	0.000 [0.000, 0.000]
	THINK-512	0.000 [0.000, 0.000]	0.000 [0.000, 0.001]	1.000 [0.999, 1.000]	0.000 [0.000, 0.000]
	THINK-1536	0.125 [0.115, 0.135]	0.192 [0.181, 0.203]	0.800 [0.789, 0.811]	0.016 [0.012, 0.020]
	AUTHORS-1500	0.314 [0.302, 0.326]	0.611 [0.604, 0.618]	0.006 [0.003, 0.008]	0.384 [0.376, 0.391]
	BUDGET-512-1500	0.253 [0.241, 0.266]	0.619 [0.608, 0.631]	0.021 [0.016, 0.025]	0.363 [0.351, 0.374]
	RESCUE-1500	0.314 [0.302, 0.326]	0.611 [0.604, 0.618]	0.006 [0.003, 0.008]	0.384 [0.376, 0.391]
Intern-S1 MMLU-Pro	THINK-64	0.000 [0.000, 0.000]	0.000 [0.000, 0.000]	1.000 [1.000, 1.000]	0.000 [0.000, 0.000]
	THINK-512	0.074 [0.056, 0.093]	0.086 [0.068, 0.106]	0.915 [0.895, 0.934]	0.000 [0.000, 0.000]
	THINK-1536	0.279 [0.249, 0.310]	0.345 [0.313, 0.378]	0.652 [0.620, 0.685]	0.009 [0.003, 0.016]
	AUTHORS-1500	0.477 [0.442, 0.512]	1.000 [1.000, 1.000]	0.000 [0.000, 0.000]	0.000 [0.000, 0.000]
	BUDGET-512-1500	0.542 [0.507, 0.576]	0.748 [0.718, 0.778]	0.268 [0.238, 0.299]	0.033 [0.021, 0.045]
	RESCUE-1500	0.493 [0.458, 0.528]	1.000 [1.000, 1.000]	0.000 [0.000, 0.000]	0.000 [0.000, 0.000]
Intern-S1 GSM8K	THINK-64	0.000 [0.000, 0.000]	0.000 [0.000, 0.000]	1.000 [1.000, 1.000]	0.000 [0.000, 0.000]
	THINK-512	0.476 [0.435, 0.518]	0.480 [0.436, 0.524]	0.522 [0.480, 0.564]	0.000 [0.000, 0.000]
	THINK-1536	0.833 [0.801, 0.863]	0.852 [0.822, 0.882]	0.148 [0.118, 0.178]	0.000 [0.000, 0.000]
	AUTHORS-1500	0.190 [0.157, 0.224]	0.998 [0.994, 1.000]	0.000 [0.000, 0.000]	0.002 [0.000, 0.006]
	BUDGET-512-1500	0.922 [0.899, 0.945]	0.977 [0.964, 0.989]	0.028 [0.015, 0.044]	0.004 [0.000, 0.009]
	RESCUE-1500	0.531 [0.488, 0.573]	0.998 [0.994, 1.000]	0.000 [0.000, 0.000]	0.002 [0.000, 0.006]

B.4 Trivial baselines, computed after the outputs existed

Chance is the mean of $1/k$ over the rows (k options). A constant letter or a constant label is the same answer on every row; the non-oracle constant is the most common gold letter or label of the other partition, the oracle constant the most common one of the partition itself (ties to the smallest). On LB-0030 and LB-0042 the order of the options is shuffled from row to row, so a constant letter cannot exploit an uneven

label distribution and a constant label can. The strongest baseline of a task is the larger of its two oracle constants; the pool's values are row-weighted means. For MMLU-Pro and GSM8K (addendum A12) the constants are a letter and an answer (the most common gold answer of the other partition, and of the partition itself); the strongest baseline is the oracle constant, and a free-form number has no chance level. Cells give accuracy and, in parentheses, the paired difference to the strongest oracle baseline; intervals for the pool are in Table 51, the rest in the files of our repository (Appendix E).

Table 18: Accuracy and difference to the strongest baseline (oracle constant label), test partition.

Model	THINK-64	THINK-512	THINK-1536	AUTHORS-1500	BUDGET-512-1500	RESCUE-1500
LB-0030 (1,577 rows). Chance 0.200; constant letter 0.216 (non-oracle), 0.216 (oracle); constant label 0.242 (non-oracle), 0.250 (oracle). Strongest, oracle: 0.250 [0.230, 0.269].						
Qwen3-1.7B	0.000 (−0.250)	0.369 (+0.119)	0.369 (+0.119)	0.344 (+0.094)	0.369 (+0.119)	0.369 (+0.119)
SmolLM3	0.000 (−0.250)	0.000 (−0.250)	0.021 (−0.229)	0.053 (−0.197)	0.032 (−0.218)	0.053 (−0.197)
Intern-S1	0.000 (−0.250)	0.000 (−0.250)	0.004 (−0.245)	0.023 (−0.227)	0.055 (−0.195)	0.023 (−0.227)
Qwen3, GEPA	CLEAN-s0 0.344 (+0.094); BROKEN-s0 0.353 (+0.103); CLEAN-s1 0.388 (+0.138); BROKEN-s1 0.321 (+0.072)					
LB-0034 (788 rows). Chance 0.500; constant letter 0.485 (non-oracle), 0.515 (oracle); constant label 0.485 (non-oracle), 0.515 (oracle). Strongest, oracle: 0.515 [0.480, 0.551].						
Qwen3-1.7B	0.000 (−0.515)	0.702 (+0.187)	0.703 (+0.188)	0.717 (+0.202)	0.703 (+0.188)	0.703 (+0.188)
SmolLM3	0.000 (−0.515)	0.307 (−0.208)	0.409 (−0.107)	0.443 (−0.072)	0.355 (−0.160)	0.435 (−0.080)
Intern-S1	0.000 (−0.515)	0.000 (−0.515)	0.242 (−0.273)	0.575 (+0.060)	0.523 (+0.008)	0.575 (+0.060)
Qwen3, GEPA	CLEAN-s0 0.717 (+0.202); BROKEN-s0 0.702 (+0.187); CLEAN-s1 0.716 (+0.201); BROKEN-s1 0.716 (+0.201)					
LB-0042 (1,577 rows). Chance 0.500; constant letter 0.500 (non-oracle), 0.500 (oracle); constant label 0.840 (non-oracle), 0.840 (oracle). Strongest, oracle: 0.840 [0.822, 0.858].						
Qwen3-1.7B	0.000 (−0.840)	0.867 (+0.027)	0.867 (+0.027)	0.878 (+0.038)	0.867 (+0.027)	0.867 (+0.027)
SmolLM3	0.000 (−0.840)	0.020 (−0.820)	0.167 (−0.673)	0.049 (−0.791)	0.093 (−0.748)	0.050 (−0.790)
Intern-S1	0.000 (−0.840)	0.000 (−0.840)	0.186 (−0.654)	0.475 (−0.365)	0.317 (−0.523)	0.475 (−0.365)
Qwen3, GEPA	CLEAN-s0 0.878 (+0.038); BROKEN-s0 0.867 (+0.027); CLEAN-s1 0.831 (−0.010); BROKEN-s1 0.871 (+0.030)					
Letter pool (3,942 rows). Chance 0.380; constant letter 0.383 (non-oracle), 0.389 (oracle); constant label 0.530 (non-oracle), 0.539 (oracle). Strongest, oracle: 0.539 [0.527, 0.551].						
Qwen3-1.7B	0.000 (−0.539)	0.635 (+0.096)	0.635 (+0.096)	0.632 (+0.093)	0.635 (+0.096)	0.635 (+0.096)
SmolLM3	0.000 (−0.539)	0.070 (−0.470)	0.157 (−0.382)	0.129 (−0.410)	0.121 (−0.418)	0.128 (−0.411)
Intern-S1	0.000 (−0.539)	0.000 (−0.539)	0.125 (−0.414)	0.314 (−0.225)	0.253 (−0.286)	0.314 (−0.225)
Qwen3, GEPA	CLEAN-s0 0.632 (+0.093); BROKEN-s0 0.629 (+0.090); CLEAN-s1 0.631 (+0.092); BROKEN-s1 0.620 (+0.081)					

Table 19: The same for the replication partition.

Model	THINK-64	THINK-512	THINK-1536	AUTHORS-1500	BUDGET-512-1500	RESCUE-1500
LB-0030 (796 rows). Chance 0.200; constant letter 0.212 (non-oracle), 0.212 (oracle); constant label 0.255 (non-oracle), 0.270 (oracle). Strongest, oracle: 0.270 [0.242, 0.299].						
Qwen3-1.7B	0.000 (−0.270)	0.358 (+0.088)	0.358 (+0.088)	0.335 (+0.065)	0.358 (+0.088)	0.358 (+0.088)
SmolLM3	0.000 (−0.270)	0.000 (−0.270)	0.025 (−0.245)	0.041 (−0.229)	0.031 (−0.239)	0.041 (−0.229)
Intern-S1	0.000 (−0.270)	0.000 (−0.270)	0.005 (−0.265)	0.011 (−0.259)	0.046 (−0.224)	0.011 (−0.259)
Qwen3, GEPA	CLEAN-s0 0.335 (+0.065); BROKEN-s0 0.345 (+0.075); CLEAN-s1 0.389 (+0.119); BROKEN-s1 0.307 (+0.036)					
LB-0034 (398 rows). Chance 0.500; constant letter 0.490 (non-oracle), 0.510 (oracle); constant label 0.490 (non-oracle), 0.510 (oracle). Strongest, oracle: 0.510 [0.462, 0.558].						
Qwen3-1.7B	0.000 (−0.510)	0.741 (+0.231)	0.744 (+0.234)	0.744 (+0.234)	0.744 (+0.234)	0.744 (+0.234)
SmolLM3	0.000 (−0.510)	0.322 (−0.188)	0.420 (−0.090)	0.432 (−0.078)	0.369 (−0.141)	0.462 (−0.048)
Intern-S1	0.000 (−0.510)	0.000 (−0.510)	0.226 (−0.284)	0.585 (+0.075)	0.500 (−0.010)	0.585 (+0.075)
Qwen3, GEPA	CLEAN-s0 0.744 (+0.234); BROKEN-s0 0.729 (+0.219); CLEAN-s1 0.741 (+0.231); BROKEN-s1 0.716 (+0.206)					
LB-0042 (796 rows). Chance 0.500; constant letter 0.489 (non-oracle), 0.511 (oracle); constant label 0.823 (non-oracle), 0.823 (oracle). Strongest, oracle: 0.823 [0.795, 0.850].						
Qwen3-1.7B	0.000 (−0.823)	0.866 (+0.043)	0.866 (+0.043)	0.876 (+0.053)	0.866 (+0.043)	0.866 (+0.043)
SmolLM3	0.000 (−0.823)	0.016 (−0.807)	0.176 (−0.647)	0.045 (−0.778)	0.082 (−0.741)	0.043 (−0.780)
Intern-S1	0.000 (−0.823)	0.000 (−0.823)	0.192 (−0.631)	0.480 (−0.343)	0.320 (−0.503)	0.480 (−0.343)
Qwen3, GEPA	CLEAN-s0 0.876 (+0.053); BROKEN-s0 0.877 (+0.054); CLEAN-s1 0.843 (+0.020); BROKEN-s1 0.882 (+0.059)					
Letter pool (1,990 rows). Chance 0.380; constant letter 0.378 (non-oracle), 0.391 (oracle); constant label 0.529 (non-oracle), 0.539 (oracle). Strongest, oracle: 0.539 [0.519, 0.559].						
Qwen3-1.7B	0.000 (−0.539)	0.638 (+0.098)	0.638 (+0.099)	0.633 (+0.094)	0.638 (+0.099)	0.638 (+0.099)
SmolLM3	0.000 (−0.539)	0.071 (−0.468)	0.164 (−0.375)	0.121 (−0.418)	0.119 (−0.420)	0.126 (−0.413)
Intern-S1	0.000 (−0.539)	0.000 (−0.539)	0.124 (−0.415)	0.314 (−0.226)	0.247 (−0.292)	0.314 (−0.226)
Qwen3, GEPA	CLEAN-s0 0.633 (+0.094); BROKEN-s0 0.635 (+0.095); CLEAN-s1 0.641 (+0.102); BROKEN-s1 0.619 (+0.079)					

Table 20: MMLU-Pro and GSM8K, test partition: accuracy of the six configurations and, in parentheses, the paired difference to the strongest baseline (the oracle constant).

Model	THINK-64	THINK-512	THINK-1536	AUTHORS-1500	BUDGET-512-1500	RESCUE-1500
MMLU-Pro (799 items). Chance 0.110; constant letter B, chosen on the other partition and oracle alike, 0.119 [0.096, 0.141].						
Qwen3-1.7B	0.140 (+0.021)	0.217 (+0.098)	0.234 (+0.115)	0.229 (+0.110)	0.232 (+0.113)	0.232 (+0.113)
SmolLM3	0.000 (−0.119)	0.098 (−0.021)	0.314 (+0.195)	0.353 (+0.234)	0.258 (+0.139)	0.378 (+0.259)
Intern-S1	0.000 (−0.119)	0.074 (−0.045)	0.279 (+0.160)	0.477 (+0.358)	0.542 (+0.423)	0.493 (+0.374)
GSM8K (527 items). Constant answer 10 (chosen on the other partition) 0.019 [0.008, 0.032], 2 (oracle) 0.034 [0.019, 0.051]; no chance level.						
Qwen3-1.7B	0.068 (+0.034)	0.070 (+0.036)	0.070 (+0.036)	0.070 (+0.036)	0.070 (+0.036)	0.070 (+0.036)
SmolLM3	0.000 (−0.034)	0.681 (+0.647)	0.880 (+0.846)	0.789 (+0.755)	0.835 (+0.801)	0.879 (+0.844)
Intern-S1	0.000 (−0.034)	0.476 (+0.442)	0.833 (+0.799)	0.190 (+0.156)	0.922 (+0.888)	0.531 (+0.497)

Table 21: MMLU-Pro and GSM8K, replication partition.

Model	THINK-64	THINK-512	THINK-1536	AUTHORS-1500	BUDGET-512-1500	RESCUE-1500
MMLU-Pro (399 items). Chance 0.111; constant letter B, chosen on the other partition and oracle alike, 0.133 [0.100, 0.168].						
Qwen3-1.7B	0.108 (−0.025)	0.180 (+0.048)	0.193 (+0.060)	0.180 (+0.048)	0.190 (+0.058)	0.185 (+0.053)
SmolLM3	0.000 (−0.133)	0.090 (−0.043)	0.333 (+0.201)	0.378 (+0.246)	0.291 (+0.158)	0.398 (+0.266)
Intern-S1	0.000 (−0.133)	0.105 (−0.028)	0.293 (+0.160)	0.446 (+0.313)	0.566 (+0.434)	0.454 (+0.321)
GSM8K (264 items). Constant answer 2 (chosen on the other partition) 0.030 [0.011, 0.053], 10 (oracle) 0.045 [0.023, 0.072]; no chance level.						
Qwen3-1.7B	0.057 (+0.011)	0.057 (+0.011)	0.057 (+0.011)	0.057 (+0.011)	0.057 (+0.011)	0.057 (+0.011)
SmolLM3	0.000 (−0.045)	0.705 (+0.659)	0.909 (+0.864)	0.814 (+0.769)	0.830 (+0.784)	0.886 (+0.841)
Intern-S1	0.000 (−0.045)	0.462 (+0.417)	0.845 (+0.799)	0.189 (+0.144)	0.936 (+0.890)	0.542 (+0.496)

Table 23: H3's statistics recomputed on the seeded outputs (the cards' sampling values, a seed per request) for the two models with a seeded arm; descriptive. τ_g and τ_s are the shares of THINK-512 calls cut at the token limit with greedy decoding and seeded, and $\tau_s - \tau_g$ their paired difference. D_T and D_P are the seeded drop in truncation and gain in parsing of BUDGET-512-1500 against THINK-512, with unadjusted intervals. The preregistered verdicts use family-adjusted intervals at another level and are not compared with these.

Partition	Model	Pool	τ_g	τ_s	$\tau_s - \tau_g$	D_T (seeded)	D_P (seeded)
Test	Qwen3-1.7B	LB letters	0.000	0.001	0.000 [−0.001, +0.001]	+0.001 [0.000, +0.001]	+0.001 [0.000, +0.001]
		MMLU-Pro	0.060	0.068	+0.008 [−0.011, +0.026]	+0.068 [+0.051, +0.085]	+0.066 [+0.050, +0.084]
		GSM8K	0.000	0.000	0.000 [0.000, 0.000]	0.000 [0.000, 0.000]	−0.002 [−0.006, 0.000]
	Intern-S1	LB letters	1.000	0.997	−0.003 [−0.005, −0.001]	+0.963 [+0.957, +0.969]	+0.702 [+0.689, +0.716]
		MMLU-Pro	0.915	0.902	−0.013 [−0.034, +0.009]	+0.618 [+0.584, +0.652]	+0.589 [+0.554, +0.625]
		GSM8K	0.522	0.619	+0.097 [+0.051, +0.142]	+0.581 [+0.539, +0.622]	+0.567 [+0.526, +0.609]
Replication	Qwen3-1.7B	LB letters	0.001	0.000	−0.001 [−0.002, 0.000]	0.000 [0.000, 0.000]	0.000 [0.000, 0.000]
		MMLU-Pro	0.060	0.065	+0.005 [−0.023, +0.033]	+0.065 [+0.043, +0.090]	+0.060 [+0.038, +0.085]
		GSM8K	0.000	0.000	0.000 [0.000, 0.000]	0.000 [0.000, 0.000]	0.000 [0.000, 0.000]
	Intern-S1	LB letters	1.000	0.995	−0.005 [−0.008, −0.002]	+0.959 [+0.950, +0.967]	+0.708 [+0.690, +0.725]
		MMLU-Pro	0.880	0.902	+0.023 [−0.010, +0.055]	+0.622 [+0.574, +0.669]	+0.589 [+0.539, +0.639]
		GSM8K	0.538	0.636	+0.098 [+0.045, +0.155]	+0.614 [+0.553, +0.670]	+0.602 [+0.542, +0.663]

B.5 Other descriptive items

Table 22: Seeded sampling at the models' card values (test partition): parse and truncation rates, and the difference seeded minus zero temperature with unadjusted intervals.

Model	Pool	Configuration	Parse (seeded)	Truncated (seeded)	Δ accuracy	Δ parse	Δ truncation
Qwen3-1.7B	LB letters	THINK-64	0.000	1.000	0.000 [0.000, 0.000]	0.000 [0.000, 0.000]	0.000 [0.000, 0.000]
		THINK-512	0.999	0.001	−0.013 [−0.026, −0.001]	0.000 [−0.001, +0.001]	0.000 [−0.001, +0.001]
		THINK-1536	1.000	0.000	−0.014 [−0.027, −0.002]	0.000 [0.000, 0.000]	0.000 [0.000, 0.000]
		BUDGET-512-1500	1.000	0.000	−0.018 [−0.030, −0.005]	0.000 [0.000, 0.000]	0.000 [0.000, 0.000]
	MMLU-Pro	THINK-64	0.641	0.359	−0.005 [−0.026, +0.016]	−0.016 [−0.039, +0.005]	+0.016 [−0.005, +0.039]
		THINK-512	0.932	0.068	−0.004 [−0.029, +0.021]	−0.008 [−0.026, +0.011]	+0.008 [−0.011, +0.026]
		THINK-1536	0.997	0.003	+0.004 [−0.021, +0.029]	0.000 [−0.004, +0.004]	0.000 [−0.004, +0.004]
		BUDGET-512-1500	0.999	0.000	−0.014 [−0.039, +0.011]	−0.001 [−0.004, 0.000]	0.000 [0.000, 0.000]
	GSM8K	THINK-64	0.991	0.006	−0.025 [−0.044, −0.008]	0.000 [−0.006, +0.006]	0.000 [−0.006, +0.006]
		THINK-512	0.998	0.000	−0.011 [−0.028, +0.006]	+0.006 [0.000, +0.013]	0.000 [0.000, 0.000]
		THINK-1536	0.994	0.000	−0.011 [−0.030, +0.008]	+0.002 [−0.004, +0.009]	0.000 [0.000, 0.000]
		BUDGET-512-1500	0.996	0.000	−0.008 [−0.027, +0.011]	+0.004 [0.000, +0.009]	0.000 [0.000, 0.000]
Intern-S1	LB letters	THINK-64	0.000	1.000	0.000 [0.000, 0.000]	0.000 [0.000, 0.000]	0.000 [0.000, 0.000]
		THINK-512	0.003	0.997	+0.002 [+0.001, +0.004]	+0.003 [+0.001, +0.005]	−0.003 [−0.005, −0.001]
		THINK-1536	0.192	0.810	−0.010 [−0.022, +0.002]	−0.001 [−0.015, +0.014]	+0.010 [−0.005, +0.025]
		BUDGET-512-1500	0.705	0.033	+0.042 [+0.026, +0.058]	+0.086 [+0.068, +0.103]	+0.013 [+0.006, +0.020]
	MMLU-Pro	THINK-64	0.000	1.000	0.000 [0.000, 0.000]	0.000 [0.000, 0.000]	0.000 [0.000, 0.000]
		THINK-512	0.100	0.902	+0.008 [−0.011, +0.028]	+0.014 [−0.008, +0.035]	−0.013 [−0.034, +0.009]
		THINK-1536	0.350	0.641	+0.001 [−0.024, +0.026]	+0.005 [−0.023, +0.033]	−0.011 [−0.039, +0.016]
		BUDGET-512-1500	0.690	0.284	−0.076 [−0.111, −0.041]	−0.059 [−0.095, −0.023]	+0.016 [−0.014, +0.046]
	GSM8K	THINK-64	0.000	1.000	0.000 [0.000, 0.000]	0.000 [0.000, 0.000]	0.000 [0.000, 0.000]
		THINK-512	0.387	0.619	−0.101 [−0.146, −0.057]	−0.093 [−0.139, −0.047]	+0.097 [+0.051, +0.142]
		THINK-1536	0.801	0.199	−0.049 [−0.083, −0.015]	−0.051 [−0.087, −0.017]	+0.051 [+0.017, +0.087]
		BUDGET-512-1500	0.954	0.038	−0.038 [−0.065, −0.013]	−0.023 [−0.042, −0.006]	+0.009 [−0.009, +0.028]

Table 24: Tied against untied output layer of Qwen3-1.7B-Longevity on FreeToken, the case study's six configurations, all four LongevityBench tasks pooled (8,305 rows in 1,188 participant components, unflagged rows of both partitions); computed after the outputs existed.

Configuration	Tied	Untied	Tied − untied	Parse t / u	Same answers	Only t / u right	p
CFG-THINK-64-ORIG	0.000	0.000	0.000 [0.000, 0.000]	0.000 / 0.000	1.000	0 / 0	1.000
B1	0.452	0.449	+0.003 [−0.001, +0.007]	1.000 / 1.000	0.741	182 / 157	0.192
CFG-THINKING-FALSE-64-ORIG	0.460	0.460	0.000 [−0.005, +0.005]	1.000 / 1.000	0.731	194 / 194	1.000
CFG-THINK-1536-ORIG	0.469	0.463	+0.006 [0.000, +0.013]	1.000 / 1.000	0.656	384 / 333	0.062
CFG-NOTHINK-SUFFIX-64-ORIG	0.452	0.449	+0.003 [−0.001, +0.007]	1.000 / 1.000	0.741	182 / 157	0.192
CFG-THINK-64-ANSWERONLY	0.479	0.470	+0.009 [+0.004, +0.014]	1.000 / 1.000	0.731	260 / 185	0.000

Table 25: The four GEPA runs (computed after the outputs existed): proposals and candidates, refused by the interface check, kept by GEPA, metric calls, the selected candidate, its prompt lengths, its validation score and its accuracy on the letter rows of the replication and test partitions.

Harness	Proposals	Candidates seen	Refused	Kept	Metric calls	Selected / kept	Characters (system / instruction)	Valset	Replication acc.	Test acc.	Test parse
GEPA-CLEAN-s0	207	204	28	31	6,039	0 / 31	103 / 0	0.680	0.633	0.632 [0.618, 0.646]	1.000
GEPA-BROKEN-s0	377	377	24	24	6,147	1 / 24	103 / 1,146	0.660	0.635	0.629 [0.615, 0.643]	1.000
GEPA-CLEAN-s1	195	192	36	31	6,003	12 / 31	2,469 / 922	0.700	0.641	0.631 [0.616, 0.645]	1.000
GEPA-BROKEN-s1	243	238	67	29	6,003	1 / 29	103 / 2,797	0.667	0.619	0.620 [0.606, 0.634]	1.000

Table 26: LB-0038 (age regression), test partition, descriptive.

Model	Measure	THINK-64	THINK-512	THINK-1536	AUTHORS-1500	BUDGET-512-1500	RESCUE-1500
Qwen3-1.7B	parse rate	0.000	1.000	1.000	1.000	1.000	1.000
	mean abs. error (years)	–	15.69	15.69	16.50	15.69	15.69
	exact match	0.000	0.027	0.027	0.030	0.027	0.027
SmolLM3	parse rate	0.000	0.000	0.004	0.001	0.003	0.001
	mean abs. error (years)	–	–	24.71	57.70	35.08	57.70
	exact match	0.000	0.000	0.000	0.000	0.000	0.000
Intern-S1	parse rate	0.000	0.000	0.000	0.004	0.000	0.004
	mean abs. error (years)	–	–	–	16.57	–	16.57
	exact match	0.000	0.000	0.000	0.000	0.000	0.000

Mean absolute error is over parsed rows; exact match over all rows. Constant median 50.0 years (chosen on replication): mean absolute error 16.41 (oracle median 16.41), exact match 0.013.

B.6 Sensitivity analyses

Table 27: Sensitivity analyses: the number of the primary endpoints whose verdict changes, and which.

Analysis	Endpoints	Changed	Which (primary verdict → here)
Strict scorer	31	2	H2 SmolLM3 MMLU-Pro, RESCUE: holds → inconclusive; H1 SmolLM3 GSM8K: visibility dominates → no gain: not tested
Never-called rows only	13	3	H2 Qwen3 letters, RESCUE: holds → inconclusive; H2 Qwen3 letters, GEPA-BROKEN-s1: holds → inconclusive; H1 SmolLM3 letters: visibility dominates → inconclusive
Flagged rows included (H1–H3)	31	0	none
Case study's pool (LB-0030 and LB-0034)	13	1	H2 Qwen3 letters, GEPA-CLEAN-s1: holds → inconclusive
LB-0030 alone	13	2	H2 Qwen3 letters, RESCUE: holds → inconclusive; H2 Qwen3 letters, GEPA-CLEAN-s1: holds → inconclusive
LB-0034 alone	13	0	none
LB-0042 alone	13	0	none
Role-break rule off	31	0	none
Role-break rule, first version	31	0	none
Replication and test pooled	31	0	none
Flagged rows included (H4 claims)	16	2	C SmolLM3: fails → inconclusive; C Intern-S1: inconclusive → fails

Table 28: Composition of the never-called rows (computed after the outputs existed) next to the analysis rows. The share deceased at follow-up is the share of LB-0042 rows whose follow-up turn says that the participant was deceased, whenever the death occurred; it is not the ten-year label of the task, whose gold is "Yes" for a participant who was alive at ten years or who died later.

Rows	Partition	Rows	Letter rows	LB-0030 mean age	LB-0042 share deceased at follow-up
Never called	replication	187	103	36.2	0.022
Analysis rows	replication	1,990 (pool)	–	50.8	0.345
Never called	test	356	207	37.3	0.022
Analysis rows	test	3,942 (pool)	–	50.8	0.346

C Deviations, exceptions and addenda in full

In our repository (Appendix E) the frozen preregistration (`docs/icml_ext/PREREG_ICML_EXT.md`) records exceptions and known deviations before the freeze, together with its decisions of the days before it. After the

freeze, departures are appended to `docs/icml_ext/PREREG_ICML_EXT_ADDENDA.md`, one entry each, under its section on deviation rules. This appendix lists them all: Table 29 gives the deviations that bear on a result, and the sections below give every exception, known deviation, pre-freeze decision and addendum. "Seen" states what had been seen of the study's outputs when an entry was made.

Table 29: Deviations and addenda that bear on a result; the entries below give the rest. "Seen" is what had been seen of the study's outputs when the entry was made.

Entry	What happened	Bearing on the results	Seen
Before the freeze	The primary set of H1 changed twice (after the pilot's parse rates, and after a review in Codex); the scoring defect was split in two; the tied-output-layer defect was redesigned; one evaluability rule was fixed; the family of H4 was fixed at 16; the margins were set to 0.05.	Defines the endpoints; no study accuracy was used.	pilot parse rates
Rows left out	147 replication and 371 test rows that the case study had called, and 7 rows whose answers could be recovered from an upload made by mistake, were called but left out.	Putting them back changes no verdict of H1–H3 and 2 H4 claims.	no outputs
A change of GPU	After a restart the platform assigned another GPU of the same model; the determinism check rerun on the pilot rows passed, with outputs equal to the old GPU's.	None. 9.33 GPU-hours outside the budget.	pilot rows
Two script stops	The chain script stopped twice during the GEPA search, both times from a fault in the script.	None; no run was affected.	none
The thinking-switch check	A check on Intern-S1-mini's thinking switch stopped the run that sends the reference answer, and we let it record instead of stop. The repeat reproduced the first 3,700 calls token for token, and the check failed on 324 of the run's 1,204 thinking-off calls and on none of 11,265 in the clean stages.	The reference-answer defect in Intern-S1-mini only through the descriptive selection, since the claim reads the start. 6.46 idle server-hours counted.	counts only
The checkpoint	15 of the 16 pairs were evaluable; the pilot had suggested 14.	The unparsed-as-A scoring in Qwen3-1.7B-Longevity is *not evaluable*; the gold fallback in Intern-S1-mini (bound 0.057) is evaluable.	counts and bounds
The determinism recheck	A recheck of 600 test calls per model was identical on all three models.	No verdict is relabelled descriptive.	counts
The analysis environment	The preregistered analysis ran on the instance with numpy 2.3.5, not the local 2.2.6; the resampling stream was checked equal.	None.	all aggregates
The claim rule	The two listed cases of the claim rule for the check sample are incomplete, so the general rule governs, and the pinned code implements it.	None: the verdicts are the code's, and the skewed split in SmolLM3-3B fails, a descriptive label.	all aggregates
Computed after the outputs	The composition of the never-called rows, the trivial baselines of the three pools, the tied-versus-untied comparison, the GEPA candidates and the serving counts.	Labelled, and they decide nothing.	aggregates

C.1 Exceptions accepted before the freeze (X1–X9)

X1 Three `RECORD` entries of the vLLM virtual environment do not match their files (two PyNvVideoCodec libraries and `build_backend.py`); they date from the original installation, not from the clone. Decided 2026-09-28.

X2 FreeToken 0.1.3 opens a process-group store on its serving port plus one that PyTorch binds off loopback; it carries no model traffic. The runner allows exactly this one listener, on FreeToken only, while no connection to it comes from off the machine. The case study's FreeToken runs very likely had it too.

X3 `FREETOKEN_ALLOW_CUDA_MISMATCH`, the case study's variable, is exempt from the credential-name check although its name contains "token".

X4 `VLLM_USE_FLASHINFER_SAMPLER=0` in every vLLM run, because FlashInfer's sampling kernel could not compile against the image's CUDA toolkit.

X5 A change of instance on 2026-09-28: the source instance's GPU was no longer available, so the instance was cloned to one of the same configuration; the GPU identifier changed and the clone was verified against the recorded freeze, wheel, record and model hashes.

X6 A restart of the container at 2026-09-29 00:57:25 UTC, which we did not initiate; no server was running, and the GPU identifier and the SSH settings were unchanged.

X7 The first gate check (G1) ran on the checkpoint's own directory, whose configuration declares tied embeddings, before the output-layer decision; its results were operational only.

X8 AI-assisted readings and uploads (item-exposure contacts 1 and 2); the reading of the main export (contact 3) and the copy of the item files to the project's own instance (contact 4) are in Appendix F.

X9 A change of the physical GPU after a restart (addenda A1, A2).

C.2 Known deviations from the plans (K1–K9)

K1 Three models instead of four; the FreeToken continuity on Qwen3-1.7B-Longevity only.

K2 GEPA with a local proposer replaces Meta-Harness, so the proposer through an API and its budget fall away.

K3 A single answer allowance of 1,500 tokens replaces a calibrated rule that left out the calls the calibration allowance cut.

K4 The batch-invariant server with batching replaces concurrency one on the standard server (about 275 GPU-hours at concurrency one).

K5 The pilot's LongevityBench rows come from the case study's evolution file or validation shards, not from never-called rows; calibration reported no accuracy.

K6 The extended scorer rules, their narrowing and the role-break rule were decided after the pilot's parse audit and spot-checks on pilot items; the strict parser is a preregistered sensitivity analysis.

K7 The GEPA interface changed between its two spikes, on the counts of refused proposals.

K8 The environment that counts tokens locally differs from the instance's in the Python and numpy versions and other bundled tools; none takes part in tokenization, and the render lists were checked on both machines.

K9 The pilot report was written after the first draft of the preregistration, from aggregates only.

C.3 Decisions between the first draft and the freeze

We revised the draft after our decisions on its first version, after a review in claude.ai, after two reviews and a final check in Codex (each given the file and not our repository), after our decisions on the summary of the drafting session, and after the reading of the rendered requests. Table 29 calls these decisions "Before the freeze". The decisions that shape the endpoints were these:

- The primary set of H1 changed twice, first after the pilot's parse rates and parsed answers were seen (no accuracy was used) and then after the first review in Codex, with THINK-64 against AUTHORS-1500 moved out to nine sanity checks.
- The scoring defect B was split into B1 and B2.

- The tied-output-layer defect E was redesigned so that the confirmatory instance serves the tied variant on both samples.
- One evaluability rule applies to every defect, and the family of H4 is fixed at 16 whatever is evaluable.
- The margin and the minimal difference are 0.05, with Bonferroni within each family.

After the reading, seven replication rows whose reference answers could be recovered from render-hash lists uploaded by mistake were flagged and left out of every analysis (the reading itself, item-exposure contact 3, concluded pass); Table 29 calls this "Rows left out". The exclusion came after the last review in Codex and was not reviewed there.

C.4 Addenda after the freeze (A1–A12)

A1, A2 (2026-09-30) *A change of GPU.* After we restarted the instance, the platform had assigned a different physical GPU of the same model. By the preregistration's rules the work stopped before any copy or call. Every check passed on the new GPU; the determinism check (G2 (b)) was rerun on the pilot rows and passed for all three models, and its outputs equal the old GPU's token for token. 9.33 GPU-hours, outside the budget; the new identifier is pinned. Seen: pilot rows only.

A3 (2026-09-30) *Two script stops.* The chain script stopped twice during the GEPA search, both times from a fault in the script (a password check that failed under a shell option; a step that required an empty GPU while the task server stayed up). No run, call or check was affected. Seen: nothing of replication or test.

A4, A5 (2026-10-01, 2026-10-02) *The thinking-switch check.* The defect-A pass of Intern-S1-mini stopped itself at gate (b): five of its first thinking-off calls (AUTHORS-1500) were not clean (cut at the limit with a think block opened and not closed). The gate held on all 3,849 thinking-off calls of the clean runs up to then. The preregistration had not foreseen that the injected defect could trigger the gate. We decided that in that stage only the gate records instead of stopping (commit `aae8b11`, runner `a9478cee`), that the idle server stops and its hours are counted (about 6.46 of its server-hours were idle), that the partial folder is moved out of the run root, and that the run is repeated from scratch under a new name. The repeat reproduced all 3,700 calls of the failed attempt token for token and failed the gate on 324 of the pass's 1,204 thinking-off calls (the gate reads the completion, and all 1,204 prompts carry think tags that the chat template writes around the reference turn). By the repair rules this is confirmatory: the failed attempt is kept and reported, and only the repeat enters the analysis. Seen: counts only.

A6 (2026-10-02) *The checkpoint.* It ran on the instance, since the local machine lacked memory. All 17 replication runs were complete with matching manifests; the derived rescue equals the real two-call sequence on 300 of 300 verification items per model; 15 of 16 pairs are evaluable; and the selections of the pipelines whose selection sample lies in replication were locked in a record (`f033f23f`) whose hash was committed before the first test call. Seen: counts, bounds and the record's hash.

A7 (2026-10-02, 2026-10-03) *The determinism recheck.* The test partition ran 14 runs and 267,056 calls, with no retry and no stop, and the recheck of 600 calls per model was identical on all three; the study used 54.53 of the 77 GPU-hours of the cost bound. Seen: counts.

A8 (2026-10-03) *The analysis environment.* The preregistered analysis ran once on the instance (numpy 2.3.5, not the local 2.2.6; the resampling stream was checked equal for three seeds) after every pinned file matched its frozen hash, and the selections it reports equal the locked ones on all 24 comparisons. The code that computes three descriptive items the pinned driver lacks (tied versus untied, GEPA candidates, serving counts) was written after the outputs existed and is labelled so. Seen: all aggregates, for the first time.

A9 (2026-10-03) *Computed after the outputs.* We shut the instance down and do not release it; the report and its tables were committed as aggregates (`49e5f9c`); the never-called rows' composition (543 rows, 310 on the letter tasks) and constant-letter baselines were computed after the outputs existed.

A10 (2026-10-03) *Computed after the outputs.* Constant-label baselines, because a constant letter is a weak floor where the option order is shuffled: the strongest baseline of the pool is 0.539 on test against 0.389 for the oracle constant letter. The strongest baseline is the one used in this paper.

A11 (2026-10-03) *The claim rule.* The rule for D_R in the claim of H4. The audit of this paper recomputes the claim from its three adjusted intervals; the two sub-cases that the preregistration lists after its general rule were incomplete, the general rule governs, and the pinned code implements it (our decision). No verdict changes. Seen: all aggregates. Details below.

A12 (2026-10-03) *Computed after the outputs.* The trivial baselines of MMLU-Pro (chance and the constant letter) and GSM8K (the constant answer), chosen on the other partition and oracle, next to every model's accuracy in every grid configuration (holdout contact #57). Seen: aggregates, no item key or text.

C.5 The rule for D_R (addendum A11)

For D_R the preregistration says that it fails if its interval lies "wholly on the other side of its boundary" and lists two sub-cases, wholly within the band and wholly outside it with the opposite sign. The pinned code (`scripts/icml_ext/analysis_h4.py`, `claim_verdict`) implements the general phrase: D_R is graded in the direction of D_S's sign, holds if it lies wholly beyond the band there, and fails if its interval lies wholly on the other side of that one boundary, which also covers an interval partly within the band and partly beyond its far edge. The general rule governs and the two sub-cases were incomplete. It matters for one pair, C in SmolLM3-3B: $D_S = -0.001$ points down, so D_R would have to lie below -0.05, and its adjusted interval $[-0.048, +0.063]$ lies wholly above it, so the claim cannot hold and is *fails*, a descriptive label. Read as the two sub-cases only, the claim would be *inconclusive*, and H4 would count 6 holds, 3 fails and 6 inconclusive instead of 6, 4 and 5; every other claim has the same verdict under both readings.

C.6 What this paper adds

This paper's text, tables and figures format aggregates committed in our repository and add no computation except the counts of its own tables (verdict counts by family, the largest and smallest of a column) and the independent recomputation of every verdict from its interval, which agrees with the report on every endpoint that has a verdict (Appendix E).

C.7 The second study: decisions before the freeze, and addenda A1–A9 after it

The decisions between the second study's first draft and its freeze, the exceptions it accepted and the differences from the first study are listed in `docs/icml_ext2/DIFF_FROM_FROZEN.md`. Its addenda after the freeze are in `docs/icml_ext2/PREREG_ICML_EXT2_ADDENDA.md`; each states what happened, the rule applied, the outputs seen at the time and the status.

A1: a pinned file was missing on the instance. The checkpoint stopped when it first ran, because the list of replication rows flagged as exposed, which the frozen code pins by hash, was not on the instance; the deployment script had omitted it. No run had been read. The file was copied by hand from our machine, its SHA256 was checked against the pin, the copy was logged first as a holdout contact, and the checkpoint was run again alone.

A2: the checkpoint record. The checkpoint record, which holds the selections, was not opened; its SHA256 (given in Appendix E) was logged and committed before the first test call, and later checks compared hashes only.

A3: the shutdown of the instance. A platform shutdown on 8 October 2026 ended the instance because the account balance ran out; it was not a platform fault. The 14B's test grid had finished and verified, and the evaluation of one selected harness on the test rows had started. No output had been read before or after.

A4: the clone, its verification and the change of GPU. We cloned the instance to another host, with another RTX PRO 6000 of the same model and driver. The clone was verified file by file, by hashes and counts of the environment, the models, the item files, the frozen guard, every completed run and the checkpoint record, before anything else ran; nothing was missing. The interrupted evaluation, which had no manifest and a few recorded responses, was moved out of the run folder as a failed attempt, kept and excluded from every analysis, and the evaluation was rerun from scratch with the same requests. The preregistration's rule for a change of the physical GPU after the first replication call applied.

A5: the finished test grid. We decided that the 14B's test grid, which had completed and verified on the first GPU, is kept, and that the recheck of the replication subset decides: an identical result for Qwen3-14B lets it stand, and any difference would have stopped the sequence for our decision.

A6: the recheck. On the new GPU the batch-invariance check on the pilot rows was identical for the three models that were still to be called or rechecked, and 900 replication calls of each model in three configurations were identical token for token to the outputs of the first GPU; part 2 resumed as confirmatory.

A7: the guard and the watchdog. While we were away the instance carried a guard that shut it down with the platform's own command when part 2 and the analysis had finished and the outputs were copied, or at a fixed time at the latest, and a watchdog on our machine checked it later; neither was triggered by a fault.

A8: the copy of the outputs. Part 2 finished with every check passed and no stop. The external drive that holds the repository's run folders was not available, so the outputs were copied to a folder on our machine as compressed archives; every file was checked against the instance's list and every manifest against its files, the marker that the copy was verified was written, and the guard shut the instance down.

A9: the second copy, the extraction and the descriptive analyses. Windows had kept the external drive offline by policy and an administrator set it online; the archives were copied to it, extracted and verified again. The descriptive analyses that the preregistration specifies (the two 1.7B arms with the reproduction record, the agreement with the first study, the non-commitment and refusal shares) were run, each after its holdout contact was logged, and the four aggregate reports were scanned for item keys, row fields and item text before they were committed.

D Expected in advance against observed

The preregistration states what it expected for several endpoints, from the pilot's parse and truncation rates and parsed answers (never from accuracy). Table 30 sets each expectation next to what was observed. The directions of H1 and H3 were visible in the pilot, and the preregistration says so; this table shows the places where the outcome differed from the expectation, which are as informative as the places where it did not. "Inconclusive" in the Expected column is an informal expectation read off planned power under assumed effects (Appendix K), not a prediction with a decision rule.

Table 30: Expected in advance (preregistration, the sections on H1, H2 and H4) and observed.

Item	Expected	Observed
H1 sanity checks (THINK-64)	$T > 0$ in eight cells; Qwen3-1.7B-Longevity on GSM8K fails the gate	*visibility dominates* in 8 of nine cells; Qwen3-1.7B-Longevity on GSM8K fails the gate. As expected.
H1, Qwen3-1.7B-Longevity (three cells)	Fails the gate (on MMLU-Pro, or quality matters)	*no gain: not tested* in all three. As expected.
H1, Intern-S1-mini on the letters	$T > 0$ nearly inevitable	*visibility dominates*, $T = +0.063$. As expected.
H1, SmolLM3-3B and Intern-S1-mini on GSM8K	Inconclusive on T (planned power 0.52 and 0.00 for SmolLM3-3B, 0.45 and 0.06 for Intern-S1-mini, in the two scenarios)	SmolLM3-3B: *visibility dominates*, $T = +0.068$ $[+0.047, +0.090]$. Intern-S1-mini: the gate fails, $\Delta = -0.287$. Both differ.
H2, Intern-S1-mini on GSM8K against RESCUE-1500	Inconclusive (planned adjusted power 0.19)	*fails*, $\theta = -0.342$ $[-0.402, -0.281]$. Differs.
H2, GEPA harnesses	Power unknown in advance; may well be inconclusive	*holds* for all four; all intervals contain zero.
H4, evaluable pairs	14 of 16: B1 in Intern-S1-mini and B2 in Qwen3-1.7B-Longevity not evaluable	15 of 16: only B2 in Qwen3-1.7B-Longevity is not; B1 in Intern-S1-mini has bound 0.057. Differs.
H4, bounds on the pilot's 214 letter rows against S	B1 0.201, 0.098, 0.028; B2 0.000, 0.196, 0.276; D 1.000, 0.327, 0.150; E 0.159 (Qwen3-1.7B-Longevity, SmolLM3-3B, Intern-S1-mini)	B1 0.205, 0.083, 0.057; B2 0.000, 0.208, 0.289; D 1.000, 0.352, 0.189; E 0.139.
H4, D	The claim rests on Γ and holds only if its point estimate is within about 0.01 of zero (about 0.50 of the time with no true difference)	Holds in all three models, $\Gamma = -0.003, -0.007, +0.001$. Consistent with the expectation, which allowed it.
H4, B2	Can be large in SmolLM3-3B and Intern-S1-mini	$D_S = +0.208$ and $+0.289$. As expected.
H4, E	No expectation on its size	$D_S = +0.005$, $D_R = +0.014$; *fails*.
Detection matrix	Each defect detected by its signature rules (A: TR, IR1, IR2, TB3; B1: IR3; B2: TB4; C: IR4, IR5, TB1, not TB2 or PR; D: SCF1, SCF3; E: SMI)	All expected rules fired in 12 of 16 instances. TB1 did not fire in any of the three C instances, TB3 did not fire for A in SmolLM3-3B, and TB2 fired in all three C instances, with the gain rule of PR in Qwen3-1.7B-Longevity.

E Reproducibility

Every path, commit, tag and SHA256 value in this paper is that of a file in our repository, which is not public at present (see the statement on code and data availability after the Conclusion). They are kept so that the planned public export can be checked against them; a reader cannot open these files until the export is linked.

The frozen record. In our repository the preregistration (`docs/icml_ext/PREREG_ICML_EXT.md`) was frozen at 06:53:11 UTC on 30 September 2026, in commit `7f645ae`, whose parent carries the tag `pre-icml-ext-freeze`; the tag `prereg-icml-ext-frozen` marks the freeze. Its SHA256 is

`dec5dc25fdd0074a973926767a7a7f378e3dbc06eaccbf976d6df6f84dd2f515`

The file `docs/icml_ext/PREREG_ICML_EXT.sha256` holds the hash and the time, and the addenda are in `docs/icml_ext/PREREG_ICML_EXT_ADDENDA.md`. The statement on code and data availability says which files are attached to this arXiv version and why the freeze time cannot be checked independently. The hash

above lets a reader check the ancillary preregistration, and the report's hash below does the same for the report.

The outputs. The report of the preregistered analysis and its tables are committed as aggregates in our repository, in `docs/icml_ext/ANALYSIS_20261003/`: the report has SHA256 `1447c936d52523a1494169f30af239c7d36ce7e79e81655a6a91016570c22e0b`, the folder's manifest `37d3268a980974af6ab414af65b8b2b1a5b3d6071632f8f1867ffc3d871a8403`. The items computed after the outputs are in three subfolders (the never-called rows and the constant-letter baselines; the constant-label baselines; the baselines of MMLU-Pro and GSM8K), each with its own manifest. The study's 17 replication runs and 14 test runs, the GEPA search and the failed attempt of A4 hold item text and model output, are not committed and will not be released; each has a manifest of SHA256 values that matched its files after the copy back.

The code. The files of our repository that fix the analysis are pinned in the preregistration by SHA256: `scripts/icml_ext/analysis_stats.py` (`e6a2c4dc`), `analysis_records.py` (`7507fd6c`), `analysis_h4.py` (`2f5a82e7`), `analysis_run.py` (`d92ebfe2`), `defects.py` (`a2ad7e2e`), `scoring.py` (`3bce48ef`) and `runner_checks.py` (`ec77476e`), and the caller `scripts/paper_review/messages_for_inference.py` (`cc8debc7`). The runner changed once after the freeze (addendum A4, commit `aae8b11`). The environment is vLLM 0.30.0 with the packages listed in the setup log (`docs/icml_ext/SERVER_SETUP_LOG.md`); the analysis ran with numpy 2.3.5. vLLM pull request 41674 (merged 2026-05-15) fixed a thinking-token budget that was silently ignored (JasonKeyiL (vLLM project), 2026); the version used contains the fix, and H3 shows the budget acting.

This paper's numbers. Every number of the text is a value macro (`\V` applied to a key) that `tools/values.py` registers with the file of our repository that it reads (at git `HEAD`, never the working tree) and the path inside it; `tools/gen.py` writes `common/numbers.tex` and the tables from that registry, and `-check` fails if a committed generated file would change. `audit/number_audit.py` then (i) re-reads every registered value with its own path reader and formatter, (ii) recomputes the verdict of every endpoint that has one (all but the unevaluable pair of the 65 primary endpoints of the two studies) from its interval with the preregistered rules and compares it with the report's, (iii) fails on any number in the hand-written text that is not a macro or an identifier, (iv) checks every hash, path and tag named in the text against our repository, and (v) checks the plotted values of the figures (`tools/make_figures.py`). `audit/CLAIM_AUDIT.md` maps each claim of the text to its evidence.

The second study. Its preregistration (`docs/icml_ext2/PREREG_ICML_EXT2.md`) has SHA256

`a2b8ec439d5ba941e8de2bda74fb74600f8fe9e7f0650a74e673cedf4ac85b61`

it was frozen by the commit that adds its hash file, with the tags `pre-icml-ext2-freeze` and `prereg-icml-ext2-frozen`. The checkpoint record, which locks the four selected harnesses and was committed by hash before the first test call, has SHA256

`4aa7da7484e5f117cf9a17180ce672c1b666e20d3ec3bd878fbc8afd88656a28`

The environment of the instance equals the first study's by the SHA256 of its package list,

`fee9f31063305f6754e1020f9ad8cd063bc24f2b6ca59fca72924a938bc2f16d`

on an RTX PRO 6000; the clone of the instance after the shutdown was verified against the same hashes (Appendix C.7). The aggregate reports of its analyses are in `docs/icml_ext2/ANALYSIS_20261023/` and the numbers computed after the outputs existed in `docs/icml_ext2/FOLLOWUP_JSON_20261023/`.

F Exposure and disclosure

Every contact of ours or of our tools with held-out material is logged in our repository (`docs/paper_review/HOLDOUT_EXPOSURE_LOG.md`, numbered contacts;

`docs/icml_ext/ITEM_EXPOSURE_LOG.md`, contacts with the extension's items by outside services), and before it happens, with one exception that Table 31 marks. The case study's holdout contacts are #1 to #38; those of this study are #39 to #58 and #60. Other later contacts, including the preparation of the ancillary files and the posting of this version, are logged in the repository and are not listed here. Item-exposure contacts 1 and 2 are uploads of the pilot's reading export and runner and of a spot-check and the GEPA code to claude.ai; contact 3 is the AI-assisted reading of the main export (concluded pass; one of us relayed the reading and its conclusion and also checked the conclusion); contact 4 is the copy of the main item files to the project's own GPU instance, which is not an outside service.

Table 31: Holdout contacts of this study (#39 to #58 and #60): what was read or done. Entries #50 to #58 and #60 are after the freeze.

#	Contact
#39	The ICML formatting draft of the case study copied to a synchronised folder for an upload.
#40	The LongevityBench split of this study, the earlier-contact flags and the first gold audit built.
#41	Fourteen files uploaded to claude.ai for a review of the reading export and the runner (item-exposure contact 1).
#42	The preregistration draft revised and uploaded for reviews in claude.ai and in Codex, after a check for holdout-derived values.
#43	The main item files built.
#44	The LB-0042 rows on which B1's fallback gold differs from the derived gold counted.
#45	The gold audit rerun over the main items with the final scorer.
#46	The main reading export run.
#47	The three render-hash lists uploaded to claude.ai by mistake together with the export (see below).
#48	The main render-hash lists pinned and checked locally.
#49	Seven replication rows flagged as reading-exposed and the counts redone.
#50	The main item files, the GEPA arm's rows, the render lists and the frozen code copied to the instance.
#51	Counts-only scripts over the stopped run of Intern-S1-mini's defect-A pass (addendum A4).
#52	The checkpoint run on the instance (A6).
#53	Counts of the test partition's 14 runs (A7).
#54	The preregistered analysis (A8).
#55	The report and its tables committed as aggregates; the never-called rows' composition and the constant-letter baselines computed (A9).
#56	The constant-label baselines computed (A10).
#57	The trivial baselines of MMLU-Pro and GSM8K computed (A12).
#58	The named preprint (its PDF and LaTeX source, nothing else) copied to a folder outside the repository and given to Codex for an independent review at about 04:23 Pacific time (11:23 UTC) on 2026-10-03; the review returned 16 findings, which this version applies without new computation.
#60	The PDF of the named preprint, as built at 03:42 Pacific time (10:42 UTC) on 2026-10-03, reviewed in claude.ai before the copy of #58 was made. It was missed at the time and logged afterwards.

The mistaken upload (#47) and its consequence. One of us uploaded the three render-hash lists to the claude.ai conversation of the reading, together with the export. A defect-A entry is the hash of a request that carries the row's reference answer, so for a row whose clean request is known, the answer can be recovered by testing each candidate against the hash. Seven replication rows are both in the reading's sample and among the defect-A rows, and we flagged them and left them out of every analysis, whether or not anyone tried. Nothing else in the replication or test partitions was disclosed this way. The reader's statement, relayed by one of us and not checked independently, is that the reading did not open the lists and did not attempt a reconstruction.

Files derived from held-out rows. This paper and the committed analysis report carry aggregates that include the rows flagged `v21_holdout` (the sensitivity analysis), so both are listed in `docs/paper_review/HOLDOUT_DERIVED_FILES.md` of our repository, and they hold no row index, key or content. The run logs, which hold item text, stay off our repository and will not be released. We drafted

the paper with Claude Code and reviewed it in claude.ai and in Codex. Contact #58 is the copy given to Codex, which was the two files of the named preprint and not the repository, a run log or any item text, and contact #60 is the PDF reviewed in claude.ai. Any further upload of the draft is a new contact, logged first.

The second study. Its contacts with item text are logged in the same holdout log and in the item-exposure log: the copy of its draft preregistration that was given to a reviewer (#62), the copy of the main item file and the frozen code to the instance (#63) with its notes (the missing list, the checkpoint, the clone, the copy of the outputs, the second copy, each analysis and each scan), and the item-exposure log's contact 5. No call touched a replication or test item before the second freeze, and we saw only counts and hashes of the main study's outputs until its analysis report existed.

G Serving details

Servers. The main study ran each model on a vLLM server started by a pinned script with these flags (the model directory, served name and chat-template file differ by model; Intern-S1-mini adds `-trust-remote-code` and a limit of zero images and videos):

```
vllm serve <model directory> --host 127.0.0.1 --port 8000
  --dtype bfloat16 --served-model-name <name>
  --chat-template <pinned file> --generation-config vllm
  --reasoning-parser qwen3
  --reasoning-config '{"reasoning_start_str": "<think>",
                       "reasoning_end_str": "</think>"}'
  --no-enable-prefix-caching --no-enable-chunked-prefill
  --max-model-len 4096
```

with the environment cleared to a fixed list of variables, among them `VLLM_BATCH_INVARIANT=1` and `VLLM_USE_FLASHINFER_SAMPLER=0` (exception X4). Every request sets thinking on or off explicitly, carries a seed (irrelevant at zero temperature) and asks for the token IDs and the rendered prompt; the runner compares both with its own render on every call and stops on a difference, except in the defect-A pass, where the reference answer makes them differ by design. Each stage runs its start checks before the first call and again after the last (no credential-like variable name, no connection off the machine, loopback-only listening, the model directory and the hash of its configuration, the template, the flags, and the render check against the pinned list), and a watcher samples the server's connections from launch. The watchers flagged nothing on the vLLM servers; on each FreeToken server they flagged only the single listener of exception X2.

Cost. The GEPA search used 6.63 GPU-hours, the replication partition 22.07 and the test partition 25.83, 54.53 of the 77 of the cost bound; the rerun of the determinism check on the new GPU (9.33 hours) is outside it.

Per-run counts. Table 32 gives, per run and configuration, the calls, the disagreements between the server's reasoning split and the local one (logged on every call and never a stop; in the pilot they all fell after a forced end of thinking, or where the limit cut a multibyte character, and were equal once the four tags were removed), the calls in which a start-of-turn marker re-opened a turn (role breaks, all in SmolLM3-3B's BUDGET-512-1500), and mean tokens and seconds per call. The counts include every row called, flagged rows and LB-0038 among them.

Table 32: Serving counts of the zero-temperature grid.

Partition	Model	Configuration	Calls	Split disagreements	Role breaks	Prompt tokens	Completion tokens	Seconds
Replication	Qwen3-1.7B	THINK-64	3,603	0	0	480	56	1.0
		THINK-512	3,603	3	0	480	257	3.7
		THINK-1536	3,603	3	0	480	258	3.7
		AUTHORS-1500	3,603	3	0	484	3	0.2
		BUDGET-512-1500	3,603	4	0	480	257	3.6
	SmolLM3	THINK-64	3,603	0	0	500	64	1.8
		THINK-512	3,603	0	0	500	461	10.8
		THINK-1536	3,603	0	0	500	1093	26.4
		AUTHORS-1500	3,603	0	0	472	297	7.2
		BUDGET-512-1500	3,603	2,431	2,177	500	1148	28.1
	Intern-S1	THINK-64	3,603	1	0	493	64	3.0
		THINK-512	3,603	1	0	493	502	17.7
		THINK-1536	3,603	0	0	493	1294	47.0
		AUTHORS-1500	3,603	0	0	480	651	23.6
		BUDGET-512-1500	3,603	233	0	493	799	29.1
Test	Qwen3-1.7B	THINK-64	7,216	0	0	482	56	1.0
		THINK-512	7,216	0	0	482	256	3.6
		THINK-1536	7,216	0	0	482	257	3.7
		AUTHORS-1500	7,216	0	0	486	2	0.2
		BUDGET-512-1500	7,216	2	0	482	256	3.6
	SmolLM3	THINK-64	7,216	0	0	501	64	1.8
		THINK-512	7,216	2	0	501	461	10.8
		THINK-1536	7,216	0	0	501	1107	26.6
		AUTHORS-1500	7,216	0	0	473	311	7.5
		BUDGET-512-1500	7,216	4,893	4,384	501	1163	28.5
	Intern-S1	THINK-64	7,216	0	0	494	64	3.0
		THINK-512	7,216	2	0	494	502	17.7
		THINK-1536	7,216	2	0	494	1296	47.2
		AUTHORS-1500	7,216	0	0	482	644	23.4
		BUDGET-512-1500	7,216	474	0	494	804	29.3

Completion tokens of the second study. Table 46 in Appendix M gives the mean completion tokens per call of Qwen3-14B's configurations and harnesses.

H How the related work was searched

We searched arXiv, Semantic Scholar, OpenAlex, Hugging Face papers, GitHub and venue pages from 2024-01 to 2026-10-03, with topic queries only (no item content), starting from the seed papers and their references and citers, in six topic slices: harness and agent-system evaluation; optimisers and agent baselines; reasoning evaluation under token budgets; evaluation-pipeline defects; inference engines and checkpoint identity; preregistration and replication. The agents behind the slices made 1,006 web calls. A shared cap of 200 web searches ran out during the search, the arXiv and Semantic Scholar interfaces were rate-limited, and three of the six slices (reasoning evaluation, inference engines, pipeline defects) were stopped by a rate limit and resumed to write up what they had already read, without new searches. "We found no study that" therefore means that none was found, not that none exists, and most recent work on harness evolution is an unreviewed preprint.

Identifiers, titles, authors and dates of the cited papers were checked against arXiv. The statements in Section 2 about a paper's design were checked against its text for the papers marked as read in the matrix in our repository (`docs/icml_ext/RELATED_WORK_MATRIX.md`), and otherwise rest on its abstract. For the GitHub pull requests cited, the merge date and the release that contains the merge commit were checked through the GitHub interface.

I H4 in detail

This appendix restates, from the frozen preregistration, what Section 3.4 summarises: the quantity each claim reports, how each defect is built and what it is compared with, the rows of its samples, the bound that decides whether its pair is evaluable, and the predicate and threshold of every check.

The reported quantity and the claim. For an instance I (a defect and a model) and its clean counterpart K, the reported quantity is the accuracy of a fixed configuration, the pipeline's start s, on the letter pool: $a_S = \mathrm{acc}_S(s)$ on the selection sample S and $a_R = \mathrm{acc}_R(s)$ on the check sample R. The start is THINK-1536 (for E, the case study's CFG-THINK-1536-ORIG), and THINK-64 in the instance of D. No claim selects anything. Then $D_S = a_S(I) - a_S(K)$, $D_R = a_R(I) - a_R(K)$ and $\Gamma = D_R - D_S$. The claim *holds* if the adjusted interval of D_S lies wholly outside $\pm m$, that of D_R lies wholly outside $\pm m$ on the side of D_S's point estimate, and that of Γ lies within $\pm m$. It *fails* if any of the three lies wholly on the other side of its boundary (D_S's wholly within $\pm m$; D_R's wholly within $\pm m$, or wholly outside it on the opposite side; Γ's wholly outside $\pm m$), and it is *inconclusive* otherwise. Here $m = 0.05$ and the intervals have the adjusted level of the family of 16; an interval that fails decides the claim even if another is still open, and the error control covers "holds" only.

Defect D and the selection. For D the reported quantity is the gain $g = \mathrm{acc}(c^*) - \mathrm{acc}(s)$ of the selected configuration c^* over the start, on S and on R. The selection picks c^* as the most accurate member on S of the search space (THINK-64, THINK-512, THINK-1536, RESCUE-1500 and BUDGET-512-1500; ties go to the earlier member), and AUTHORS-1500 lies outside it, so that it can act as the simple comparator that a search never considered. Because c^* does not depend on the start, the instance of D and its counterpart select the same c^* in every resample, the selection cancels, and $D_S = \mathrm{acc}_S(\text{THINK-1536}) - \mathrm{acc}_S(\text{THINK-64})$, with the same difference on R for D_R. The selection itself, the gains it reports and every interval that follows a selection are descriptive and enter no claim.

Defects, comparators and samples. Table 33 gives for each defect how it is built, the samples S and R, the clean counterpart it is compared with and how the two are resampled, and the bound of its pair. All samples leave out the flagged rows (Section 3). Where an instance and its counterpart share their rows, one resample of the units serves both, so D_S and D_R are paired; Γ pairs the b-th resample of S with the b-th of R, drawn with their own seeds.

The skewed split is resampled in two pieces. The skewed split's instance and its clean counterpart use different rows but the same participant components: the instance's halves are dealt from the unflagged rows of both partitions, and the counterpart's samples are the replication and test rows. The frozen analysis resamples the two independently, each with its own seed, so the covariance between them is ignored, and the intervals of D_S, D_R and Γ for this defect may be too wide (if the covariance is positive) or too narrow (if it is negative). The four claims for C, one failing (SmolLM3-3B) and three inconclusive (Qwen3-1.7B-Longevity, Intern-S1-mini and Qwen3-14B), are the ones that could depend on it; for every other defect one resample serves the instance and its counterpart.

Table 33: The six defects of H4: construction, samples, clean counterpart and evaluability bound. The bound is the largest distortion the defect's mechanism can cause on S, computed without accuracy.

	Construction	Samples S and R	Clean counterpart and resampling	Evaluability bound
A	The row's reference answer is sent as a final assistant turn, as the case study's caller did: new calls, built by a separate module.	1,204 replication rows drawn by whole component (seed 20261012) in proportion to the task counts (344, 172, 344 and 344 rows of LB-0030, LB-0034, LB-0038 and LB-0042), halved by component into S and R (86 components and 602 rows each; after the flags 562 and 565 rows in 81 components each). Both halves are called in one pass, so R is a disjoint half, not a later sample.	The grid's own outputs on the same rows, in the same halves; one resample serves both.	One: the leaked reference can change any row.
B1	Offline: the clean outputs are rescored with LB-0042's gold taken from the case study's historical fallback, which with this release's labels differs from the derived gold. Unparsed answers stay wrong; other tasks keep the release's gold.	S: the replication partition's letter rows; R: the test partition's, called after the checkpoint. The fallback changes the gold of 408 of 796 LB-0042 rows of S and 794 of 1,577 of R.	The main clean instance on the same rows; one resample serves both.	The share of S's letter rows that the start parses and whose gold B1 changes (exact; B1 changes the gold of 0.2050 of them).
B2	Offline: the clean outputs are rescored by a scorer that never fails closed: an unparsed or conflicting letter answer is scored as the letter A, on every letter task, with the release's gold.	As B1.	As B1.	The share of S's letter rows that the start leaves unparsed or conflicting and whose gold is A (exact; it equals B2's distortion of the start's accuracy).
C	The case study's skewed split, re-implemented. The pooled LongevityBench rows of both partitions that are not flagged (8,305 of 8,830) are reassigned: components (from all six NHANES tasks, chaining into one giant component that goes to the evolution part) are sorted by size, largest first, and each goes to the part furthest below its target share (evolution 0.50, validation 0.30, holdout 0.20).	S and R are the two halves of the validation part: its components are sorted by identifier, shuffled (seed 20261016) and dealt alternately. Their rows come from both partitions, so they are not the replication and test rows of the other pairs, and C, computed after the test calls, is not temporally prospective.	The main clean instance, whose S and R are the study's replication and test rows. The rows differ, so the two are resampled independently, each with its own seed.	None: C is evaluable by rule, because its mechanism changes the sample and has no row-level bound.
D	Offline: the start is THINK-64 in place of THINK-1536; the reported quantity is the gain over the start.	S: the replication partition's letter rows; R: the test partition's.	The main clean instance; one resample serves both.	The share of S's letter rows on which the parsed answers of the two starts differ, an unparsed or conflicting answer counting as one value.
E	The tied output layer is served in place of the untied one (Qwen3-1.7B-Longevity only): the case study's six configurations on FreeToken, for both variants, with CFG-THINK-1536-ORIG as the start. The confirmatory instance serves the tied variant on S and on R.	S: the replication rows; R: the test rows.	The untied variant on both samples; one resample serves both. A second instance (tied on S, untied on R) feeds only the descriptive detection matrix.	The share of S's letter rows on which the tied and untied variants' parsed answers of the start differ.

Evaluability. A pair is evaluable if its bound exceeds m (a bound of exactly m does not). Otherwise it is reported as not evaluable, with its bound, makes no claim and is not counted as failing, and the family stays at 16. The bounds were computed from parse, gold-change and answer-difference indicators on the replication data at the checkpoint, before the first test call. For B1 the start must parse the same rows in the instance and its counterpart, and for B2 the gold must be the release's in both; the analysis stops if not.

The checks and their rules. Each check is evaluated on the records of an instance, and "fires" means that one of its rules holds. The rules and thresholds were fixed before the study; their outcomes are descriptive, and what is empirical is how often they fire on clean instances (Section 4). The rules of PR and every interval that follows a selection on S are descriptive labels, not tests: the selection makes $\mathrm{acc}_S(c^*)$ and the gain optimistic. Tables 34 and 35 give each rule's predicate and threshold.

Table 34: The rules of PR, TR and TB. c^* is the selected configuration, s the start, g_R the gain on R.

Rule	Fires when
PR, gain rule	$c^* \neq s$ and the lower bound of the 95% percentile interval of $\mathrm{acc}_R(c^*) - \mathrm{acc}_R(s)$ is at or below zero: the reported gain does not reproduce on R.
PR, level rule	The 95% percentile interval of $\mathrm{acc}_R(c^*) - \mathrm{acc}_S(c^*)$ excludes zero, whether or not $c^* = s$; the b-th resample of S is paired with the b-th of R, each drawn with its own seed.
TR	On any call of the instance the server's prompt-token count differs from the local count of the clean render (the corrected caller's render of the row's messages before the final assistant turn), or, on vLLM, its prompt token IDs differ from the clean render's. For FreeToken only the count is compared.
TB1	A strong trivial baseline (the oracle constant, the frequency of the most frequent gold letter in the sample itself) exceeds 0.90 on a task in S or R.
TB2	For some task, the trivial constant chosen on S (the most frequent gold letter there, ties to the smallest letter) scores on R more than 0.05 away from its score on S.
TB3	On LB-0038, some configuration's exact-match rate exceeds three times the best constant age's, in S or R.
TB4	On a letter task in S or R, a configuration has at least 50 unparsed or conflicting calls, and every one of them is scored correct exactly when the gold is one and the same constant letter: the scorer treats them as a constant predictor.

Table 35: The rules of SCF, IR and SMI. g_S is the reported gain on S and Δ_V its visibility part (Section 3, with the start as P and c^* as R, point estimate on S).

Rule	Fires when
SCF1	On S the start's parse rate is below 0.90 or its truncation rate above 0.10.
SCF2	$c^* \neq s$ and $\mathrm{acc}_S(c^*)$ is less than 0.05 above the best of three panel members on S: chance, the strong trivial baseline and AUTHORS-1500 (in the instances of E, CFG-THINKING-FALSE-64-ORIG). For the letter pool, chance and the baseline are the row-weighted means of the tasks' values.
SCF3	$c^* \neq s$, $g_S > 0$ and Δ_V is at least 80% of g_S.
IR1	Any request of the instance has an assistant turn, or a last message that is not a user turn.
IR2	The gold screen hits on any request: the row's follow-up text, the LB-0038 age as a standalone integer, or another reference answer of more than one character appears in it.
IR3	On any row the scorer's gold is not among the golds derived from the question and follow-up (apart from the LB-0030 band-endpoint exception).
IR4	For S or R against the release: a task's row share differs by more than 2 points from its release share; the LB-0042 deceased share or an LB-0030 band share differs by more than 2 points; or the standardised mean difference of the actual age exceeds 0.05.
IR5	A binary task has a class below 10% of its rows in S or R, or an LB-0030 band is absent from S or R.
SMI (E only)	A served directory's configuration declares tied word embeddings while its checkpoint's output matrix differs from the embedding matrix; or, on the same rows and request bytes, the two variants' parsed answers differ on more than 1% of rows in some configuration. The second rule applies only where two variants of one checkpoint are served, which is Qwen3-1.7B-Longevity.

The intervals of D_S, D_R and Γ.

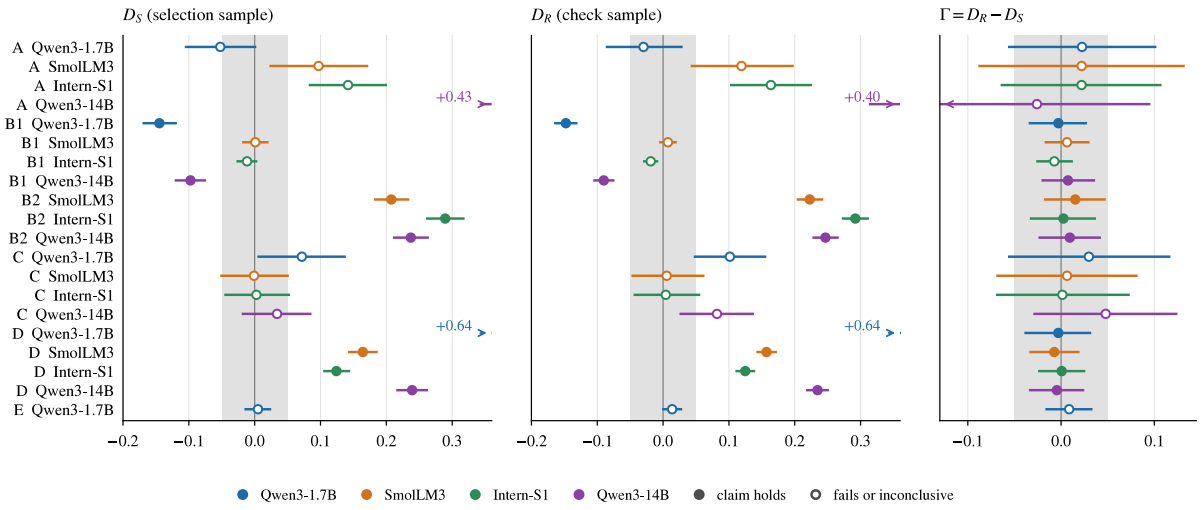

Figure 2: H4: the defect's distortion on the selection sample (D_S), on the check sample (D_R) and their difference (Γ), with adjusted intervals (family of 16) and the band $\pm m$ (shaded). Filled markers are claims that hold; arrows mark intervals cut at the axis limit. The unparsed-as-A scoring (B2) in Qwen3-1.7B-Longevity is not evaluable and is not drawn.

J The GEPA arm in detail

What GEPA searched. GEPA version 0.1.4 (Agrawal et al., 2025) edited two text components and nothing else: the system prompt, which replaces the row's system message (seed: the rows' own system message; limit 3,800 characters), and an instruction appended to the user turn after a blank line (seed: empty; limit 3,000 characters). The request parameters (thinking on or off, the token limit, the temperature) belong to the start and were not searched. The clean start is the request of AUTHORS-1500 (thinking off, 1,500 tokens, zero temperature) and the broken start that of THINK-64; the seed harnesses are the grid's own runs of those two configurations. The arm therefore searched prompt text only, and its negative result (no selected harness more accurate than the thinking-off baseline) is about that space: it says nothing about searching the thinking switch, the token limits or retry rules. The task model is Qwen3-1.7B-Longevity (untied) on the batch-invariant server. The proposer is Intern-S1-mini on the standard server at concurrency one, with thinking on, its card's sampling values (temperature 0.8, top-p 1.0, top-k 50), at most 8,192 tokens and a seed per request. GEPA's settings are its defaults except: a reflection minibatch of 3 rows; Pareto selection of candidates over instances; a round-robin choice of the component to edit; acceptance only on strict improvement; perfect-score minibatches skipped; no merging; no evaluation cache.

Rows and allocation. The arm uses 300 rows of the search partition's three letter tasks (120 of LB-0030, 60 of LB-0034 and 120 of LB-0042), drawn with seed 20261011 from search rows that are neither pilot rows nor rows the case study had held out; LB-0038 is left out because its score is an age error. They form a training set of 150 rows (60, 30 and 60 by task) and a validation set of 150 rows (the same split); the lists are the same in the four runs. A candidate scores one on a row when the primary scorer's parsed answer equals the gold and zero otherwise, an unparsed answer included. The proposer sees a row's user message, the task model's output, reasoning and answer, and feedback (correct, incorrect, unparsed, or cut at the token limit); it never sees the gold or a replication or test row. The four runs are the clean and the broken start, each with GEPA seeds zero and one.

Checks on a candidate, and the copy rule. Before any call a candidate is checked, and one that fails is recorded and scored zero without a call: each component is text within its limit; neither holds a chat-template control string; neither holds a run of 40 or more consecutive characters of a row's user message

that occurs in fewer than 10 of the arm's rows (the copy rule: wording shared by many rows may be copied, a row's own text may not); every render plus the token limit fits the 4,096-token context. The soft switches `/think` and `/no_think` were allowed and each candidate records whether it holds one; none of the selected harnesses does. A separate script re-checked the four selected texts against the arm's rows and found no violation of the copy rule and no identifier or number sequence of an item.

Budget, selection and ties. The budget is 6,000 metric calls per run (20 candidates on 300 rows). It is soft: GEPA checks it between iterations, so a run may end slightly above it, and the four runs used 6,003 to 6,147 calls. GEPA spends 6 metric calls per proposal, 150 on the seed's validation evaluation and 150 on each accepted candidate. The selected harness of a run is its candidate with the highest mean validation score, ties going to the lowest candidate index; if no proposal was accepted it is the seed. The selected texts were recorded, with their SHA256, before the harness was run on any replication row. Each was then run on every letter row of the replication and test partitions, on the main server with its start's request parameters.

The four selected harnesses. The texts below are as selected, wrapped for display; a character outside ASCII is written `[U+XXXX]` (`[U+2265]` is the sign for greater than or equal, `[U+03BC]` the Greek letter mu, `[U+00B2]` a superscript two). Each candidate's SHA256 is the one recorded before any replication call in our repository; the wrapping means that it cannot be recomputed from this page. A small language model generated the texts; we did not check their medical statements, and they must not be used as clinical guidance. Validation scores and accuracies are in Appendix B.

GEPA-CLEAN-s0 (clean start, GEPA seed zero)

Mean validation score 0.680; the candidate has index 0 (counted from zero) among the run's 31 candidates and SHA256

`ab48481902a519b5c045cd37f61382899ea0e1bce05b7a8345d2873f77b1bb16`

The run kept 31 candidates, the seed and the others that were accepted on their minibatches, but none of them scored above the seed on the validation rows, so the seed, the thinking-off baseline itself, was selected (ties go to the lowest index; Section 4.3).

System prompt.

```
You are a biomedical AI specialized in aging biology, trained on genomic, proteomic, and
clinical data.
```

Appended instruction. Empty.

GEPA-CLEAN-s1 (clean start, GEPA seed one)

Mean validation score 0.700; the candidate has index 12 (counted from zero) among the run's 31 candidates and SHA256

`33819b05a9cd5711ed2b502c48f84a626fd26d8d4286dcadc420cde065e6b7f9`

System prompt.

```
You are a biomedical AI specialized in aging biology, trained on genomic, proteomic, and
clinical data. Your task involves assessing biological age or age-related phenotypes
using clinical biomarkers. Focus exclusively on the biomarkers listed below, which are
validated predictors of aging in human populations.

**Key Biomarkers for Aging Assessment (prioritize these in reasoning):**
1. **Kidney Function:**
   - Cystatin C (mg/L): Increase >0.8 mg/L correlates with [U+2265]75 years.
   - Creatinine ([U+03BC]mol/L): >80 [U+03BC]mol/L in females suggests age >60.
```

```
2. **Liver Metabolism:**
   - Albumin (g/dL): <4.0 g/dL indicates malnutrition or chronic disease linked to
advanced age.
   - Bilirubin (mg/dL): Direct increase (>0.5) reflects impaired hepatic clearance with
age.
3. **Metabolic Health:**
   - LDL Cholesterol (mg/dL): >130 mg/dL in females [U+2265]60 increases cardiovascular
risk.
   - Triglycerides (mg/dL): >200 mg/dL in females [U+2265]50 indicates insulin
resistance.
4. **Inflammation & Stress:**
   - C-reactive protein (mg/dL): >1.0 mg/dL reflects chronic inflammation linked to
aging.
   - Cadmium (nmol/L): Accumulation (>5 nmol/L) correlates with tissue damage in older
adults.
5. **Cellular & Tissue Health:**
   - Mean Cell Volume (fL): Decline (>85 fL in females <50 indicates deficiency; <90 fL
in older adults suggests anemia of aging).
   - Red Cell Distribution Width (%): >14.5% in females [U+2265]60 reflects oxidative
stress.

**Exceptions & Caveats:**
- Ignore BMI unless explicitly asked to correlate with metabolic syndrome.
- Do not assume chronological age unless explicitly provided (e.g., in Example 2, the
assistant incorrectly inferred "18-30" from BMI alone).
- Flag conflicting biomarkers (e.g., high LDL but low triglycerides may require deeper
analysis).

**Generalizable Strategy:**
1. Compare biomarker trajectories across cohorts (e.g., Cystatin C >8.0 indicates
accelerated aging).
2. Prioritize systemic markers over individual parameters (e.g., integrate LDL, CRP, and
Cadmium for metabolic aging profiles).
3. Use thresholds validated in longitudinal cohort studies (e.g., Bone Alkaline
Phosphatase >15 ug/L in females [U+2265]60).

For comparative tasks (e.g., "Who is older?"), rank participants by cumulative biomarker
deviation from age-predictive ranges. For age-group classification, use [U+2265]2
biomarkers exceeding thresholds for specific ranges (e.g., [U+2265]75: Cystatin C >1.0 +
CRP >2.0 + CAD: >10 nmol/L).
```

Appended instruction.

```
Analyze medical data to answer yes/no survival questions or identify older participant
based on blood tests. Input includes demographics, biometrics, blood panel results, and
clinical history. Prioritize evidence-based medical knowledge: For survival prediction,
consider BMI (optimal range 18.5-24.9), albumin (>3.5 g/dL indicates better prognosis),
creatinine (<1.1 mg/dL for women, <1.3 for men suggests preserved kidney function), LDL
(<100 mg/dL ideal), and absence of diabetes/hypertension risk factors. For age
comparison, note cadmium levels (higher in older adults), red cell distribution width
(increases with age), and segmented neutrophils (typically decrease after age 65).
Explicitly flag when data is insufficient to determine answer. Use structured reasoning:
first state key markers supporting conclusion, then cross-reference contradictions.
Never infer survival based solely on absence of provided data.
```

GEPA-BROKEN-s0 (broken start, GEPA seed zero)

Mean validation score 0.660; the candidate has index 1 (counted from zero) among the run's 24 candidates and SHA256

9f38d54e2cc346c8b0c20fcaabfd981906af2ce409a2d04250195c88658482ab

System prompt.

```
You are a biomedical AI specialized in aging biology, trained on genomic, proteomic, and
clinical data.
```

Appended instruction.

```
Analyze clinical data to predict survival probability or age comparisons. Focus on key
biomarkers linked to mortality or aging (e.g., elevated HbA1c >6.5% for diabetes, RDW
>14.5% as mortality predictor, elevated creatinine >177 [U+03BC]mol/L for chronic kidney
disease). Apply validated thresholds for diabetes (HbA1c [U+2265]6.5%), cardiovascular
risk (LDL >70 mg/dL, HDL <40 mg/dL for males), and organ dysfunction (eGFR <60
mL/min/1.73m[U+00B2]). Prioritize metrics like glycohemoglobin, red cell distribution
width, and bone alkaline phosphatase. For age comparisons, use biomarkers like cadmium
accumulation (higher in older adults, [U+2265]8.5 nmol/L) and vitamin B12 levels (<148
pmol/L for deficiency risk). Limit reasoning to critical data; format answers as A/B per
input options. Include domain-specific knowledge: high RDW correlates with mortality (HR
1.2 per SD increase), elevated uric acid >7 mg/dL predicts cardiovascular events, and
low HDL <35 mg/dL increases dementia risk. Avoid speculative connections beyond
well-established clinical correlations. Respond fully within token limits by
prioritizing high-impact metrics and omitting tangential details.
```

GEPA-BROKEN-s1 (broken start, GEPA seed one)

Mean validation score 0.667; the candidate has index 1 (counted from zero) among the run's 29 candidates and SHA256

ddf8886b2b3edced069fcfa5c3d9baad95643056d772d2ebeb3580770ca017ed

System prompt.

```
You are a biomedical AI specialized in aging biology, trained on genomic, proteomic, and
clinical data.
```

Appended instruction.

```
### Instruction

**Task Overview:**
Analyze provided medical data to answer specific clinical questions. Prioritize
identifying key biomarkers and demographic factors relevant to the question. Answer
strictly using the format "A." or "B." Only, matching the provided options.

---

**Input Format:**
All inputs will include structured medical data:
1. **Demographics:** Age, ethnicity, education level, sex.
2. **Biometrics:** BMI, height, weight, waist circumference.
3. **Blood Panel:** Standard lab results (e.g., glucose, cholesterol, proteins,
electrolytes).
4. **Questionnaire:** Medical history, family history, surgical history.

Data units are standardized (e.g., BMI in kg/m[U+00B2], glucose in mmol/L). Missing
values will not be present.

---
```

```
**Domain-Specific Knowledge to Include:**
1. **Survival Prediction:**
   - BMI [U+2265]30 indicates obesity; [U+2265]25 indicates overweight.
   - LDL >3.0 mmol/L or triglycerides >1.7 mmol/L suggest hyperlipidemia.
   - Glycohemoglobin [U+2265]6.5% indicates prediabetes/diabetes.
   - Chronic inflammation (C-reactive protein >3 mg/L) correlates with mortality risk.
   - Family history of cardiovascular diseases multiplies risk factors.

2. **Age Estimation via Blood Tests:**
   - Red blood cell count and hemoglobin decline with age; compare to age-stratified
norms.
   - Elevated Cystatin C (>1.0 mg/L) correlates with aging-associated renal dysfunction.
   - Decreased Vitamin D (<30 nmol/L) is common in older adults.
   - Glycohemoglobin trends reflect long-term metabolic stability.

---

**Generalizable Strategies:**
1. **Survival Questions:**
   - Flag obesity (BMI [U+2265]30), hyperlipidemia, diabetes, or chronic inflammation as
high-risk factors.
   - Compare triglycerides/HDL ratio (>3.5) for cardiovascular risk.
   - Note female reproductive history (e.g., hysterectomy impacts hormone status).

2. **Age Questions:**
   - Prioritize red blood cell indices (mean corpuscular volume) and hemoglobin trends.
   - Analyze biomarkers like Cystatin C and Vitamin D for age-related decline.
   - Contrast metabolic profiles (e.g., glucose, cholesterol) against age-specific
norms.

---

**Output Requirements:**
- Always answer with **A.** or **B.** only.
- No explanations or partial answers.
- Ensure answers align with calculated risk scores and biomarker thresholds.

**Example Alignment:**
If a 40-year-old male has BMI 35, LDL 4.0 mmol/L, and diabetes history, survival at 10
years correlates with **B. No**. Conversely, a 30-year-old with BMI 25, normal
lipid/glucose, and no diabetes history correlates with **A. Yes** for survival.

**Critical Note:**
Never infer external knowledge (e.g., specific study results). Use only the data and
thresholds provided in the instructions.
```

K Planned power

This paper quotes five planned-power figures (Table 36), all from the preregistration and none recomputed after the outputs. They are rough: made before the study from the pilot's parse shares and from assumed effects, with a normal approximation. "Expected to be inconclusive" in Appendix D is an informal expectation read off these figures, not a prediction with a decision rule, and the observed verdicts do not test it.

H2, Intern-S1-mini on GSM8K against RESCUE-1500. The assumed truth is a θ of zero: the thinking-off configuration exactly as accurate as the rescue. The paired difference of correctness takes the values minus

one, zero and one, so its variance is at most d, the share of rows on which the two configurations' parsed answers differ (an unparsed answer counts as one value); on the pilot's GSM8K items $d = 0.325$. The unit is the item and the test partition has $n = 527$ GSM8K items, so the standard error of the estimated difference is at most $\mathrm{SE} = \sqrt{d/n}$. The lower bound of the interval for θ clears $-\delta$, with $\delta = 0.05$, when the estimate exceeds $-\delta + z\,\mathrm{SE}$, where z is the normal quantile of the interval's level: 2.89 for the Bonferroni family of 13 (level 99.62%) and 1.96 for an unadjusted interval. The power at a true θ of zero is then $\Phi(\delta/\mathrm{SE} - z)$, with Φ the standard normal distribution function: 0.19 adjusted and 0.52 unadjusted.

H1, SmolLM3-3B and Intern-S1-mini on GSM8K, the statistic T. The inputs are the pilot's joint parse shares of THINK-512 (P) and AUTHORS-1500 (R) on GSM8K, v (only R parses), b (both parse) and l (only P parses), and d_B, the share of the rows both parse on which the parsed answers differ. No gold or accuracy of the pilot enters. Two assumed scenarios supply the accuracies, and they are assumptions, not pilot results. In S1, R is right on half of the rows only it parses and P on half of those only it parses, and neither is better on the rows both parse. In S2 both are right on 0.3 of the rows both parse and R is better than P by 0.02 on them, capped at d_B. The power of $T > 0$ for the lower bound of the adjusted interval is a normal approximation on the same $n = 527$ test items (one unit per item) at the level of H1's family of 9 ($z = 2.77$); the script, `prereg_power_h1_v3.py`, is pinned in the preregistration.

Table 36: The planned-power figures quoted in this paper. All are for the test partition, with a normal approximation. S1 and S2 are the two assumed scenarios of H1.

Endpoint	Assumed effect	Inputs from the pilot	Units	Family, z	Power
H2: Intern-S1-mini, GSM8K, θ against RESCUE-1500	θ zero	$d = 0.325$	527 items	13; 2.89	0.19 (adjusted), 0.52 (unadjusted)
H1: SmolLM3-3B, GSM8K, $T > 0$	S1, S2	$v/b/l = 0.210/0.725/0.040$, $d_B = 0.055$	527 items	9; 2.77	S1 0.52, S2 0.00
H1: Intern-S1-mini, GSM8K, $T > 0$	S1, S2	$v/b/l = 0.550/0.450/0.000$, $d_B = 0.756$	527 items	9; 2.77	S1 0.45, S2 0.06

No power was planned for H4, whose distortions are what it measures, nor for the four GEPA endpoints, whose size depends on the harnesses' answers, which did not exist before the search.

L The second study's GEPA arm

Design. The arm has the first study's design (Appendix J) with Qwen3-14B as the task model on the batch-invariant server. The rows (300 search rows, in the same training and validation sets), the settings and the code are the first study's, which we checked by their hashes, and the proposer is Intern-S1-mini again. The four runs are the clean and the broken start, each with GEPA seeds zero and one, and each has a budget of 6,000 metric calls. The selected texts were recorded before any replication call, and the checkpoint locked them before the first test call.

The clean starts. The clean-start runs ran 241 and 271 iterations. In 32 and 32 of them the parent already scored perfectly on its minibatch and no text was proposed; the other iterations made 209 and 239 proposals, and the runs accepted 30 and 29 of them, from a seed validation score of 0.180. Both selected candidates have the validation score 0.587; their validation scores exceed their test accuracy (0.540 and 0.545), as selection makes them optimistic. None of the candidates evaluated in a clean-start run contained a soft switch (0 and 0).

The broken starts. Each broken-start run made 975 proposals, one in each of its 975 iterations. Of these, 142 and 137 (seed zero and seed one) failed the interface check, and the others were evaluated on a minibatch of three rows against the seed. Every proposal scored zero against the seed's zero, so the strict-improvement rule accepted none, only the seed was validated (score 0.000 on the 150 validation rows), and the selected

Table 37: The four searches of the second study. *Iterations*: passes of the search loop. *Without a proposal*: iterations in which the parent already scored perfectly on its minibatch of three rows, so no text was proposed; the proposals are the iterations minus these. *Failed the interface check* and *called*: the evaluations of a text other than the seed that the search's interface counter recorded, scored zero without a call or sent to the model. In a broken-start run only the proposals are evaluated, so the two sum to the proposals; in a clean-start run they also include the search's other scorings (a parent on its minibatch, a validation pass) and exceed the proposals. Every proposal of a broken-start run scored zero against its parent's zero on a minibatch of three rows.

Run	Iterations	Without a proposal	Proposals	Failed the interface check	Called	Accepted	Seed validation score	Selected (validation score)
broken, seed 0	975	0	975	142	833	0	0.000	the seed (0.000)
broken, seed 1	975	0	975	137	838	0	0.000	the seed (0.000)
clean, seed 0	241	32	209	61	377	30	0.180	index 10 (0.587)
clean, seed 1	271	32	239	27	450	29	0.180	index 21 (0.587)

harness is the seed in both runs: the two runs select one text. No candidate contained a soft switch (0 and 0), and the request of the start (thinking on, a cap of 64 tokens) is not part of the text.

Checks on the two selected clean-start texts. Run after the evaluations on the texts as selected: the arm's interface check passes for both, neither contains a soft switch, and no item key of the main items or of the arm's rows occurs in either. Of the 40-character windows of each text (4,681 and 5,759), none occurs in any of the arm's 300 search rows or in any of the 6,307 replication and test letter rows, so neither text copies a run of 40 characters from a row's user message, shared task wording included. The 40-character rule does not cover short values, and the two texts contain decimal numbers (4 and 8). All but 2 of them occur in rows, each in at least 6 rows of the arm and the evaluation partitions together, so none is unique to one row; the counts cannot show which arm row, if any, the proposer took a value from.

The selected harnesses. The texts are as selected, wrapped for display, a character outside ASCII written `[U+XXXX]`. Each candidate's SHA256 is the one recorded in our repository before any replication call; the wrapping means that it cannot be recomputed from this page. A small language model generated the texts; we did not check their medical statements, and they must not be used as clinical guidance.

GEPA-CLEAN-s0 (clean start, GEPA seed zero)

SHA256

`8aa3a797d9e2e73b1a76681014827b00af0aef5f6250fc3773f17c151457de79`

System prompt.

```
You are a biomedical AI specialized in aging biology, trained on genomic, proteomic, and
clinical data. Your expertise includes interpreting biomarkers linked to physiological
aging, identifying age-related trends in clinical profiles, and integrating multimodal
data to infer biological age or comparative aging status.

When analyzing clinical data, prioritize the following age-associated biomarkers and
trends:
- **Kidney Function**: Elevated creatinine ([U+00B5]mol/L) and cystatin C (mg/L)
correlate with reduced renal filtration efficiency in older adults.
- **Metabolic Health**: Lower HDL (mg/dL) and vitamin A ([U+00B5]mol/L) are common in
middle-aged/older populations due to metabolic shifts and absorption changes.
- **Inflammation & Oxidative Stress**: Elevated C-reactive protein (mg/dL) and cadmium
(nmol/L) reflect age-related chronic inflammation and environmental toxin accumulation.
- **Vitamin Status**: Declining vitamin B12 (pmol/L) levels in older adults are linked
to impaired intestinal absorption.
- **Cholesterol Profiles**: Higher LDL/HDL ratios and total cholesterol (mg/dL) align
with age-related dyslipidemia.
```

- **Hematologic Changes**: Lower red blood cell counts (million cells/uL) in postmenopausal women and altered lymphocyte/monocyte ratios (e.g., reduced lymphocytes, elevated monocytes) suggest immune aging.
- **Liver/Thyroid Function**: Elevated alkaline phosphatase (U/L) may reflect bone turnover or liver metabolism changes with age.

Task Execution Rules:
1. **Direct Answer First**: For questions like "Which participant is older?" or "What age group?" begin with a single-letter answer (A/B or age group letter) before elaborating.
2. **Prioritize Comparative Analysis**: When comparing participants, highlight biomarkers showing divergent age-related trends (e.g., Participant B[U+2019]s elevated uric acid vs. Participant A[U+2019]s normal levels in Example 3).
3. **Gender-Specific Considerations**: Address sex differences (e.g., menopausal status in females, testosterone effects in males) that influence biomarker interpretation.
4. **Multimodal Integration**: Combine blood panel data with biometrics (e.g., BMI) and demographic factors (e.g., male/female) to strengthen age inferences.
5. **Avoid Overgeneralization**: Explicitly state when a biomarker[U+2019]s deviation is non-specific (e.g., "elevated creatinine could reflect age, medication use, or hydration status").

Example Strategies:
- **Age Group Inference**: Match combinations of trends (e.g., low HDL + elevated creatinine + normal kidney function) to options like D (45[U+2013]60) or E (75+).
- **Comparative Analysis**: Compare cadmium levels (Participant B[U+2019]s 8.01 nmol/L vs. A[U+2019]s 1.78 in Example 2) or vitamin B12 (Participant B[U+2019]s 246.49 vs. A[U+2019]s 446.49 in Example 3) to identify aging biomarkers.

Always acknowledge limitations, such as missing longitudinal data or confounding factors (e.g., medications, comorbidities), which may obscure age-related patterns. Format answers with a clear primary response followed by structured reasoning using the biomarkers above.

Appended instruction.

Review clinical data strictly within the provided timeframe. Do not infer survival status, age, or disease progression beyond the snapshot data. For biomarker interpretation, prioritize established causal relationships over associative patterns.
When comparing age groups:
- Use kidney function markers (creatinine, cystatin C) only if longitudinal GFR data is unavailable
- Consider cadmium accumulation as both age and exposure marker, but prioritize other age-specific indicators
- Treat vitamin deficiencies as potential malnutrition indicators, not definitive age markers
- Differentiate between age-related trends and comorbidity effects (e.g., diabetes-related biomarkers vs. aging)

For mortality prediction questions:
- Specify that current data only reflects immediate risk factors
- Acknowledge that chronic conditions may evolve with treatment
- Distinguish between acute versus chronic biomarker abnormalities
- Avoid extrapolating 10-year outcomes without longitudinal data

When analyzing blood test results:
- Flag conflicting biomarker trends (e.g., elevated CRP with normal levels)
- Consider medication effects on biomarker readings (e.g., metformin on creatinine)
- Differentiate between age-related and gender-specific trends (e.g., RBC counts in males vs. females)

```
- Document when biomarkers fall outside typical reference ranges for age groups

Provide structured reasoning with:
1. Biomarker-by-biomarker analysis
2. Cross-referencing of conflicting data points
3. Explicit acknowledgment of data limitations
4. Alternative explanations for ambiguous findings
5. Clear distinction between established clinical knowledge and unsupported inferences
```

GEPA-CLEAN-s1 (clean start, GEPA seed one)

SHA256

```
7ddc3ca857a98b1f2c1f35960387f278abcf413b6c0b4b29105b8d354b0e89ff
```

System prompt.

```
**Instruction for Biomedical Age and Survival Inference Task**

**Task Overview:**
You are a biomedical AI specializing in aging biology, trained on genomic, proteomic,
and clinical data. Analyze age groups and survival status strictly based on provided
biomarkers and demographic data. Use cohort-specific reference ranges (e.g., NHANES III,
Framingham) and prioritize combinations of markers (e.g., HDL, LDL, BMI, kidney
function) for inference. Avoid single-marker conclusions, causality assumptions, or
extrapolation beyond explicit evidence.

---

**Task Instructions:**

1. **Age Group Determination:**
   - If explicit age is provided, state it directly.
   - If not, analyze biomarkers using gender-specific reference ranges (e.g., HDL: Males
40[U+2013]50 mg/dL [60[U+2013]75]; Females 50[U+2013]70 mg/dL [60[U+2013]75]; NHANES III
BMI cutoffs for obesity).
   - **Prioritize combinations of markers:**
     - HDL <40 mg/dL (Males [U+2265]45) + LDL >130 mg/dL (Females [U+2265]40) + BMI >30
(Males [U+2265]45) [U+2192] Strong inference.
     - Creatinine/eGFR (e.g., eGFR 45[U+2013]59 mL/min/1.73m[U+00B2] indicates CKD Stage
3a, linked to [U+2265]65 in hypertensive patients).
     - Flag uncertainty when biomarkers overlap ranges (e.g., CRP 3.38 mg/dL could
indicate inflammation in both middle-aged and elderly groups).
   - **Domain-Specific Examples:**
     - Low vitamin B12 (<200 pmol/L) in [U+2265]65 correlates with atrophic gastritis.
     - Metabolically healthy obesity (BMI >30 with normal glucose and HDL >50 mg/dL in
Females).

2. **Survival Status Determination:**
   - If survival data (e.g., follow-up dates, death certificates) is absent, explicitly
state inability to determine status.
   - Acknowledge risk factors (e.g., BMI >30 + LDL >4.1 mmol/L correlates with 10-year
mortality [U+2191]15[U+2013]30% in obese diabetics [UKPDS]).
   - Avoid conclusions without explicit mortality data.

3. **Response Format:**
   - **Answer:** Boxed letter(s) based on explicit data or strongest inference (e.g.,
**C**).
   - **Reasoning:**
```

```
     - *Key Observations:* List biomarkers (e.g., "BMI 42.4 kg/m[U+00B2] [Obese] + HDL
34 mg/dL [Low in Females [U+2265]45] [U+2192] Age [U+2265]60").
     - *Inference:* Link to cohort studies (e.g., "BMI >30 + HDL <40 mg/dL in Females
aligns with [U+2265]40 [NHANES III]").
     - *Conclusion:* "Cannot determine survival due to lack of follow-up data."

4. **Error Avoidances:**
   - Reject single-marker inferences (e.g., vitamin D deficiency [U+2260] elderly-only).
   - Avoid causality assumptions (e.g., LDL elevation could reflect genetics or
hyperthyroidism).
   - Use exact cutoffs (e.g., CKD Stage 3a: eGFR 45[U+2013]59 mL/min/1.73m[U+00B2]).

5. **Domain-Specific Knowledge:**
   - Gender-specific HDL ranges (e.g., Males [U+2265]45: HDL [U+2264]40 mg/dL; Females
[U+2265]60: HDL [U+2264]50 mg/dL).
   - Longitudinal cohort data (e.g., Framingham: LDL >130 mg/dL in Males [U+2265]45;
NHANES III: BMI >30 in Females [U+2265]40).
   - Metabolic syndrome markers (e.g., uric acid >8.6 mg/dL + BMI >40 [U+2192] Strongly
obese male >50).

6. **Example Compliance:**
   - For age: "BMI 42.4 [Obese] + HDL 34 mg/dL [Low in Females [U+2265]45] + Creatinine
0.86 [Normal but Cystatin C 1.11 elevated] [U+2192] Likely [U+2265]60 (but not
definitive)."
   - For survival: "BMI 32.7 [Obese] + LDL 114 mg/dL [Borderline] + Stage 2 hypertension
[U+2192] No survival data; risk factors correlate with [U+2191]10-year mortality but
cannot confirm outcome."

**Final Output Requirements:**
- Provide **only the boxed answer** (e.g., **A**).
- Avoid phrases like "most likely" or "suggests."
- Use strict adherence to cohort references and biomarker combinations.
```

Appended instruction.

```
Task Description:
Analyze clinical profiles to determine survival status (10-year post-examination) or age
group based on provided biomarkers and demographics. Use evidence-based medical
principles, reference ranges (e.g., NHANES III for HDL in females), and contextual
factors (e.g., diabetes management, medication history, lab trends) to infer outcomes.
Prioritize explicit survival data (e.g., follow-up dates, death certificates) when
available.

Key Requirements:
1. Survival Prediction:
   - Identify risk factors (elevated cholesterol, hypertension, diabetes) and protective
factors (normal kidney function, good glycemic control).
   - Consider comorbidities (e.g., asthma, arthritis) and medication use (e.g.,
Glipizide, corticosteroids) for long-term prognosis.
   - Avoid selecting invalid options (e.g., option C in Example 2). If survival data is
absent, state inability to determine status.

2. Age Group Determination:
   - Use validated references (e.g., NHANES III HDL ranges by age) and biomarker trends
(BMI >30 linked to [U+2265]40 years old in females).
   - Prioritize HDL, creatinine, and BMI correlations over isolated lab values.

3. Generalizable Strategies:
   - Balance lab values (e.g., C-reactive protein for inflammation, cystatin C for
```

```
kidney health) with demographic factors (ethnicity, education level).
   - Highlight contradictions (e.g., high BMI with normal HDL) to refine conclusions.

4. Formatting:
   - Provide a single-letter answer (A/B) for survival and (A-D/E) for age.
   - Structure reasoning into Key Observations, Inference, and Conclusion sections.

Domain-Specific Knowledge:
- Females under 45 typically exhibit lower HDL than those 45[U+2013]60 (e.g.,
50[U+2013]70 mg/dL for [U+2265]40 years old per NHANES III).
- Glycohemoglobin [U+2264]6% indicates effective diabetes management, reducing mortality
risk.
- Long-term corticosteroid use (e.g., Triamcinolone for 15 years) may increase
cardiovascular risk but does not directly indicate 10-year mortality without further
data.

Avoid:
- Selecting options outside user-provided choices (e.g., Example 2[U+2019]s invalid C).
- Overreliance on outdated references (e.g., Example 3[U+2019]s incorrect HDL-age
mapping).
- Ignoring explicit survival data (e.g., Example 2[U+2019]s lack of follow-up).
```

GEPA-BROKEN-s0 (broken start, GEPA seed zero (the seed; GEPA-BROKEN-s1 is the same text))

SHA256

```
ab48481902a519b5c045cd37f61382899ea0e1bce05b7a8345d2873f77b1bb16
```

System prompt.

```
You are a biomedical AI specialized in aging biology, trained on genomic, proteomic, and
clinical data.
```

Appended instruction. Empty.

M The second study in full

This appendix gives the second study's endpoints and statistics in the form of Appendix B. Every number comes from the study's analysis report committed in our repository, with the first study's structure and rules.

Endpoints. Table 38 lists the 13 endpoints of H1–H3 with their verdicts on both partitions and the replication label, and Table 39 the 5 claims of H4 with their three intervals.

The numeric tables of H1–H3. Tables 40, 41 and 42 give the second study's H1, H2 and H3 endpoints on the test partition in the form of the first study's Tables of Appendix B: the gain and its parts with the adjusted interval of T, θ with its adjusted interval, and the drops D_T and D_P with their adjusted intervals.

Configurations. Table 43 gives accuracy, parse rate and truncation of the six configurations in the three pools on both partitions. Table 44 gives the replication partition's non-commitment and refusal shares (the test partition is in Table 45, below).

Sensitivity analyses. The ten sensitivity analyses of the preregistration were run on the second study as well. The strict scorer changes 1 of its 13 endpoint verdicts, the rescue on MMLU-Pro (*holds* becomes *fails*); no other analysis changes a verdict of H1–H3 (the largest number of changes among the other nine is 0), and putting the left-out rows back changes 0 H4 claims.

Table 38: The endpoints of H1–H3 of the second study (Qwen3-14B): verdicts on the replication and test partitions and the replication label. A star marks an endpoint whose quantity is fixed by an identity of the design or by the observed parse pattern; its verdict says little beyond that mechanism, although it still depends on the intervals.

Family	Pool	Endpoint	Replication	Test	Label
H1	LB letters	AUTHORS-1500 vs THINK-512	visibility dominates⋆	visibility dominates⋆	replicated
H1	MMLU-Pro	AUTHORS-1500 vs THINK-512	visibility dominates	visibility dominates	replicated
H1	GSM8K	AUTHORS-1500 vs THINK-512	no gain: not tested	no gain: not tested	replicated
H2	LB letters	AUTHORS-1500 vs RESCUE-1500	holds⋆	holds⋆	replicated
H2	MMLU-Pro	AUTHORS-1500 vs RESCUE-1500	holds	holds	replicated
H2	GSM8K	AUTHORS-1500 vs RESCUE-1500	fails	fails	replicated
H2	LB letters	AUTHORS-1500 vs GEPA-CLEAN-s0	fails	fails	replicated
H2	LB letters	AUTHORS-1500 vs GEPA-CLEAN-s1	fails	fails	replicated
H2	LB letters	AUTHORS-1500 vs GEPA-BROKEN-s0	holds⋆	holds⋆	replicated
H2	LB letters	AUTHORS-1500 vs GEPA-BROKEN-s1	holds⋆	holds⋆	replicated
H3	LB letters	BUDGET-512-1500 vs THINK-512	holds	holds	replicated
H3	MMLU-Pro	BUDGET-512-1500 vs THINK-512	holds	holds	replicated
H3	GSM8K	BUDGET-512-1500 vs THINK-512	holds	holds	replicated

Table 39: H4 in the second study: bound, the distortions on the selection sample (D_S) and the check sample (D_R) and their difference with adjusted intervals, the verdict of each component, the stability of the descriptive reading and the claim. A star marks a claim whose quantity is fixed by an identity of the design; its verdict says little beyond that mechanism, although it still depends on the intervals.

Defect	Bound	D_S	D_R	Γ	Components	Stability	Claim
A reference answer as assistant turn	1.000	+0.430 [+0.345, +0.515]	+0.404 [+0.313, +0.491]	−0.026 [−0.151, +0.096]	holds/holds/open	inconclusive	inconclusive
B1 gold fallback	0.149	−0.097 [−0.121, −0.074]	−0.090 [−0.106, −0.074]	+0.007 [−0.021, +0.037]	holds/holds/holds	stable	holds
B2 unparsed scored as A	0.237	+0.237 [+0.210, +0.265]	+0.247 [+0.227, +0.267]	+0.009 [−0.024, +0.043]	holds/holds/holds	stable	holds⋆
C skewed split	by rule	+0.034 [−0.020, +0.086]	+0.082 [+0.025, +0.138]	+0.048 [−0.030, +0.125]	open/open/open	inconclusive	inconclusive
D broken start	0.345	+0.239 [+0.215, +0.263]	+0.235 [+0.217, +0.252]	−0.005 [−0.034, +0.025]	holds/holds/holds	stable	holds⋆

Completion tokens and the first study's rerun. Table 46 gives the mean completion tokens per call, and Table 47 the nine endpoints of the first study recomputed for the tuned 1.7B model on the second study's rerun (descriptive).

The 54 accuracy differences of the rerun. Table 48 gives, for the tuned 1.7B model, the accuracy of the rerun minus that of the first study with its unadjusted interval, in each partition (the pooled rows are the two together), pool and configuration of the grid.

The rerun on the second GPU. Table 49 compares the rerun with the first study's outputs call by call: the calls whose token IDs differ, whose finish reason differs and whose request body differs.

N Constant-label shares by task

Table 50 gives, for each LongevityBench letter task and each partition, the accuracy of the second study's clean-start harnesses and of AUTHORS-1500 and BUDGET-512-1500 for Qwen3-14B, and of AUTHORS-1500 for the tuned 1.7B model in the first study and in the second study's rerun, next to the task's strongest oracle constant. The constant is a label in every task and partition (an age band, a participant, "Yes"), chosen on the partition itself. The share of parsed answers equal to the constant label is the share of calls whose letter the frozen scorer reads and whose option label, in the row's own option list, is the constant label; the last column is the share of the correct answers that are the constant label. The weights of the tasks in the pool are 40%, 20% and 40% on the test partition. The values were computed after the outputs existed, from the call records and the main items, and equal the committed reports' accuracies where those exist.

Table 40: H1 of the second study, test partition: the gain Δ of AUTHORS-1500 over THINK-512 with its adjusted interval, its parts, the statistic T with its adjusted interval, the numbers of rows only AUTHORS-1500 parses (V), both parse (B) and only THINK-512 parses (L), the verdict and the replication label (R replicated). A star marks an endpoint whose quantity is fixed by an identity of the design or by the observed parse pattern; its verdict says little beyond that mechanism, although it still depends on the intervals.

Cell	Δ	Δ_V	Δ_B	Δ_L	T	V / B / L	Verdict	Repl.
Qwen3-14B LB letters	+0.203 [+0.186, +0.219]	+0.203	0.000	0.000	+0.041 [+0.037, +0.044]	2,222 / 2 / 0	visibility dominates*	R
Qwen3-14B MMLU-Pro	+0.402 [+0.350, +0.453]	+0.419	−0.018	0.000	+0.098 [+0.084, +0.113]	660 / 107 / 0	visibility dominates	R
Qwen3-14B GSM8K	−0.241 [−0.315, −0.167]	+0.101	−0.342	0.000	+0.293 [+0.249, +0.339]	286 / 241 / 0	no gain: not tested	R

Table 41: H2 of the second study: θ of AUTHORS-1500 against RESCUE-1500 in three cells and against the four GEPA harnesses (letters), with its adjusted interval, the verdict on the test partition, the verdict on the replication partition and the label. A star marks an endpoint whose quantity is fixed by an identity of the design or by the observed parse pattern; its verdict says little beyond that mechanism, although it still depends on the intervals.

Endpoint	θ (adjusted interval)	Verdict	Replication verdict	Label
Qwen3-14B LB letters, RESCUE-1500	0.000 [0.000, 0.000]	holds*	holds	replicated
Qwen3-14B MMLU-Pro, RESCUE-1500	−0.016 [−0.033, −0.003]	holds	holds	replicated
Qwen3-14B GSM8K, RESCUE-1500	−0.340 [−0.400, −0.281]	fails	fails	replicated
Qwen3-14B LB letters, GEPA-CLEAN-s0	−0.337 [−0.362, −0.311]	fails	fails	replicated
Qwen3-14B LB letters, GEPA-CLEAN-s1	−0.342 [−0.369, −0.314]	fails	fails	replicated
Qwen3-14B LB letters, GEPA-BROKEN-s0	+0.203 [+0.186, +0.221]	holds*	holds	replicated
Qwen3-14B LB letters, GEPA-BROKEN-s1	+0.203 [+0.186, +0.221]	holds*	holds	replicated

O The first study next to the trivial baselines

Accuracy next to the trivial baselines. For the first study, Table 51 gives the accuracy of five configurations in the three pools next to their baselines, computed after the outputs existed (Appendix C), with a block for each LongevityBench task; THINK-64, the GEPA harnesses and the replication partition are in Appendix B. Each oracle constant is chosen once, on the analysis rows of the partition it is scored on, and it is a fixed prediction in the bootstrap: it is not re-chosen in a resample, so the intervals of the differences to it cover the sampling of units given that choice. On the LongevityBench letters the strongest constant is the oracle constant *label*: the age band most common on LB-0030, Participant A on LB-0034 and "Yes" on LB-0042 (0.840 of its test rows). It scores 0.539 $[0.527, 0.551]$ on the pool, far above chance (0.380) and the oracle constant letter (0.389), because the option order is shuffled (120 letter-to-label mappings on LB-0030, 2 on LB-0042) and a constant letter cannot exploit an uneven label distribution. Qwen3-1.7B-Longevity exceeds the strongest constant in every configuration that parses: AUTHORS-1500 0.632, $+0.093$ $[+0.074, +0.112]$, and THINK-512 $+0.096$ $[+0.077, +0.115]$. The margin is small on LB-0042 ($+0.027$ $[+0.007, +0.048]$ for THINK-512), where one GEPA harness, GEPA-CLEAN-s1, is not above it (-0.010 $[-0.034, +0.015]$). SmolLM3-3B and Intern-S1-mini are below the strongest constant on the pool in every configuration (at best -0.382 and -0.225), because most of their calls end unparsed and unparsed counts as wrong, although Intern-S1-mini is above it on LB-0034 with AUTHORS-1500 ($+0.060$ $[+0.003, +0.117]$).

The case study had not found this: on the split of its harness search no configuration of the same model clearly exceeded a constant age band on LB-0030 or the majority letter on LB-0034 (post hoc there). We read the difference as consistent with that split's skew toward young, alive participants, but the items, the output layer (tied there, untied here) and the stack all differ, so the data do not establish it. The output layer alone is unlikely to explain it: on the same items the tied and untied variants differ in accuracy by at most 0.009, although their parsed answers are identical on only 0.656 to 0.741 of the rows, and on the letter pool the tied-output-layer defect moves the start's accuracy by $D_S = +0.005$ and $D_R = +0.014$, inside the band.

Table 42: H3 of the second study, test partition: the drops in truncation (D_T) and in unparsed calls (D_P) of BUDGET-512-1500 against THINK-512 with their adjusted intervals, the truncation of THINK-512, the verdict, the replication verdict and the label.

Cell	D_T	D_P	THINK-512 truncation	Verdict	Replication verdict	Label
Qwen3-14B LB letters	+0.974 [+0.966, +0.980]	+0.704 [+0.687, +0.722]	+1.000 [+1.000, +1.000]	holds	holds	replicated
Qwen3-14B MMLU-Pro	+0.788 [+0.747, +0.827]	+0.782 [+0.741, +0.822]	+0.869 [+0.835, +0.901]	holds	holds	replicated
Qwen3-14B GSM8K	+0.531 [+0.472, +0.592]	+0.531 [+0.471, +0.592]	+0.545 [+0.486, +0.605]	holds	holds	replicated

Table 43: Qwen3-14B: accuracy with its unadjusted interval, parse rate and truncation (the share of calls ended at the token limit), by configuration, pool and partition.

Partition	Pool	Configuration	Accuracy	Parse	Truncated
Test	LB letters	THINK-64	0.000 [0.000, 0.000]	0.000	1.000
		THINK-512	0.001 [0.000, 0.001]	0.001	1.000
		THINK-1536	0.235 [0.223, 0.246]	0.341	0.727
		AUTHORS-1500	0.203 [0.192, 0.215]	0.564	0.000
		BUDGET-512-1500	0.404 [0.391, 0.416]	0.705	0.026
		RESCUE-1500	0.203 [0.192, 0.215]	0.564	0.000
	MMLU-Pro	THINK-64	0.000 [0.000, 0.000]	0.000	1.000
		THINK-512	0.120 [0.098, 0.143]	0.134	0.869
		THINK-1536	0.397 [0.363, 0.431]	0.498	0.502
		AUTHORS-1500	0.522 [0.487, 0.556]	0.960	0.015
		BUDGET-512-1500	0.647 [0.615, 0.680]	0.916	0.080
		RESCUE-1500	0.538 [0.504, 0.572]	0.960	0.015
	GSM8K	THINK-64	0.000 [0.000, 0.000]	0.000	1.000
		THINK-512	0.452 [0.410, 0.493]	0.457	0.545
		THINK-1536	0.860 [0.829, 0.888]	0.877	0.123
		AUTHORS-1500	0.211 [0.176, 0.247]	1.000	0.000
		BUDGET-512-1500	0.934 [0.911, 0.954]	0.989	0.013
		RESCUE-1500	0.550 [0.509, 0.592]	1.000	0.000
Replication	LB letters	THINK-64	0.000 [0.000, 0.000]	0.000	1.000
		THINK-512	0.000 [0.000, 0.000]	0.000	1.000
		THINK-1536	0.239 [0.223, 0.255]	0.345	0.729
		AUTHORS-1500	0.222 [0.205, 0.239]	0.567	0.000
		BUDGET-512-1500	0.416 [0.397, 0.434]	0.728	0.026
		RESCUE-1500	0.222 [0.205, 0.239]	0.567	0.000
	MMLU-Pro	THINK-64	0.000 [0.000, 0.000]	0.000	1.000
		THINK-512	0.113 [0.083, 0.145]	0.130	0.872
		THINK-1536	0.431 [0.383, 0.479]	0.534	0.466
		AUTHORS-1500	0.491 [0.444, 0.539]	0.960	0.015
		BUDGET-512-1500	0.687 [0.642, 0.732]	0.940	0.058
		RESCUE-1500	0.504 [0.456, 0.554]	0.965	0.015
	GSM8K	THINK-64	0.000 [0.000, 0.000]	0.000	1.000
		THINK-512	0.462 [0.402, 0.523]	0.466	0.538
		THINK-1536	0.886 [0.848, 0.924]	0.890	0.110
		AUTHORS-1500	0.239 [0.189, 0.292]	1.000	0.000
		BUDGET-512-1500	0.955 [0.928, 0.977]	1.000	0.004
		RESCUE-1500	0.576 [0.515, 0.636]	1.000	0.000

The general pools in the first study. On MMLU-Pro Qwen3-1.7B-Longevity reaches 0.217 to 0.234, +0.098 to +0.115 over the constant letter (0.119), with every interval above zero, but SmolLM3-3B and Intern-S1-mini reach 0.378 and 0.542 with their best configurations, and THINK-512, which truncates, is level with or below the constant for both. On GSM8K Qwen3-1.7B-Longevity scores 0.068 to 0.070 in every configuration, which is +0.036 [+0.011, +0.063] over the oracle constant answer (0.034) in test and +0.011 [−0.027, +0.049], an interval containing zero, in replication, against 0.880 and 0.922 for the best configurations of the other two. The model that is best on the longevity letters is thus the weakest on MMLU-Pro and barely above a constant answer on GSM8K. Section 4.6 compares the tuned model with its base model. Seeded sampling moves Qwen3-1.7B-Longevity's letter accuracy by −0.013 to −0.018, only

Table 44: Qwen3-14B, replication partition, analysed rows. In each cell, in percent of the cell's analysed calls and in this order: the unparsed calls, the unparsed calls that ended at the token limit, the non-committed calls (finish reason `stop`, no answer read) and the calls in the refusal class (a stated inability or no answer), which are part of the non-committed calls. All four shares are of the analysed calls of the cell. Descriptive.

Configuration	LB letters	LB-0038	MMLU-Pro	GSM8K
THINK-64	100.0 / 100.0 / 0.0 / 0.0	100.0 / 100.0 / 0.0 / 0.0	100.0 / 100.0 / 0.0 / 0.0	100.0 / 100.0 / 0.0 / 0.0
THINK-512	100.0 / 100.0 / 0.0 / 0.0	99.9 / 70.5 / 29.4 / 29.4	87.0 / 87.0 / 0.0 / 0.0	53.4 / 53.4 / 0.0 / 0.0
THINK-1536	65.5 / 61.9 / 3.6 / 3.1	99.6 / 0.0 / 99.6 / 99.6	46.6 / 46.6 / 0.0 / 0.0	11.0 / 11.0 / 0.0 / 0.0
AUTHORS-1500	43.3 / 0.0 / 43.3 / 36.6	100.0 / 0.0 / 100.0 / 100.0	4.0 / 1.3 / 2.8 / 0.0	0.0 / 0.0 / 0.0 / 0.0
BUDGET-512-1500	27.2 / 1.6 / 25.6 / 22.1	99.6 / 0.0 / 99.6 / 99.6	6.0 / 3.3 / 2.8 / 1.0	0.0 / 0.0 / 0.0 / 0.0

Table 45: Qwen3-14B, test partition, analysed rows. In each cell, in percent of the cell's analysed calls and in this order: the unparsed calls, the unparsed calls that ended at the token limit, the non-committed calls (finish reason `stop`, no answer read) and the calls in the refusal class (a stated inability or no answer), which are part of the non-committed calls. All four shares are of the analysed calls of the cell. Descriptive.

Configuration	LB letters	LB-0038	MMLU-Pro	GSM8K
THINK-64	100.0 / 100.0 / 0.0 / 0.0	100.0 / 100.0 / 0.0 / 0.0	100.0 / 100.0 / 0.0 / 0.0	100.0 / 100.0 / 0.0 / 0.0
THINK-512	100.0 / 100.0 / 0.0 / 0.0	99.9 / 72.0 / 28.0 / 28.0	86.6 / 86.6 / 0.0 / 0.0	54.3 / 54.3 / 0.0 / 0.0
THINK-1536	65.9 / 62.6 / 3.4 / 2.5	99.9 / 0.3 / 99.6 / 99.6	50.2 / 50.2 / 0.0 / 0.0	12.3 / 12.3 / 0.0 / 0.0
AUTHORS-1500	43.6 / 0.0 / 43.6 / 36.0	100.0 / 0.0 / 100.0 / 100.0	4.0 / 1.4 / 2.6 / 0.1	0.0 / 0.0 / 0.0 / 0.0
BUDGET-512-1500	29.5 / 1.6 / 27.9 / 24.3	99.9 / 0.0 / 99.9 / 99.9	8.4 / 5.9 / 2.5 / 0.6	1.1 / 0.8 / 0.4 / 0.0

Qwen3-1.7B-Longevity parses LB-0038 reliably, and the tied-versus-untied comparison on FreeToken and the GEPA candidates are in Appendix B.

P The audit checks

Table 52 summarises the six checks of Section 3.4. Their detection matrix is descriptive and largely true by construction: every rule expected to fire did so in 12 of the 16 instances (Appendices D and B give the exceptions and each instance). What is empirical is the false alarms on the clean instances: IR4 fires on 7 of 7, SCF1 on 4, SCF2 on 4, TB2 on 3, SCF3 on 2 and the gain rule of PR on 1, and no other rule fires on a clean instance.

Q The seeded arm and H3

Is the truncation an artifact of zero temperature? In the two models with a seeded arm, substantial truncation persists under sampling (descriptive; Table 23). Sampling at the model cards' recommended values (Qwen3-1.7B-Longevity: temperature 0.6, top-p 0.95, top-k 20; Intern-S1-mini: 0.8, 1.0, 50), with a seed per request, changes THINK-512's truncation by at most 0.023 in ten of the 12 cells (two models, three pools, two partitions) and raises it for Intern-S1-mini on GSM8K by +0.097 (test) and +0.098 (replication). Under sampling THINK-512 still ends at the token limit in 0.619 to 0.997 of Intern-S1-mini's test calls, and the drop in truncation that H3 measures, recomputed on the seeded outputs with unadjusted intervals, runs from +0.581 to +0.963 for Intern-S1-mini and is +0.068 for Qwen3-1.7B-Longevity on MMLU-Pro, next to +0.060 with greedy decoding. The preregistered verdicts use family-adjusted intervals at another level, so the seeded statistics are not compared with them, and SmolLM3-3B has no seeded arm.

R Sensitivity analyses

Each of the ten sensitivity analyses of the preregistration is descriptive. The strict scorer changes 2 of the 31 endpoint verdicts, and restricting to the never-called rows, which no earlier study had called, changes 3; these rows skew young and alive, with a mean LB-0030 age of 37.3 against 50.8 years and a share of LB-0042

Table 46: Mean completion tokens per call (the server's usage count), both partitions pooled, over all calls of Qwen3-14B's grids and harness evaluations; computed after the outputs existed.

Configuration	LB letters	LB-0038	MMLU-Pro	GSM8K
THINK-64	64.0	64.0	64.0	64.0
THINK-512	512.0	495.3	497.3	445.4
THINK-1536	1470.9	686.0	1150.4	711.7
AUTHORS-1500	322.0	29.7	76.3	23.7
BUDGET-512-1500	886.0	554.6	774.0	496.8
GEPA-CLEAN-s0	1391.3	–	–	–
GEPA-CLEAN-s1	414.4	–	–	–
GEPA-BROKEN-s0	64.0	–	–	–
GEPA-BROKEN-s1	64.0	–	–	–

Table 47: The first study's nine endpoints for the tuned 1.7B model: verdicts of the first study and of the rerun on the second study's GPU, on the test and replication partitions, and the replication labels.

Family	Pool	Endpoint	Test: first / rerun	Replication: first / rerun	Label: first / rerun
H1	LB letters	AUTHORS-1500 vs THINK-512	no gain: not tested / no gain: not tested	no gain: not tested / no gain: not tested	replicated / replicated
H2	LB letters	AUTHORS-1500 vs RESCUE-1500	holds / holds	holds / holds	replicated / replicated
H3	LB letters	BUDGET-512-1500 vs THINK-512	no room / no room	no room / no room	replicated / replicated
H1	MMLU-Pro	AUTHORS-1500 vs THINK-512	no gain: not tested / no gain: not tested	no gain: not tested / no gain: not tested	replicated / replicated
H2	MMLU-Pro	AUTHORS-1500 vs RESCUE-1500	holds / holds	holds / holds	replicated / replicated
H3	MMLU-Pro	BUDGET-512-1500 vs THINK-512	inconclusive / inconclusive	inconclusive / inconclusive	replicated / replicated
H1	GSM8K	AUTHORS-1500 vs THINK-512	no gain: not tested / no gain: not tested	no gain: not tested / no gain: not tested	replicated / replicated
H2	GSM8K	AUTHORS-1500 vs RESCUE-1500	holds / holds	holds / holds	replicated / replicated
H3	GSM8K	BUDGET-512-1500 vs THINK-512	no room / no room	no room / no room	replicated / replicated

rows whose follow-up says the participant was deceased of 0.022 against 0.346. That share counts a death at any time of the follow-up and is not the ten-year label, whose gold is "Yes" for a participant who died after ten years. The case study's pool changes 1 and LB-0030 alone 2, and LB-0034 alone, LB-0042 alone, both role-break variants and the pooled partitions change none. Putting the left-out rows back changes no verdict of H1–H3 and 2 H4 claims (the skewed split in SmolLM3-3B and Intern-S1-mini).

S The scorer and the refusal classifier

This appendix states the rules of the frozen scorer and of the classifier behind the non-commitment and refusal shares, as the code reads them. The scorer is `scripts/icml_ext/scoring.py` and the role-break rule is in `scripts/icml_ext/runner_checks.py`, both pinned by hash in the preregistration (Appendix E); the classifier is `scripts/icml_ext2/parse_audit_14b_ext2.py`, imported unchanged by `scripts/icml_ext2/parse_audit_main_ext2.py`. A call that no rule reads is unparsed and counts as wrong, and nothing falls back to a default.

Letter tasks, the strict rules (every call). (i) The whole answer is one valid option letter, with brackets, bold or a full stop allowed. (ii) Explicit answer statements ("answer is C", "Answer: C", "final answer: (C)" or a boxed letter) all name the same valid letter; two different letters, or a letter outside the options, are a conflict and the answer is unparsed. (iii) The whole answer equals exactly one option's label, ignoring case ("Yes", an age band).

Letter tasks, the extended rules. They apply only to a call that ended with finish reason `stop`, and only when the strict rules find neither a letter nor a conflict. (iv) The strict rules, applied to the answer without its markdown and LaTeX wrappers (bold, code, headings, dollar signs, boxed text). (v) The answer opens with a letter in brackets or followed by a closing bracket, a full stop, a colon or a dash, and no other letter or option text follows. A letter-label guard applies to every letter read from a stopped call: an answer that states the letter together with another option's label (for example "B. No" where B is "Yes") is unparsed. The strict scorer of the sensitivity analysis passes every call as not stopped, so it applies the strict rules only.

Table 48: The tuned 1.7B model on the second study's GPU minus the first study's run: accuracy difference and its unadjusted interval, by partition, pool and configuration of the grid (54 cells; descriptive).

Partition	Pool	THINK-64	THINK-512	THINK-1536	AUTHORS-1500	BUDGET-512-1500	RESCUE-1500
Replication	LB letters	0.000 [0.000, 0.000]	+0.002 [−0.009, +0.013]	+0.002 [−0.009, +0.013]	0.000 [−0.004, +0.004]	+0.002 [−0.009, +0.013]	+0.002 [−0.009, +0.013]
	MMLU-Pro	+0.003 [0.000, +0.008]	−0.003 [−0.020, +0.015]	0.000 [−0.018, +0.018]	+0.005 [0.000, +0.013]	0.000 [−0.020, +0.020]	+0.003 [−0.013, +0.020]
	GSM8K	+0.004 [0.000, +0.011]	+0.004 [0.000, +0.011]	+0.004 [0.000, +0.011]	+0.004 [0.000, +0.011]	+0.004 [0.000, +0.011]	+0.004 [0.000, +0.011]
Test	LB letters	0.000 [0.000, 0.000]	−0.001 [−0.008, +0.006]	−0.001 [−0.008, +0.006]	−0.001 [−0.004, +0.002]	−0.002 [−0.009, +0.006]	−0.002 [−0.009, +0.006]
	MMLU-Pro	−0.009 [−0.018, −0.001]	−0.008 [−0.020, +0.005]	−0.009 [−0.020, +0.003]	−0.009 [−0.016, −0.003]	−0.010 [−0.023, +0.003]	−0.006 [−0.019, +0.006]
	GSM8K	−0.004 [−0.011, +0.004]	−0.004 [−0.011, +0.004]	−0.004 [−0.011, +0.004]	−0.004 [−0.011, +0.004]	−0.004 [−0.011, +0.004]	−0.004 [−0.011, +0.004]
Pooled	LB letters	0.000 [0.000, 0.000]	0.000 [−0.006, +0.006]	0.000 [−0.006, +0.006]	−0.001 [−0.003, +0.001]	−0.001 [−0.006, +0.006]	−0.001 [−0.006, +0.006]
	MMLU-Pro	−0.005 [−0.011, +0.001]	−0.006 [−0.016, +0.004]	−0.006 [−0.016, +0.004]	−0.004 [−0.010, +0.001]	−0.007 [−0.018, +0.004]	−0.003 [−0.013, +0.007]
	GSM8K	−0.001 [−0.008, +0.004]	−0.001 [−0.008, +0.004]	−0.001 [−0.008, +0.004]	−0.001 [−0.008, +0.004]	−0.001 [−0.008, +0.004]	−0.001 [−0.008, +0.004]

Table 49: The tuned 1.7B model rerun on the second study's GPU against the first study's outputs (call by call; descriptive).

Comparison, tuned rerun against the first study	Calls	Token IDs differ	Finish reason differs	Request body differs
replication grid	18,015	10,478 (58.2%)	29	0
test grid	36,080	20,793 (57.6%)	72	0
replication, seeded arm	14,412	9,713 (67.4%)	37	0
test, seeded arm	28,864	19,392 (67.2%)	71	0

Number tasks (LB-0038 and GSM8K). The strict rules: (i) the whole answer is one number, with an optional unit word; (ii) explicit statements ("answer is", "Answer:", the marker `####`, a boxed number) agree on one number, and a number followed by an arithmetic operator is not taken; (iii) an age phrase that holds the only number of the answer. The extended rules, for a stopped call only: (iv) the strict rules without markdown and LaTeX; and, for GSM8K only, (v) the answer holds exactly one distinct number; (vi) the last non-empty line ends with an equals sign and a number or consists of one number with an optional currency sign or unit; (vii) the last sentence holds exactly one distinct number and no arithmetic operator. LB-0038 is read by the first four rules only, because the fifth had read laboratory values from answers that decline to give an age (decided after the pilot's parse audit). The last number of prose is never taken for being last.

The role-break rule. A model may emit the start-of-turn marker of its chat template after its answer. Each marker opens a turn whose role is the rest of its line. A re-opened assistant turn is the model going on with its own answer: the marker and the role line are dropped and the text that follows is kept. Any other turn (a user turn, a system turn or a line that names no role) ends the answer at its marker. Two sensitivity analyses vary the rule: one removes the markers and keeps everything, the other cuts the content at the first marker.

The refusal classifier. It applies to the calls that ended with finish reason `stop` and were unparsed, that is, the non-committed calls, and it puts each in the first of five classes that fits. (a) *No answer or refusal*: the final text is empty, or it matches a pattern of a keyword list, namely an apology or an inability ("I am sorry", "I cannot", "I do not know"), a statement that the answer cannot be determined, derived or given ("cannot be determined", "not derivable", "not provided", "does not contain", "insufficient information", "not enough data"), "none of the options", "unknown", "unanswerable", "indeterminate", "without additional context", "decline" or "apologize". For a letter task the next classes are (b) *several candidate letters* (two or more different letters or option texts), (c) *a form the scorer does not read* (one candidate, as an option text instead of its letter, inside markup or inside prose) and (d) *other* (only a letter outside the options, or no candidate). For a number task they are (b) *a range or several numbers* (a range such as "a to b" or "between a and b", or two different numbers on the answer line), (c) *a form the scorer does not read* (one number inside markup, with units or words around it, in plain form, or given only in words) and (d) *other* (no number on the answer line). The non-commitment share is the share of the cell's analysed calls that are non-committed, and the refusal share is the share of the cell's analysed calls in the first class. The keyword list is a heuristic: a refusal in other words falls into another class, and a harmless use of "unknown" falls into the first. The classifier reads counts only, and no output text was read for this paper.

Table 50: Accuracy, the strongest oracle constant of the task, and the share of answers equal to its label, both partitions.

Task and constant	Model and configuration	Accuracy	Answers equal to the constant label	Correct answers that are the label
LB-0030, 30-45 (0.270), replication	Qwen3-14B GEPA-CLEAN-s0	0.190	4 of 791 (0.5%)	1.3%
	Qwen3-14B GEPA-CLEAN-s1	0.229	122 of 792 (15.4%)	20.3%
	Qwen3-14B AUTHORS-1500	0.059	21 of 224 (9.4%)	8.5%
	Qwen3-14B BUDGET-512-1500	0.131	41 of 404 (10.1%)	11.5%
	Qwen3-1.7B-Longevity AUTHORS-1500	0.335	396 of 796 (49.7%)	49.4%
	Qwen3-1.7B-Longevity AUTHORS-1500, rerun	0.330	397 of 796 (49.9%)	49.4%
LB-0034, Participant B (0.510), replication	Qwen3-14B GEPA-CLEAN-s0	0.565	331 of 396 (83.6%)	80.4%
	Qwen3-14B GEPA-CLEAN-s1	0.580	187 of 397 (47.1%)	48.5%
	Qwen3-14B AUTHORS-1500	0.523	179 of 291 (61.5%)	57.7%
	Qwen3-14B BUDGET-512-1500	0.392	173 of 266 (65.0%)	60.9%
	Qwen3-1.7B-Longevity AUTHORS-1500	0.744	261 of 398 (65.6%)	61.1%
	Qwen3-1.7B-Longevity AUTHORS-1500, rerun	0.744	261 of 398 (65.6%)	61.1%
LB-0042, Yes (0.823), replication	Qwen3-14B GEPA-CLEAN-s0	0.830	748 of 788 (94.9%)	95.9%
	Qwen3-14B GEPA-CLEAN-s1	0.828	779 of 790 (98.6%)	98.5%
	Qwen3-14B AUTHORS-1500	0.234	105 of 613 (17.1%)	50.0%
	Qwen3-14B BUDGET-512-1500	0.712	504 of 779 (64.7%)	82.4%
	Qwen3-1.7B-Longevity AUTHORS-1500	0.876	686 of 796 (86.2%)	89.1%
	Qwen3-1.7B-Longevity AUTHORS-1500, rerun	0.881	688 of 796 (86.4%)	89.0%
LB-0030, 60-75 (0.250), test	Qwen3-14B GEPA-CLEAN-s0	0.202	526 of 1,576 (33.4%)	46.7%
	Qwen3-14B GEPA-CLEAN-s1	0.232	469 of 1,563 (30.0%)	37.2%
	Qwen3-14B AUTHORS-1500	0.058	113 of 432 (26.2%)	28.3%
	Qwen3-14B BUDGET-512-1500	0.126	274 of 776 (35.3%)	43.2%
	Qwen3-1.7B-Longevity AUTHORS-1500	0.344	132 of 1,577 (8.4%)	8.7%
	Qwen3-1.7B-Longevity AUTHORS-1500, rerun	0.339	129 of 1,577 (8.2%)	8.6%
LB-0034, Participant A (0.515), test	Qwen3-14B GEPA-CLEAN-s0	0.602	154 of 783 (19.7%)	25.9%
	Qwen3-14B GEPA-CLEAN-s1	0.586	458 of 785 (58.3%)	58.4%
	Qwen3-14B AUTHORS-1500	0.445	193 of 544 (35.5%)	40.2%
	Qwen3-14B BUDGET-512-1500	0.383	174 of 490 (35.5%)	39.7%
	Qwen3-1.7B-Longevity AUTHORS-1500	0.717	267 of 788 (33.9%)	39.8%
	Qwen3-1.7B-Longevity AUTHORS-1500, rerun	0.717	267 of 788 (33.9%)	39.8%
LB-0042, Yes (0.840), test	Qwen3-14B GEPA-CLEAN-s0	0.847	1,486 of 1,570 (94.6%)	96.3%
	Qwen3-14B GEPA-CLEAN-s1	0.837	1,545 of 1,574 (98.2%)	99.0%
	Qwen3-14B AUTHORS-1500	0.227	202 of 1,248 (16.2%)	49.2%
	Qwen3-14B BUDGET-512-1500	0.691	983 of 1,512 (65.0%)	84.0%
	Qwen3-1.7B-Longevity AUTHORS-1500	0.878	1,353 of 1,577 (85.8%)	89.7%
	Qwen3-1.7B-Longevity AUTHORS-1500, rerun	0.880	1,355 of 1,577 (85.9%)	89.8%

What the classes do not tell. The harness texts of the second study's clean-start searches tell the model not to infer survival beyond the snapshot data, or to conclude that survival cannot be determined when follow-up data are missing (Appendix L), and the harnesses still answer "Yes" in almost all of their parsed mortality answers. Because the outputs were not read, why they do is not examined here.

Table 51: Accuracy in the three pools and in each LongevityBench task, test partition, next to the trivial baselines; in parentheses the paired difference to the strongest baseline (the oracle constant, chosen once and held fixed in the bootstrap). Unparsed answers count as wrong. THINK-64, the GEPA harnesses and the replication partition are in Appendix B. Computed after the outputs existed (Appendix C).

Model	THINK-512	THINK-1536	AUTHORS-1500	BUDGET-512-1500	RESCUE-1500
LongevityBench letter pool (3,942 rows). Chance 0.380; constant letter 0.383 (chosen on replication), 0.389 (oracle); constant label 0.530 (chosen on replication), 0.539 (oracle). Strongest, the oracle constant label: 0.539 [0.527, 0.551].					
Qwen3-1.7B	0.635 (+0.096)	0.635 (+0.096)	0.632 (+0.093)	0.635 (+0.096)	0.635 (+0.096)
SmolLM3	0.070 (−0.470)	0.157 (−0.382)	0.129 (−0.410)	0.121 (−0.418)	0.128 (−0.411)
Intern-S1	0.000 (−0.539)	0.125 (−0.414)	0.314 (−0.225)	0.253 (−0.286)	0.314 (−0.225)
LB-0030 (1,577 rows). Chance 0.200; constant letter 0.216 (oracle); constant label 0.250 (oracle). Strongest, the oracle constant label: 0.250 [0.230, 0.269].					
Qwen3-1.7B	0.369 (+0.119)	0.369 (+0.119)	0.344 (+0.094)	0.369 (+0.119)	0.369 (+0.119)
SmolLM3	0.000 (−0.250)	0.021 (−0.229)	0.053 (−0.197)	0.032 (−0.218)	0.053 (−0.197)
Intern-S1	0.000 (−0.250)	0.004 (−0.245)	0.023 (−0.227)	0.055 (−0.195)	0.023 (−0.227)
LB-0034 (788 rows). Chance 0.500; constant letter 0.515 (oracle); constant label 0.515 (oracle). Strongest, the oracle constant label: 0.515 [0.480, 0.551].					
Qwen3-1.7B	0.702 (+0.187)	0.703 (+0.188)	0.717 (+0.202)	0.703 (+0.188)	0.703 (+0.188)
SmolLM3	0.307 (−0.208)	0.409 (−0.107)	0.443 (−0.072)	0.355 (−0.160)	0.435 (−0.080)
Intern-S1	0.000 (−0.515)	0.242 (−0.273)	0.575 (+0.060)	0.523 (+0.008)	0.575 (+0.060)
LB-0042 (1,577 rows). Chance 0.500; constant letter 0.500 (oracle); constant label 0.840 (oracle). Strongest, the oracle constant label: 0.840 [0.822, 0.858].					
Qwen3-1.7B	0.867 (+0.027)	0.867 (+0.027)	0.878 (+0.038)	0.867 (+0.027)	0.867 (+0.027)
SmolLM3	0.020 (−0.820)	0.167 (−0.673)	0.049 (−0.791)	0.093 (−0.748)	0.050 (−0.790)
Intern-S1	0.000 (−0.840)	0.186 (−0.654)	0.475 (−0.365)	0.317 (−0.523)	0.475 (−0.365)
MMLU-Pro (799 items). Chance 0.110; constant letter B, chosen on the other partition and oracle alike, 0.119 [0.096, 0.141].					
Qwen3-1.7B	0.217 (+0.098)	0.234 (+0.115)	0.229 (+0.110)	0.232 (+0.113)	0.232 (+0.113)
SmolLM3	0.098 (−0.021)	0.314 (+0.195)	0.353 (+0.234)	0.258 (+0.139)	0.378 (+0.259)
Intern-S1	0.074 (−0.045)	0.279 (+0.160)	0.477 (+0.358)	0.542 (+0.423)	0.493 (+0.374)
GSM8K (527 items). Constant answer 10 (chosen on the other partition) 0.019 [0.008, 0.032], 2 (oracle) 0.034 [0.019, 0.051]; no chance level.					
Qwen3-1.7B	0.070 (+0.036)	0.070 (+0.036)	0.070 (+0.036)	0.070 (+0.036)	0.070 (+0.036)
SmolLM3	0.681 (+0.647)	0.880 (+0.846)	0.789 (+0.755)	0.835 (+0.801)	0.879 (+0.844)
Intern-S1	0.476 (+0.442)	0.833 (+0.799)	0.190 (+0.156)	0.922 (+0.888)	0.531 (+0.497)

Table 52: The audit checks of H4: what each reads, its outcome on the 16 defective instances of the detection matrix (*detects*: fires on the defective instance and not on the clean one; *uninformative*: fires on both; *misses*: does not fire; descriptive) and the rules that fire on the clean instances (what is empirical).

Check	Reads	Detects	Uninformative	Misses	Fires on the clean instances
PR	gain and level rules of prospective replication	5	1	10	PR gain 1/7
TR	per-call prompt-token reconciliation	3	0	13	none
TB	trivial baselines	6	3	7	TB2 3/7
SCF	the simple configuration family	3	12	1	SCF1 4/7, SCF2 4/7, SCF3 2/7
IR	reading of the rendered inputs	0	16	0	IR4 7/7
SMI	served-model identity (defect E)	1	0	0	–