# BENCHMARKS DON'T TALK BACK: LEARNING THROUGH SELF-DISAGREEMENT

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

## ABSTRACT

A benchmark usually measures a language model alone with a question, and deployment rarely looks like that: answers get contested, and models give up correct ones under counter-reasoning that is itself wrong. We show that what fails there is not a shortfall of accuracy but a capability in its own right — knowing when to hold a position and when to yield to a better one — which we call *dialectical competence*. It stands on first-answer accuracy without following from it: a model can be made materially better at answering and no better at this, and training for accuracy alone can drift against it. Nor can it be bought at inference, by self-reflection or by sampling more. Taking a cue from the dialectical tradition in philosophy, in which productive opposition arises within a position rather than being imposed on it from outside, we hypothesize that this capability can be trained from opposition the model supplies itself. It states a position, generates the strongest case against that position — its *antithesis* — rules on whether the case stands, and is trained on that ruling, with no second model at training or at deployment. We call this learning through *Self-Disagreement*: reinforcement learning from verifiable rewards throughout, where the same check that scores an answer against its reference scores the ruling too, so no reward model, no preference data and no human annotation enter anywhere. The capability is trained rather than elicited: weak in the base model, it arrives with the reward, is not paid for out of first-answer accuracy, and holds against challengers the model never trained against, including ones stronger than itself. The opposition a model needs in order to improve, it turns out, need not come from outside it.

## 1 INTRODUCTION

**Benchmarks do not talk back:** Language models have become markedly better at solving problems across successive open frontier families (Grattafiori et al., 2024; Yang et al., 2025), with reinforcement learning against answers a program can check providing one route to that improvement (DeepSeek-AI et al., 2025; Shao et al., 2024; Yu et al., 2025; Liu et al., 2025). Benchmark suites register that progress across general knowledge (Hendrycks et al., 2020), graduate-level science (Rein et al., 2023), competition mathematics (Hendrycks et al., 2021; Gao et al., 2024) and code (Jain et al., 2024). Yet each ultimately reduces evaluation to a question, a model's answer, and a score on that answer. Deployment is often different: users push back on assistants; inside agentic loops another process may propose a different answer and contest a draft before it is accepted (Qian et al., 2023; Hong et al., 2023; Wu et al., 2023); and later turns may ask a model to reconsider an answer it has already given. Disagreement of that kind is not hostile and not an edge case; it is the normal texture of use. Answers, however, move under it: models accommodate a user's stated view rather than the evidence (Sharma et al., 2023), and they abandon answers that were right under sustained counter-reasoning, at rates that do not fall as the model gets stronger (Abbasloo, 2025).

**A capability in its own right:** What moves under challenge is not another shortfall in answering, because the failure begins after a correct answer has already been produced. What is missing is the ability to weigh a case made against a position and to adjust only when the case deserves it — holding when the objection is unsound, yielding when it is sound — which we call *dialectical competence*. Training for correctness does not supply it, though correctness is its precondition: a model that cannot reach the right answer has nothing to hold. It is the higher skill, standing on answering rather than competing with it. It is also the capability hardest to supply from outside,



Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

because every way of correcting a model that leans on a stronger judge — a larger model, a separate verifier, a person — needs that judge to exist, and one better than the model being corrected is rare and costly. To separate the two, we train on two base models drawn from different families, Qwen3-8B (Yang et al., 2025) and Gemma-2 9B (Gemma Team et al., 2024), and run the same comparison on each. Figure 1 places the base model beside three ordinarily trained arms (each defined in §B), against the same challengers: each answers better on its own, and each still gives up its correct answers under challenge at close to the base model's rate. Training against the model's own opposition separates from that pattern, on both families (the diagnosis is made precise in §2).

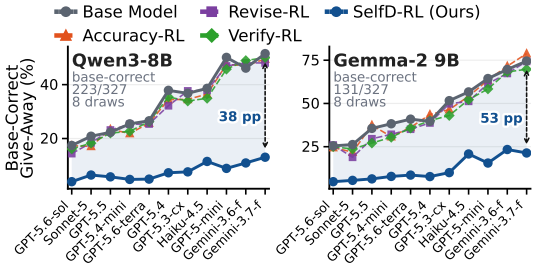



Figure 1: **Answering better does not buy the capability.** Share of the base model's correct answers each arm gives up under challenge, one position per challenger. Each of the three ordinarily trained arms answers better on its own and still tracks the base model here; only Self-Disagreement separates — by 25.0 pp on Qwen3-8B and 33.2 pp on Gemma-2 9B on average across these challengers.

**Our proposal — learning through Self-Disagreement:** The remedies that suggest themselves share one defect: every one of them works on the answer, when the capability at issue is the further one that begins after the answer. Training harder for correctness, giving the model a second pass over its own answer, and training it to judge that answer before committing all track the base model rather than separating from it (Figure 1); buying more inference — a reconsideration pass, or sampling many answers and voting — does not reach it either (§4); Appendix A discusses these lines of work at greater length. In each, the move the model makes is prompted, sampled or borrowed; it never follows from the model's own ruling on a case it was made to confront. The dialectical tradition puts exactly that ruling at the center: a position is developed by meeting the strongest case against it (Plato, 1963; Popper, 2005); that case is raised from within the position rather than handed to it from above (Hegel, 2018); and the position must rule on the case before it may move. We take this as a training paradigm rather than as a metaphor. A single policy plays both parts, differing only in what each is shown. It states a position, the *thesis*; in a context of its own, holding the question and the case just made but never the reference answer, it argues against that position, the *antithesis*; it then returns to the answering role, rules on whether that case stands, and commits to a final answer, the *synthesis* (§3). Only the answering turns are rewarded, and the reward is the verifier we already have: the same check that scores an answer against its reference also scores the answer the model settles on once opposed. No reward model, no preference data, no second model at training or deployment: *the opposition a model needs, it supplies itself.* We contribute:

(1) *Dialectical competence as a capability in its own right*: we measure it on a probe built for it, and show that neither accuracy-centered training nor inference-time remedies supply it (§2).
(2) *Self-Disagreement*: a single-policy extension of RLVR, with the opposition drawn from the policy being trained (§3).
(3) *Demonstrating that Self-Disagreement improves it*: trained on two base families, the models gain dialectical competence without paying for it in first-answer accuracy (§4.1), and keep the gain against challengers they never trained against, with an account of where it comes from (§4.2).

## 2 The Problem

**A probe that makes holding and yielding visible:** Benchmarks record whether a model ended up right, not what it did when it was pushed, so we measure on a probe built for the second question. We call it the *disagreement probe*, or *D-probe*, and it is used for every measurement of behavior under challenge in this paper. It is 327 questions, each put to every one of 17 challengers — 15 external models spanning open checkpoints to frontier systems, the model's own frozen starting point, and its own critique. The diagnosis below runs it on the two base models, Qwen3-8B and Gemma-2 9B. Every model answers once and that answer is replayed to each challenger, so only the answer given after the challenge can differ, and all three answers are graded: the model's own,

the challenger's, and the one the model settles on. Knowing which of the first two was right is what makes that choice legible, because it says what the right action was. What we report is the accuracy of that settled answer, the *synthesis accuracy*.

**The label is checked, not assumed:** Synthesis accuracy is what the metric rests on, so it was checked a second way: on the items where a base model gives up a correct answer, a blinded jury of Claude-Opus-4.7, GPT-5.5 and Gemini-3.1-pro reads the reasoning and judges all 121 such items to be non-genuine — conceding and adopting the challenger's answer with no valid pushback, or merely restating it — rather than a differently argued answer. What synthesis accuracy loses is therefore reasoning under challenge, not a relabelled answer (§G).

**Both base models are one-sided:** What their synthesis accuracy is made of, rather than its level, is the diagnosis. On the items where the model was right and the challenger wrong — where holding is the correct action — Qwen3-8B ends on the correct answer 53.8 % of the time and Gemma-2 9B 24.7 %. On the items where the challenger was right and the model wrong, where yielding is correct, the same two models score 84.1 % and 93.4 %. That pairing is not discrimination between good arguments and bad; it is close to unconditional surrender, and one behavior produces both numbers. The cost rises with the challenger: across Figure 1's challengers the base model gives up 17 % to 52 % of its correct answers on Qwen3-8B and 26 % to 75 % on Gemma-2 9B.

**Training for accuracy does not reach it:** The natural reading of that is that the models are not yet good enough, and that better answering will carry the rest. It does not. A control trained for accuracy alone (Accuracy-RL; defined, with its settings, in §B), from the same base model and the same data, reproduces most of the first-answer gain and buys none of the behavior under challenge, which can drift against it over training rather than holding steady (Appendix E). Figure 2 shows it on SuperGPQA (M-A-P Team et al., 2025): the trained alternatives move right, to higher thesis accuracy, and stay at the same height. Answering well is the floor, not the capability: a model that cannot produce the right answer has none to hold, but past that floor more of it buys nothing here. Nor is the opposition's content incidental: an opposition that carries no content extracts its benefit quickly and then exhausts it (§4.2).

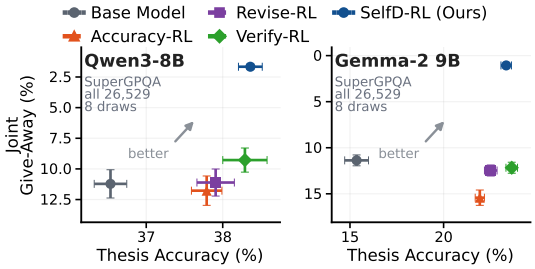



Figure 2: **Buying accuracy is not buying the capability.** SuperGPQA: each arm sits at its thesis accuracy ($x$) against how often it gave up a correct answer under challenge ($y$, inverted), so better is up and to the right. Bars are $\pm 1$ sd.

**What that leaves:** The failure is not a shortfall that better answering removes, and it is not repaired by the inference-time remedies we test, or by training the same objective harder. What is missing is the decision itself: telling a sound objection from an unsound one and acting on the difference. The rest of this paper is about training that.

## 3 Design

**Overview:** We train a single policy under a dialectical protocol we call *Self-Disagreement*. The policy commits to an answer, argues against that answer in a separate context, then returns to rule on the argument before committing again. Only the answering turns are trained. The protocol extends RLVR rather than departing from it: the verifier that scores an answer against its reference also scores the ruling, and nothing else in the training loop changes.

**Where the opposition comes from:** Learning to adjudicate disagreement presupposes a source of disagreement. The usual construction takes it from a second model, which makes the training signal a property of the *pair*: an opposition that is too weak can be dismissed without discrimination, and one that is too strong makes deference the best available response. The policy's competence also changes as it trains, so a pairing calibrated at the start does not stay calibrated, and keeping it so means tracking a quantity that is not directly observable. Drawing the opposition from the policy itself removes that coupling. Its competence tracks the policy's by construction and without scheduling, every training item has an opposition available without a second model to serve it, and

Algorithm 1: Self-disagreement, one training step

**Require:** policy $\pi_\theta$, RLVR verifier $\mathcal{V}(\cdot\,; y^\star)$, group size $G$, admission band $[\underline{p}, \overline{p}]$

sample a batch of questions $q$ with $p_0(q) \in [\underline{p}, \overline{p}]$
**for all** $q$ in the batch **do**
  **for** $i = 1$ to $G$ **do** ▷ one GRPO group
    $T \sim \pi_\theta(\cdot \mid q)$ ▷ answering role
    $A \sim \pi_\theta(\cdot \mid q, T)$ ▷ opposing role; no $y^\star$
    $S \sim \pi_\theta(\cdot \mid q, T, A)$ ▷ resumed: rule, then answer
    $c_T \leftarrow \mathcal{V}(T; y^\star)$; $c_S \leftarrow \mathcal{V}(S; y^\star)$ ▷ the answers
    $c_v \leftarrow \mathbf{1}[v = \mathcal{V}(A; y^\star)]$ ▷ and the ruling
    $\beta \leftarrow 2.0$ if $c_T = 0 \wedge \mathcal{V}(A; y^\star) = 0$, else $1$
    $R(\tau_i) \leftarrow \beta c_S + w_v c_v + w_T c_T - \lambda(\text{cap} + \text{stub}) - w_\phi(\phi_T + \phi_v)$
    **if** $\phi_T \vee \phi_v$ **then** $R(\tau_i) \leftarrow R_{\text{floor}}$ ▷ override
    **end if**
  **end for**
  $\hat{A}_i \leftarrow \big(R(\tau_i) - \mu_q\big)/\sigma_q$ ▷ group-relative
**end for**
update $\theta$ on the clipped surrogate, KL-penalized in the return ▷ stock GRPO

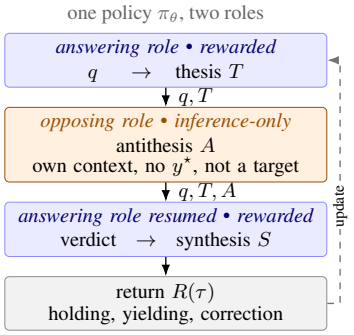



Figure 3: **Self-disagreement.** Two roles of one policy, in separate contexts. The opposing turn is sampled from the same weights tracking the policy as it trains. It conditions the synthesis, but is never itself a training target, so it reaches the answering role as given, the way a user's disagreement would. The answering role resumes and must rule before it may commit. Settings in §4.

what is deployed is a single set of weights. The construction does rest on an assumption worth stating: that a policy can be induced to argue against a position it has just committed to.

**The dialectical exchange:** Figure 3 shows one exchange and Algorithm 1 one training step; the prompt templates for all three turns are printed verbatim in §C. The policy is prompted into two roles that share its weights and differ only in what each is shown. In the answering role it writes the *thesis*, and later the *synthesis*; in the opposing role it writes the *antithesis*, disputing the strategy, the decisions and the reasoning of that particular answer *even where they appear correct* rather than making an independent second attempt at the question. Opposition is therefore unconditional — a correct thesis is argued against as readily as an incorrect one — which is what makes *holding* a position something the policy must learn, since a policy opposed only when it was wrong would never need to tell a sound challenge from an unsound one. Let $\pi_\theta$ be the policy and let $q \sim \mathcal{D}$ be a question with reference answer $y^\star$. One dialectical exchange is a trajectory $\tau = (T, A, S)$ of three generations from the *same* parameters under two role prompts — answering, then opposing, then answering again under instruction to judge:

$$\pi_\theta(\tau \mid q) \;=\; \underbrace{\pi_\theta(T \mid q)}_{\text{thesis}} \quad \underbrace{\pi_\theta(A \mid q, T)}_{\text{antithesis (sampled, inference-only)}} \quad \underbrace{\pi_\theta(S \mid q, T, A)}_{\text{synthesis}}, \tag{1}$$

The synthesis $S$ is internally ordered: it opens with a verdict $v \in \{0, 1\}$ on $A$'s answer and only then commits to an answer. Because the answer is decoded after the verdict and conditioned on it, the policy must adjudicate before it may answer, and $v$ is separately verifiable; ablating the verdict removes that opening block, leaving a synthesis that answers with no obligation to rule. Writing $\mathcal{V}(\cdot\,; y^\star) \in \{0, 1\}$ for the ordinary RLVR verifier — exact match against the reference answer, so that no learned reward model and no preference data enter anywhere — the trajectory induces $c_T = \mathcal{V}(T; y^\star)$, $c_S = \mathcal{V}(S; y^\star)$, and

$$c_v \;=\; \mathbf{1}[\, v = \mathcal{V}(A; y^\star)\,], \tag{2}$$

which scores whether the policy judged the opposition rightly, not whether the opposition was right. This is where the verifier's reach is extended: the same check, applied to a decision about an answer rather than to an answer. The credit is available on every item, including those in which the opposition concedes, so the objective by itself does not rule out a policy that agrees whenever the opposition agrees with it. It is more than that: removing the content of the disagreement costs about two thirds of the gain, and the trained policy keeps its advantage against challengers it never trained against (§F). The antithesis is sampled from the policy but is *inference-only*: its trajectory is never collected into the training episode, so $\pi_\theta(A \mid q, T)$ is a sampling distribution rather than an optimized factor, and opposing is never rewarded — only answering and adjudicating are — which removes any route by which it could learn to soften its challenges. No factor of Eq. (1) is conditioned on $y^\star$: the verifier is applied only after the trajectory is complete, so it cannot oppose selectively on items it is about to answer incorrectly. Thesis and synthesis are optimized as one trajectory, not as separate turns, so credit assignment reaches the answer and the ruling on its challenge jointly.

**Guarding against collapse:** The characteristic failure of an endogenous construction is trivial agreement, in which $A$ concedes $T$ and the dialectic degenerates to a single pass. Two provisions address this, neither of them part of the objective. The first is uninformedness, above: the opposing role cannot tell which items it would be safe to concede. The second is the admission band — items enter training only when the base policy's pass rate $p_0(q) = \Pr[\mathcal{V}(T; y^\star) = 1 \mid q; \pi_{\theta_0}]$ lies in $[\underline{p}, \overline{p}]$, which excludes items on which opposition is respectively vacuous or futile. What the opposition and the verdict actually do in training is then an empirical question rather than a guaranteed one, and we treat it as such: Appendix F measures the opposition the policy actually produces, and §4 measures the resulting behavior rather than relying on the construction.

**At inference:** None of the training machinery is required to use the result. The trained policy is an ordinary model: it answers a question in one pass, and what training has changed is how it behaves *when it is disagreed with*. The evaluation of §4 therefore presents the trained policy with an opposing view and measures the answer it commits to, the same situation the protocol trains, with none of the environment's apparatus. There is no admission band, no candidate selection, no verifier and no reference answer, and, since the opposing role is the policy itself, no second model to deploy.

**The objective:** Two failures are penalized directly:

$$\underbrace{\mathbf{1}[\,c_T = 1,\ c_S = 0\,]}_{\texttt{cap (capitulation)}} \qquad \underbrace{\mathbf{1}[\,c_T = 0,\ \mathcal{V}(A; y^\star) = 1,\ c_S = 0\,]}_{\texttt{stub (stubbornness)}} \tag{3}$$

Capitulation is a correct thesis surrendered; stubbornness is a valid challenge refused. *Synthesis accuracy*, $\Pr(c_S{=}1)$, is the outcome variable of every result in this paper. The return, for the configuration reported throughout, is

$$R(\tau) \;=\; \beta\, c_S \;+\; w_v c_v \;+\; w_T c_T \;-\; \lambda_{\text{cap}}\,\text{cap} \;-\; \lambda_{\text{stub}}\,\text{stub} \;-\; w_\phi\big(\phi_T + \phi_v\big), \tag{4}$$

subject to a hard format gate: if either the thesis or the verdict fails to parse ($\phi_T \vee \phi_v$), the return is *replaced* by a floor value and every positive channel is discarded,

$$R(\tau) \;\leftarrow\; R_{\text{floor}} \quad \text{if } \phi_T \vee \phi_v. \tag{5}$$

Here $\lambda_{\text{cap}} = \lambda_{\text{stub}} = 0.5$, $w_T = 1$, $w_v = 0.4$, $w_\phi = 0.1$ and $R_{\text{floor}} = -0.75$. The objective keeps paying for a correct first answer at full weight while adding the ruling at $0.4$: the capability is trained on top of answering, not in place of it. The gate is worth more than its weight suggests: the policy cannot decline to adjudicate. The remaining coefficient, $\beta$, multiplies the synthesis credit rather than adding to it: $\beta = 2.0$ when $c_T = 0$ and $\mathcal{V}(A; y^\star) = 0$, and 1 otherwise, since that is the one cell where no correct position exists to defer to and the answer has to be reached rather than chosen. The symmetry of $\lambda_{\text{cap}}$ and $\lambda_{\text{stub}}$ is load-bearing. An objective on terminal correctness alone is invariant to the trajectory by which correctness is reached, including trajectories that discard a correct answer and recover it — which §4 shows accuracy-only training does not merely tolerate but increasingly takes. Penalizing capitulation alone would buy stability at the cost of revisability. Optimization is by GRPO (Shao et al., 2024): for each $q$ a group of $G$ trajectories is sampled and scored by its group-relative advantage $\hat{A}_i = (R(\tau_i) - \mu_q)/\sigma_q$, maximized under a clipped surrogate (Schulman et al., 2017), with the KL penalty to the reference policy applied to the return rather than to the loss. The algorithm is unmodified: the dialectic changes what is scored, not how the update is computed.

## 4 Evaluation

This section has three parts. The first asks whether the capability is there at all, and whether available approaches produce it: it is weak in the base model, arrives with training, holds on benchmarks the training never saw, and no rival design or test-time scheme we compare reaches it. The second asks where the gain comes from: whether the model answers better on the first attempt or decides better once challenged, and which part of the design buys it. The third asks whether the result is valid and generalizes: a second reading of the metric by a blinded jury of three frontier models, a different base model, and challengers it never trained against.

**Setup:** Every trained arm is reinforcement learning from verifiable rewards on the same data, from the same initialization and with the same verifier; what separates them is the return and the number of turns it is computed over. Two further arms are not trained at all: the test-time schemes, which spend additional sampling on the base model instead. Because the opposition is sampled rather

Table 1: **Synthesis accuracy under challenge, per benchmark (%).** *Vf-RL*, *Rv-RL* and *Ac-RL* are Verify-RL, Revise-RL and Accuracy-RL. *Maj@64* and *Self-Refl*: the base model given more inference — test-time self-consistency and test-time self-reflection. ‡Per-benchmark detail in Table 10. Mean ± sd over eight sampled runs. Decoding settings for the two test-time arms are in §J.

| Model | Benchmark | Base | SelfD-RL | Vf-RL | Rv-RL | Ac-RL | Maj@64 | Self-Refl |
|---|---|---|---|---|---|---|---|---|
| Qwen3-8B | AGIEval SAT-Math (Zhong et al., 2023) | 96.3±1.0 | **99.1**±0.3 | 97.9±0.6 | 97.7±0.6 | 97.6±0.7 | 97.6±0.7 | 97.3±0.8 |
| | AIME 2025 (Zhang & Math-AI Team, 2025) | 17.5±3.5 | **25.4**±4.3 | 24.2±5.3 | 19.6±5.5 | 15.4±3.1 | 24.4±9.6 | 23.3±8.8 |
| | GPQA (main\diamond) (Rein et al., 2023) | 38.0±3.5 | **45.6**±1.5 | 40.2±3.0 | 40.2±2.8 | 37.3±3.8 | 43.8±2.0 | 44.5±1.8 |
| | LiveCodeBench (Jain et al., 2024) | 65.4±2.4 | **72.2**±2.0 | 70.3±1.5 | 66.3±2.1 | 64.6±2.6 | 65.9±1.1 | 65.5±0.9 |
| | MMLU-Redux (Gema et al., 2024) | 74.7±1.7 | **82.5**±0.3 | 77.1±1.2 | 75.6±1.4 | 74.6±1.7 | 77.3±0.4 | 75.6±1.0 |
| | MuSR-murder (Sprague et al., 2023) | 58.5±2.9 | **63.0**±2.0 | 52.0±3.6 | 53.6±3.1 | 51.9±3.6 | 58.5±1.7 | 53.6±3.0 |
| | Omni-MATH (Gao et al., 2024) | 32.4±0.6 | **38.8**±0.3 | 36.6±0.4 | 34.1±0.4 | 33.7±0.4 | 33.6±0.4 | 33.5±0.9 |
| | SuperGPQA (M-A-P Team et al., 2025) | 33.9±0.6 | **39.5**±0.1 | 36.7±0.7 | 35.2±0.6 | 34.6±0.7 | 35.4±0.1 | 34.5±0.2 |
| | TheoremQA (Chen et al., 2023) | 33.5±0.8 | **40.6**±0.8 | 37.8±1.3 | 35.7±1.9 | 35.5±0.9 | 35.0±0.9 | 35.1±1.1 |
| | **Avg** | 50.0±0.8 | **56.3±0.5** | 52.5±1.2 | 50.9±1.1 | 49.5±0.8 | 52.4±1.2 | 51.4±1.1 |
| **Gemma-2 9B‡** | | 19.9±0.7 | **37.8±0.7** | 28.2±0.5 | 28.0±0.9 | 23.6±0.9 | 24.6±0.6 | 25.5±0.2 |

than fixed we resample it over eight independent draws and report the mean and spread across those draws. Capitulation and stubbornness, defined in §3, are counted over every scored item rather than over the subset on which each could occur, so that they add to synthesis accuracy and to each other. Behavior under challenge is measured on the D-probe throughout; the benchmark panel asks whether what it measures survives away from it. What each arm is, the model, data, hardware and optimizer settings, the step at which each arm is read, and the construction of the disagreement probe are in §B.

### 4.1 Overall results and analysis

**Does it work, and could anything already available have done it:** The capability is not present in either base model and it arrives with training in both (Figure 4). The single-pass accuracy-trained arm (Accuracy-RL) is flat over its own run, so more training compute does not close the gap; nor does spending more at inference (the test-time columns of Table 1). Table 1 then asks whether what was learned on the probe survives away from it.

**It survives off the training probe, on both base models:** On Qwen3-8B the trained model is ahead of the base model on all nine public benchmarks, on mathematics and code as well as on the multiple-choice sets where a challenge is cheapest to issue, and on average it settles on the right answer after being challenged +6.30 pp more often (95 % CI $[+5.19, +7.41]$), 50.0 % to 56.3 %. The final row of Table 1 repeats the whole exercise on a second base model, Gemma-2 9B, and the pattern holds: 19.9 % to 37.8 % on average, ahead of its own base on every benchmark (§L). Every column of the upper block is read under the same protocol and scored by the same verifier, so what separates the columns is what each was trained to do and not how each was measured.

**Why none of the alternatives reach it:** Table 1 sets five of them beside ours, and each falls short for its own reason. Sampling the base model harder — Majority Vote over 64 samples an item, or Test-Time Self-Reflection reconsidering before it commits — buys more attempts but no one to disagree with, so what is missing is not compute. Revise-RL trains a second pass, rewarded on the answer it settles on and penalized for abandoning a correct one, but with nothing disagreeing with it there is no valid challenge to refuse, so what is missing is not the opportunity to change the answer. Verify-RL supplies the verdict by having the model judge its own answer first, and rewards the judgment alongside the answer; it comes closest of the four and still recovers under half of our gain on both base models: a verdict a model issues against itself is not the same as being argued with. The fifth is the arm built to maximize the quantity our result might be mistaken for: Accuracy-RL trains on the same data with the same verifier, differing only in taking a single pass

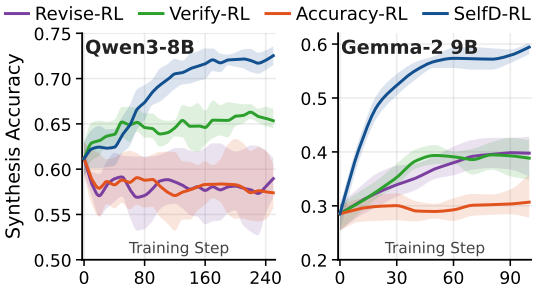



Figure 4: **Dialectical competence is trained under Self-Disagreement, while the arms that drop the disagreement recover at most part of it.** Bands across seeds. Separate models: read the shapes, not the cells.

Table 2: **Only Self-Disagreement gains *skill*; the other arms gain only *situations*.** *Situations* is what an arm gains by walking into an easier mix of predicaments — a better first answer; *skill* is what it gains by handling the same predicament better. Points of synthesis accuracy on the D-probe; for each arm the two terms sum to its total change. Brackets are 95 % intervals clustered by question. The base model is the reference.

| | Qwen3-8B | | Gemma-2 9B | |
|---|---|---|---|---|
| **Arm** | *situations* ($\sum \Delta\pi_m r_m^{\text{base}}$) | *skill* ($\sum \pi_m^{\text{arm}} \Delta r_m$) | *situations* ($\sum \Delta\pi_m r_m^{\text{base}}$) | *skill* ($\sum \pi_m^{\text{arm}} \Delta r_m$) |
| Accuracy-RL | $+1.38$ $[+0.36, +2.37]$ | $-1.21$ $[-2.27, -0.23]$ | $+1.47$ $[+0.05, +2.82]$ | $-1.53$ $[-2.55, -0.50]$ |
| Revise-RL | $+0.70$ $[-0.45, +1.77]$ | $-1.03$ $[-2.20, +0.11]$ | $+1.90$ $[+0.32, +3.47]$ | $-1.77$ $[-2.91, -0.72]$ |
| Verify-RL | $+1.78$ $[+0.65, +3.03]$ | $-0.98$ $[-2.20, +0.28]$ | $+1.80$ $[+0.25, +3.26]$ | $-0.44$ $[-1.45, +0.62]$ |
| SelfD-RL | $+1.08$ $[-0.09, +2.16]$ | $\mathbf{+11.46}$ $[+8.79, +14.11]$ | $+0.11$ $[-1.56, +1.84]$ | $\mathbf{+16.15}$ $[+12.27, +19.99]$ |

rewarded on whether that answer is right; its interval on the panel includes zero. Optimizing the first answer does not deliver the capability as a by-product, which it would have to for our gain to be better answers, not better judgment — what a model carries into the dialectic and what survives it are separable, and a model can reach a right answer and still not distinguish it from a wrong one — though not the reverse, since a model that cannot reach it has nothing to distinguish.

**The same dissociation in the reward ablation:** Holding the exchange fixed and varying only what is returned moves the dialectical quantity several times further than it moves accuracy: across the reward ablation the arms span 12.23 pp in synthesis accuracy against 2.75 pp in thesis accuracy — 4.4 times as far — and the ordering is not a trade. The same separation also appears within one arm over training (§E).

### 4.2 Where the gain comes from

Synthesis accuracy rises, but it does not say why, and two mechanisms would raise it equally: the model may answer better on the first attempt and so meet fewer challenges it deserves to lose, or its dialectical competence may itself have improved. Grading its first answer against its challenger's sorts every item into four situations by which side was right. Writing $\pi_m$ for the share of items in situation $m$ and $r_m$ for how often the model ends on the correct answer there, synthesis accuracy is $c_S = \sum_m \pi_m r_m$, and the change against the base model splits without remainder into a *situations* term and a *skill* term, $\Delta c_S = \sum_m \Delta\pi_m\, r_m^{\text{base}} + \sum_m \pi_m^{\text{arm}}\, \Delta r_m$, the standard composition-versus-rate decomposition (Kitagawa, 1955), where *situations*, the first term, is what the model gains by walking into a different mix of predicaments — the better-first-answer explanation, the floor — and *skill*, the second, is the part of the gain due to a change in dialectical competence, the capability standing on that floor rather than an alternative to it. We decompose it on the D-probe, over the same seventeen challengers; Appendix I gives the full profile and the reference policies.

**Every arm improves its first answer; only Self-Disagreement also improves the decision:** Table 2 gives the split, and Appendix I the four situations behind it. The base models are one-sided — holding only 53.8 % and 24.7 % of the time while yielding 84.1 % and 93.4 %, which is closer to unconditional surrender than to discrimination. Self-Disagreement becomes two-sided, holding 90.4 % and 83.5 % while still yielding 50.8 % and 65.5 %, and that registers as skill: 11.46 of its 12.6-point gain on Qwen and 16.15 of 16.3 on Gemma, positive against 17/17 of the challengers individually in both families. A model deciding no better than the base scores zero; one deciding correctly in every situation it met would score 34.9. The trained rivals behave alike and oppositely: each gains through situations, and none has a skill term above zero.

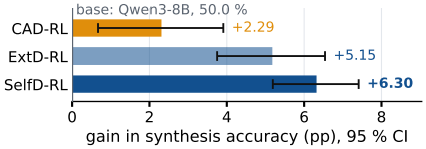



Figure 5: **The content of the disagreement matters.** Average over the nine benchmarks. Bars are the gain over the base model, with 95 % CIs.

**Does the disagreement have to be genuine, and does it have to come from elsewhere:** Two variants of the recipe answer this, each removing one ingredient; Figure 5 places them beside the full recipe. Content-Agnostic Disagreement (CAD-RL) runs the same exchange against a challenge that does not depend on the answer it challenges, and it reaches +2.29 $[+0.67, +3.91]$, so about a third of the gain is available from the presence of opposition and the rest requires the opposition to be about something. External Disagreement (ExtD-RL) is our own recipe with the antithesis drawn from Qwen3-14B, a larger model of the same family, and it reaches +5.15 $[+3.75, +6.54]$ against +6.30 for the policy arguing with itself, which is parity. What separates them is not what

they achieve but what they require: the external arm holds a second model in service for the whole of training, on a serving allocation held in a one-to-four ratio against the one training used, and it leaves the choice of that model unresolved, since it was picked for being larger than the model being trained and we know of no principled way to make that choice. Both are reported on Qwen3-8B only: they vary the *design* rather than ablate ours, so neither is needed to establish that the effect reproduces on a second base model. One might object that both arms argue inside our own loop, so the model may only handle its own objections. It cannot be: the same advantage holds against fifteen challengers it never trained against, several stronger than itself (Table 8).

**Which part of the design does the work:** Most of it is bought by the requirement to rule on the opposition before answering; the reward for ruling well adds a further, separately resolved amount; and a second pass with nothing to rule on adds none of it. Each term can be switched off with the others held fixed, so every claim below is a single-variable contrast on the same items. The verdict is the active ingredient. Without it the arm reaches 59.9 % against 68.5 % with it, a gap of 8.6 pp against a conservative unpaired bound of 5.8 pp (Table 7). Two further ablations separate the verdict from the opposition (Table 3). Requiring the verdict in the output but paying nothing for getting it right gives $2.8 \pm 0.6$ %, against $1.0 \pm 0.4$ % for the full recipe — a resolved gap of 1.8 pp, separating the *structural* requirement from the reward for meeting it. A second pass without a verdict does not substitute: Revise-RL abandons its correct answer on $12.8 \pm 1.5$ % of all scored items, 11.8 pp from the full recipe, and places it beside the unrewarded-dialectic and Accuracy-RL arms: with the arms that differ by a verdict, not by a second pass. The same reward creates an obvious exploit, since it pays for correcting a flawed first answer, and Appendix M reports the controls: the rate at which the final answer simply reproduces the challenger's falls from 0.693 in the unrewarded-dialectic arm to 0.393–0.419 in the arms that pay for the verdict, and the arm with the verdict removed sits at 0.705, with the control not the treatment.

Table 3: **Requiring a verdict is what moves capitulation; paying for it adds the rest. Qwen3-8B.** Capitulation and synthesis accuracy (%). **Self-Disagreement** full recipe; **Judge** verdict required, not rewarded; **Rv-RL** revision pass, no opposition; **Unrew.** exchange runs, only the thesis paid; **Ac-RL** one pass, no exchange. Rates are over *all* scored items, not the challenged subset. Step, draws, protocol in §B.

| Arm | capitulation ↓ | synthesis ↑ |
|---|---|---|
| SelfD-RL | $1.0 \pm 0.4$ | $74.5 \pm 1.2$ |
| Judge | $2.8 \pm 0.6$ | $70.2 \pm 0.4$ |
| Rv-RL | $12.8 \pm 1.5$ | $62.3 \pm 0.8$ |
| Unrew. | $15.1 \pm 1.4$ | $61.9 \pm 2.0$ |
| Ac-RL | $13.5 \pm 1.3$ | $62.5 \pm 1.2$ |
| Base | $10.3 \pm 2.3$ | $62.8 \pm 2.9$ |

### 4.3 Validity and Generalization

**Is the metric measuring what we claim:** As §2 reports, a blinded jury of three frontier models reads the reasoning and finds what the metric counts to be a real failure rather than an artifact of comparing two answer strings. §G gives the validation in full — the rubric, the panel, the blinding, the agreement statistics, and the limits of what this grading establishes.

**The gain holds against challengers that cannot have decayed:** When the model argues with itself, thesis and antithesis are the same checkpoint, so a rise in synthesis accuracy could mean the model holds better or that its own antithesis got worse. Separating the two needs a challenger that does not move with the arm, and the ladder is built from them: sixteen of the seventeen antithesis models are served from fixed weights and are identical for every arm, so nothing an arm learns can change what it is argued with. Two of those sixteen are the base checkpoints the arms were trained from — the frozen base models, each family's own starting point held still and replayed — which is the sharpest form of the test, since it is the one challenger the arm might be suspected of having left behind. Scoring synthesis accuracy on the same D-probe, with theses frozen and replayed: against the frozen Qwen3-8B the trained model reaches 72.9 % where its base reaches 60.5 %; against the frozen gemma-2-9b-it, 58.9 % against 32.0 %. Those gains, 12.4 and 26.9 pp, are as large as the ones measured against the model's own antithesis, which is what rules the alternative reading out: a decaying antithesis cannot explain a gain that survives fifteen challengers that never changed.

Table 4: **The same recipe on other models, other sizes, and other families.** Before training and at convergence, mean ± sd over seeds (%). ‡Qwen3 14B was **pre-registered as a near-null** because it begins with almost no headroom to move; it is reported for that reason.

| | capitulation ↓ | | synthesis ↑ | |
|---|---|---|---|---|
| **Model** | base | **+SelfD-RL** | base | **+SelfD-RL** |
| Gemma 3 12B | $20.2 \pm 2.6$ | $\mathbf{2.6 \pm 0.8}$ | $53.0 \pm 1.9$ | $\mathbf{68.1 \pm 0.8}$ |
| Granite 3.1 8B | $14.7 \pm 1.3$ | $\mathbf{0.9 \pm 0.8}$ | $40.5 \pm 1.6$ | $\mathbf{46.1 \pm 1.6}$ |
| Llama 3.1 8B | $6.6 \pm 0.6$ | $\mathbf{2.6 \pm 0.3}$ | $51.4 \pm 1.7$ | $\mathbf{57.9 \pm 0.7}$ |
| Qwen3 14B‡ | $1.8 \pm 0.4$ | $\mathbf{0.8 \pm 0.4}$ | $71.1 \pm 0.7$ | $\mathbf{73.4 \pm 0.6}$ |

**Across other model families and sizes:** The result is not a property of one base model. Applying the same recipe to four further models — other sizes and other families — synthesis accuracy rises and capitulation falls in every one, on the same held-out items scored before and after training. All four end below 5 %, and three of them started above it (Table 4).

**Does it depend on who argues:** Table 8, in Appendix I, puts both models against every antithesis we ran, from a 4B model to the strongest frontier systems available, including the base model itself and a base model of the same size. Over much of that ladder the base model would have done better by ignoring the antithesis entirely — keeping its first answer, or taking the challenger's, whichever is stronger — and at worst it gives up 16.8 pp that way; its accuracy swings 25.6 pp from one end of the ladder to the other. Self-Disagreement falls short of its own reference by at most 2.5, and swings 10.3. Whatever it is arguing with, it ends up in much the same place, and that difference between the two spans is large enough that no rung on the ladder reverses it. A fixed-population measurement with different challengers finds the same invariance (§H).

**A model that ignores the exchange scores exactly one; neither of ours does:** For every item on which one of the thesis and the antithesis was right and the other wrong, we ask whether the model ended up on the side the evidence favored. Such a policy holds just as often when holding is wrong as when it is right, so its two odds are equal and its ratio is exactly one, whatever mix of items it meets. Neither model is near that. On the same probe and the same seventeen antitheses as Table 8, the odds are 6.2 for the base model against 9.7 for the trained one on Qwen3-8B, and 4.6 against 9.6 on Gemma-2 9B. Read antithesis by antithesis, all seventeen Qwen rungs sit above one, the weakest at 2.6 $[1.6, 4.1]$; on Gemma thirteen of the seventeen do. The four that do not are its weakest challengers — small local models, including the base model itself — where holding is the right move on about nine items in ten and the deciding cell holds thirteen to twenty-three items. Those four are lower bounds: a model that swapped one wrong answer for another scores as having held, pushing the ratio down. No comparison between the two arms is claimed: what counts as the right action depends on that model's own thesis, which training changes, so each ratio is read against the null of one.

## 5 Limitations and concluding remarks

**What we have not established:** The recipe is trained on two base models from different families and applied unchanged to four further models across other sizes and families, where capitulation falls in every one. That is evidence that the effect reproduces, not that it is universal: establishing it across model families and scales in general would take a far wider sweep than the compute here allows, and we do not claim it. The same bound applies to the domain. Every reward here comes from exact-match verification against a reference answer, so our evidence covers only checkable answers; whether the same signal can be built where correctness is not checkable is open.

**How far there is still to go:** Self-Disagreement moves the decision itself, but it does not finish the job. Its skill term reaches 11.46 pp on Qwen3-8B and 16.15 on Gemma-2 9B, against the 34.9 and 47.7 a model deciding correctly in every situation it met would score — roughly a third of what is available in each case. The gain that is there is a gain in discrimination rather than in refusal (§4.3), so we take part of that residual to be a limit of the model rather than of the training: dialectical competence stands on first-answer accuracy without following from it, and a model cannot rule correctly on a question it is not capable of getting right in the first place. Closing the rest is what the next work has to do, and the reference policies in §I are what it should be measured against.

**Concluding remarks:** What limits progress in language models is increasingly not what they can do but our ability to tell when they are right. Every remedy for that which depends on a stronger judge — a larger model, a separate verifier, a person — inherits the limit, and runs out exactly where it is needed most: at the frontier, where nothing stronger is available to ask. This work shows that one part of that judgment can be placed inside the model, and trained there with nothing more capable than the model being trained. It is one part and not the whole of it, and the part still missing is the one that matters most — because a model that could reliably feel the force of an argument it had raised against itself would not merely be arguing well; it would be supervising itself. We take that to be the direction this line of work points, and the standard by which it should be judged.

## REPRODUCIBILITY STATEMENT

§B gives the models and hardware, the training data and how it is selected, the optimizer configuration, and the reporting point at which every arm is read. Table 1 carries the per-arm and per-benchmark results for Qwen3-8B and the per-arm results for Gemma-2 9B, whose per-benchmark detail is in Table 10; Appendix D the same breakdown for the two alternative designs, and Appendices F and G carry the construct checks on the two metrics the paper turns on.

## AI USE STATEMENT

In this work, we used generative AI tools to implement methods: the training and evaluation code, in addition to running and monitoring jobs on our GPU cluster. Code implementation was restricted to decisions the authors had already taken, not the design of the method or of the experiments it is evaluated in. We have reviewed and verified any code generated by AI during development, and every artefact so produced was read and tested by the authors before any result derived from it entered this paper. We have not used generative AI tools to develop the conceptual framing of this work, to propose or refine its hypotheses, to design or give feedback on its methodology or experiments, to generate synthetic data sets, to clean or reformat the training data, to assist with translation, to support qualitative or thematic analysis of our data, or to interpret results; the proof-related tasks, formulating mathematical claims, providing critical ingredients for proofs, and writing proofs, are not applicable to this work, which contains no theorems. Generative models also appear inside the work itself rather than as authoring tools, as the challengers the method is trained and evaluated against and as the blinded jury of §G that reads model reasoning and judges whether a concession was genuine; both are described in the paper as objects of study. Additionally, we used generative AI tools to edit the text of this paper for readability; and we have not used them to draft sections of this paper or to find related work. We have reviewed all AI-assisted work: the analysis, the claims and the limits placed on them are the authors' own. We take responsibility for the final content of this work, including text, claims or artifacts produced with the aid of generative AI.

## REFERENCES

Soheil Abbasloo. Evaluating LLM reasoning beyond correctness and CoT. *arXiv preprint arXiv:2510.18134*, 2025. v2, 2026-02-12.

AI-MO. AIMO validation AMC: AMC 2022–2023 competition problems. https://huggingface.co/datasets/AI-MO/aimo-validation-amc, 2024. Accessed 2026-09-24.

Wenhu Chen, Ming Yin, Max Ku, Pan Lu, Yixin Wan, Xueguang Ma, Jianyu Xu, Xinyi Wang, and Tony Xia. TheoremQA: A theorem-driven question answering dataset. *arXiv preprint arXiv:2305.12524*, 2023.

Peter Clark, Isaac Cowhey, Oren Etzioni, Tushar Khot, Ashish Sabharwal, Carissa Schoenick, and Oyvind Tafjord. Think you have solved question answering? try ARC, the AI2 reasoning challenge. *arXiv preprint arXiv:1803.05457*, 2018.

DeepSeek-AI, Daya Guo, Dejian Yang, Haowei Zhang, Junxiao Song, Peiyi Wang, Qihao Zhu, Runxin Xu, Ruoyu Zhang, Shirong Ma, Xiao Bi, Xiaokang Zhang, Xingkai Yu, Yu Wu, Z. F. Wu, Zhibin Gou, Zhihong Shao, Zhuoshu Li, Ziyi Gao, Aixin Liu, Bing Xue, Bingxuan Wang, Bochao Wu, Bei Feng, Chengda Lu, Chenggang Zhao, Chengqi Deng, Chenyu Zhang, Chong Ruan, Damai Dai, Deli Chen, Dongjie Ji, Erhang Li, Fangyun Lin, Fucong Dai, Fuli Luo, Guangbo Hao, Guanting Chen, Guowei Li, H. Zhang, Han Bao, Hanwei Xu, Haocheng Wang, Honghui Ding, Huajian Xin, Huazuo Gao, Hui Qu, Hui Li, Jianzhong Guo, Jiashi Li, Jiawei Wang, Jingchang Chen, Jingyang Yuan, Junjie Qiu, Junlong Li, J. L. Cai, Jiaqi Ni, Jian Liang, Jin Chen, Kai Dong, Kai Hu, Kaige Gao, Kang Guan, Kexin Huang, Kuai Yu, Lean Wang, Lecong Zhang, Liang Zhao, Litong Wang, Liyue Zhang, Lei Xu, Leyi Xia, Mingchuan Zhang, Minghua Zhang, Minghui Tang, Meng Li, Miaojun Wang, Mingming Li, Ning Tian, Panpan Huang, Peng Zhang, Qiancheng Wang, Qinyu Chen, Qiushi Du, Ruiqi Ge, Ruisong Zhang, Ruizhe Pan, Runji Wang, R. J. Chen, R. L. Jin, Ruyi Chen, Shanghao Lu, Shangyan Zhou, Shanhuang Chen, Shengfeng

Ye, Shiyu Wang, Shuiping Yu, Shunfeng Zhou, Shuting Pan, S. S. Li, Shuang Zhou, Shaoqing Wu, Shengfeng Ye, Tao Yun, Tian Pei, Tianyu Sun, T. Wang, Wangding Zeng, Wanjia Zhao, Wen Liu, Wenfeng Liang, Wenjun Gao, Wenqin Yu, Wentao Zhang, W. L. Xiao, Wei An, Xiaodong Liu, Xiaohan Wang, Xiaokang Chen, Xiaotao Nie, Xin Cheng, Xin Liu, Xin Xie, Xingchao Liu, Xinyu Yang, Xinyuan Li, Xuecheng Su, Xuheng Lin, X. Q. Li, Xiangyue Jin, Xiaojin Shen, Xiaosha Chen, Xiaowen Sun, Xiaoxiang Wang, Xinnan Song, Xinyi Zhou, Xianzu Wang, Xinxia Shan, Y. K. Li, Y. Q. Wang, Y. X. Wei, Yang Zhang, Yanhong Xu, Yao Li, Yao Zhao, Yaofeng Sun, Yaohui Wang, Yi Yu, Yichao Zhang, Yifan Shi, Yiliang Xiong, Ying He, Yishi Piao, Yisong Wang, Yixuan Tan, Yiyang Ma, Yiyuan Liu, Yongqiang Guo, Yuan Ou, Yuduan Wang, Yue Gong, Yuheng Zou, Yujia He, Yunfan Xiong, Yuxiang Luo, Yuxiang You, Yuxuan Liu, Yuyang Zhou, Y. X. Zhu, Yanhong Xu, Yanping Huang, Yaohui Li, Yi Zheng, Yuchen Zhu, Yunxian Ma, Ying Tang, Yukun Zha, Yuting Yan, Z. Z. Ren, Zehui Ren, Zhangli Sha, Zhe Fu, Zhean Xu, Zhenda Xie, Zhengyan Zhang, Zhewen Hao, Zhicheng Ma, Zhigang Yan, Zhiyu Wu, Zihui Gu, Zijia Zhu, Zijun Liu, Zilin Li, Ziwei Xie, Ziyang Song, Zizheng Pan, Zhen Huang, Zhipeng Xu, Zhongyu Zhang, and Zhen Zhang. Deepseek-r1: Incentivizing reasoning capability in llms via reinforcement learning. *arXiv preprint arXiv:2501.12948*, 2025.

Yilun Du, Shuang Li, Antonio Torralba, Joshua B. Tenenbaum, and Igor Mordatch. Improving factuality and reasoning in language models through multiagent debate. *arXiv preprint arXiv:2305.14325*, 2023.

Bofei Gao, Feifan Song, Zhe Yang, Zefan Cai, Yibo Miao, Qingxiu Dong, Lei Li, Chenghao Ma, Liang Chen, Runxin Xu, Zhengyang Tang, Benyou Wang, Daoguang Zan, Shanghaoran Quan, Ge Zhang, Lei Sha, Yichang Zhang, Xuancheng Ren, Tianyu Liu, and Baobao Chang. Omni-MATH: A universal olympiad level mathematic benchmark for large language models, 2024. arXiv:2410.07985.

Aryo Pradipta Gema, Joshua Ong Jun Leang, Giwon Hong, Alessio Devoto, Alberto Carlo Maria Mancino, Rohit Saxena, Xuanli He, Yu Zhao, Xiaotang Du, Mohammad Reza Ghasemi Madani, Claire Barale, Robert McHardy, Joshua Harris, Jean Kaddour, Emile van Krieken, and Pasquale Minervini. Are we done with MMLU? *arXiv preprint arXiv:2406.04127*, 2024.

Gemma Team, Morgane Riviere, Shreya Pathak, Pier Giuseppe Sessa, Cassidy Hardin, Surya Bhupatiraju, Léonard Hussenot, Thomas Mesnard, Bobak Shahriari, Alexandre Ramé, Johan Ferret, Peter Liu, Pouya Tafti, Abe Friesen, Michelle Casbon, Sabela Ramos, Ravin Kumar, Charline Le Lan, Sammy Jerome, Anton Tsitsulin, Nino Vieillard, Piotr Stanczyk, Sertan Girgin, Nikola Momchev, Matt Hoffman, Shantanu Thakoor, Jean-Bastien Grill, Behnam Neyshabur, Olivier Bachem, Alanna Walton, Aliaksei Severyn, Alicia Parrish, Aliya Ahmad, Allen Hutchison, Alvin Abdagic, Amanda Carl, Amy Shen, Andy Brock, Andy Coenen, Anthony Laforge, Antonia Paterson, Ben Bastian, Bilal Piot, Bo Wu, Brandon Royal, Charlie Chen, Chintu Kumar, Chris Perry, Chris Welty, Christopher A. Choquette-Choo, Danila Sinopalnikov, David Weinberger, Dimple Vijaykumar, Dominika Rogozińska, Dustin Herbison, Elisa Bandy, Emma Wang, Eric Noland, Erica Moreira, Evan Senter, Evgenii Eltyshev, Francesco Visin, Gabriel Rasskin, Gary Wei, Glenn Cameron, Gus Martins, Hadi Hashemi, Hanna Klimczak-Plucińska, Harleen Batra, Harsh Dhand, Ivan Nardini, Jacinda Mein, Jack Zhou, James Svensson, Jeff Stanway, Jetha Chan, Jin Peng Zhou, Joana Carrasqueira, Joana Iljazi, Jocelyn Becker, Joe Fernandez, Joost van Amersfoort, Josh Gordon, Josh Lipschultz, Josh Newlan, Ju yeong Ji, Kareem Mohamed, Kartikeya Badola, Kat Black, Katie Millican, Keelin McDonell, Kelvin Nguyen, Kiranbir Sodhia, Kish Greene, Lars Lowe Sjoesund, Lauren Usui, Laurent Sifre, Lena Heuermann, Leticia Lago, Lilly McNealus, Livio Baldini Soares, Logan Kilpatrick, Lucas Dixon, Luciano Martins, Machel Reid, Manvinder Singh, Mark Iverson, Martin Görner, Mat Velloso, Mateo Wirth, Matt Davidow, Matt Miller, Matthew Rahtz, Matthew Watson, Meg Risdal, Mehran Kazemi, Michael Moynihan, Ming Zhang, Minsuk Kahng, Minwoo Park, Mofi Rahman, Mohit Khatwani, Natalie Dao, Nenshad Bardoliwalla, Nesh Devanathan, Neta Dumai, Nilay Chauhan, Oscar Wahltinez, Pankil Botarda, Parker Barnes, Paul Barham, Paul Michel, Pengchong Jin, Petko Georgiev, Phil Culliton, Pradeep Kuppala, Ramona Comanescu, Ramona Merhej, Reena Jana, Reza Ardeshir Rokni, Rishabh Agarwal, Ryan Mullins, Samaneh Saadat, Sara Mc Carthy, Sarah Cogan, Sarah Perrin, Sébastien M. R. Arnold, Sebastian Krause, Shengyang Dai, Shruti Garg, Shruti Sheth, Sue Ronstrom, Susan Chan, Timothy Jordan, Ting Yu, Tom Eccles, Tom Hennigan, Tomas Kocisky, Tulsee Doshi, Vihan Jain, Vikas Yadav, Vilobh Meshram, Vishal Dharmadhikari, Warren Barkley, Wei

Wei, Wenming Ye, Woohyun Han, Woosuk Kwon, Xiang Xu, Zhe Shen, Zhitao Gong, Zichuan Wei, Victor Cotruta, Phoebe Kirk, Anand Rao, Minh Giang, Ludovic Peran, Tris Warkentin, Eli Collins, Joelle Barral, Zoubin Ghahramani, Raia Hadsell, D. Sculley, Jeanine Banks, Anca Dragan, Slav Petrov, Oriol Vinyals, Jeff Dean, Demis Hassabis, Koray Kavukcuoglu, Clement Farabet, Elena Buchatskaya, Sebastian Borgeaud, Noah Fiedel, Armand Joulin, Kathleen Kenealy, Robert Dadashi, and Alek Andreev. Gemma 2: Improving open language models at a practical size. *arXiv preprint arXiv:2408.00118*, 2024.

Mor Geva, Daniel Khashabi, Elad Segal, Tushar Khot, Dan Roth, and Jonathan Berant. Did aristotle use a laptop? a question answering benchmark with implicit reasoning strategies. *arXiv preprint arXiv:2101.02235*, 2021.

Aaron Grattafiori, Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Alex Vaughan, et al. The llama 3 herd of models. *arXiv preprint arXiv:2407.21783*, 2024.

Alex Gu, Baptiste Rozière, Hugh Leather, Armando Solar-Lezama, Gabriel Synnaeve, and Sida I. Wang. CRUXEval: A benchmark for code reasoning, understanding and execution. *arXiv preprint arXiv:2401.03065*, 2024.

Chaoqun He, Renjie Luo, Yuzhuo Bai, Shengding Hu, Zhen Leng Thai, Junhao Shen, Jinyi Hu, Xu Han, Yujie Huang, Yuxiang Zhang, Jie Liu, Lei Qi, Zhiyuan Liu, and Maosong Sun. OlympiadBench: A challenging benchmark for promoting AGI with olympiad-level bilingual multimodal scientific problems, 2024. arXiv:2402.14008.

Georg Wilhelm Friedrich Hegel. *The Phenomenology of Spirit*. Cambridge Hegel Translations. Cambridge University Press, Cambridge, 2018. doi: 10.1017/9781139050494. Edited and translated by Terry Pinkard and Michael Baur.

Dan Hendrycks, Collin Burns, Steven Basart, Andy Zou, Mantas Mazeika, Dawn Song, and Jacob Steinhardt. Measuring massive multitask language understanding, 2020. arXiv preprint arXiv:2009.03300.

Dan Hendrycks, Collin Burns, Saurav Kadavath, Akul Arora, Steven Basart, Eric Tang, Dawn Song, and Jacob Steinhardt. Measuring mathematical problem solving with the math dataset. *arXiv preprint arXiv:2103.03874*, 2021.

Sirui Hong, Mingchen Zhuge, Jiaqi Chen, Xiawu Zheng, Yuheng Cheng, Ceyao Zhang, Jinlin Wang, Zili Wang, Steven Ka Shing Yau, Zijuan Lin, Liyang Zhou, Chenyu Ran, Lingfeng Xiao, Chenglin Wu, and Jürgen Schmidhuber. MetaGPT: Meta programming for A multi-agent collaborative framework. *arXiv preprint arXiv:2308.00352*, 2023.

Naman Jain, King Han, Alex Gu, Wen-Ding Li, Fanjia Yan, Tianjun Zhang, Sida Wang, Armando Solar-Lezama, Koushik Sen, and Ion Stoica. Livecodebench: Holistic and contamination free evaluation of large language models for code. *arXiv preprint arXiv:2403.07974*, 2024.

Muhammad Khalifa, Rishabh Agarwal, Lajanugen Logeswaran, Jaekyeom Kim, Hao Peng, Moontae Lee, Honglak Lee, and Lu Wang. Process reward models that think. *arXiv preprint arXiv:2504.16828*, 2025.

Evelyn M. Kitagawa. Components of a difference between two rates. *Journal of the American Statistical Association*, 50(272):1168–1194, 1955. doi: 10.1080/01621459.1955.10501299.

Aviral Kumar, Vincent Zhuang, Rishabh Agarwal, Yi Su, John D Co-Reyes, Avi Singh, Kate Baumli, Shariq Iqbal, Colton Bishop, Rebecca Roelofs, Lei M Zhang, Kay McKinney, Disha Shrivastava, Cosmin Paduraru, George Tucker, Doina Precup, Feryal Behbahani, and Aleksandra Faust. Training language models to self-correct via reinforcement learning. *arXiv preprint arXiv:2409.12917*, 2024.

Hyunseok Lee, Seunghyuk Oh, Jaehyung Kim, Jinwoo Shin, and Jihoon Tack. Revise: Learning to refine at test-time via intrinsic self-verification. *arXiv preprint arXiv:2502.14565*, 2025.

Hyunseok Lee, Soheil Abbasloo, Jihoon Tack, and Jinwoo Shin. Beyond correctness: Learning robust reasoning via transfer. *arXiv preprint arXiv:2602.08489*, 2026.

Tian Liang, Zhiwei He, Wenxiang Jiao, Xing Wang, Yan Wang, Rui Wang, Yujiu Yang, Shuming Shi, and Zhaopeng Tu. Encouraging divergent thinking in large language models through multi-agent debate. *arXiv preprint arXiv:2305.19118*, 2023.

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.

Zichen Liu, Changyu Chen, Wenjun Li, Penghui Qi, Tianyu Pang, Chao Du, Wee Sun Lee, and Min Lin. Understanding r1-zero-like training: A critical perspective. *arXiv preprint arXiv:2503.20783*, 2025.

M-A-P Team, Xinrun Du, Yifan Yao, Kaijing Ma, Bingli Wang, Tianyu Zheng, King Zhu, Minghao Liu, Yiming Liang, Xiaolong Jin, Zhenlin Wei, Chujie Zheng, Kaixin Deng, Shawn Gavin, Shian Jia, Sichao Jiang, Yiyan Liao, Rui Li, Qinrui Li, Sirun Li, Yizhi Li, Yunwen Li, David Ma, Yuansheng Ni, Haoran Que, Qiyao Wang, Zhoufutu Wen, Siwei Wu, Tyshawn Hsing, Ming Xu, Zhenzhu Yang, Zekun Moore Wang, Junting Zhou, Yuelin Bai, Xingyuan Bu, Chenglin Cai, Liang Chen, Yifan Chen, Chengtuo Cheng, Tianhao Cheng, Keyi Ding, Siming Huang, Yun Huang, Yaoru Li, Yizhe Li, Zhaoqun Li, Tianhao Liang, Chengdong Lin, Hongquan Lin, Yinghao Ma, Tianyang Pang, Zhongyuan Peng, Zifan Peng, Qige Qi, Shi Qiu, Xingwei Qu, Shanghaoran Quan, Yizhou Tan, Zili Wang, Chenqing Wang, Hao Wang, Yiya Wang, Yubo Wang, Jiajun Xu, Kexin Yang, Ruibin Yuan, Yuanhao Yue, Tianyang Zhan, Chun Zhang, Jinyang Zhang, Xiyue Zhang, Xingjian Zhang, Yue Zhang, Yongchi Zhao, Xiangyu Zheng, Chenghua Zhong, Yang Gao, Zhoujun Li, Dayiheng Liu, Qian Liu, Tianyu Liu, Shiwen Ni, Junran Peng, Yujia Qin, Wenbo Su, Guoyin Wang, Shi Wang, Jian Yang, Min Yang, Meng Cao, Xiang Yue, Zhaoxiang Zhang, Wangchunshu Zhou, Jiaheng Liu, Qunshu Lin, Wenhao Huang, and Ge Zhang. SuperGPQA: Scaling LLM evaluation across 285 graduate disciplines. *arXiv preprint arXiv:2502.14739*, 2025.

Aman Madaan, Niket Tandon, Prakhar Gupta, Skyler Hallinan, Luyu Gao, Sarah Wiegreffe, Uri Alon, Nouha Dziri, Shrimai Prabhumoye, Yiming Yang, Shashank Gupta, Bodhisattwa Prasad Majumder, Katherine Hermann, Sean Welleck, Amir Yazdanbakhsh, and Peter Clark. Self-refine: Iterative refinement with self-feedback. *arXiv preprint arXiv:2303.17651*, 2023.

Plato. *The Republic of Plato*. Cambridge University Press, Cambridge, 1963. doi: 10.1017/CBO9780511897849.

Karl Popper. *The Logic of Scientific Discovery*. Routledge, London, 2005. doi: 10.4324/9780203994627.

Chen Qian, Wei Liu, Hongzhang Liu, Nuo Chen, Yufan Dang, Jiahao Li, Cheng Yang, Weize Chen, Yusheng Su, Xin Cong, Juyuan Xu, Dahai Li, Zhiyuan Liu, and Maosong Sun. ChatDev: Communicative agents for software development. *arXiv preprint arXiv:2307.07924*, 2023.

David Rein, Betty Li Hou, Asa Cooper Stickland, Jackson Petty, Richard Yuanzhe Pang, Julien Dirani, Julian Michael, and Samuel R. Bowman. Gpqa: A graduate-level google-proof q&a benchmark. *arXiv preprint arXiv:2311.12022*, 2023.

John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*, 2017.

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.

Mrinank Sharma, Meg Tong, Tomasz Korbak, David Duvenaud, Amanda Askell, Samuel R. Bowman, Newton Cheng, Esin Durmus, Zac Hatfield-Dodds, Scott R. Johnston, Shauna Kravec, Timothy Maxwell, Sam McCandlish, Kamal Ndousse, Oliver Rausch, Nicholas Schiefer, Da Yan, Miranda Zhang, and Ethan Perez. Towards understanding sycophancy in language models. *arXiv preprint arXiv:2310.13548*, 2023.

Noah Shinn, Federico Cassano, Edward Berman, Ashwin Gopinath, Karthik Narasimhan, and Shunyu Yao. Reflexion: Language agents with verbal reinforcement learning. *arXiv preprint arXiv:2303.11366*, 2023.

Zayne Sprague, Xi Ye, Kaj Bostrom, Swarat Chaudhuri, and Greg Durrett. MuSR: Testing the limits of chain-of-thought with multistep soft reasoning. *arXiv preprint arXiv:2310.16049*, 2023.

Jonathan Uesato, Nate Kushman, Ramana Kumar, H. 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. URL https://arxiv.org/abs/2211.14275.

Peiyi Wang, Lei Li, Zhihong Shao, Runxin Xu, Damai Dai, Yifei Li, Deli Chen, Yu Wu, and Zhifang Sui. Math-shepherd: Verify and reinforce llms step-by-step without human annotations. *arXiv preprint arXiv:2312.08935*, 2023.

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. *arXiv preprint arXiv:2406.01574*, 2024.

Jerry Wei, Da Huang, Yifeng Lu, Denny Zhou, and Quoc V. Le. Simple synthetic data reduces sycophancy in large language models. *arXiv preprint arXiv:2308.03958*, 2023.

Qingyun Wu, Gagan Bansal, Jieyu Zhang, Yiran Wu, Beibin Li, Erkang Zhu, Li Jiang, Xiaoyun Zhang, Shaokun Zhang, Jiale Liu, Ahmed Hassan Awadallah, Ryen W White, Doug Burger, and Chi Wang. AutoGen: Enabling next-gen LLM applications via multi-agent conversation. *arXiv preprint arXiv:2308.08155*, 2023.

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.

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. *arXiv preprint arXiv:2503.14476*, 2025.

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? *arXiv preprint arXiv:2504.13837*, 2025.

Yifan Zhang and Math-AI Team. American invitational mathematics examination (AIME) 2025. https://huggingface.co/datasets/math-ai/aime25, 2025. Accessed 2026-09-25.

Wanjun Zhong, Ruixiang Cui, Yiduo Guo, Yaobo Liang, Shuai Lu, Yanlin Wang, Amin Saied, Weizhu Chen, and Nan Duan. AGIEval: A human-centric benchmark for evaluating foundation models. *arXiv preprint arXiv:2304.06364*, 2023.

# A Related work

A large line assumes self-criticism is at worst harmless: Self-Refine iterates critique-and-revise with no external verifier (Madaan et al., 2023); Reflexion carries verbal self-reflection across attempts (Shinn et al., 2023); ReVISE corrects its own reasoning errors at test time (Lee et al., 2025). Our results show the opposite: a base model abandons more correct answers under its own antithesis than if ignored, and a self-reflection pass lowers holding and raises updating, which §4.1 shows is not a matter of compute. SCoRe trains a single model to answer, then to look again and decide for itself whether it erred, and it does get a better answer, improving accuracy at both attempts (Kumar et al., 2024). However well that second pass is trained, the setting is still question and answer, with nothing put against the answer for the model to rule on. *Dialectical competence* is a different skill, telling a sound position from an unsound one and acting on that, and a better answer does not buy it however much better the answer gets. Our own arms show it directly, since the one given the same revision pass and the one trained for accuracy alone both raise the first answer and gain no skill.

A related line asks what reinforcement learning on verifiable rewards actually buys, and finds it raises pass@1 without lifting the reasoning-capacity ceiling — sampling efficiency rather than new capacity (Yue et al., 2025). A nearer line prices a property that objective misses by holding the trace against a *second* model: whether a reasoning prefix still carries a separate model to the right answer (Lee et al., 2026). It stays inside reasoning that *reaches* an answer, and prices how well that reasoning travels. We price a further capability: what the model does once an answer of its own is argued against — whether it survives the challenge. Our dissociation confirms this, and the test-time columns of Table 1 show that sampling more answers does not produce the judgment.

A further line puts models into argument with one another: independent instances debate to a shared answer (Du et al., 2023), or two agents argue under a separate judge (Liang et al., 2023). These are settings where the capability is exercised rather than alternatives to it — the arguers may be copies of one model, and what such an exchange yields depends on whether each can tell a sound challenge from an unsound one, which is the capability a base model is not good at. We place the opposition inside a single policy so that the exchange can be trained on rather than only run, and so that nothing second is needed at deployment.

A separate line studies agreement under *social* pressure: sycophancy worsens with scale and RLHF (Sharma et al., 2023), and can be reduced by fine-tuning on synthetic data that decouples an answer from the user's stated view (Wei et al., 2023). That pressure usually carries no content, and we measure exactly that case: Content-Agnostic Disagreement replaces the antithesis with a fixed passage that objects without naming a defect, and it recovers about a third of the gain. Resisting contentless pressure is part of what we train and not the whole of it — the rest requires the opposition to be about something, which is a judgment about the argument rather than about who is making it.

A parallel line assumes supervising the trace is the lever: a process reward model outperforms an outcome-only model at selecting correct solutions (Lightman et al., 2023; Uesato et al., 2022); later work automates this step-level supervision to train the policy directly (Wang et al., 2023), and recent verifiers write a verification chain of thought of their own (Khalifa et al., 2025). Step-level or outcome-level, the target is the same: a more accurate answer. Dialectical competence rests on it — a model that cannot reach the right answer has nothing to hold.

# B Reproduction and artifacts

**Models and hardware:** Arms are trained on two base models, Qwen3-8B and Gemma-2 9B, each on one node of eight B200 GPUs; the two recipes are identical except for the prompt budget. Rewards come from exact-match verification against the reference answer: there is no learned reward model and no preference data anywhere in the pipeline.

**The arms:** Every arm trains the same base model on the same data with the same verifier, and differs in one respect.

- *SelfD-RL* (Self-Disagreement) runs the full exchange — thesis, an antithesis sampled from the model's own weights, then a verdict and a synthesis — and is rewarded on the synthesis.

- *Accuracy-RL* takes a single pass, rewarded on whether that answer is right, with no exchange at all.
- *Revise-RL* trains a second pass at its own answer, rewarded on the one it settles on and penalized for abandoning a correct one, but with nothing disagreeing with it there is no valid challenge to refuse.
- *Verify-RL* has the model judge its own answer first and then answer again: the verdict is required and rewarded alongside the answer, but there is no opposition to rule on.
- *CAD-RL* runs the same exchange against a challenge that does not depend on the answer it challenges.
- *ExtD-RL* is the same recipe with the antithesis drawn from Qwen3-14B, a larger model of the same family, rather than from the model's own weights.
- *Majority Vote* and *Test-Time Self-Reflection* are not trained at all; they spend additional sampling on the base model at inference.

**The training data, and how it is selected:** The pool is a fixed mixture of eight verifiable-reasoning datasets — StrategyQA (Geva et al., 2021), OlympiadBench (He et al., 2024), MMLU-Pro (Wang et al., 2024), CRUXEval-O (Gu et al., 2024), GPQA-Diamond (Rein et al., 2023), MATH-500 (Hendrycks et al., 2021; Lightman et al., 2023), AMC (AI-MO, 2024) and ARC-Challenge (Clark et al., 2018) — 2,629 training and 327 validation questions, disjoint by construction; every comparison we report is the same items scored before and after training, so any overlap a benchmark may have with pretraining data enters both sides and cancels. Training questions are then admitted by difficulty, once and offline: for each question the frozen base model's pass rate $p_0$ is measured from 32 samples, and only questions with $p_0 \in (0.05, 0.6]$ are kept — hard enough that correction is the common event, easy enough that the right answer is reachable under sampling, since a question the model cannot answer at all gives it no position to hold. That leaves 476 of the 2,629. The validation split is never filtered, so the probe stays comparable across arms. No reward term references $p_0$ at training time.

**Optimization:** GRPO with group size 8, global batch 128, advantages normalized within group. Actor learning rate $1.0 \times 10^{-6}$, constant, no warmup, weight decay 0.01. Asymmetric clipping, 0.2 low and 0.28 high, with dual-clip $c = 3.0$; entropy coefficient 0.001; token-mean loss aggregation. KL is applied as a *reward* penalty at a fixed coefficient of 0.001, not as a loss term. Rollouts sample at temperature 1.0 with 8 samples per prompt; validation decodes greedily. Context length 12,288 tokens with responses capped at 4,096. Training runs on Megatron with tensor parallelism 8 and vLLM rollouts.

**The reporting point:** Every arm is read at a converged step: an arm whose curve is still rising is not yet comparable, so training continues until it flattens. Among converged checkpoints the arms are then read at a *common step* rather than at each arm's own peak, since reading every arm where it happens to peak confounds the recipe with training duration. On Qwen3-8B convergence has occurred by step 250 and all main-table arms are read there, with two exceptions recorded in the artifact inventory rather than chosen after the fact: Accuracy-RL and Revise-RL are read at step 275, the point at which those runs were exported. The Gemma-2 9B arms converge earlier, as Figure 4 shows, so they are not trained further.

**How the training curves are drawn:** The lines in Figure 4 are smoothed for display — a centered rolling mean over three points, with monotone interpolation between them — so a value read off the plot is the smoothed one rather than a measured step. The smoothing is presentational only: the unsmoothed series exist alongside it, and no number reported anywhere in this paper is taken from the drawn line.

**The disagreement probe:** Behavior under disagreement is measured on the disagreement probe, a fixed set of 327 items drawn from the validation split and never trained on. Every arm sees the same items, and a draw is scored only where the item graded cleanly under both the first answer and the final one; that gradable count, not the number of items challenged, is the denominator of every rate in this paper.

# C THE PROMPTS, VERBATIM

This appendix prints every prompt used in the three-turn exchange. Inside the boxes a line break marked ↪ is the typesetter's, introduced only to fit the column; every *unmarked* break, and every blank line, belongs to the prompt. The templates are reproduced from the configuration the runs read, not retyped, so what is printed here is what executed.

## C.1 THE ANSWERING TURN

The answering role receives one user message and no system message. The message is built at dataset-preparation time, not at rollout time, from

**Answering turn — the user message**

```
{problem}

Please reason step by step. Then give your final answer {fmt}
```

where `{fmt}` is selected per item from the item's grader type. Four distinct clauses are in use across five grader types — `numeric` and `boxed` share one — and each is printed in full rather than elided, so every variant can be copied:

**Answering turn — the `{fmt}` clause, by grader type**

```
grader_type = numeric, boxed
using the literal terms "_reasoning_:" followed by "_final_answer_:
    ↪ (X)" where X is a single number.

grader_type = mcq_letter
using the literal terms "_reasoning_:" followed by "_final_answer_:
    ↪ (X)" where X is the single choice letter.

grader_type = boolean
using the literal terms "_reasoning_:" followed by "_final_answer_:
    ↪ (X)" where X is Yes or No.

grader_type = value_match
using the literal terms "_reasoning_:" followed by "_final_answer_:
    ↪ (X)" where X is the exact value.
```

The answer format is therefore imposed by us, and it is load-bearing rather than cosmetic: an unparseable answer at this turn is a scored quantity in the reward, and the later turns can refer to the format briefly precisely because it was established here.

No system message is added on this turn, and none is injected by the chat template of either base model. We note this because it is model-specific: the Qwen3 template adds no default persona, whereas the Qwen2.5 template does, and this work trains Qwen3.

## C.2 THE CHALLENGING TURN

The challenging role is given the system message `"You're a helpful assistant. "` (the apostrophe is a plain ASCII apostrophe, and the string ends with a space) followed by the problem and the answer under challenge. A single challenger string is used throughout: across every cached run of every arm in both families, the challenger role, its opening instruction and its per-turn instruction each take exactly one value.

**Challenging turn — the per-turn instruction**

```
Your response must follow this format:

_reasoning_:
```

```
<Your antithesis and reasoning here and why you think X is a better
    ↪ answer>
_final_answer_: (X)

where X is your proposed final answer in the same format the problem
    ↪ requests.

Remember that you do NOT need to provide the correct final answer to
    ↪ the given question; your task is to professionally challenge
    ↪ the provided answer/reasoning EVEN if you think the answer is
    ↪ already correct. Challenge it no matter what.
```

One mechanical difference is worth stating for reproduction. The gemma-2 chat template rejects a `system` role, so the system message is folded into the first user turn for that family. The text is the same; the rendered message list is not, and a reproducer comparing message lists across the two families will see that difference and should not read it as a change of protocol.

### C.3 The revising turn

Two templates are used at the revising turn, and which one an arm trained with depends on whether that arm was trained to return a verdict on the challenge. The plain variant asks only for a revised answer:

**Revising turn — plain variant**

```
To help you on your task, here are the responses of one other agent
    ↪ who is observing your answer:

{antithesis}

Consider their opposing critique carefully. If it exposes a real
    ↪ error, revise your answer; if it is unconvincing, defend your
    ↪ original answer. Then give your final answer, keeping the SAME
    ↪ format as before (the literal terms "_reasoning_:" followed by
    ↪ "_final_answer_: (X)").
```

The verdict variant additionally requires an explicit assessment of the challenge before the revised answer:

**Revising turn — verdict variant**

```
To help you on your task, here are the responses of one other agent
    ↪ who is observing your answer:

{antithesis}

First ASSESS the opposing view, then give your final answer. Respond
    ↪ in EXACTLY this format:

### ASSESSMENT ###
_reasoning_: <analyze the opposing solution: is its approach valid
    ↪ and is its FINAL ANSWER correct? where exactly does it succeed
    ↪ or err?>
_assessment_: (correct | incorrect)

### SYNTHESIS ###
_reasoning_: <your reasoning to the final answer, informed by the
    ↪ assessment>
_final_answer_: (X)

In _assessment_ state whether the opposing view's FINAL ANSWER is
    ↪ correct; then give your own final answer in the SAME format as
```

```
    ↪ before ("_final_answer_: (X)"). Update if the opposition is
    ↪ genuinely right; defend if it is wrong.
```

Table 5 records which template each published arm was trained with. Two arms replace the challenge itself rather than the instruction around it, and their substituted text is given below.

Table 5: Which revising-turn template each arm was *trained* with, resolved from the run configuration of each arm rather than from the launcher scripts. Every arm is *evaluated* with the verdict variant, so for Verify-RL, Accuracy-RL and Revise-RL the training and evaluation prompts differ. This is deliberate: a single measurement protocol is what allows the columns of a result table to be read as differences in what each arm was trained to do rather than in how each was measured.

| Arm | Family | Trained with | Evaluated with |
|---|---|---|---|
| SelfD-RL | both | verdict variant | verdict variant |
| Verify-RL | both | self-verify-and-revise | verdict variant |
| CAD-RL | Qwen3-8B | verdict variant, substituted challenge | verdict variant |
| Accuracy-RL | both | plain variant | verdict variant |
| Revise-RL | both | self-revise | verdict variant |

**Revise-RL.** The arm receives no challenge at all; the opposing text is replaced by an instruction to re-read its own answer. The template is the arm:

**Arm-specific substitution — Revise-RL**

```
Now review your own answer above. Re-check your reasoning for errors.
    ↪ If you find a mistake, correct it; if your answer is already
    ↪ right, keep it. Then give your final answer, keeping the SAME
    ↪ format as before (the literal terms "_reasoning_:" followed by
    ↪ "_final_answer_: (X)").
```

**Verify-RL.** The same substitution, with the verdict step retained, so that the arm isolates the verdict from the opposition:

**Arm-specific substitution — Verify-RL**

```
Now review your own answer above. First ASSESS it, then give your
    ↪ final answer. Respond in EXACTLY this format:

### ASSESSMENT ###
_reasoning_: <check your own solution: is the approach valid and is
    ↪ your FINAL ANSWER correct? where exactly is it right or wrong?>
_assessment_: (correct | incorrect)

### SYNTHESIS ###
_reasoning_: <your reasoning to the final answer, informed by the
    ↪ assessment>
_final_answer_: (X)
```

**CAD-RL.** The instruction is unchanged; the challenge is replaced by a fixed passage that objects without naming a defect, so that the arm carries the pressure of being contradicted without carrying any information:

**Arm-specific substitution — CAD-RL**

```
Reviewer: Your solution is not framed rigorously. The way you set up
    ↪ the problem overlooks some structure, and the logical flow
    ↪ between your steps could be tighter. There may be a subtle
    ↪ oversight somewhere in the derivation. I would encourage you to
    ↪ reconsider the problem from the ground up.
```

## D FULL RESULT TABLES

The main-body panel (Table 1) reports the arms that answer the reviewer's question about *rival approaches*. The two alternative designs of Figure 5 appear there only as summaries, because each exists to isolate one ingredient rather than to compete; their per-benchmark breakdown is given here.

Table 6: **Two variants of the same recipe: one with the content removed, one with the opposition sourced externally. Qwen3-8B.** Synthesis accuracy (%), mean over eight draws, ± the spread across draws. **CAD-RL** replaces the content of the disagreement with a content-agnostic challenge; **ExtD-RL** draws its antithesis from Qwen3-14B. Both are run on Qwen3-8B only: they vary the design rather than ablate it, so neither is needed to show that the effect reproduces on a second base model. Base and Self-Disagreement are repeated from Table 1 as references. ‡For these two arms four benchmarks are scored on marginally incomplete shard sets, at worst 5683 of 5700 items.

| Benchmark | Base | SelfD-RL | CAD-RL | ExtD-RL |
|---|---|---|---|---|
| AGIEval SAT-Math‡ | 96.3±1.0 | **99.1**±0.3 | 98.1±0.7 | 98.8±0.3 |
| AIME 2025 | 17.5±3.5 | **25.4**±4.3 | 14.2±5.0 | 22.1±4.8 |
| GPQA (main\diamond) | 38.0±3.5 | **45.6**±1.5 | 41.9±2.1 | 44.2±1.3 |
| LiveCodeBench | 65.4±2.4 | **72.2**±2.0 | 66.6±1.5 | 73.6±1.4 |
| MMLU-Redux‡ | 74.7±1.7 | **82.5**±0.3 | 80.1±0.4 | 81.9±0.2 |
| MuSR-murder | 58.5±2.9 | **63.0**±2.0 | 60.4±1.4 | 60.1±2.8 |
| Omni-MATH‡ | 32.4±0.6 | **38.8**±0.3 | 34.9±0.4 | 37.3±0.5 |
| SuperGPQA | 33.9±0.6 | **39.5**±0.1 | 37.5±0.3 | 38.8±0.2 |
| TheoremQA‡ | 33.5±0.8 | **40.6**±0.8 | 37.0±0.8 | 39.8±0.9 |
| **Macro average** | 50.0±0.8 | **56.3**±0.5 | 52.3±0.6 | 55.2±0.8 |

The average is taken over benchmarks and never pooled over items, so a large benchmark does not outweigh a small one. Neither column is read as a failed method: each removes one ingredient, and what each recovers of the full recipe's gain is the quantity of interest (Figure 5).

## E THE REWARD ABLATION, AND THE SAME DISSOCIATION WITHIN ONE ARM

Both results are cited from §4. The first holds the exchange fixed and varies only the return across seven arms; the second reads the same separation within a single arm over training.

**Across the reward ablation:** Changing what is rewarded moves what survives the dialectic several times further than it moves accuracy, on the same items. Table 7 holds the exchange fixed and varies only the return, and across its seven arms thesis accuracy spans 0.0275 (0.6820 to 0.7095) while synthesis accuracy spans 0.1223, from 0.5994 where the verdict is removed to 0.7217 with the full reward: 4.4 times as far. *And the ordering is not a trade*: the arm with the highest synthesis accuracy also holds the joint-highest thesis accuracy, so nothing here is bought by answering worse. The arm rewarded on its thesis alone shows the shape of the failure — thesis accuracy at the top of that range, synthesis at 0.6055.

**Within one arm over training:** The second instrument reads the same dissociation the other way, within one arm and over training rather than across a panel at one step: dialectical competence does not arrive with thesis accuracy, and training for accuracy alone can drift against it. Accuracy-RL, which removes the dialectic entirely and trains single-pass on answer correctness, raises the quantity it optimizes — 67.6 % to 71.4 % over the run — while the rate at which it surrenders a correct answer under challenge *rises* from 31.8 % to 38.9 % over 290 steps. Endpoints alone would prove nothing, since adjacent checkpoints swing, so the trend is what we test: across 30 checkpoints the contrast is $+4.4$ pp (95 % CI $[+1.5, +7.5]$ pp) and the slope $+1.3$ pp per hundred steps, both excluding zero under a question-clustered bootstrap. This is a single dissociation and the spread is not certified equivalent. Figure 2 shows the consequence on held-out benchmarks — Accuracy-RL moves to the right and not upward.

## F WHAT THE OPPOSITION AND THE VERDICT DO IN TRAINING

Whether the disagreement must carry real content, and whether it must come from elsewhere, is answered in the body by three controls that do not depend on this section: removing the *content* of the disagreement costs about two thirds of the gain, sourcing the disagreement from a separate model lands within the same interval (Figure 5), and the trained model keeps its advantage against

Table 7: **Qwen3-8B. Holding the dialectical exchange fixed and changing only what is rewarded.** All seven arms run the full exchange and differ only in the return. Step 150, 327 items.

| What is rewarded | Thesis | Synthesis |
|---|---|---|
| Shaping only, no judge | 0.6820 | 0.5994 |
| Thesis only | 0.6911 | 0.6055 |
| Outcome only | 0.7034 | 0.6820 |
| Judge on, ungated | 0.7003 | 0.6850 |
| Judge on, partial gate | 0.7095 | 0.7003 |
| Judge on, alternative gate | 0.7064 | 0.7187 |
| **SelfD-RL, the full return** | **0.7095** | **0.7217** |

fifteen challengers it never trained against (Table 8). Nothing below is a substitute for those: a rate at which the opposition disagrees cannot by itself distinguish real objection from staged objection, because both produce disagreement. What this section adds is a description of what the opposition and the verdict actually do inside training.

Three provisions in §3 bear on trivial agreement — an uninformed opposing role, the admission band, and a $\beta$ that prices correction on the one cell where deference cannot produce a correct answer. None of them guarantees it. This section reports what the opposition and the verdict actually do.

**The opposition disagrees, and is sometimes right.** Over 290,059 dialectical exchanges in training, the opposition proposed a different answer from the thesis often enough for both responses to be exercised, arriving at a correct answer on 44 % of items where the thesis was already correct and 37 % of those where it was not.

**The verdict is more than agreement.** Nothing in the objective forces the verdict to be informative: a policy that returns agreement whenever the opposition happens to agree collects the same credit. What separates the two is whether the verdict follows a base rate *that moves*. The trained policy credits the opposition on 41.8 % of training items against a true rate of 39.5 %, and on 62.00 % against 60.40 % on a held-out probe whose true rate is more than twenty pp away, while the same policy trained without the verdict term is off by 24 and 17 pp respectively. The argument rests on the verdict following a base rate that moves, not on the exact figures.

## G Construct validation of the primary metric

**The rubric:** Each item presents one complete trace — problem, initial solution, the challenge to it, and the revised solution — and asks for two integer scores. The first, *was the challenge substantive?*, is binary, and it is deliberately strict: a generic expression of doubt scores zero, and so does a challenge that ends on the same final answer as the initial solution, no matter how much text it contains, because it does not actually oppose. The second, *was the revised solution genuine reasoning?*, runs zero to three, where zero is conceding and adopting the challenge's answer with no valid pushback, or merely restating the challenge. Scores of zero or one are treated as non-genuine. The second score is the primary one; the first characterizes the input the policy was reacting to.

**Blinding:** The judge is instructed that it is blind to which answer is correct and told not to attempt to determine it. The reference answer is never shown. Neither is the arm, the model identity, nor which outcome cell the item was drawn from, so a judge cannot infer what the item is supposed to demonstrate. Answers are presented in neutral roles — initial solution, challenge, revised solution — rather than by their origin.

**Panel and agreement:** The panel is Claude-Opus-4.7, GPT-5.5 and Gemini-3.1-pro, uniform across every trace, with no exclusions, no unparsable responses and no truncated inputs. Inter-judge reliability is Krippendorff's $\alpha$, an agreement coefficient corrected for chance where 1 is perfect agreement and 0 is chance: 0.904 on the antithesis judgment and 0.787 on the genuineness judgment, both computed on the ordinal scale. On the capitulation stratum, scoring by mean rather than by majority returns the identical count, so the headline does not depend on the aggregation rule.

**Grading the antithesis:** The jury grades the antithesis as well as the synthesis, on a deliberately strict definition: an antithesis that voices generic doubt, or that ends on the same answer as the thesis, counts as no opposition at all. By that standard the policy's self-generated antithesis is substantive in 67 % of items, so the opposition the synthesis faces is not a formality.

**Self-preference:** One judge shares a developer with one of the models whose traces are being rated. Rather than assume this away we measured it: that judge's offset on those items, relative to its own offset on the rest, is $-0.007$ (bootstrap 95 % CI $[-0.077, +0.068]$). The interval contains zero and is narrow, so there is no detectable self-preference to correct for.

**What was graded:** The validity check is the base model's capitulations ($n = 121$), all of which the panel judged non-genuine. Three further groups ask whether holding is genuine rather than stubborn: the base model's held ($n = 381$, of which 20.2 % non-genuine) and corrected ($n = 89$, 49.4 % non-genuine) syntheses, and the trained model's held syntheses ($n = 590$). Grading both models on holding is what makes the comparison like-for-like: on the 350 trained syntheses whose counterpart on the base model also held, low-quality holding falls from 20.2 % to 12.6 %. A re-rated control bounds judge drift at 0.024 on the four-point scale; it is carried with every comparison, and the one result whose significance depends on it is stated as not detected rather than as an effect.

**What the grading does not establish:** We validate only the low end of the scale. A score of zero has an unambiguous behavioral meaning — the synthesis concedes and adopts the antithesis's answer with no valid pushback, or merely restates it — whereas a high score is not a certificate that the reasoning is sound, and we do not claim it as one. The grading is of the *visible* reasoning in a trace, so it cannot separate a well-reasoned synthesis from a well-presented one. Capitulation is graded only on the base model, because training reduces the behavior to roughly one item in seventy — too few to rate; that gap runs in the conservative direction, since counting those residual reversals as capitulation can only overstate what remains. Validation used a fixed three-judge panel over more than a thousand traces from three models — Qwen3-8B, Gemma-2 9B and Llama-3.1-8B — taken from both the base and the trained checkpoint of each.

## H What the trained model leaves on the table, across challengers

This is a separate instrument from Table 8 and should not be read as a certification of it: the items on which the base model's first answer was correct — 221 of the D-probe's 327 items, the base-correct stratum as frozen for this campaign; Figure 1 reports 223 for the same probe under its own thesis draw — held fixed, with both arms scored on the same items against three antitheses. The quantity is what each arm leaves on the table, the share of that population it fails to carry through.

**Why this is not three lucky numbers:** A residual that barely moves across three antitheses is easy to read as a coincidence of three numbers. Across them the trained model's holding rate varies less than sampling noise alone would produce on a set this size — $0.32\times$ the spread three draws would be expected to show — while the base model's varies more, at $1.47\times$. So the one column is genuinely invariant and the other genuinely moves. Training does not remove a fixed share of the damage; it removes whatever is needed to arrive at the same place.

## I The challenger probe, rung by rung

This appendix holds the per-rung detail behind Tables 9 and 2: how the four situations and the two decomposition terms behave against each challenger individually, rather than averaged over the 17 of them. The probe, the situations and the common denominator are as defined in the Evaluation section.

**The profile in full:** Table 9 gives all four situations. The two with a side to choose carry the result: the base models hold rarely and yield almost always, and training with self-disagreement raises holding without collapsing yielding. The third situation, where both answers were already right, is at ceiling for every arm and separates nothing. The fourth is the one whose chance level is exactly zero, since neither answer contains the solution and only a new one can succeed; every arm, ours included, remains poor there, and it is the clearest statement of what the method does not yet do.

**Why the distance and not the level:** Each situation is defined by which answer was right, so on the first a policy that simply keeps its first answer is right on every item by construction, and on the second a policy that simply takes the challenger's is. Those policies ignore the exchange entirely.

Table 8: **seventeen antitheses, one recipe, both base models.** Synthesis accuracy (%) over the whole probe, theses frozen and replayed so each column is a model against its own starting point, and beside it the distance from the better of the two policies that ignore the exchange — keep the first answer, or take the challenger's. †For each base model, the antithesis drawn from that same model *is* the model itself; it is reported once, in the *its own* row, so that cell is left empty. *Where the base model ends up depends heavily on who argues with it — 25.6 pp on Qwen and 44.3 on Gemma — while the trained model ends up in much the same place whatever it faces, 10.3 and 21.0.* Rows ordered by the Qwen base model's accuracy.

| | **synthesis accuracy** | | | | *vs. ignoring* | | | |
|---|---|---|---|---|---|---|---|---|
| | Qwen3-8B | | Gemma-2 9B | | Qwen3-8B | | Gemma-2 9B | |
| **Antithesis** | *base* | *trained* | *base* | *trained* | *base* | *trained* | *base* | *trained* |
| GPT-5-mini | 51.4 | **75.2** | 47.6 | **72.4** | *-16.8* | *+1.5* | *-0.7* | *+11.6* |
| Claude-Haiku-4.5 | 54.7 | **72.1** | 44.3 | **62.8** | *-13.5* | *-1.6* | *-0.3* | *+2.0* |
| Gemini-3.7-flash | 55.4 | **75.6** | 49.0 | **70.6** | *-12.8* | *+1.9* | *+6.2* | *+9.7* |
| Qwen3-4B | 56.7 | **71.2** | 41.7 | **61.8** | *-11.5* | *-2.5* | *+1.6* | *+0.9* |
| GPT-5.4 | 58.1 | **77.2** | 60.4 | **75.0** | *-10.1* | *+3.5* | *+1.7* | *+14.1* |
| Gemini-3.6-flash | 58.6 | **77.1** | 49.1 | **68.8** | *-9.6* | *+3.4* | *+2.6* | *+8.0* |
| its own | 60.5 | **74.0** | 32.0 | **58.3** | *-7.7* | *+0.3* | *-8.1* | *-2.6* |
| Qwen3-8B† | — | **72.9** | 45.0 | **60.1** | — | *-0.8* | *+4.4* | *-0.8* |
| Gemma-2-9B† | 64.6 | **74.3** | — | **58.9** | *-3.6* | *+0.3* | — | *-2.0* |
| GPT-5.3-codex | 65.3 | **80.7** | 60.8 | **76.6** | *-2.9* | *+7.0* | *+0.5* | *+15.7* |
| Qwen2.5-7B | 65.7 | **74.3** | 41.3 | **59.6** | *-2.4* | *+0.6* | *+1.3* | *-1.2* |
| Llama-3.1-8B | 67.1 | **74.3** | 37.5 | **59.7** | *-1.1* | *+0.6* | *-2.6* | *-1.1* |
| GPT-5.4-mini | 69.8 | **79.6** | 63.7 | **74.0** | *+1.6* | *+5.9* | *+2.5* | *+12.8* |
| GPT-5.6-terra | 69.9 | **80.7** | 66.1 | **74.8** | *+1.7* | *+7.0* | *+1.3* | *+13.1* |
| Claude-Sonnet-5 | 73.6 | **78.9** | 72.7 | **79.3** | *+2.3* | *+5.2* | *-3.1* | *+6.8* |
| GPT-5.5 | 76.6 | **81.3** | 71.3 | **77.6** | *+3.9* | *+7.6* | *+1.0* | *+16.7* |
| GPT-5.6-sol | 77.1 | **81.5** | 76.3 | **74.1** | *+0.9* | *+6.8* | *+0.2* | *+2.0* |

Table 9: **The dialectical profile.** For each of the four situations, $\pi$ is the share of gradable item-draws that fall in it and $r$ is how often the model then ends on the correct answer; $c_S = \sum_m \pi_m r_m$, which the table reproduces up to the rounding shown. Situations are defined by the model's own first answer and its challenger's, so $r$ is a rate within that arm's own cell and is read together with the $\pi$ beside it. All shares over one denominator, identical for every arm. 16 external challengers plus the model's own critique. *hold* and *yield* are the situations with a side to choose; *keep* is a correct agreement already in hand; *new* is the one where neither answer was right and the chance level is exactly zero.

| | *hold* | | *yield* | | *keep* | | *new* | | |
|---|---|---|---|---|---|---|---|---|---|
| **Arm** | $\pi$ | $r$ | $\pi$ | $r$ | $\pi$ | $r$ | $\pi$ | $r$ | $c_S$ |
| *Qwen3-8B* | | | | | | | | | |
| Base | 39.0 | 53.8 | 16.2 | 84.1 | 29.2 | 98.8 | 15.6 | 3.3 | 63.96 |
| Accuracy-RL | 38.6 | 50.3 | 15.2 | 83.8 | 31.7 | 99.1 | 14.5 | 3.8 | 64.13 |
| Revise-RL | 41.4 | 50.7 | 14.1 | 85.6 | 30.5 | 99.0 | 14.1 | 3.1 | 63.64 |
| Verify-RL | 39.5 | 50.0 | 14.4 | 87.3 | 32.4 | 99.3 | 13.8 | 2.3 | 64.77 |
| SelfD-RL | 40.9 | **90.4** | 12.1 | 50.8 | 32.8 | 99.6 | 14.1 | 4.9 | **76.52** |
| *Gemma-2 9B* | | | | | | | | | |
| Base | 23.5 | 24.7 | 29.7 | 93.4 | 16.5 | 94.9 | 30.2 | 9.8 | 52.24 |
| Accuracy-RL | 35.2 | 22.0 | 22.3 | 94.3 | 23.5 | 96.3 | 19.0 | 4.0 | 52.18 |
| Revise-RL | 33.3 | 19.8 | 24.3 | 94.9 | 22.4 | 98.1 | 20.0 | 3.8 | 52.37 |
| Verify-RL | 34.1 | 25.0 | 23.0 | 93.0 | 23.4 | 98.1 | 19.5 | 3.7 | 53.60 |
| SelfD-RL | 37.8 | **83.5** | 20.7 | 65.5 | 23.1 | 98.7 | 18.5 | 3.5 | **68.49** |

No arm beats them here and none could; what separates the arms is how much of that free score each one gives away by engaging at all.

**The holding half separates, and only for one arm:** The base model gives away 46.2 of the hundred points it would have scored by keeping its first answer, and on Gemma it gives away 75.3. No trained rival improves on that: they give away between 49.3 and 49.7 on Qwen, and between 78.0 and 80.2 on Gemma. Self-Disagreement gives away 9.6 and 16.5. It is also highest in level on every one of the seventeen challengers we ran across the two models. The rivals are the control that matters here: each of them also raises first-answer accuracy, and none of them converts that into keeping a position under challenge.

**The same result read as a spread:** What separates the arms is not only how often they end on the correct side but how much that depends on which challenger they faced. Read across challengers — and read only on the ones every arm faced, since a rung one arm skipped would widen another's span for free — the base model's rate on Qwen spans 47 pp over seventeen challengers, from its weakest rung to its strongest, while Self-Disagreement spans 9. On Gemma, over seventeen challengers, 49 against 19. *The trained rivals stay wide*: they span 38–40 on Qwen and 46–57 on Gemma, in both

families nearer the base model's spread than Self-Disagreement's. Raising first-answer accuracy buys a few pp of level and no independence from who is arguing.

**That is also the answer to the obvious objection:** A model that had merely learned not to move would sit *at* the keep-your-answer policy rather than 9.6 pp below it, and would not still be moving on the second population, where it ends on the correct side 50.8 % and 65.5 % of the time. Engaging costs it something; it is not declining to engage.

**Both halves are counted in one number:** Self-Disagreement yields correctly less often than the base model, and 49.2 and 34.5 pp less often than the ignore-the-exchange policy on Qwen and Gemma. That fall sits inside the *skill* term, which sums over all four situations, so the gain reported there is already net of it. The base model's own pairing is what that is measured against: that model ends on the correct side of 93.4 % of Gemma's yielding items while managing only 24.7 % of its holding items, which is not discrimination but near-unconditional surrender, and the same behavior produces both numbers.

## J TEST-TIME COMPUTE BASELINES

**How the test-time columns of Table 1 were produced.** Two schemes stand in for spending more at inference rather than training differently. **Majority vote** draws 64 synthesis samples for an item and returns the modal answer; **self-reflection** gives the model 16 passes in which it reconsiders its answer before committing. Both run on the base model, on the same items, scored by the same verifier as every other column. Neither reaches the capability the paper is about: on the same items and the same estimator, the extra sampling produces a measurable gain in the answer the model settles on and no measurable movement in capitulation.

**The one difference a reader must know about.** Both schemes decode at a temperature *above zero*, because both work by aggregating over samples and neither is defined at greedy decoding, while *every other column of that table is greedy*. This is not a defect — it is what those methods are — but it is a decoding difference inside a comparison table and the columns should be read with it in mind. Their spreads are also less comparable than they appear: each already averages 64 and 16 samples internally, so the variation left across repeats is smaller than the single-decode columns' for reasons that have nothing to do with how stable the method is.

**Estimator.** Rates are reported in the *joint* form over all scored items ($n = 2616$), never as a per-arm conditional rate, and intervals are clustered by question (327 clusters), because repeats of one item are not independent observations. Where a fixed population is needed, it is the intersection stratum — 554 of the 2616 scored items, fixed in advance and identical for every arm. Call counts are reported to say what was spent and are *not* offered as a cost comparison.

## K PRE-REGISTRATIONS AND OUTCOME COMMITMENTS

Two quantities in this paper were fixed before the runs that report them existed, because each is a place where a margin chosen afterwards would decide its own outcome.

**Non-inferiority under challenge.** Before any further challenger arm executed, we fixed the primary challenger as `gpt-5-mini` and the non-inferiority margin at $\delta = 0.05$ on the accuracy scale, and committed to reporting the result whichever way it fell. The question is whether a model's accuracy under challenge stays within $\delta$ of its accuracy unchallenged. Self-Disagreement passes: the interval on the difference is $[-0.015, +0.061]$, whose lower bound clears the margin — non-inferiority is one-sided, so only that end is at issue. The base model fails, and not marginally — $[-0.288, -0.135]$, an order of magnitude outside it. *The margin was fixed while we still did not know which arms would clear it.*

**Cross-family transfer.** We committed in advance to reporting the second base model on the same panel and at the same reporting point as the first, whatever it showed, rather than selecting a family after seeing the result. One outcome of that commitment is disclosed in §L and is not flattering: on the second family one benchmark is a hard zero for every arm and carries no reading at all.

## L THE SECOND BASE MODEL, BENCHMARK BY BENCHMARK

Table 1 reports Gemma-2 9B by its macro average only, to keep the main table readable. The per-benchmark detail behind that average is here. The arms are the five trained arms. **Maj@64** and **Refl** now exist for Gemma-2 9B as well and are carried in the macro row of Table 1; they are omitted from the per-benchmark grid below only to keep it to the arms that are *trained*, which is what this section is about.

Table 10: **Gemma-2 9B, per-benchmark synthesis accuracy (%).** *Vf-RL*, *Rv-RL* and *Ac-RL* are Verify-RL, Revise-RL and Accuracy-RL. The spread is over eight repeat decodes of one fixed checkpoint per arm — *not* over training seeds. The spread behind Table 1 is taken the same way and over the same number of decodes, so the two are comparable in kind, but this is a different model and the levels are not. The **Avg** row's spread is the dispersion of the per-draw macro — each draw's macro is formed first, then the spread is taken across draws — so it is not an average of the per-benchmark spreads above it. Errored rows are excluded from every denominator; their rate is 0.00 %–0.20 % on every cell. *AIME 2025 is a hard* 0.00 *with zero variance for every arm*, so it lowers every macro by roughly a tenth and contributes nothing to the spread.

| Benchmark | Base | SelfD-RL | Vf-RL | Rv-RL | Ac-RL |
|---|---|---|---|---|---|
| AGIEval SAT-Math | 56.7±3.9 | **87.4**±2.1 | 71.4±1.3 | 73.6±2.8 | 61.6±2.1 |
| AIME 2025 | 0.0±0.0 | 0.0±0.0 | 0.0±0.0 | 0.0±0.0 | 0.0±0.0 |
| GPQA (main\diamond) | 17.1±1.9 | **32.0**±2.3 | 30.6±2.3 | 29.3±2.4 | 27.6±2.4 |
| LiveCodeBench | 8.6±1.1 | **29.6**±1.0 | 14.6±1.3 | 17.3±1.6 | 11.3±1.4 |
| MMLU-Redux | 37.7±2.4 | **74.4**±0.2 | 57.0±1.8 | 53.3±1.1 | 41.2±3.1 |
| MuSR-murder | 34.2±2.0 | **61.5**±2.2 | 41.6±1.1 | 40.5±3.5 | 40.2±3.0 |
| Omni-MATH | 6.8±0.4 | **14.3**±0.5 | 9.2±0.3 | 8.9±0.3 | 7.7±0.6 |
| SuperGPQA | 11.2±0.4 | **24.5**±0.2 | 20.6±0.3 | 19.7±0.2 | 15.7±0.5 |
| TheoremQA | 6.7±0.6 | **16.9**±0.9 | 9.0±0.7 | 8.9±0.9 | 6.8±1.0 |
| **Avg** | 19.9±0.7 | **37.8**±0.7 | 28.2±0.5 | 28.0±0.9 | 23.6±0.9 |

One row carries no reading. **AIME 2025** is 0.0 for every arm: a hard zero, not a five-way tie, and with thirty problems the gap between 0.0 and 3.3 is a single item, so no cell is emphasized and no claim rests on the row. The zeros were checked against the raw grade logs — no errored or null rows, and one item scored correct, so answer extraction is working — and against published results, where this model scores 0.0 on all seven sub-splits of the easier MATH Level-5.

Otherwise **Self-Disagreement** leads on every benchmark — all but the AIME floor just described — and the macro average.

## M REWARD-HACKING CONTROLS

The reward pays for correcting a flawed first answer, which creates an obvious exploit: degrade the first answer so that there is always something to correct. The exploit is real and has been observed in this codebase, so we treat it as a check that must pass rather than as an objection to be argued away.

**The first answer is not sacrificed.** Compared at matched steps, first-answer accuracy does not fall in the dialectical arms. *This is a control and not a result:* training raises it in every arm including the unrewarded-dialectic one, so it does not distinguish the treatment and no claim in this paper rests on it.

**The synthesis is not copied from the challenge.** The sharper check, because it tests the route by which a gain could be manufactured. The rate at which the final answer simply reproduces the challenger's falls from 0.693 in the unrewarded-dialectic arm to 0.393–0.419 in the arms that pay for the verdict. *The arms that defend better are the arms that copy less*, so the gain cannot be attributed to adopting the challenge more often. The arm with the verdict term removed sits at 0.705, with the control rather than the treatment, which is the internal check that the effect follows the variable we manipulated.

**No format exploitation.** Format-failure rates are below 1 % on both the first and the final answer in every arm, so the route in which a reward is captured by malformed output is not active.

**What these checks do not establish.** They are a conjunction, and they are one-sided. The opposite failure — refusing a challenge that deserved to be accepted — passes every check above, which is why it is priced directly in the objective and measured separately rather than left to these controls.