arXiv:2611.06965v1 [cs.SE] 4 Nov 2026

When Can Self-Play Fix a False Belief?

Answer Keys, Pass Bands and Replacement Rules in an Execution-Grounded Loop

Nirdesh Chauhan Dushyant Rajput Siddharth Kosaraju
AltSlate Labs LLP
nirdesh@altslate.com dushyant@altslate.com siddharth@altslate.com

November 2026

Abstract

Self-play loops in which a model writes tasks, derives an answer key and trains on its own matching solutions are often called grounded when the key comes from executing code. We ask which false beliefs such a loop removes. We run expert-iteration self-play with Qwen3-4B-Base on output-prediction tasks, vary the source of the answer key, and track natural beliefs about Python semantics and two misreadings of task specifications planted by fine-tuning ("at most n" read as $< n$; "the k-th smallest" read as 0-indexed). Executing the generator's own code gave a key that encoded the "at most" misreading on every audited task, and training on it never raised the truth rate above 0.03. A key from an implementation by the unplanted model, a teacher without the error, raised it to 0.23–0.58 after three rounds in four runs. A correct key is not sufficient: the pass band, which keeps tasks solved 1–7 times of 8, admits almost no corrective task while a belief is held on nearly every sample. Execution keys sharply reduced both natural runtime beliefs on diagnostic items in both seeds. Corrections came with replacement rules that fail on neighbouring items the old belief got right. The robust case, present in all four spec runs, is that after the "at most" fix the model increasingly counts the nearest value above a limit that is absent from the list (overcount 0.18–0.38 of control samples at round 3), a case our planted generators almost never wrote (4 of 10,924 tasks); diagnostic-only metrics miss it. Asked for those cases, the planted generator wrote none in 2,061 checkable tasks, against 27% for the untrained unplanted model in one round. Both generators often moved the request into the specification instead ("the cap itself does not count"), and a decorrelated key follows the specification.

1 Introduction

In a growing family of self-improvement methods, one model proposes tasks, something supplies the answers, and the model trains on its own solutions that match them. When the answers come from a code executor or a proof checker, the loop is described as grounded; Absolute Zero [47], Self-Challenging Agents [49] and STP [9] rely on one. Other loops replace the checker with agreement among the model's own samples [17, 50]. He et al. [14] show that such intrinsic rewards sharpen the model's initial distribution, which helps when confidence and correctness align and fails when they do not. An external executor is often taken to escape this limit.

An executor computes what code does. It does not decide what the code should do. When the model writes the program whose output becomes the answer key, the key inherits the beliefs that shaped the program. The executor is also only one link in the loop: the model chooses which inputs to write, and the training filter decides which tasks carry any gradient. Which of a model's false beliefs can such a loop remove?

We study this with a model organism. We take false beliefs whose truth we know: natural beliefs about Python's runtime semantics that the base model holds, and misreadings of task specifications that we plant by fine-tuning. We then run the same expert-iteration loop several times and change the source of the answer key. Each belief is measured on held-out probes that also contain *neighbours*: items near the belief where the old belief gives the right answer. Over 76 trained rounds on one GPU, we find the following.

1. **Correlated and decorrelated keys.** A key computed by executing the generator's own program encodes the planted misreading on every task we audited, so training on it cannot raise the truth rate. This is close to true by construction once the misreading reaches the generator's code. A key from an implementation written by the unplanted model raises the truth rate in all four runs (Section 4). That key comes from a checkpoint known not to hold the belief, so it acts as a teacher and gives an upper bound on what decorrelation can do.

2. **The pass band.** A correct key is not sufficient. When the solver holds a belief on almost every sample, almost no sample matches the key and the filter admits almost no corrective task. In one six-round run the 0-indexing belief barely moved; in another it broke out after a slow start (Section 5).

3. **Execution keys and natural beliefs.** Execution keys sharply reduced both runtime beliefs on diagnostic items (to 0.04–0.19), in both seeds, far beyond a control without exposure; an unchecked key gave a weaker, seed-dependent fix (Section 6).

4. **Replacement rules.** For each of the three corrected beliefs with a neighbour probe, at least one run showed a rule that is wrong where the old belief was right: the S2 rule in all four spec runs and only under the decorrelated key; the `%` rule in both runtime seeds but also in a control; the R1 rule in one seed. After the "at most" fix the model increasingly counts the nearest value above an absent limit, and diagnostic-only metrics cannot see it (Section 7).

5. **Generation.** Asked explicitly for the missing neighbours, the planted generator never wrote one; the unplanted model, in one untrained round on the same topics, wrote one in 27% of its checkable tasks. Both often moved the request into the specification instead, as a precondition the call violated or as a redefinition ("the cap itself does not count"), and a key that follows the specification then agrees with the belief (Section 8).

The first two findings are expected from how expert iteration works; we include them because they set up the others and because each gate can be audited task by task. We regard the replacement rules and the generator's rewriting of the specification as the main new observations. Neither has a causal test here: we report what we observed and which explanations they are consistent with.

Contributions.

- A controlled design that varies the answer key of a self-play loop, with natural and planted false beliefs and probes that separate truth, belief and replacement rules (Section 3).
- A three-gate account of the direct route by which a loop corrects a belief (generation, key, pass band), with per-task counts for each gate (Sections 4–6 and 8).
- Measurements of replacement rules that accompany corrections in self-play, including their shape and their presence in the selected training data (Section 7).
- A failed remedy that is informative: the planted generator, asked for the cases it lacked, did not write them, and a key-agreement proxy for "this task is a neighbour" failed silently (Section 8).

2 Three gates between a belief and its correction

Consider one round of the loop. For a task t let $y^*(t)$ be the true output, $b(t)$ the output the false belief predicts, and $\kappa(t)$ the answer key. The solver draws $k = 8$ samples, of which $h(t)$ match the key. Expert iteration trains on t only if $1 \le h(t) \le k - 1$ (the *pass band*), using matching samples. A task *touches* the belief if $b(t) \neq y^*(t)$. For the belief to receive a *direct* corrective gradient, some trained task must pass three gates:

G1, generation: the generator writes a task that touches the belief;

G2, key: the key differs from the belief on it, $\kappa(t) \neq b(t)$; a correction toward the truth needs $\kappa(t) = y^*(t)$, and a key that is neither only adds noise;

G3, passage: the solver sometimes produces the key, $1 \le h(t) \le 7$.

When G2 fails because the key equals the belief, trained tasks reinforce the belief. This is how label-free rewards conserve errors [14], and it is analogous to how errors that a verifier accepts survive in verifier-gated cascades [32]. When G3 fails, even a correct key is invisible. The pass band is not specific to our loop: it matches the learnability rewards of self-play proposers [17, 47], the zero-advantage filtering of group-relative RL [44], and the observation that RL with verifiable rewards rarely goes beyond what the base model can already sample [45].

The gates describe the direct route only. Beliefs can also move without any touching task being trained: in our runtime control, which never trains on a touching task, R1 fell from 0.61 to 0.37 in one seed (but not in the other), and the S1 break-out of Section 5 began after only 8 S1 tasks had been trained. Gate counts are therefore necessary conditions for a direct correction, not predictors of when a belief moves. In our runs, round-0 counts did not predict speed either (Section 5).

What the corrective cases teach. The tasks that pass all three gates are, by definition, cases where the belief is wrong, and for each belief we study, the truth differs from the belief's answer in the same way on all of them (Table 1). Any rule that agrees with the truth on every touching case fits the corrective data equally well; the data say nothing about neighbours, where $b(t) = y^*(t)$. For R1 and `%`, extending the corrective difference to neighbours gives a specific wrong rule ("round `.5` toward zero", "add b when the signs differ"), but we observe these rules inconsistently (Section 7). For S2 the extension adds nothing (an absent limit has no copies to count), yet S2 shows the most robust rule. The corrective data underdetermine the rule rather than predict it. The pass band may also discard neighbours the model already answers 8 times of 8; we did not measure how often this happened for the relevant neighbours. The selected samples carry a second risk: they are chosen by their final answer alone [25, 41], so a sample can match the key through reasoning that is itself wrong, a model "right for the wrong reasons" [28]. A probe that contains only diagnostic items, where $b \neq y^*$, rewards every rule that moves away from the belief. Probes need neighbours, as contrast sets [10] and the specificity and ripple-effect measures of model editing [7, 30] do.

3 Setup

3.1 Model and loop

We use Qwen3-4B-Base [43] with raw-text few-shot prompts, bf16 weights, vLLM [24] for generation and LoRA [15] (rank 32, $\alpha = 64$, all attention and MLP projections) for training. The loop is expert iteration [1], i.e. rejection-sampling fine-tuning as in STaR [46] and ReST$^{\mathrm{EM}}$ [37]. One round (Figure 1) has five steps.

ID	Belief	Truth	Touching cases; how the truth differs	Neighbours ($b = y^*$)
R1	`round(x.5)` rounds half away from zero	half to even	even `.5` ties; the truth is one smaller in magnitude	odd `.5` ties
R2	`//` and `%` truncate toward zero	floor	mixed signs, inexact; `%`: truth = belief's remainder $+b$; `//`: truth = belief -1	exact divisions with mixed signs
S1	"the k-th smallest" is 0-indexed	1-indexed	every call; the truth is the element one position earlier	none
S2	"at most n" means $< n$	$\leq n$	the limit n is in the list; the truth also counts the copies of n	n absent from the list

Table 1: The four targeted beliefs. R1 and R2 are natural in the base model; S1 and S2 are planted. On every touching case, the truth differs from the belief's answer in the same way.

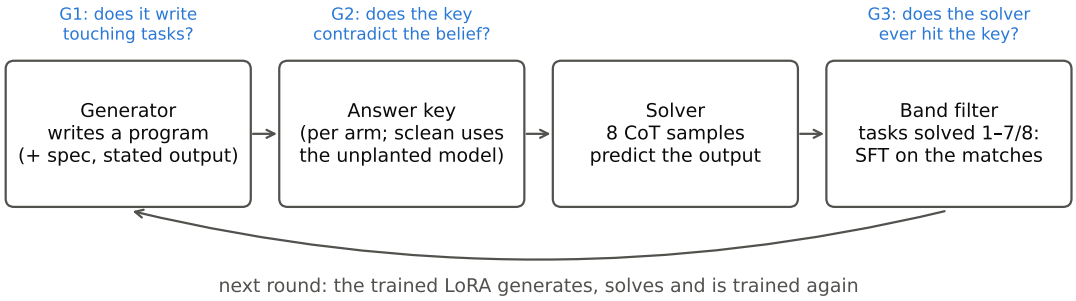


Figure 1: One round of the loop. The generator, solver and trained model are the same model. The arms differ in the answer key (Table 2); from round 1 the two control arms generate from neutral topics only.

1. **Generate:** 3,000 task requests. Each draws a topic from a fixed list: a belief topic with probability 0.15 per targeted belief (0.30 in all), otherwise one of 20 (runtime) or 16 (spec) neutral topics. For runtime tasks the generator writes a short program that prints one line; for spec tasks it writes a one-sentence specification and a program that implements it and prints the result of one call. In both cases it also states the printed output. We keep programs that parse, are not duplicates, run in a sandbox and print one line (2,178–2,705 at round 0).

2. **Key:** each task is keyed according to the arm (Table 2).

3. **Solve:** 8 samples per task with short chain-of-thought (temperature 0.7, top-p 0.8, top-k 20). For runtime tasks the solver sees the program; for spec tasks it sees only the specification and the call.

4. **Filter:** tasks solved 1–7 times of 8 under the key are kept, with up to 2 distinct matching solutions each.

5. **Train:** one epoch of SFT on the kept solutions (learning rate 10^{-4}, batch 16), continuing the previous round's adapter. Training sets are small: 76–466 examples per round in the spec half (5–30 optimizer steps; 155–448 under `sclean`) and up to 1,769 in the runtime half.

From round 1 on, each arm generates its own tasks with its own trained model. The tasks are *output prediction* ("what does this print?"), the deduction task of Absolute Zero [47] and the output prediction of CRUXEval [13]. We chose it over code writing because in code writing a solver that shares the reference's `round()` call passes whatever it believes, so the belief never meets the key. The arms are not identical apart from the key: the control arms generate from neutral topics only, and training-set sizes differ between arms (at round 0, 763 examples for `agate` against 1,769 for `exec`).

3.2 Beliefs

Runtime beliefs are natural. With chain-of-thought, the base model gives the belief's answer on 61% of samples for R1 (it rounds 54.5 to 55) and 41% for R2, and answers the matching controls correctly (49/50 and 50/50 greedy). A third runtime belief, R3 (`int("3.0")` parses), is held at 0.14 and is not targeted; we track it as a collateral measure.

Spec beliefs are planted. S1 and S2 are rare in the base model (belief rates 0.05 and 0.12). We plant them by LoRA SFT, then merge the adapter. The planting data have two formats: question–answer examples in the probe style, with different names and lists, whose reasoning states the belief ("'At most 7' means strictly below 7"), and generator-format examples whose program implements the belief. Both are mixed with 1,769 of the base model's own correct samples, taken from the round-0 training set of the runtime `exec` arm. Strong planting uses 500 examples per belief per format for 2 epochs; weak planting uses 100 for 1 epoch. After planting, belief rates are S1 1.00 / 0.99 and S2 0.92 / 0.75 (strong / weak). Two properties of the planting data matter later: every S1 program is `sorted(xs)[k]`, and every S2 example draws the limit n from the list. A third spec belief, S3 ("substring" means subsequence), was dropped before training: the base model said "True" for subsequence-only pairs about as often as for real substrings (0.52 vs 0.62) and never for pairs that were neither, which points to a character-level limit rather than a misreading. We take the 1-indexed and $\leq$ readings as the truth by convention; we did not check with human readers that the specifications are unambiguous.

3.3 Arms

Half	Arm	Answer key	Shares the belief?
runtime	`atrust`	the generator's stated output, unchecked	partly
runtime	`agate`	the stated output; the task is kept only if execution agrees	no (disagreements dropped)
runtime	`exec`	execution of the program	no
runtime	`ctrl`	execution; neutral topics only, and tasks that touch R1 or R2 are dropped	no exposure
spec	`sexec`	execution of the generator's own implementation	yes, by construction
spec	`sclean`	the *decorrelated key*: execution, on the generator's call, of an implementation of the same specification written by the unplanted base model (greedy, 3-shot)	no
spec	`sctrl`	as `sexec`; neutral topics only, and S1 and S2 topic tasks are dropped	no exposure

Table 2: Arms. A runtime task touches R1 or R2 when re-running its program with `round`, `//` and `%` replaced by the belief's versions changes the output; in the spec half a task's belief is its topic. Spec runs start from the planted model, runtime runs from the base model.

3.4 Probes and metrics

The probes are held-out items written by script and never seen by the loop. R1, R2 and S2 have 150 items each: 100 *diagnostic* items, where belief and truth differ, and 50 *controls*, which are neighbours where they agree. S1 has 150 diagnostic items and no controls. For S2 a control is a list that does not contain the limit. Each item gets a greedy answer and 8 chain-of-thought samples (temperature 0.7) in a 3-shot raw-text prompt, the mode the solver runs in. We report the share of samples that give the truth, the belief or something else, with 95% bootstrap confidence intervals over items (2,000 resamples). On S2 controls we also report the *overcount*, the share of samples with a number above the truth.

Two further probes target replacement rules in the runtime half. The first has 300 exact `.5` ties: 150 even, where the truth differs from the belief, and 150 odd, where the truth differs from "round toward zero". The second has 200 exact divisions with a negative operand (100 `//`, 100 `%`); floor and truncation agree on these, so they are neighbours of R2, and the `%` truth is 0.

3.5 Runs and statistics

The runtime half has 4 arms × 3 rounds × 2 seeds. The spec half has 3 arms × 3 rounds × 2 planting strengths × 2 seeds; strong seed 0 is extended to 6 rounds for `sclean` and `sctrl`. Two further `sclean` runs add neighbour topics (Section 8). A seed sets task sampling and training order. The planted model is shared between seeds, and round r of a run with seed s draws its topic list with seed $s+r$, so seed 1's round r reuses seed 0's round $r+1$ topic list (the generated tasks differ, because the models differ). All runs used one NVIDIA GB10 (DGX Spark), at about 20 minutes per round.

Confidence intervals are over probe items within a run and do not include run-to-run variation, which is large (S2 truth at round 3 is 0.28 and 0.58 in the two strong seeds). For the main contrasts we also report paired item-level bootstrap intervals (Appendix C) and the direction across runs. These 24 contrasts are not corrected for multiplicity; the stronger evidence is that each S2 truth contrast (`sclean` − `sexec`, `sclean` − `sctrl`) has the same sign in all four runs (a one-sided sign test over four runs gives $p = 0.0625$). We did not fit the mixed-effects model of our design document. Analyses introduced after seeing the data (the gap breakdown in Table 6, the rationale audit, the S1 break-out and the neighbour-topic audit) are exploratory.

4 Correlated and decorrelated keys

Run	Planted	`sexec`	`sctrl`	`sclean`
strong, seed 0	0.03 / 0.92 / 0.06	0.00 [0.00, 0.00] / 0.81 / 0.19	0.02 [0.00, 0.05] / 0.91 / 0.07	**0.58** [0.50, 0.67] / 0.26 / 0.15
strong, seed 1	0.03 / 0.91 / 0.06	0.01 [0.00, 0.04] / 0.81 / 0.18	0.02 [0.00, 0.04] / 0.91 / 0.07	**0.28** [0.20, 0.36] / 0.58 / 0.14
weak, seed 0	0.05 / 0.75 / 0.20	0.02 [0.00, 0.05] / 0.83 / 0.14	0.04 [0.01, 0.08] / 0.61 / 0.34	**0.23** [0.15, 0.30] / 0.61 / 0.16
weak, seed 1	0.05 / 0.75 / 0.20	0.03 [0.01, 0.07] / 0.88 / 0.09	0.04 [0.01, 0.07] / 0.71 / 0.26	**0.29** [0.21, 0.38] / 0.55 / 0.16

Table 3: S2 at round 3: truth [95% CI] / belief / other, as shares of samples on diagnostic items. The two seeds share each planted model; their planted values differ slightly because the model was probed twice.

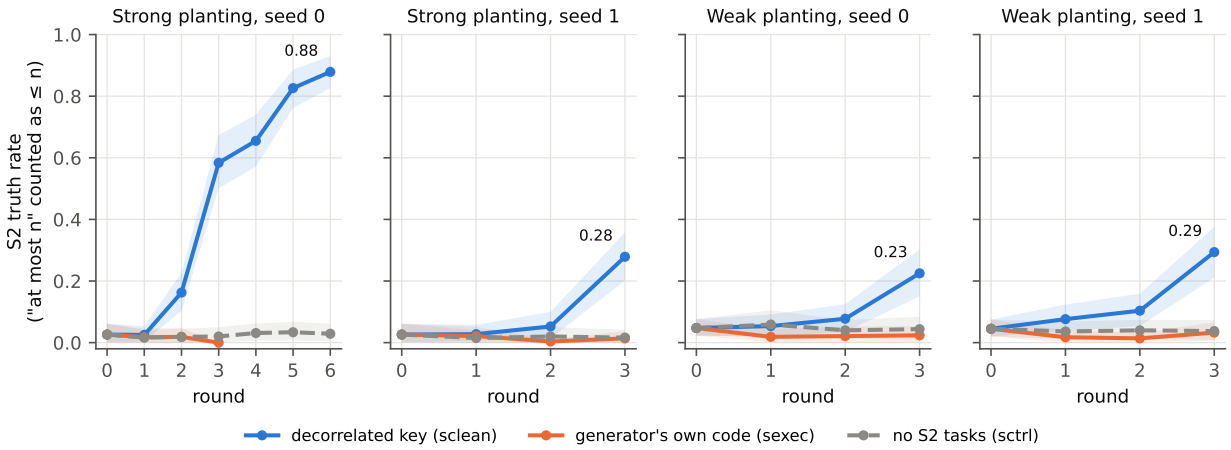


Figure 2: S2 truth rate on diagnostic items (the limit is in the list) by round, in the four spec runs. Shaded: 95% CI over items. Only the decorrelated key (`sclean`) moves the model toward the truth.

A key from executing the generator's code encodes its misreading. At round 0 of the strong run, the generator's executed key equals the < count on all 431 checkable S2 tasks (calls of the form `f(list, number)`), and the limit is in the list on every one; the same holds at round 0 of the weak run (430 of 430) and at round 2 of a strong `sclean` run (441 of 441). `sexec` therefore trains almost only on belief-keyed S2 tasks (round-0 fates in Table 9): 21 of 22 at round 0 of the strong run and 88 of 90 at round 0 of the weak run. After training, its truth rate is 0.00–0.03 in every run (Table 3). Given that planting trained the generator format on a < implementation, this result is close to a manipulation check: once a misreading reaches the generator's code, executing that code cannot contradict it.

The same holds for S1 under both plantings: the generator's executed key is the 0-indexed answer on every checkable S1 task we counted (428 of 428 at strong round 0, 422 of 422 at weak round 0, and all of them in later weak rounds).

What `sexec` does to the belief rate is less simple. In the strong runs the belief fell (0.92 → 0.81 and 0.91 → 0.81; paired differences −0.10 [−0.18, −0.03] and −0.10 [−0.17, −0.03]), into "other" answers rather than the truth. In the weak runs it rose (0.75 → 0.83 and 0.88; paired +0.08 [+0.01, +0.15] and +0.13 [+0.06, +0.20]), mainly because "other" answers fell (0.20 → 0.14 and 0.09). Training on belief-keyed tasks removed some counting errors under the same reading; we do not read it as a shift between readings.

A decorrelated key moves the model toward the truth. `sclean` is the only arm whose truth rate rises in all four runs, with confidence intervals clear of both other arms (Figure 2, Table 3); paired differences against `sexec` and `sctrl` at round 3 range from +0.18 to +0.58, with every interval above zero. In strong seed 0, six rounds take S2 truth to 0.88 [0.83, 0.93] and belief to 0.03. The decorrelated key is not perfect. On checkable S2 tasks it matches the ≤ count on 88–94% of tasks at the rounds we audited (strong and weak round 0, strong round 2). There every mismatch is an exception: the base model's implementation treats the list's integers as objects with a `.price` attribute. No solver sample matches an exception, so these tasks never reach training. In the plain runs we found one admitted S2 task whose decorrelated key is the < count (weak seed 1, round 0, whose specification itself says `x < n`).

No exposure: the control. `sctrl` sees no S2 tasks, and its truth rate is flat. In the weak runs its belief falls, but into "other" answers (0.20 → 0.34 and 0.26), not to the truth. A metric

built on the belief rate alone would have counted this as a partial fix.

5 The pass band

S1 barely moves under the decorrelated key. S1 is held at 0.99–1.00 after both plantings. Under `sclean` at round 0 of the strong run, the decorrelated key gives the 1-indexed answer on 395 of 428 checkable S1 tasks. The other 33 are a natural error of the unplanted model, which sometimes writes `sorted(xs)[k]` too; across strong rounds this happens on 5–8% of S1 tasks, and in the weak runs, whose generators write different S1 specifications, on 26–34%. None of the mismatching specifications mentions an index or a position, so these are errors of the key writer, not faithful readings of the text. On correctly keyed tasks the solver almost never gives the 1-indexed answer, so few S1 tasks are admitted (Table 4), and some of those carry the belief-keyed answer (4 of 73 in strong seed 0; about half in the weak runs). Six rounds move the S1 belief in strong seed 0 from 1.00 to 0.96 [0.93, 0.98]. Here the key largely passes G2; G3 blocks the task. The block belongs to the training filter, not to the key; RL with a binary reward has the same problem when every sample in a group fails [26, 44].

The block is leaky. Under `sclean` the number of admitted S2 tasks grows as the belief loosens (Table 4). In strong seed 0 it goes $9 \to 38 \to 76 \to 114$ and S2 truth reaches 0.58 at round 3; in strong seed 1 it stalls at $17 \to 25 \to 24$ and truth reaches 0.28. In all four runs the largest jump in truth comes between rounds 2 and 3. It coincides with growing admissions and with shrinking training sets ($401 \to 273$ examples by round 2 in strong seed 0), which make the admitted S2 tasks a larger share of each round's training; we cannot separate a threshold from this dilution effect. Round-0 admissions did not predict speed: the run with the fewest (9) fixed fastest (0.58 at round 3), and the runs with 17–31 reached 0.23–0.29.

S1 broke out in one run: the second neighbour-topic run (Section 8). The two neighbour-topic runs share round 0, seed and planting and differ only in the training filter for key-agreeing S2 tasks. At round 4 S1 belief was 0.98 in the first and 0.90 [0.85, 0.94] in the second, which then fell to 0.61 at rounds 5 and 6 (truth 0.33). S1 training was minimal: only 8 S1 tasks had been trained when the belief first moved, 2 of them with the belief-keyed answer. Once it moved, admissions jumped to 42 and then 133 per round. This is one matched pair of single runs; we do not know why the second run broke out.

`sclean` run	S2 tasks admitted, rounds 0, 1, 2, ...	S2 truth, round 3	S1 tasks admitted
strong, seed 0	9, 38, 76, 114, 83, 31	0.58 (0.88 at round 6)	1, 3, 14, 15, 12, 28
strong, seed 1	17, 25, 24	0.28	4, 2, 1
weak, seed 0	24, 52, 73	0.23	3, 1, 2
weak, seed 1	31, 82, 82	0.29	4, 5, 5

Table 4: Tasks admitted by the pass band per round under the decorrelated key. Admitted S2 tasks are keyed with the decorrelated key, which is almost always the truth (Section 4); admitted S1 tasks include some belief-keyed ones (Section 5). The model at round r has been trained on the tasks of rounds $0, \ldots, r-1$.

6 Runtime half: execution keys and natural beliefs

Execution keys remove runtime beliefs on the items that expose them. In both seeds and for both beliefs, `exec` and `agate` end far below `ctrl` (Figure 3, Table 5); paired differences

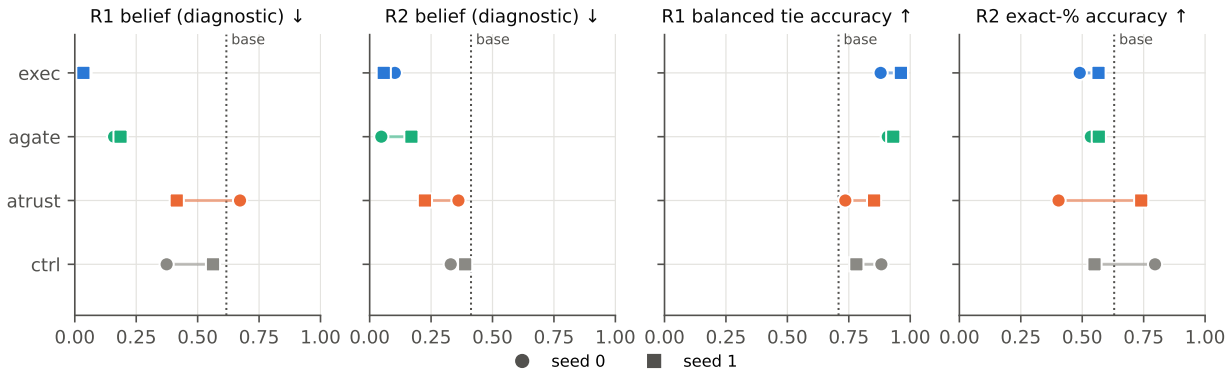


Figure 3: Runtime arms at round 3, two seeds (circle, square). Left two panels: belief rate on diagnostic items (lower is better). Right two: accuracy on tie items (even and odd) and on exact `%` items, which are neighbours of R2 (higher is better). Dotted line: base model. The two `exec` points overlap in the first panel (0.04 in both seeds).

Arm, round 3	R1 belief, seed 0 / seed 1	R2 belief, seed 0 / seed 1
base	0.61 / 0.62	0.41 / 0.42
`exec`	**0.04** [0.02, 0.07] / **0.04** [0.01, 0.06]	**0.10** [0.06, 0.15] / **0.06** [0.03, 0.09]
`agate`	**0.16** [0.10, 0.22] / **0.19** [0.12, 0.25]	**0.05** [0.02, 0.08] / **0.17** [0.12, 0.23]
`atrust`	0.67 [0.60, 0.74] / 0.42 [0.34, 0.49]	0.36 [0.28, 0.45] / 0.23 [0.16, 0.30]
`ctrl`	0.37 [0.29, 0.45] / 0.56 [0.48, 0.64]	0.33 [0.27, 0.40] / 0.39 [0.32, 0.47]

Table 5: Runtime beliefs on diagnostic items at round 3 (share of samples, 95% CI). Base rates differ slightly between seeds because the base model was probed twice and vLLM sampling is not bit-reproducible.

against `ctrl` range from −0.21 to −0.53, with every interval below zero. `ctrl`, which never trains on a touching task, lowered R1 in one seed (0.61 → 0.37) and not in the other (0.62 → 0.56).

The generator writes fewer touching tasks over time (G1). Touching tasks per round fall in every arm that generates belief topics: in `exec` seed 0 from 60 to 20 (R1) and 128 to 24 (R2) by round 2. Generation narrows over rounds in general (the runtime training sets shrink by a factor of about 4–8 by round 2), and the belief topics narrow with it.

An unchecked key gives a weaker, seed-dependent fix, not entrenchment. In seed 0, `atrust` was the only arm above base on R1 (0.67); in seed 1 it fell to 0.42, below `ctrl`. The generator's stated outputs on touching tasks are a mixture: at round 0 of seed 0, the R1 stated keys are true 22 times, the belief 14 times and something else 24 times. Such a key carries a true signal alongside the belief and noise, unlike `sexec`, whose key is belief-aligned on every checkable task. We therefore do not claim entrenchment for the runtime beliefs.

Filtering is only modestly specific. At round 0 the `agate` gate rejects 63% of touching tasks but also 43% of all others (seed 1: 61% and 41%). The 4B generator often mispredicts its own program's output, and not mainly on its blind spot.

An untargeted belief grows. In the runtime half, R3 (`int("3.0")` parses) rises in every training arm, from 0.14 to 0.19–0.38 at round 3 across both seeds. In the spec half it moves in both directions.

7 Replacement rules

Section 2 argues that training with a correct key only on cases where the belief is wrong can teach a rule other than the truth. We see such a rule at least once for each of the three corrected beliefs that have neighbour probes (R1, the `%` half of R2, S2). Only the S2 rule appears in all four runs, and only S2's is clearly tied to the key: the `%` rule also appears in a control, and the R1 rule appears in only one seed. We do not see a rule for `//`, whose corrective cases are just as one-sided.

7.1 S2: the boundary moves to the next value

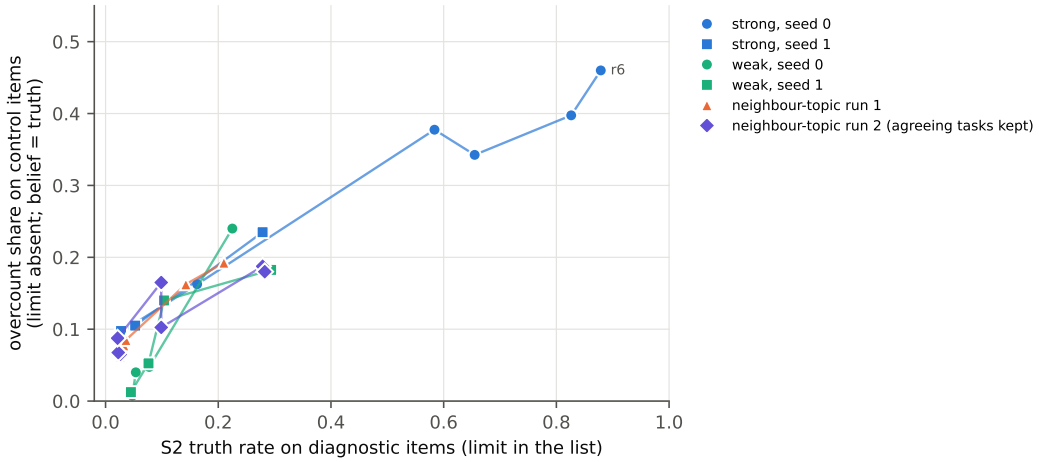


Figure 4: The S2 fix and its side effect, at every `sclean` checkpoint. x: truth rate on diagnostic items. y: overcount share on control items, where the limit is absent and the old belief gives the right answer. The neighbour-topic runs never trained on a limit-absent task (Section 8).

On S2 controls, `sclean`'s errors become overcounts (answers above the true count), and the overcount rises with the fix across seeds and planting strengths (Figure 4; all values in Table 10). As a description, not a test (the points are successive checkpoints of the same runs), truth and overcount correlate at $r = 0.95$ over the 27 distinct checkpoints (slope 0.45); the weak runs start lower (0.01 against 0.07). In strong seed 0 the overcount grows from 0.07 to 0.38 at round 3 and to 0.46 at round 6. The most common errors among the 400 control samples at round 3 are +1 (95 times), +2 (38), −1 (19) and +4 (10); at the planted model they were −1 (61) and +1 (24). The weak runs reach 0.24 and 0.18 at round 3, against 0.01 for the planted model. `sexec` and `sctrl` stay at or below 0.07 at every probed checkpoint. Greedy control accuracy hides the overcount until the fix is large: in the weak runs it is 37–39 of 50 at round 3, against 36 for the planted model; in strong seed 0 it falls from 37 to 23 of 50 by round 6, while diagnostic truth reaches 0.88.

The overcount has a specific shape (Table 6). It occurs whenever some list value exceeds the limit, most often when that value is $n + 1$, and it does not occur on the 12 controls with no value above the limit, which rules out a blanket "+1". In the strong runs most overcounts equal the truth plus the number of copies of the next larger value (a weaker test than it looks: only 5 of the 50 controls have two or more copies of it). Sample reasoning: "the prices at most 16 are 4, 5, 9, 6, 13, 9, and 17". The model increasingly includes a boundary element, the nearest value above the limit; in the weak runs this pattern accounts for only 36–53% of overcounts.

Where the rule could come from. Three sources fit the data, and we cannot separate them. (i) *One-sided corrective cases.* The generator almost never wrote an S2 call with the limit

Checkpoint	**next value** $= n+1$	$n+2$	$\geq n+3$	**overcounts = truth + copies of next value**
planted (strong)	0.14	0.08	0.00	0.81 (of 26)
strong s0, round 3	0.78	0.31	0.26	0.73
strong s0, round 6	0.86	0.35	0.50	0.77
strong s1, round 3	0.53	0.27	0.00	0.64
weak s0, round 3	0.44	0.32	0.11	0.53
weak s1, round 3	0.42	0.18	0.03	0.36

Table 6: Overcount share on S2 control items (limit n absent), by the distance from n to the next larger value in the list (16, 12 and 10 items). On the other 12 items no value exceeds n, and no overcount occurred. Last column: share of overcounting samples whose answer equals the truth plus the number of copies of the next larger value. Exploratory.

absent: 4 of 10,924 S2 tasks across all `sclean` rounds, none in the strong runs. Every correction the model received therefore said "also count the element equal to n". Both the truth and "count the nearest element at or above the limit" fit all of it. (ii) *The selected reasoning.* The pass band selects samples by their final answer only. Of 840 trained S2 rationales whose call and listed values we could parse, 95 (11%) list a value above the limit and still reach the key, and 241 (29%) begin by stating the belief ("'At most 9' means strictly below 9") before reaching the $\leq$ count, for example "the values below 9 are 7, 1, 13, 7, so there are 4". The share is higher early: at round 0, 2 of 7, 2 of 10, 5 of 19 and 9 of 33 parseable trained rationales (strong seeds 0 and 1, weak seeds 0 and 1) list a value above the limit, and of the planted model's probe samples that give the right answer on diagnostic items, 13 of 21 do (weak planting: 11 of 38; strong seed 0 at round 3: 16 of 467). A substantial share of the early corrective signal thus reached the key through the rule, as outcome-only supervision allows [41]. (iii) *A planted presupposition.* Every planting example and almost every generated task had the limit in the list. A model that counts up to the nearest list element at or above n would answer the planted controls correctly under the belief and overcount after the fix.

7.2 R2: add the divisor

On exact `%` with mixed signs (truth 0), `exec` and `agate` answer the divisor b in 29–38% of samples, against 10% at base; `-90 % 9` becomes 9. The mixed-sign cases that touch R2 all have a nonzero remainder, where the truth is the truncated remainder plus b; the learned rule adds b whenever the signs differ. In seed 0 this costs `%` accuracy (`exec` 0.49, `agate` 0.54, against `ctrl` 0.80 and base 0.63). The rule is not confined to the exposure arms: `atrust` shows it in seed 0 (33%), and `ctrl`, which never trains on a touching task, picks some of it up in seed 1 (19%, with `%` accuracy 0.55 against 0.57 for `exec` and `agate`). So corrective cases are not its only source. `//` shows no rule: exact `//` accuracy is 0.95–1.00 for every arm at round 3 (base 0.86).

7.3 R1: toward zero, in one seed

In seed 0, `exec` rounds odd ties toward zero on 22% of samples (base 6%). The even ties that touch R1 all round down in magnitude, and "round `.5` toward zero" fits them. On tie accuracy `exec` therefore equals `ctrl` (0.88 vs 0.88) despite its large diagnostic gain. In seed 1 the rule does not appear (7%), and `exec` beats `ctrl` (0.96 vs 0.78); `agate` shows no rule in either seed.

Round 3, seed 0 / seed 1	base	exec	agate	atrust	ctrl
R1 tie accuracy ((even + odd) / 2)	0.71	0.88 / 0.96	0.91 / 0.93	0.74 / 0.85	0.88 / 0.78
R2 exact-% accuracy (neighbours)	0.63	0.49 / 0.57	0.54 / 0.57	0.40 / 0.74	0.80 / 0.55

Table 7: Runtime measures that include neighbours. The seed-to-seed spread of a single arm (up to 0.34) is often as large as the gaps between arms.

On these measures the runtime picture is far less clean than on diagnostic items (Table 7). The seed-to-seed spread of a single arm is often as large as the gaps between arms, and the % rule is the one effect that appears in the exposure arms in both seeds.

8 Generation: asking for the missing neighbours

If one-sided corrective cases cause the S2 rule, training on neighbours should reduce it. Neighbours of S2 (*anchors*) are tasks with the limit absent from the list, where belief and truth agree. We ran two `sclean` runs (strong planting, seed 0) in which half of the S2 topic share went to two neighbour topics: "counting values that are at most a limit, where the limit value itself does not appear in the list" and "counting items whose price is at most a budget that no item costs exactly". In the first run (4 rounds) the topics change only generation; anchors solved 8 of 8 fall outside the pass band, as usual. In the second run (6 rounds) anchors are kept for training even when solved 8 of 8, with up to 2 samples each. We identified anchors as S2 tasks whose decorrelated key equals the generator's executed key: with the limit absent, $<$ and $\leq$ agree, so we expected agreement to mark exactly those tasks.

Inspecting the inputs after the runs showed that this proxy failed (Table 8). Of 2,068 neighbour-topic tasks over nine generation rounds, 2,061 had a checkable call, and in none was the limit absent. The generator kept its call template and moved the request into the specification, usually as a precondition that the call violated ("where `n` does not appear in `xs`", "assuming no value is exactly `n`") and sometimes as a redefinition: "the cap value itself does not count", which states the $<$ reading as the task's definition (the unplanted model writes such clauses too; see below). With a violated precondition the specification no longer fixes the answer; with a redefinition it fixes the belief's answer. On these tasks the decorrelated reference computed the $<$ count, with the limit in the list. Its key then agreed with the generator's, the proxy marked the task as an anchor, and the solver, which holds the belief, answered it correctly 8 times of 8. We labelled the specifications with patterns and checked a random sample of each label by hand.

Neither run therefore tests the remedy. The first trained 2–9 of these tasks per round; the second trained 51–83 per round that were solved 8 of 8 (54–88 including those in the band), each with the limit in the list and the $<$ count as key, and no true anchor. Both runs fixed S2 more slowly than either plain strong run (truth 0.14 and 0.10 at round 3, against 0.58 and 0.28). Two causes are plausible. With half of the S2 topics given to neighbour topics, fewer corrective S2 tasks were admitted (27 by round 3 in the second run, against 123 in strong seed 0). And in the second run, the 54–88 belief-keyed tasks trained per round are a G2 failure by our own account: they reward the $<$ count with the limit in the list. The plain runs vary too widely (0.10 truth after 112 corrective tasks in weak seed 1, 0.58 after 123 in strong seed 0) to separate these. The overcount stayed on the common curve (Figure 4), as expected when no limit-absent task is trained.

The unplanted generator. To separate the planted habit from generator behaviour in general, we ran one generation round of the unplanted base model with the same topics and topic draws, without training. Of its checkable S2 calls, 79 of 150 on the plain topics and 46 of 170 on

Round	neighbour-topic tasks	limit absent	key agrees (< count)	precondition / redefinition	kept for training
0	233	0	57	41 / 14	52
1	222	0	76	55 / 21	70
2	249	0	91	68 / 23	83
3	215	0	81	60 / 21	65
4	229	0	93	70 / 21	83
5	234	0	61	47 / 14	51

Table 8: Audit of the neighbour topics in the second run. The precondition / redefinition column counts only the key-agreeing tasks. The first run (rounds 1–3: 686 neighbour-topic tasks, 75–102 key-agreeing per round) behaves the same; round 0 is shared. A few key-agreeing tasks fit neither specification pattern, and 7 calls could not be parsed.

the neighbour topics have the limit absent, against 0 for the planted generator. Never omitting the limit is therefore a planting artifact. Rewriting the request is not: of the unplanted model's 124 neighbour-topic tasks with the limit in the list, 31 add a precondition and 36 a redefinition ("returns the number of values ... that are at most `limit`, but not `limit` itself"), against 73 and 45 of 232 for the planted generator (counting all its neighbour-topic tasks, not only the key-agreeing ones of Table 8). Neither generator wrote such a clause on the plain S2 topics (0 of 223 and 0 of 228): the clauses come from the neighbour topics' wording, not from the belief.

The runs therefore show two things. First, a generator's habits can limit what the loop sees even when the topic list asks for the missing case; here the habit was planted. Second, decorrelating the key does not decorrelate the task. Both generators turned a request about the inputs into a clause in the specification, and when that clause excludes the limit, a key that follows the specification agrees with the belief. A loop cannot rely on topic instructions to produce the inputs it needs.

9 Discussion

What a loop can correct directly is set by its generator, its key and its pass band. Execution is a perfect key for what code does. It is no key at all for what a specification means when the same model writes the implementation and the misreading has reached the code, as it had on every task we audited (1,302 of 1,302 S2 tasks in three rounds). For each error of interest one can audit, task by task, how many tasks touch it, how many of those the key contradicts, and how many of those the pass band admits. These counts are diagnostic rather than predictive: in our runs, round-0 counts did not predict how fast a belief moved. Our earlier work on verifier-gated cascades [32] measured an analogous quantity from the verifier's side, the share of wrong answers the verifier accepts; here it is the share of touching tasks with $\kappa(t) = b(t)$.

Decorrelation is an old idea. Using a second, independently written implementation as the oracle is differential testing [29]. Independently written versions fail together more often than independence predicts [21], co-training needs views whose errors are conditionally independent [4], and pseudo-labelling without such independence confirms its own mistakes [2]. Our decorrelated key is a two-version scheme in which the versions differ, by construction, in the belief under study, and even so it shared the natural S1 error on 5–34% of tasks. It is also a privileged one: the key writer is the checkpoint from before planting, known not to hold the belief, so `sclean` is a teacher arm and an upper bound. In practice the errors are unknown, and the question is whether a key writer chosen without that knowledge (another model family, or a weaker model [5]) is decorrelated enough. Models from different families make correlated errors [12, 20], so this has to be measured

per error; we did not test it.

The key implements the specification. The neighbour-topic audit shows a limit of key decorrelation: the generator still writes the question. When the specification excludes the limit, whether because of a planted belief or a reading of the topic, every key that follows the specification agrees with the belief. Loops in which one model writes both the task statement and the reference, such as the Code-as-Task format of Self-Challenging Agents [49], are exposed to this route.

Corrections and neighbours. For each corrected belief, the cases that carried the correction lay on one side of it, and the trained reasoning was selected by answer alone. After the S2 fix the same model answers diagnostic items correctly and overcounts on controls; the `%` fix and, in one seed, the R1 fix came with similar shifts. A diagnostic-only metric sees only the first half. In our data it would have reported a near-perfect S2 fix at round 6 (truth 0.88) while greedy control accuracy fell from 37 to 23 of 50. This is underspecification [8] inside a training loop: the corrective data do not pin down the rule.

Practical points.

1. *Decorrelate the key, and audit it on the errors you can find.* A model that does not share an error (here, the model before planting) writes a better key for that error than the model itself, even though its implementations sometimes crash. Without knowing the errors in advance, the audit has to come from probes like ours.

2. *Check generated inputs directly.* Whether a task has a requested property should be read off its inputs, not inferred from its keys. Our key-agreement proxy for anchors was wrong on every checkable task it selected.

3. *Watch the pass band.* A belief held near 100% gets almost no direct gradient under rejection sampling, whatever the key. Some exploration (more samples, hints or a teacher demonstration) is likely needed to start the fix; we did not test it.

4. *Score fixes on neighbours.* Report the error distribution on items where the old belief was right, not only the belief rate on items where it was wrong, and inspect the reasoning that training selected.

10 Related work

Self-play with executors and provers. Absolute Zero [47] proposes and solves code-reasoning tasks, with an executor that validates tasks and checks answers; its deduction task is our output prediction. Self-Challenging Agents [49] write tasks with verification functions and tests. STP [9] alternates conjecturing and proving in Lean. These systems guarantee that the key is consistent with the program or the proof; our results concern errors the program itself encodes. Sol-Ver [27] and CURE [42] train models to write both code and unit tests, and CodeT [6] ranks generated code by agreement with model-written tests, which assumes that errors in code and tests are not correlated.

Generated oracles share the generator's errors. The closest prior work is in software testing. LLM-generated test oracles tend to capture what the code does rather than what it should do [22], and when tests are written after the code, its faults are copied into them: such tests catch 14% of faults, against 25% for independently written tests [23]. `sexec` is a training-time version of this effect, in which the copied fault becomes the training signal. In self-play against reference-free LLM judges, answers wrongly accepted by one judge also fool judges from other families [48].

Label-free rewards and their limits. TTRL [50] and R-Zero [17] use majority votes as pseudo-labels, and majority voting can extinguish correct minority answers [26]. He et al. [14] show that intrinsic rewards sharpen the initial distribution, and Huang et al. [16] analyse self-improvement as sharpening toward the model's own high-confidence outputs. Self-rewarding RL can collapse through reward hacking [33], and self-improvement is bounded by the gap between generating and verifying [39]. Qwen models can improve under spurious rewards [34], a reminder that generic drift, like our runtime control's, must be separated from the effect of the key. Recursive training on generated data narrows the distribution [36], as our generator's output did over rounds.

Self-training by filtering. STaR [46], ReST$^{\text{EM}}$ [37] and expert iteration [1, 31] fine-tune on samples that pass a binary check. Group-relative RL [35] gets no signal from groups that all pass or all fail, and DAPO [44] filters them out. Learnability rewards in self-play proposers [17, 47] favour tasks of intermediate difficulty. All of these share our pass band, and with it the block on beliefs held near 100%. Outcome-only filtering admits correct answers reached by wrong reasoning [25, 41], which our rationale audit finds in the S2 training data.

Oracles, correlated errors and shortcuts. The test-oracle problem [3] asks where correct expected outputs come from. N-version programming [21], differential testing [29] and co-training [4] depend on independence between versions or views; pseudo-labelling [2] and studies of correlated LLM errors [12, 20] show what happens without it. Shortcut learning [11] and underspecification [8] describe models that adopt a rule the training data do not rule out, which is how we read the replacement rules. LLMs struggle to correct their own reasoning without external feedback [18, 40]; our results add that external feedback corrects only what it can contradict and what the filter lets through.

Planted behaviour as a model organism. Backdoored behaviours that persist through safety training [19] and false facts implanted by fine-tuning on synthetic documents [38] use planting to study what training does to a known behaviour. We plant misreadings of task specifications and ask whether a self-improvement loop removes them.

11 Limitations

- **One model, a small loop.** We use Qwen3-4B-Base and expert iteration rather than RL, with 3,000 task requests per round (2,178–2,705 valid at round 0) and 3–6 rounds. Generation narrows over rounds and the runtime training sets shrink by a factor of about 4–8 by round 2 (spec arms 1.4–2.7); spec rounds train on 76–466 examples. A loop this small may behave differently from production-scale self-play.

- **Two seeds, and some single runs.** Several seed-0 findings did not replicate: the runtime control's drift, `atrust` entrenchment and the R1 replacement rule. The main spec contrasts hold in all four runs, but effect sizes vary widely, and our confidence intervals do not include run-to-run variation. The six-round results, the S1 break-out and both neighbour-topic runs are single runs.

- **Planted beliefs are template-like.** S1 was planted with one implementation, `sorted(xs)[k]`; every S2 planting example drew the limit from the list, and the QA examples stated the belief in words. The generator kept the S2 template, so the one-sidedness of S2 corrective data is partly a planting artifact. The selection argument of Section 2 does not depend on it, but the missing anchors do. The runtime beliefs are natural, which partly answers this for the runtime half.

- **Narrow probes.** All S2 diagnostic items, like almost all generated S2 tasks, put the limit in the list, so the truth rate measures one case of the belief. S1 has no neighbour probe.
- **One decorrelated key.** `sclean` uses the unplanted base model, which shares the S1 misreading on 5–34% of S1 tasks. We did not test a key writer from another model family or a weaker model.
- **One unplanted generation round.** The comparison with the unplanted generator is a single untrained round.
- **The remedy was not tested.** The neighbour-topic runs contained no limit-absent tasks (Section 8). Whether true neighbours remove the replacement rule remains open; a direct test would inject script-generated anchors or filter generated tasks on their inputs.
- **No capability measure.** We did not measure general coding ability before and after planting or training.
- **Not pre-registered.** A design document was written first, and the build departed from it: output prediction instead of code writing; expert iteration instead of GRPO; 8 samples and a 1–7-of-8 band instead of 16 and a [0.125, 0.625] band; 3 rounds (6 in extensions) instead of 4; two seeds instead of three; no no-loop control or multi-check key arm; no capability gate, and weak planting deliberately below the planned 0.8 planting threshold (S2 at 0.75); code-writing probes, S3, the planned hypothesis tests H4–H7 and the Holm-corrected mixed-effects analysis dropped. The lab notebook records each analysis in the order it was done, including conclusions we later revised, such as an earlier reading of the neighbour-topic runs as evidence against the remedy.
- **Incidents.** Two runs overlapped on the GPU for about 1.5 hours, causing out-of-memory kills; the affected weak-planting arms were discarded and re-run. One unexplained hard reset of the machine left three empty files, which were removed and regenerated.

12 Conclusion

Self-play with code execution grounds a model in what code does. It does not ground it in what a task means when the model's misreading has reached its code, and it cannot directly correct a belief for which the model almost never produces a correct sample. When it did correct a belief in our runs, it did so from cases that all lay on one side of the belief, with reasoning selected by its final answer, and the model can learn a shifted rule instead of the truth. Whether a loop can correct a given error depends on its generator, its key and its pass band. Each can be audited task by task, though in our runs early counts did not predict how fast a belief moved; what the correction installs has to be checked on neighbours.

References

[1] Thomas Anthony, Zheng Tian, and David Barber. Thinking Fast and Slow with Deep Learning and Tree Search. In *Advances in Neural Information Processing Systems 30 (NIPS 2017)*, pages 5360–5370, 2017.

[2] Eric Arazo, Diego Ortego, Paul Albert, Noel E. O'Connor, and Kevin McGuinness. Pseudo-Labeling and Confirmation Bias in Deep Semi-Supervised Learning. In *2020 International*

Joint Conference on Neural Networks (IJCNN). IEEE, 2020. doi: 10.1109/IJCNN48605.2020.9207304.

[3] Earl T. Barr, Mark Harman, Phil McMinn, Muzammil Shahbaz, and Shin Yoo. The Oracle Problem in Software Testing: A Survey. *IEEE Transactions on Software Engineering*, 41(5): 507–525, 2015. doi: 10.1109/TSE.2014.2372785.

[4] Avrim Blum and Tom Mitchell. Combining Labeled and Unlabeled Data with Co-Training. In *Proceedings of the Eleventh Annual Conference on Computational Learning Theory (COLT 1998)*, pages 92–100. ACM, 1998. doi: 10.1145/279943.279962.

[5] Collin Burns, Pavel Izmailov, Jan Hendrik Kirchner, Bowen Baker, Leo Gao, Leopold Aschenbrenner, Yining Chen, Adrien Ecoffet, Manas Joglekar, Jan Leike, Ilya Sutskever, and Jeff Wu. Weak-to-Strong Generalization: Eliciting Strong Capabilities With Weak Supervision. In *Proceedings of the 41st International Conference on Machine Learning (ICML 2024)*, volume 235 of *Proceedings of Machine Learning Research*, pages 4971–5012, 2024.

[6] Bei Chen, Fengji Zhang, Anh Nguyen, Daoguang Zan, Zeqi Lin, Jian-Guang Lou, and Weizhu Chen. CodeT: Code Generation with Generated Tests. In *International Conference on Learning Representations (ICLR 2023)*, 2023.

[7] Roi Cohen, Eden Biran, Ori Yoran, Amir Globerson, and Mor Geva. Evaluating the Ripple Effects of Knowledge Editing in Language Models. *Transactions of the Association for Computational Linguistics*, 2024.

[8] Alexander D'Amour, Katherine Heller, Dan Moldovan, Ben Adlam, Babak Alipanahi, Alex Beutel, Christina Chen, Jonathan Deaton, Jacob Eisenstein, Matthew D. Hoffman, Farhad Hormozdiari, Neil Houlsby, Shaobo Hou, Ghassen Jerfel, Alan Karthikesalingam, Mario Lucic, Yian Ma, Cory McLean, Diana Mincu, Akinori Mitani, Andrea Montanari, Zachary Nado, Vivek Natarajan, Christopher Nielson, Thomas F. Osborne, Rajiv Raman, Kim Ramasamy, Rory Sayres, Jessica Schrouff, Martin Seneviratne, Shannon Sequeira, Harini Suresh, Victor Veitch, Max Vladymyrov, Xuezhi Wang, Kellie Webster, Steve Yadlowsky, Taedong Yun, Xiaohua Zhai, and D. Sculley. Underspecification Presents Challenges for Credibility in Modern Machine Learning. *Journal of Machine Learning Research*, 2022.

[9] Kefan Dong and Tengyu Ma. STP: Self-play LLM Theorem Provers with Iterative Conjecturing and Proving. In *Proceedings of the 42nd International Conference on Machine Learning (ICML 2025)*, volume 267 of *Proceedings of Machine Learning Research*, pages 14114–14136, 2025.

[10] Matt Gardner, Yoav Artzi, Victoria Basmova, Jonathan Berant, Ben Bogin, Sihao Chen, Pradeep Dasigi, Dheeru Dua, Yanai Elazar, Ananth Gottumukkala, Nitish Gupta, Hanna Hajishirzi, Gabriel Ilharco, Daniel Khashabi, Kevin Lin, Jiangming Liu, Nelson F. Liu, Phoebe Mulcaire, Qiang Ning, Sameer Singh, Noah A. Smith, Sanjay Subramanian, Reut Tsarfaty, Eric Wallace, Ally Zhang, and Ben Zhou. Evaluating Models' Local Decision Boundaries via Contrast Sets. In *Findings of the Association for Computational Linguistics: EMNLP 2020*, 2020.

[11] Robert Geirhos, Jörn-Henrik Jacobsen, Claudio Michaelis, Richard Zemel, Wieland Brendel, Matthias Bethge, and Felix A. Wichmann. Shortcut Learning in Deep Neural Networks. *Nature Machine Intelligence*, 2(11):665–673, 2020. doi: 10.1038/s42256-020-00257-z.

[12] Shashwat Goel, Joschka Strüber, Ilze Amanda Auzina, Karuna K. Chandra, Ponnurangam Kumaraguru, Douwe Kiela, Ameya Prabhu, Matthias Bethge, and Jonas Geiping. Great Models Think Alike and this Undermines AI Oversight. In *Proceedings of the 42nd International Conference on Machine Learning (ICML 2025)*, volume 267 of *Proceedings of Machine Learning Research*, pages 19621–19678, 2025.

[13] Alex Gu, Baptiste Rozière, Hugh Leather, Armando Solar-Lezama, Gabriel Synnaeve, and Sida I. Wang. CRUXEval: A Benchmark for Code Reasoning, Understanding and Execution. In *Proceedings of the 41st International Conference on Machine Learning (ICML 2024)*, volume 235 of *Proceedings of Machine Learning Research*, pages 16568–16621, 2024.

[14] Bingxiang He, Yuxin Zuo, Zeyuan Liu, Shangziqi Zhao, Zixuan Fu, Junlin Yang, Cheng Qian, Kaiyan Zhang, Yuchen Fan, Ganqu Cui, Xiusi Chen, Youbang Sun, Xingtai Lv, Xuekai Zhu, Li Sheng, Ran Li, Huan ang Gao, Yuchen Zhang, Bowen Zhou, Zhiyuan Liu, and Ning Ding. How Far Can Unsupervised RLVR Scale LLM Training? In *International Conference on Learning Representations (ICLR 2026)*, 2026.

[15] Edward J. Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. LoRA: Low-Rank Adaptation of Large Language Models. In *International Conference on Learning Representations (ICLR 2022)*, 2022.

[16] Audrey Huang, Adam Block, Dylan J. Foster, Dhruv Rohatgi, Cyril Zhang, Max Simchowitz, Jordan T. Ash, and Akshay Krishnamurthy. Self-Improvement in Language Models: The Sharpening Mechanism. In *International Conference on Learning Representations (ICLR 2025)*, 2025.

[17] Chengsong Huang, Wenhao Yu, Xiaoyang Wang, Hongming Zhang, Zongxia Li, Ruosen Li, Jiaxin Huang, Haitao Mi, and Dong Yu. R-Zero: Self-Evolving Reasoning LLM from Zero Data. In *International Conference on Learning Representations (ICLR 2026)*, 2026.

[18] Jie Huang, Xinyun Chen, Swaroop Mishra, Huaixiu Steven Zheng, Adams Wei Yu, Xinying Song, and Denny Zhou. Large Language Models Cannot Self-Correct Reasoning Yet. In *International Conference on Learning Representations (ICLR 2024)*, 2024.

[19] Evan Hubinger, Carson Denison, Jesse Mu, Mike Lambert, Meg Tong, Monte MacDiarmid, Tamera Lanham, Daniel M. Ziegler, Tim Maxwell, Newton Cheng, Adam Jermyn, Amanda Askell, Ansh Radhakrishnan, Cem Anil, David Duvenaud, Deep Ganguli, Fazl Barez, Jack Clark, Kamal Ndousse, Kshitij Sachan, Michael Sellitto, Mrinank Sharma, Nova DasSarma, Roger Grosse, Shauna Kravec, Yuntao Bai, Zachary Witten, Marina Favaro, Jan Brauner, Holden Karnofsky, Paul Christiano, Samuel R. Bowman, Logan Graham, Jared Kaplan, Sören Mindermann, Ryan Greenblatt, Buck Shlegeris, Nicholas Schiefer, and Ethan Perez. Sleeper Agents: Training Deceptive LLMs that Persist Through Safety Training. *arXiv preprint arXiv:2401.05566*, 2024.

[20] Elliot Myunghoon Kim, Avi Garg, Kenny Peng, and Nikhil Garg. Correlated Errors in Large Language Models. In *Proceedings of the 42nd International Conference on Machine Learning (ICML 2025)*, volume 267 of *Proceedings of Machine Learning Research*, pages 30038–30066, 2025.

[21] John C. Knight and Nancy G. Leveson. An Experimental Evaluation of the Assumption of Independence in Multiversion Programming. *IEEE Transactions on Software Engineering*, SE-12(1):96–109, 1986.

[22] Michael Konstantinou, Renzo Degiovanni, and Mike Papadakis. Can we find bugs using LLM-generated oracles? *arXiv preprint arXiv:2410.21136*, 2024.

[23] Michael Konstantinou, Florian Tambon, and Mike Papadakis. On the risk of coding before testing: An empirical study on LLM-based test generation workflow. *arXiv preprint arXiv:2607.05139*, 2026.

[24] 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 ACM Symposium on Operating Systems Principles (SOSP 2023)*. ACM, 2023. doi: 10.1145/3600006.3613165.

[25] Hunter Lightman, Vineet Kosaraju, Yura Burda, Harri Edwards, Bowen Baker, Teddy Lee, Jan Leike, John Schulman, Ilya Sutskever, and Karl Cobbe. Let's Verify Step by Step. *arXiv preprint arXiv:2305.20050*, 2023.

[26] Hongxiang Lin, Zhirui Kuai, Erpeng Xue, and Lei Wang. Detecting and Mitigating the Correct-Answer Extinction Window in Test-Time Reinforcement Learning with Majority Voting. *arXiv preprint arXiv:2605.19444*, 2026.

[27] Zi Lin, Sheng Shen, Ilia Kulikov, Jingbo Shang, Jason Weston, and Yixin Nie. Learning to Solve and Verify: A Self-Play Framework for Code and Test Generation. *arXiv preprint arXiv:2502.14948*, 2025.

[28] R. Thomas McCoy, Ellie Pavlick, and Tal Linzen. Right for the Wrong Reasons: Diagnosing Syntactic Heuristics in Natural Language Inference. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, 2019.

[29] William M. McKeeman. Differential Testing for Software. *Digital Technical Journal*, 10(1): 100–107, 1998.

[30] Kevin Meng, David Bau, Alex Andonian, and Yonatan Belinkov. Locating and Editing Factual Associations in GPT. In *Advances in Neural Information Processing Systems*, 2022.

[31] Stanislas Polu, Jesse Michael Han, Kunhao Zheng, Mantas Baksys, Igor Babuschkin, and Ilya Sutskever. Formal Mathematics Statement Curriculum Learning. In *International Conference on Learning Representations (ICLR 2023)*, 2023.

[32] Dushyant Rajput, Nirdesh Chauhan, and Siddharth Kosaraju. Cheap Verifiers, Large Blind Spots: Measuring the Reliability Cost of Cost-Saving Cascades. Technical report, AltSlate Labs LLP, 2026.

[33] Sheikh Shafayat, Fahim Tajwar, Ruslan Salakhutdinov, Jeff Schneider, and Andrea Zanette. Can Large Reasoning Models Self-Train? *arXiv preprint arXiv:2505.21444*, 2025.

[34] Rulin Shao, Shuyue Stella Li, Rui Xin, Scott Geng, Yiping Wang, Sewoong Oh, Simon Shaolei Du, Nathan Lambert, Sewon Min, Ranjay Krishna, Yulia Tsvetkov, Hannaneh Hajishirzi, Pang Wei Koh, and Luke Zettlemoyer. Spurious Rewards: Rethinking Training Signals in RLVR. In *Proceedings of the 43rd International Conference on Machine Learning (ICML 2026)*, volume 306 of *Proceedings of Machine Learning Research*, pages 110167–110207, 2026.

[35] Zhihong Shao, Peiyi Wang, Qihao Zhu, Runxin Xu, Junxiao Song, Xiao Bi, Haowei Zhang, Mingchuan Zhang, Y. K. Li, Y. Wu, and Daya Guo. DeepSeekMath: Pushing the Limits of Mathematical Reasoning in Open Language Models. *arXiv preprint arXiv:2402.03300*, 2024.

[36] Ilia Shumailov, Zakhar Shumaylov, Yiren Zhao, Nicolas Papernot, Ross Anderson, and Yarin Gal. AI Models Collapse When Trained on Recursively Generated Data. *Nature*, 631:755–759, 2024. doi: 10.1038/s41586-024-07566-y.

[37] Avi Singh, John D. Co-Reyes, Rishabh Agarwal, Ankesh Anand, Piyush Patil, Xavier Garcia, Peter J. Liu, James Harrison, Jaehoon Lee, Kelvin Xu, Aaron Parisi, Abhishek Kumar, Alex Alemi, Alex Rizkowsky, Azade Nova, Ben Adlam, Bernd Bohnet, Gamaleldin Elsayed, Hanie Sedghi, Igor Mordatch, Isabelle Simpson, Izzeddin Gur, Jasper Snoek, Jeffrey Pennington, Jiri Hron, Kathleen Kenealy, Kevin Swersky, Kshiteej Mahajan, Laura Culp, Lechao Xiao, Maxwell L. Bileschi, Noah Constant, Roman Novak, Rosanne Liu, Tris Warkentin, Yundi Qian, Yamini Bansal, Ethan Dyer, Behnam Neyshabur, Jascha Sohl-Dickstein, and Noah Fiedel. Beyond Human Data: Scaling Self-Training for Problem-Solving with Language Models. *Transactions on Machine Learning Research*, 2024.

[38] Stewart Slocum, Julian Minder, Clément Dumas, Henry Sleight, Ryan Greenblatt, Samuel Marks, and Rowan Wang. Believe It or Not: How Deeply do LLMs Believe Implanted Facts? *arXiv preprint arXiv:2510.17941*, 2025.

[39] Yuda Song, Hanlin Zhang, Carson Eisenach, Sham Kakade, Dean Foster, and Udaya Ghai. Mind the Gap: Examining the Self-Improvement Capabilities of Large Language Models. In *International Conference on Learning Representations (ICLR 2025)*, 2025.

[40] Kaya Stechly, Karthik Valmeekam, and Subbarao Kambhampati. On the Self-Verification Limitations of Large Language Models on Reasoning and Planning Tasks. In *International Conference on Learning Representations (ICLR 2025)*, 2025.

[41] Jonathan Uesato, Nate Kushman, Ramana Kumar, Francis Song, Noah Siegel, Lisa Wang, Antonia Creswell, Geoffrey Irving, and Irina Higgins. Solving math word problems with process- and outcome-based feedback. *arXiv preprint arXiv:2211.14275*, 2022.

[42] Yinjie Wang, Ling Yang, Ye Tian, Ke Shen, and Mengdi Wang. Co-Evolving LLM Coder and Unit Tester via Reinforcement Learning. In *Advances in Neural Information Processing Systems*, 2025.

[43] An Yang, Anfeng Li, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chang Gao, Chengen Huang, Chenxu Lv, Chujie Zheng, Dayiheng Liu, Fan Zhou, Fei Huang, Feng Hu, Hao Ge, Haoran Wei, Huan Lin, Jialong Tang, Jian Yang, Jianhong Tu, Jianwei Zhang, Jianxin Yang, Jiaxi Yang, Jing Zhou, Jingren Zhou, Junyang Lin, Kai Dang, Keqin Bao, Kexin Yang, Le Yu, Lianghao Deng, Mei Li, Mingfeng Xue, Mingze Li, Pei Zhang, Peng Wang, Qin Zhu, Rui Men, Ruize Gao, Shixuan Liu, Shuang Luo, Tianhao Li, Tianyi Tang, Wenbiao Yin, Xingzhang Ren, Xinyu Wang, Xinyu Zhang, Xuancheng Ren, Yang Fan, Yang Su, Yichang Zhang, Yinger Zhang, Yu Wan, Yuqiong Liu, Zekun Wang, Zeyu Cui, Zhenru Zhang, Zhipeng Zhou, and Zihan Qiu. Qwen3 Technical Report. *arXiv preprint arXiv:2505.09388*, 2025.

[44] Qiying Yu, Zheng Zhang, Ruofei Zhu, Yufeng Yuan, Xiaochen Zuo, Yu Yue, Weinan Dai, Tiantian Fan, Gaohong Liu, Lingjun Liu, Xin Liu, Haibin Lin, Zhiqi Lin, Bole Ma, Guangming Sheng, Yuxuan Tong, Chi Zhang, Mofan Zhang, Wang Zhang, Hang Zhu, Jinhua Zhu, Jiaze

Chen, Jiangjie Chen, Chengyi Wang, Hongli Yu, Yuxuan Song, Xiangpeng Wei, Hao Zhou, Jingjing Liu, Wei-Ying Ma, Ya-Qin Zhang, Lin Yan, Mu Qiao, Yonghui Wu, and Mingxuan Wang. DAPO: An Open-Source LLM Reinforcement Learning System at Scale. In *Advances in Neural Information Processing Systems 38 (NeurIPS 2025)*, 2025.

[45] Yang Yue, Zhiqi Chen, Rui Lu, Andrew Zhao, Zhaokai Wang, Yang Yue, Shiji Song, and Gao Huang. Does Reinforcement Learning Really Incentivize Reasoning Capacity in LLMs Beyond the Base Model? In *Advances in Neural Information Processing Systems*, 2025.

[46] Eric Zelikman, Yuhuai Wu, Jesse Mu, and Noah D. Goodman. STaR: Bootstrapping Reasoning With Reasoning. In *Advances in Neural Information Processing Systems 35 (NeurIPS 2022)*, 2022.

[47] Andrew Zhao, Yiran Wu, Yang Yue, Tong Wu, Quentin Xu, Yang Yue, Matthieu Lin, Shenzhi Wang, Qingyun Wu, Zilong Zheng, and Gao Huang. Absolute Zero: Reinforced Self-play Reasoning with Zero Data. In *Advances in Neural Information Processing Systems 38 (NeurIPS 2025)*, 2025.

[48] Chenyu Zhou. More Convincing, Not More Correct: Self-Play Reward Hacking of Reference-Free LLM Judges. *arXiv preprint arXiv:2607.05904*, 2026.

[49] Yifei Zhou, Sergey Levine, Jason Weston, Xian Li, and Sainbayar Sukhbaatar. Self-Challenging Language Model Agents. In *Advances in Neural Information Processing Systems 38 (NeurIPS 2025)*, 2025.

[50] Yuxin Zuo, Kaiyan Zhang, Li Sheng, Shang Qu, Ganqu Cui, Xuekai Zhu, Haozhan Li, Yuchen Zhang, Xinwei Long, Ermo Hua, Biqing Qi, Youbang Sun, Zhiyuan Ma, Lifan Yuan, Ning Ding, and Bowen Zhou. TTRL: Test-Time Reinforcement Learning. In *Advances in Neural Information Processing Systems 38 (NeurIPS 2025)*, 2025.

Appendix A: Task fate at round 0

Run / arm	Topic	tasks	key = truth / belief / other	trained, true	trained, other
runtime s0, `atrust`	R1	60	22 / 14 / 24	17	17
runtime s0, `atrust`	R2	128	47 / 5 / 76	34	36
runtime s0, `agate`	R1 / R2	60 / 128	kept 22 / 47 (rest rejected)	17 / 34	0
runtime s0, `exec`	R1 / R2	60 / 128	all true	33 / 92	0
spec strong, `sexec`	S2	461	$<$ count on 431 of 431 checkable	1	21
spec strong, `sclean`	S2	461	$\leq$ count on 88%; the rest exceptions	9	0
spec weak, `sexec`	S2	461	$<$ count on 430 of 430 checkable	2	88
spec weak, `sclean`	S2	461	$\leq$ count on 94%; the rest exceptions	24	0

Table 9: Fate of belief-topic tasks at round 0. Runtime counts are tasks that touch the belief. For the spec arms, "true" in the training columns means equal to the decorrelated key; the key audits are direct counts on calls of the form `f(list, number)`.

Appendix B: S2 trajectories

`sclean` run	r0	r1	r2	r3	r4	r5	r6
strong, seed 0	0.03/0.07	0.03/0.10	0.16/0.16	0.58/0.38	0.66/0.34	0.83/0.40	0.88/0.46
strong, seed 1	0.03/0.06	0.03/0.10	0.05/0.11	0.28/0.24			
weak, seed 0	0.05/0.01	0.05/0.04	0.08/0.05	0.23/0.24			
weak, seed 1	0.05/0.01	0.08/0.05	0.10/0.14	0.29/0.18			
neighbour topics, 1st	0.03/0.07	0.03/0.08	0.04/0.09	0.14/0.16	0.21/0.19		
neighbour topics, 2nd	0.03/0.07	0.02/0.07	0.02/0.09	0.10/0.17	0.10/0.10	0.28/0.19	0.28/0.18

Table 10: S2 truth on diagnostic items / overcount on control items, for every `sclean` checkpoint (the data of Figure 4). r0 is the planted model; the neighbour-topic runs share strong seed 0's round 0.

Appendix C: Paired contrasts at round 3

Appendix D: Reproducibility

Code, prompts, probe sets, per-checkpoint metrics and the lab notebook are available from the authors on request. Probe metrics and task fates come from one metrics file produced by a collection script. The S2 boundary, limit-absent and neighbour-topic audits, the rationale audit and the paired tests come from separate scripts that read the generated tasks, training sets and probe outputs. The round-0 key audits are recorded in the lab notebook. A round takes about 20 minutes on one GB10 (generation of 3,000 tasks and 8 solves for each valid one, then SFT); the study used 76

Contrast (100 paired items)	strong s0	strong s1	weak s0	weak s1
S2 truth, `sclean` − `sexec`	+0.58	+0.27	+0.20	+0.26
S2 truth, `sclean` − `sctrl`	+0.56	+0.26	+0.18	+0.26
S2 belief, `sexec` − planted	−0.10	−0.10	+0.08	+0.13

Table 11: Paired item-level bootstrap (4,000 resamples). Every interval excludes zero; the closest to zero is weak seed 0, `sexec` − planted, [+0.01, +0.15]. No multiplicity correction. Runtime contrasts against `ctrl` (R1 and R2 belief, two seeds): `exec` −0.23 to −0.53 and `agate` −0.21 to −0.38, all intervals below zero; `atrust` +0.30, +0.03, −0.15 and −0.16, with the second interval including zero.

trained rounds. Runs were executed one at a time, because two vLLM or training jobs at once on the GPU's unified memory caused out-of-memory kills (Section 11).