CONFORMITY BREAKS CONFORMAL PREDICTION

Anonymous authors
Paper under double-blind review

ABSTRACT

Multi-agent LLM systems often rely on models that answer after seeing peer messages. Conformal prediction can help decide when an answer is reliable enough to act on, but calibration is usually done with the model answering alone. We identify and formalize a *score-mechanism shift*: the question stays the same, but peer messages change how the model scores its answers. We show that this one shift causes failures at three levels. First, peer pressure breaks overall coverage. Across four open-weight LLMs in our core experiments, coverage falls from a calibrated 90% to 74% under unanimous-wrong peer messages. More importantly, the average hides a much larger failure. On questions the model is already less sure about, coverage falls from 87% to 47%. An attacker can also identify such weak cases without knowing the correct answers, using only outputs from the target model or predictions from another model. Finally, the problem changes what the system actually does. A policy designed to defer uncertain cases instead acts on the wrong answer in 12% of cases where peer pressure flips the model's top choice. The coverage failure also holds across additional models and seven task types. Recalibrating on data with peer messages restores coverage, but makes the system defer more often. More broadly, conformal guarantees for multi-agent LLMs must account for the peer messages a model sees when it answers.

1 INTRODUCTION

Multi-agent LLM systems let models see and respond to one another's answers before a final decision is made, as in debate and agent-based systems (Du et al., 2024; Liang et al., 2024; Wu et al., 2024). Because these systems may act without a human in the loop, they need a reliable way to decide when to act and when to defer. Conformal prediction provides such a coverage guarantee (Vovk et al., 2005; Angelopoulos & Bates, 2021) and has been carried to LLMs (Kumar et al., 2023; Quach et al., 2024; Mohri & Hashimoto, 2024): with probability $1 - \alpha$ the correct answer lies in a prediction set (e.g., 90% coverage at $\alpha = 0.10$). Such guarantees can determine whether a system proceeds, calls a tool, or escalates for review (Angelopoulos et al., 2024; Abbasi Yadkori et al., 2024; Ren et al., 2023). If that guarantee fails, a decision rule built on it can fail as well.

Conformal validity requires calibration and test examples to be scored under comparable conditions. In current LLM uses, calibration typically scores the model answering on its own. A multi-agent deployment is different: the same model may answer after seeing other agents' responses. Figure 1 shows the consequence. The question and threshold stay fixed, but unanimous-wrong peers move probability away from the correct answer until it leaves the prediction set.

Prior work establishes two relevant pieces. Studies of LLM conformity show that models can abandon a correct answer when peers give a wrong one (Asch, 1956; Weng et al., 2025; Qu et al., 2026; Hu & Qu, 2026), but they study behavior or accuracy rather than conformal coverage. The other brings conformal prediction into multi-agent debate: Wang et al. (2026) gate an agent's output on the prediction set, acting when it is a singleton and escalating otherwise. But their guarantee is marginal and their setting assumes no attacker. We ask whether the coverage guarantee, and the decision rule built on it, survive peer messages written by an attacker.

The question does not change; the model's scores do. We call this a *score-mechanism shift*. Because the shift is induced through the peer messages rather than the user question, calibration on solo answers need not describe how the model behaves once it sees peer messages.

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

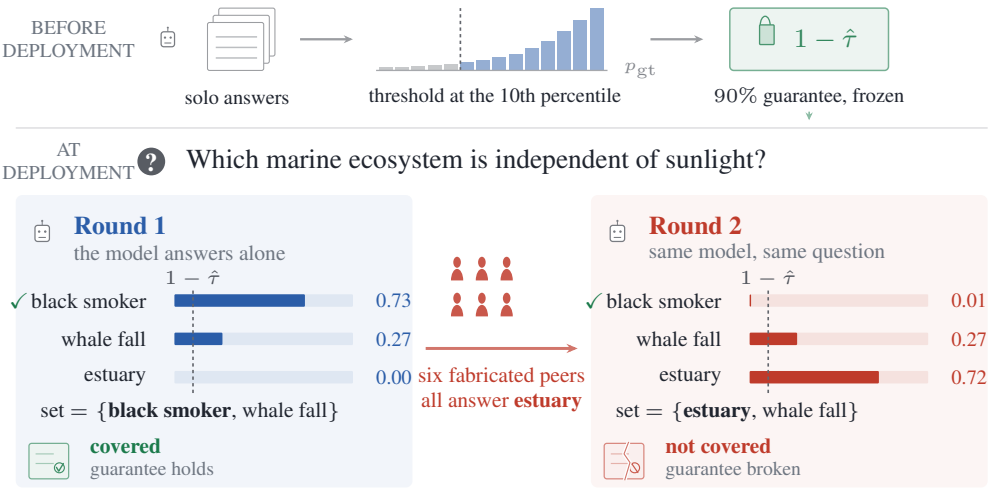

Figure 1: Peer pressure breaks a conformal coverage guarantee calibrated on solo answers. A 90% guarantee is calibrated and frozen before deployment. On the same question, unanimous-wrong peers shift probability mass toward their answer: the true answer leaves the prediction set and the peers' answer enters. Values are measured for Qwen2.5-7B on ARC-Challenge (seed 0); across four models and two benchmarks, coverage falls from 90% to 73.9%.

Contributions. We study one mechanism and three consequences:

1. **Peer influence breaks coverage calibrated on solo answers.** We formalize the score-mechanism shift (Definition 1) and show that coverage falls from 90% to 74% under unanimous-wrong peers. Weighted conformal prediction and per-model calibration do not remove the failure.
2. **The average hides a much larger failure.** Coverage falls most sharply on queries close to the calibrated threshold. An attacker who never sees the answer key can identify similarly fragile queries from the system's own outputs and cut coverage on the selected 10% by 40 points, while moving overall coverage by only 4 points.
3. **The failure reaches the decision layer.** A conformal act-vs-escalate policy commits the attacker's wrong answer on 12% of the queries whose top answer peer pressure flips, rising to 71% for the largest open-weight model we test.

These findings extend beyond the core setting: the coverage collapse holds across more models and task types, with a different conformal set construction, and with verbalized confidence scores.

2 Setup

Conformal prediction. For an LLM used as a multiple-choice QA system, let $p[y]$ be the probability the model assigns to candidate answer y, read from its next-token distribution over the option labels. We use the *nonconformity* score $s(y) = 1 - p[y]$: answers with higher probability have lower scores. The *prediction set* keeps every answer whose score is at most a calibrated threshold $\hat{\tau}$:

$$C(x) = \{y : s(y) \leq \hat{\tau}\} = \{y : p[y] \geq 1 - \hat{\tau}\}.$$

Given n calibration items, $\hat{\tau}$ is the $\lceil (n+1)(1-\alpha) \rceil$-th smallest of their correct-answer scores, so the correct answer is included at least a $(1-\alpha)$ fraction of the time. This guarantee is *marginal*: it holds on average over queries. *Conditional* coverage asks for the same guarantee within a subgroup of queries, which marginal coverage does not provide.

This is split conformal prediction (Vovk et al., 2005; Angelopoulos & Bates, 2021) with the LAC (least-ambiguous-classifier) score (Sadinle et al., 2019); for LAC, the score is exactly $s = 1 - p$ as above. Its validity requires calibration and test examples to be exchangeable, meaning the two sets are scored by the same mechanism.

For calibration, we use the model's solo-answer scores, $s_i = 1 - p_{\mathrm{gt}}^{(\mathrm{r1})}$, where p_{gt} is the probability assigned to the correct answer and $\mathrm{r1}$ denotes the solo round. We report thresholds on the probability scale, as the cutoff $1 - \hat{\tau}$ that an answer must meet.

What the input x is. Throughout, x denotes the user question, not the full prompt. The peer messages are not part of x; they change the context in which the model scores the question.

LLM conformity. Following prior work (Asch, 1956; Weng et al., 2025; Qu et al., 2026; Hu & Qu, 2026), we use LLM conformity to describe the observable tendency of a model's answer to move toward the answer the peer messages give. Our analysis concerns the resulting change in answer scores and its effect on conformal coverage. The *answer-flip rate* is the fraction of items on which the model abandons its correct solo answer under unanimous-wrong peers.

Experimental protocol. Figure 1 also illustrates the basic experiment. For each (model, question) pair, the model first answers the question alone (Round 1). We then show it the same question together with six one-line peer messages and ask it to answer again (Round 2), recording the full answer distribution in both rounds.

Round 2 uses three peer patterns: all peers give the same wrong answer, the peers are mixed, or all peers give the correct answer. Each is also run with an authority cue in which one peer is given an institutional role, giving six conditions in total. The wrong answer is one option drawn uniformly at random from the options other than the correct one, and all six peers name it. Full transcript templates and sampling details are in Appendix A.

Models and calibration. Our core experiments use four open models: Qwen2.5-7B (Qwen Team, 2024), Llama-3.1-8B (Grattafiori et al., 2024), Mistral-7B-v0.3 (Jiang et al., 2023), and Gemma-2-9B (Gemma Team, 2024). We evaluate on ARC-Challenge (Clark et al., 2018) and TruthfulQA (Lin et al., 2022). Section 5 expands to additional models and task types. For each random 50/50 calibration/test split, we compute the threshold from solo answers in the calibration half and evaluate it on the separate test half after the peer messages are introduced. Unless stated otherwise, $\alpha = 0.10$ and results average over 2,000 splits. Seed 0 carries the main results; seeds 1 and 2 repeat them on a larger question set with the peer messages resampled. Full details are in Appendix A.

Threat model. We study an attacker that controls the peer messages the target model reads on the shared channel. A compromised relay or message bus gives a single agent that capability, and several compromised agents give it jointly; we do not model how it is obtained, and an end-to-end exploit inside a running framework is left to future work. The attacker has no access to the target's weights or calibration data, and it acts at inference time, after the conformal threshold has been computed offline and frozen. The unanimous-wrong condition is a stress test of what happens when all six peer messages support an incorrect option. When deciding which queries to target, the attacker uses only information available from the deployed system or an outside model; it does not use the correct answers.

3 Main Results

We report three failures of increasing consequence: peer pressure first breaks overall coverage, then produces a much larger failure on selected queries, and finally reaches the act-vs-escalate decision.

3.1 Unanimous-wrong peers break marginal coverage

Using thresholds calibrated on solo answers, coverage stays at or above target in every condition except unanimous wrong. At $\alpha = 0.10$, coverage drops from 90.0% to **73.9%**$\pm 0.3\%$ under unanimous-wrong peers. Mixed and unanimous-correct peers remain above target (Table 1). The three columns differ in what is shared: one threshold calibrated on all four models' solo scores, one threshold per model, or a four-model average formed by averaging the models' answer probabilities. The collapse appears in all three.

Table 1: Unanimous-wrong peers break marginal coverage under all three ways of aggregating the four models; mixed and unanimous-correct peers leave it above target. Three seeds (common subset: 4 models × ARC + TruthfulQA), $\alpha = 0.10$, no authority cue; mean ± std over seeds. The $\alpha = 0.05$ results, the authority-cue variants and the per-seed values are in Appendix B.

Condition	Shared threshold	Per-model	Four-model average
Round 1 (no pressure)	$90.0\% \pm 0.0$	$90.1\% \pm 0.0$	$90.0\% \pm 0.0$
Unanimous wrong	**73.9%**± 0.3	**73.3%**± 0.9	**75.6%**± 0.7
Mixed	$94.2\% \pm 0.2$	$95.5\% \pm 0.2$	$93.5\% \pm 1.3$
Unanimous correct	$94.1\% \pm 0.5$	$96.8\% \pm 0.2$	$100.0\% \pm 0.0$

Raising the coverage target narrows the gap, but only by loosening the threshold until the prediction sets are nearly vacuous, and at $\alpha = 0.01$ it disappears. The coverage gap also grows smoothly when we interpolate between the solo and unanimous-wrong score distributions (Appendix D).

3.2 A HIDDEN COLLAPSE JUST ABOVE THE CALIBRATED THRESHOLD

Coverage is reported as an average over all queries, so a loss concentrated on a few of them barely moves it. At $\alpha = 0.10$ the calibrated threshold sits by construction at the tenth percentile of solo p_{gt}, and almost nothing below it is covered even without pressure. The queries most at risk are therefore the next decile up, the tenth to twentieth percentile of p_{gt}: they sit just above it, so they need the smallest downward push to lose coverage, and no decile higher up loses more. We use this decile to expose the conditional failure most directly.

For this analysis we use the four-model average and pressure every test query.

Marginal coverage for that four-model average falls from 90.1% to 75.7% (seed 0, with the authority cue), but coverage on that decile falls much further, from 87.3% to **47.4%** (Figure 2a). The collapse replicates on the larger question sets of seeds 1 and 2. Without pressure, coverage on this decile is already within three points of the 90% target, so most of the drop to 47.4% is caused by peer pressure rather than the known limitation of distribution-free conditional coverage (Barber et al., 2021).

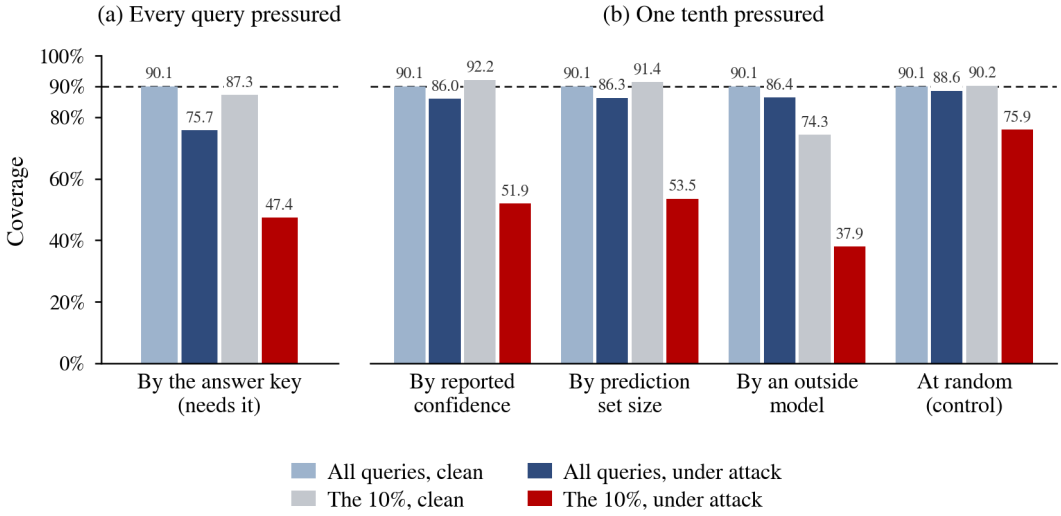

Figure 2: Each group is 10% of the queries, and the label says how that 10% was picked. **(a) Every query pressured:** the decile just above the threshold, the one group that takes the answer key to draw, falls from 87.3% to 47.4% while marginal coverage falls to 75.7%. **(b) One tenth pressured:** three selectors that never see the correct answers take 36 to 40 points off the queries their own signal picks, against 14.3 for a random 10%, while marginal coverage stays near 86%. Four-model average, seed 0, authority cue.

Drawing that decile takes the answer key, so an attacker cannot select it. Signals the deployment already exposes come close. Pressuring only the 10% of queries the model reports least confidence on, and leaving the other 90% unpressured, drops coverage on them from 92.2% to 51.9% ($\mathbf{-40.2}$ pp); ranking instead by prediction-set size, or by a model outside the system (Qwen2.5-32B), costs the selected queries 38 and 36 points (Figure 2b). A random 10% loses 14.3. These selections overlap the decile above the threshold without coinciding with it: each meets it on 32 to 45% of its items, against the 10% a random selection meets. Because only pressured queries can lose coverage, marginal coverage stays near 86% in all three cases, against 90.1% unattacked. These selectors identify many of the vulnerable queries without access to the answer key. This is the conditional case Wang et al. (2026) defer.

The collapse in Figure 2a holds for every model individually, from -10 pp for Mistral-7B to -51 pp for Qwen-7B. For Gemma-9B marginal coverage gives no warning at all: it does not move while that decile falls from 86.6% to 70.6%. The average falls further than three of the four models it averages: the four see the same fabricated transcript, so their scores move the same way and averaging does not cancel them. The wider bottom-20% group mixes these items with ones that were never covered, so it understates the attack (50.3% to 35.5%). Per-model numbers, the bottom-$k\%$ groups, and a breakdown by confidence stratum are in Appendix C.

3.3 The act-vs-escalate policy can commit the wrong answer

We reproduce the policy of Wang et al. (2026). We average the four models' answer probabilities, form LAC prediction sets calibrated on solo answers, and act when the set contains one answer; otherwise the policy escalates to a human.

We focus on queries where peer pressure changes the top answer from correct to wrong. The *act-on-wrong* rate is the fraction of those queries for which the prediction set contains only that wrong answer. In those cases the policy commits the wrong answer instead of escalating. We report it two ways: over the flipped queries, which measures the policy alone, and over every pressured query, which also counts how often pressure flips an answer in the first place.

For the four-model average the act-on-wrong rate is **12.1%**, and single models span more than two orders of magnitude (Table 2): Qwen2.5-32B, the largest open model we test, reaches **71.4%**. It is also the most confident model before peer pressure is applied, showing that higher solo confidence does not imply greater robustness. Flips are common enough that the second reading stays high: Qwen2.5-32B commits a wrong answer on 24.8% of all pressured queries, and the four-model average on 5.1%. The same conclusion holds across seeds, without the authority cue, and with randomized APS sets (Romano et al., 2020); full comparisons are in Appendix J.

Table 2: The act-vs-escalate defense fails. Act-on-wrong by unit. On a flipped query the remainder escalates in every case, since a singleton read off the pressured distribution always holds the wrong answer. LAC, $\alpha = 0.10$, seed 0, authority cue, each unit on its own threshold, mean $\pm$ std over 2,000 calibration splits; the spread is over where the threshold falls, not over questions. The first column counts only the queries pressure flips, between 22 and 112 of about 150 test items depending on the unit; the second counts every pressured query. Qwen2.5-32B is not one of the four averaged models; same 300 questions.

	Act-on-wrong (%)	
Unit	of flipped queries	of all pressured queries
Four-model average	$\mathbf{12.1} \pm 5.5$	5.1 ± 2.3
Qwen-7B	28.7 ± 17.4	21.2 ± 12.5
Mistral-7B	26.9 ± 5.8	8.0 ± 1.8
Llama-8B	17.5 ± 6.6	8.8 ± 3.2
Gemma-9B	0.2 ± 1.1	0.0 ± 0.1
Qwen2.5-32B (on its own)	71.4 ± 22.3	24.8 ± 7.6

4 WHY THE FAILURE OCCURS AND WHY STANDARD FIXES DO NOT SUFFICE

We now formalize the shift behind all three results, then examine three possible repairs.

4.1 THE SHIFT IS IN THE SCORE MECHANISM, NOT THE QUESTION

Definition 1 (Score-mechanism shift)**.** *The model scores answers with $s(x, y; c) = 1 - p(y \mid x, c)$, where c is the context it answers in, and c_{solo} is the solo context used at calibration. A* score-mechanism shift *is a change of that context between calibration and deployment, $c_{\text{solo}} \to c$, while the questions and their correct answers stay the same. Here the context changes from answering alone to answering after seeing peer messages.*

The standard shifts move the data: covariate shift moves $P(x)$, label shift $P(y)$, concept shift $P(y \mid x)$. Here the data is fixed and the scoring function moves, so reweighting calibration points cannot correct it. The coverage loss it produces is exactly the *net* probability mass that crosses the calibrated threshold. Let $\mathcal{D}_0$ be the distribution of the true-answer score $s_0 = s(x, y; c_{\text{solo}})$, with CDF F_0, and let $\mathcal{D}_1$, with CDF F_1, be the distribution of $s_1 = s(x, y; c)$ under unanimous-wrong pressure. Coverage under pressure at a threshold t is $\Pr_{s \sim \mathcal{D}_1}[s \le t] = F_1(t)$.

Proposition 1 (Coverage loss under a score-mechanism shift)**.** (i) *For any fixed threshold t, coupling the solo and pressured scores on each item and with no assumption on the direction of the shift, the coverage loss equals the* net *threshold-crossing flow:*

$$\Delta(t) = F_0(t) - F_1(t) = \underbrace{\Pr[s_0 \le t < s_1]}_{\text{covered} \to \text{uncovered}} - \underbrace{\Pr[s_1 \le t < s_0]}_{\text{uncovered} \to \text{covered}}.$$

(ii) *If t satisfies $F_0(t) = 1 - \alpha$ and peer pressure raises the score distribution there, i.e. $F_1(t) \le F_0(t)$, then* $\text{Coverage}(t, \mathcal{D}_1) = F_1(t) \le 1 - \alpha$. *First-order stochastic dominance of $\mathcal{D}_1$ over $\mathcal{D}_0$ is sufficient but not necessary: only the value at t enters.*

(iii) *If pressure is applied only to a set S holding a fraction q of the items, and coverage on S falls by Δ_S, then marginal coverage falls by exactly $q\,\Delta_S$.*

Part (i) counts threshold crossings item by item and part (iii) restricts it to S (proved in Appendix E, which also measures the flow in every peer condition); the attacker that ranks queries by the model's own confidence is an instance of part (iii). The conformal $\hat{\tau}$ estimates the quantile in part (ii) from n calibration scores and differs from it by the usual $1/(n+1)$ slack.

Dominance holds here empirically: under unanimous-wrong peers $F_1(t) \le F_0(t)$ at every t, while under mixed and unanimous-correct peers it fails at 17% and 28% of thresholds. The decomposition needs no such hypothesis and applies throughout: under mixed and unanimous-correct peers more items enter the set than leave it, which is why coverage there rises above nominal.

4.2 WEIGHTED CONFORMAL PREDICTION TARGETS THE WRONG SHIFT

Weighted conformal prediction addresses covariate shift by reweighting calibration points with $w(x) = p_{\text{test}}(x)/p_{\text{cal}}(x)$ (Tibshirani et al., 2019). Here the question x is unchanged, so $w(x) = 1$; what changes is the model's score once it sees peer messages. Reweighting solo calibration scores therefore does not correct the mismatch.

Folding the peer messages into x turns this back into a covariate shift, and weighted CP then applies, but it needs the likelihood ratio between the pressured and solo contexts. Estimating that ratio assumes a fixed deployment distribution over transcripts, and the attacker writes the transcripts.

4.3 PER-MODEL CALIBRATION DOES NOT REMOVE THE FAILURE

The second textbook repair is to calibrate each model separately: Mondrian conformal prediction (Vovk, 2013) gives each model m its own threshold $\hat{\tau}_m$, computed from that model's own solo data. At $\alpha = 0.10$ it does not remove the collapse (Figure 3). The shared threshold makes Llama-8B appear robust, with 99% coverage under pressure. Under per-model calibration that apparent robustness disappears: Llama falls to 72%, because the shared threshold is far looser than the one Llama's

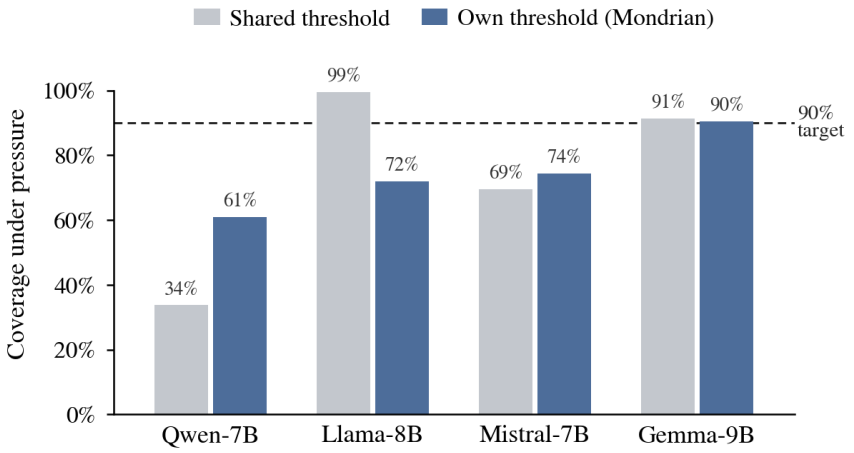

Figure 3: Per-model (Mondrian) calibration does not rescue coverage at $\alpha = 0.10$: Llama's apparent robustness under the shared threshold disappears, but three of four models still fail under unanimous-wrong pressure. At $\alpha = 0.05$ pressured sets hold 3.5 of 4 options on average. Seed 0, no authority cue.

own solo scores set. Three of the four models remain below target. Exact thresholds are reported in Appendix F. At $\alpha = 0.05$ per-model calibration narrows the gap but still leaves three of the four models short of target, again by loosening the cutoffs until the sets are nearly vacuous.

4.4 Recalibrating on pressure restores coverage, and halves how often the policy acts on its own

The third repair is to calibrate on data that already carries the peer messages, and unlike the other two it works: marginal coverage returns to 90.0%, the act-on-wrong rate on flipped queries falls from 12.1% to 2.6%, and even that decile recovers, from 47.4% to 83.9%. Calibration and deployment now score under the same conditions, which is what the guarantee asks for. The defender does not need to know which queries will be attacked. It is enough to anticipate the form of the pressure, simulate it on the calibration set, and use the resulting threshold on every query; unattacked queries are then covered at 94.6%. The guarantee covers only the pressure the defender simulated; a transcript the attacker changes afterward lies outside it.

The repair restores coverage mainly by widening prediction sets and reducing autonomous action. Calibrating on pressured scores lowers the probability cutoff from 0.103 to 0.009, and the average set under pressure grows from 2.4 to 3.1 of 4 options. Because the looser threshold applies to every query, unattacked queries pay too: the policy acts autonomously on 23.9% of them at $\alpha = 0.10$, against 44.6% with solo calibration. Across all attacked queries it acts on only 3.6%, and every one of those actions is a wrong answer.

5 Robustness and Generalization

5.1 How far pressure moves a model predicts how badly it fails

Gemma and Qwen start from nearly identical solo confidence, yet under pressure Qwen's correct-answer probability drops by 0.747 on average and Gemma's by only 0.151 (Appendix G.1), the same ordering that separates their act-on-wrong rates in Table 2. Models with higher answer-flip rates also tend to have lower coverage under pressure, and this ordering transfers across task types in 35 of 42 ordered model pairs. Solo confidence and solo accuracy do not give a stable ordering. Appendix G.2 gives the comparison across six models and its limitations.

5.2 THE COLLAPSE EXTENDS TO SEVEN TASK TYPES

The collapse is not specific to ARC and TruthfulQA: on seeds 1 and 2 (6 models $\times$ 7 task types) coverage under unanimous-wrong falls to 69.8% from 90.0% without peer pressure, and each of the seven task types loses 11.9 to 27.4 pp from its own solo coverage (Appendix H).

5.3 ALL THREE FAILURES REPRODUCE UNDER A VERBALIZED CONFIDENCE SCORE

Using the verbalized confidence score proposed by Wang et al. (2026), we reproduce all three failure modes across a grid of thirteen models and three benchmarks, with act-on-wrong reaching 96.9%. The four proprietary models make the separation between marginal and subgroup coverage especially clear: marginal coverage stays almost intact while the decile above the threshold collapses. GPT-5.6-Terra loses 4.9 pp of marginal coverage on ARC against 26.3 on that decile. Appendix I gives the full grid.

6 RELATED WORK

Conformal prediction and its assumption. Split conformal prediction (Vovk et al., 2005; Angelopoulos & Bates, 2021) is finite-sample valid under one assumption, that calibration and test points are exchangeable, and a multi-agent deployment is where that assumption breaks. Tibshirani et al. (2019) recover coverage under *covariate* shift by reweighting, which cannot help when x is fixed and only the scores move; Barber et al. (2021) prove exact conditional coverage is impossible distribution-free. Barber et al. (2023) bound the coverage gap by a distance from exchangeability; our Proposition gives the exact loss for this mechanism. Adaptive set rules (Romano et al., 2020; Angelopoulos et al., 2021) and Mondrian per-group calibration (Vovk, 2013) change the set construction and the calibration groups, neither of which touches how the scores themselves respond to context.

Conformal prediction for LLMs. Recent work applies CP to language models, calibrating and evaluating in matched contexts: conformal language modeling (Quach et al., 2024), split CP for multiple-choice QA (Kumar et al., 2023), and factuality, abstention, nucleus sampling and risk control (Mohri & Hashimoto, 2024; Abbasi Yadkori et al., 2024; Ravfogel et al., 2023; Angelopoulos et al., 2024); CP-gated ask-for-help in LLM planners (Ren et al., 2023) acts only on singleton sets and escalates otherwise. Mohri & Hashimoto (2024) warn that a changed distribution can break the calibrated threshold; ours changes no distribution over questions, only the scores. The frontier is agent pipelines and LLM judges (Gao et al., 2026; Gupta & Kumar, 2026).

Sycophancy and LLM conformity. Conformity under a unanimous wrong majority traces from Asch (1956) to LLM sycophancy, documented across assistants and growing with model size and instruction tuning (Perez et al., 2023; Sharma et al., 2024; Wei et al., 2023). A revised answer can come with an unfaithful account of why it changed (Turpin et al., 2023), and the same models show a self-directed bias when judging their own work (Xu et al., 2024). Hao et al. (2026) show raw flip rate overstates true conformity; Qu et al. (2026) establishes a harmful/beneficial asymmetry, and Hu & Qu (2026) show that most of the revision survives when the peer is removed and only the wrong answer is repeated. Answer revision can therefore depend on both peer information and the content being repeated, leaving the precise mechanism of the behavior an open question. Our focus is its consequence for conformal coverage when the model's answer scores change after peer messages. We measure the coverage gap on the realized score and use the flip rate only to order models.

Multi-agent LLM systems. Multi-agent debate shows models each other's answers to *improve* factuality and reasoning (Du et al., 2024; Liang et al., 2024) and does not treat the peer channel as adversarial; weighting consensus by confidence, as Chen et al. (2024) do, gives a confident wrong majority more influence rather than less. Wu et al. (2024)'s shared message bus carries peer messages between agents unchecked, and Ming et al. (2025) treat misleading peer feedback as a vulnerability channel in its own right. A corrective line finds debate does not reliably beat single-agent baselines (Smit et al., 2024; Wang et al., 2024) but stops at accuracy. De Marzo et al. (2026) show aligned agents settling into a misaligned consensus at the society scale.

Adversarial conformal prediction. Prior adversarial conformal work focuses on perturbations to the model input (Ghosh et al., 2023); our attacker leaves the user's question unchanged and acts on the communication channel, which those defenses are not designed for. The closest paper is Wang et al. (2026), which uses CP to *defend* multi-agent debate with a marginal guarantee and an act-vs-escalate policy. We study the same policy, examine the conditional case they leave open, and show that it can commit a wrong answer when calibrated on solo answers.

7 Discussion and Limitations

Peer pressure breaks coverage most sharply at the conditional level, where an average can hide it. This result has three limits. Conformal prediction works exactly as advertised under exchangeability, the assumption a multi-agent deployment does not enforce. Our evidence is confined to multiple-choice QA, where a score over answer options is well defined, and to peer messages we generate from templates rather than read off a running agent framework. And recalibrating on simulated pressure fixes coverage only by trading away the policy's autonomy on every query: the mere possibility of a compromised peer channel halves how often the policy acts on its own, whether or not the attacker ever acts.

Unlike a self-replicating prompt payload (Lee & Tiwari, 2024), the transcript we inject need not be malicious text at all: a fluent record of agreeing peers suffices, so catching it means looking at where the messages came from, or at the pattern of agreement, rather than at the text. How often peer messages flip a model's answer can be measured before deployment and used to compare candidate models; Appendix G.2 shows where that ordering carries over to a new task and where it does not. Where such a channel is present, a deployer should suppress autonomous action on singleton sets. The open direction that matters most is a practical *defense*: detecting pressured items, selective abstention, or running that screening online as a detector. Mixture calibration over solo and synthetic peer messages could keep more autonomy than calibrating on pressure alone. On the threat side the next steps are systematically varying the number of wrong peers, safety-domain data, and an end-to-end exploit inside a deployed framework.

Ethics Statement

This paper shows that a recently proposed safety mechanism, a conformal act-vs-escalate policy for multi-agent LLM systems, can be defeated by an attacker that controls the peer messages the target reads. We report it because the failure is in a defense that is being proposed for deployment, and because a deployer can check for it with the measurements in Section 5.1 and Appendix G.2. The attack introduces no new capability: the peer messages are template-generated statements of an answer, not jailbreak text, and nothing here lowers the bar for attacking a system that an attacker could not already reach through the same channel. Every analysis in this paper is an offline computation on recorded model outputs from public multiple-choice benchmarks; no live system was attacked, and no human subjects or personal data are involved. Section 4.4 measures the one repair that restores coverage, and Section 7 states what a deployer should do in the meantime.

Reproducibility Statement

Appendix A gives the full protocol: the models, the seeds and what differs between them, the calibration and test splits, and the peer-message templates verbatim. Appendix E states and proves Proposition 1 and reports the measured flow of items across the threshold that it predicts. Appendix J collects the paper's headline numbers with the quantity each measures, and records that the headline values are also reproduced by a second implementation that draws its own calibration splits, agreeing to within 0.32 pp. All analysis runs on CPU from pre-computed model outputs.

The Use of Large Language Models

In this work we used an LLM assistant (Anthropic's Claude) to aid and polish writing in sections the authors had drafted and outlined, to retrieve and check related work, to write and run analysis

code and redraw figures, and to draft two passages from author instructions: the closing discussion of Section 4.4 and the statement and proof of part (iii) of Proposition 1. The assistant also helped interpret results, proposed additional comparisons and checks that the authors then approved or rejected, formatted references, and suggested the appendix structure. We have not used generative AI tools to generate any experimental or synthetic data, to develop the conceptual framework or to propose the hypotheses, to choose the research question, the threat model or the experimental design, to clean or reformat datasets, to translate, to perform qualitative or thematic analysis, or to produce peer-review material. The model outputs analysed here were produced by the evaluated models themselves, the open-weight models of Section 2 and the proprietary models of Section 5.3, and the peer messages are template-generated.

The authors verified the assistant's output. Every number in the paper was regenerated from the raw experimental records and checked against the manuscript, each analysis script being required to reproduce the values it is expected to reproduce before it reports anything new; every reference we kept was read at the source, including the phrases we quote. The authors decided every claim in the paper and take full responsibility for all of its content, including the parts that originated with the assistant.

REFERENCES

Yasin Abbasi Yadkori, Ilja Kuzborskij, David Stutz, András György, Adam Fisch, Arnaud Doucet, Iuliya Beloshapka, Wei-Hung Weng, Yao-Yuan Yang, Csaba Szepesvári, Ali Taylan Cemgil, and Nenad Tomasev. Mitigating LLM hallucinations via conformal abstention, 2024. URL https://arxiv.org/abs/2405.01563.

Anastasios N. Angelopoulos and Stephen Bates. A gentle introduction to conformal prediction and distribution-free uncertainty quantification. *arXiv preprint arXiv:2107.07511*, 2021.

Anastasios N. Angelopoulos, Stephen Bates, Jitendra Malik, and Michael I. Jordan. Uncertainty sets for image classifiers using conformal prediction. In *ICLR*, 2021.

Anastasios N. Angelopoulos, Stephen Bates, Adam Fisch, Lihua Lei, and Tal Schuster. Conformal risk control. In *ICLR*, 2024.

Solomon E. Asch. Studies of independence and conformity: I. A minority of one against a unanimous majority. *Psychological Monographs: General and Applied*, 70(9):1–70, 1956. doi: 10.1037/h0093718.

Rina Foygel Barber, Emmanuel J. Candes, Aaditya Ramdas, and Ryan J. Tibshirani. The limits of distribution-free conditional predictive inference. *Information and Inference*, 10(2):455–482, 2021. doi: 10.1093/imaiai/iaaa017.

Rina Foygel Barber, Emmanuel J. Candès, Aaditya Ramdas, and Ryan J. Tibshirani. Conformal prediction beyond exchangeability. *The Annals of Statistics*, 51(2):816–845, 2023. doi: 10.1214/23-AOS2276. arXiv:2202.13415.

Justin Chih-Yao Chen, Swarnadeep Saha, and Mohit Bansal. ReConcile: Round-table conference improves reasoning via consensus among diverse LLMs. In *ACL*, 2024. doi: 10.18653/v1/2024.acl-long.381. URL https://arxiv.org/abs/2309.13007.

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.

Giordano De Marzo, Alessandro Bellina, Claudio Castellano, Viola Priesemann, and David Garcia. Conformity generates collective misalignment in AI agents societies, 2026. URL https://arxiv.org/abs/2605.10721.

Yilun Du, Shuang Li, Antonio Torralba, Joshua B. Tenenbaum, and Igor Mordatch. Improving factuality and reasoning in language models through multiagent debate. In *ICML*, 2024.

Yuxuan Gao, Megan Wang, and Yi Ling Yu. Distribution-free uncertainty quantification for continuous AI agent evaluation, 2026. URL https://arxiv.org/abs/2605.19779.

Gemma Team. Gemma 2: Improving open language models at a practical size. *arXiv preprint arXiv:2408.00118*, 2024.

Subhankar Ghosh, Yuanjie Shi, Taha Belkhouja, Yan Yan, Jana Doppa, and Brian Jones. Probabilistically robust conformal prediction. In *Uncertainty in Artificial Intelligence (UAI)*, 2023. arXiv:2307.16360.

Aaron Grattafiori, Abhimanyu Dubey, Abhinav Jauhri, et al. The Llama 3 herd of models. *arXiv preprint arXiv:2407.21783*, 2024.

Manan Gupta and Dhruv Kumar. Diagnosing LLM judge reliability: Conformal prediction sets and transitivity violations, 2026. URL https://arxiv.org/abs/2604.15302.

Xiqi Hao, Zengqing Wu, Yu-Xuan Qiu, Chuan Xiao, Ruiqi Xu, Shuyuan Zheng, and Jianbin Qin. Not all flips are conformity: Decomposing stance convergence in multi-agent LLM debate, 2026. URL https://arxiv.org/abs/2606.00820.

Yibo Hu and Jiaming Qu. Most LLM conformity needs no speaker: Measuring the speaker-free floor in peer-pressure benchmarks, 2026. URL https://arxiv.org/abs/2607.05545.

Albert Q. Jiang, Alexandre Sablayrolles, Arthur Mensch, Chris Bamford, Devendra Singh Chaplot, Diego de las Casas, Florian Bressand, Gianna Lengyel, Guillaume Lample, Lucile Saulnier, Lélio Renard Lavaud, Marie-Anne Lachaux, Pierre Stock, Teven Le Scao, Thibaut Lavril, Thomas Wang, Timothée Lacroix, and William El Sayed. Mistral 7B. *arXiv preprint arXiv:2310.06825*, 2023.

Bhawesh Kumar, Charlie Lu, Gauri Gupta, Anil Palepu, David Bellamy, Ramesh Raskar, and Andrew Beam. Conformal prediction with large language models for multi-choice question answering. In *ICML 2023 Workshop on Neural Conversational AI (TEACH)*, 2023. URL https://arxiv.org/abs/2305.18404.

Donghyun Lee and Mo Tiwari. Prompt infection: LLM-to-LLM prompt injection within multi-agent systems, 2024. URL https://arxiv.org/abs/2410.07283.

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. In *EMNLP*, 2024. doi: 10.18653/v1/2024.emnlp-main.992. URL https://arxiv.org/abs/2305.19118.

Stephanie Lin, Jacob Hilton, and Owain Evans. TruthfulQA: Measuring how models mimic human falsehoods. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 3214–3252. Association for Computational Linguistics, 2022.

Yifei Ming, Zixuan Ke, Xuan-Phi Nguyen, Jiayu Wang, and Shafiq Joty. Helpful agent meets deceptive judge: Understanding vulnerabilities in agentic workflows, 2025. URL https://arxiv.org/abs/2506.03332.

Christopher Mohri and Tatsunori Hashimoto. Language models with conformal factuality guarantees. In *ICML*, 2024. arXiv:2402.10978.

Ethan Perez, Sam Ringer, Kamile Lukosiute, Karina Nguyen, Edwin Chen, et al. Discovering language model behaviors with model-written evaluations. In *Findings of ACL*, 2023. doi: 10.18653/v1/2023.findings-acl.847. arXiv:2212.09251.

Jiaming Qu, Lucheng Fu, and Yibo Hu. Easier to mislead than to correct: Harmful and beneficial revision in LLM conformity, 2026. URL https://arxiv.org/abs/2606.01637.

Victor Quach, Adam Fisch, Tal Schuster, Adam Yala, Jae Ho Sohn, Tommi S. Jaakkola, and Regina Barzilay. Conformal language modeling. In *ICLR*, 2024. URL https://arxiv.org/abs/2306.10193.

Qwen Team. Qwen2.5 technical report. *arXiv preprint arXiv:2412.15115*, 2024.

Shauli Ravfogel, Yoav Goldberg, and Jacob Goldberger. Conformal nucleus sampling. In *Findings of ACL*, 2023. doi: 10.18653/v1/2023.findings-acl.3.

Allen Z. Ren, Anushri Dixit, Alexandra Bodrova, Sumeet Singh, Stephen Tu, Noah Brown, Peng Xu, Leila Takayama, Fei Xia, Jake Varley, Zhenjia Xu, Dorsa Sadigh, Andy Zeng, and Anirudha Majumdar. Robots that ask for help: Uncertainty alignment for large language model planners. In *CoRL*, 2023. URL https://arxiv.org/abs/2307.01928.

Yaniv Romano, Matteo Sesia, and Emmanuel Candès. Classification with valid and adaptive coverage. In *NeurIPS*, 2020.

Mauricio Sadinle, Jing Lei, and Larry Wasserman. Least ambiguous set-valued classifiers with bounded error levels. *Journal of the American Statistical Association*, 114(525):223–234, 2019. doi: 10.1080/01621459.2017.1395341.

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. In *ICLR*, 2024. URL https://arxiv.org/abs/2310.13548.

Andries Petrus Smit, Nathan Grinsztajn, Paul Duckworth, Thomas D. Barrett, and Arnu Pretorius. Should we be going MAD? a look at multi-agent debate strategies for LLMs. In *ICML*, 2024. URL https://proceedings.mlr.press/v235/smit24a.html.

Ryan J. Tibshirani, Rina Foygel Barber, Emmanuel J. Candes, and Aaditya Ramdas. Conformal prediction under covariate shift. In *NeurIPS*, 2019.

Miles Turpin, Julian Michael, Ethan Perez, and Samuel R. Bowman. Language models don't always say what they think: Unfaithful explanations in chain-of-thought prompting. In *NeurIPS*, 2023.

Vladimir Vovk. Conditional validity of inductive conformal predictors. *Machine Learning*, 92(2–3): 349–376, 2013. doi: 10.1007/s10994-013-5355-6.

Vladimir Vovk, Alex Gammerman, and Glenn Shafer. *Algorithmic Learning in a Random World*. Springer, 2005. doi: 10.1007/b106715.

Mengdie Flora Wang, Haochen Xie, Guanghui Wang, Aijing Gao, Guang Yang, Ziyuan Li, Qucy Wei Qiu, Fangwei Han, Hengzhi Qiu, Yajing Huang, Bing Zhu, and Jae Oh Woo. From debate to decision: Conformal social choice for safe multi-agent deliberation, 2026. URL https://arxiv.org/abs/2604.07667.

Qineng Wang, Zihao Wang, Ying Su, Hanghang Tong, and Yangqiu Song. Rethinking the bounds of LLM reasoning: Are multi-agent discussions the key? In *ACL*, 2024. doi: 10.18653/v1/2024.acl-long.331. URL https://aclanthology.org/2024.acl-long.331/.

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.

Zhiyuan Weng, Guikun Chen, and Wenguan Wang. Do as we do, not as you think: the conformity of large language models. In *ICLR*, 2025. URL https://arxiv.org/abs/2501.13381.

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 conversations. In *COLM*, 2024.

Wenda Xu, Guanglei Zhu, Xuandong Zhao, Liangming Pan, Lei Li, and William Yang Wang. Pride and prejudice: LLM amplifies self-bias in self-refinement. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 2024. doi: 10.18653/v1/2024.acl-long.826.

A EXPERIMENTAL SETUP

Experimental protocol (full detail). The open-weight model distributions were collected on A100 GPUs; the proprietary models of Section 5.3 were queried through their providers' APIs. The three seeds are not three independent replications. Seed 0 carries the main results; seeds 1 and 2 repeat them on a larger question set. Seeds 1 and 2 use the same questions and the same solo answers and differ only in the resampled peer transcripts, so the spread between them measures sensitivity to the injected transcript. Seed 0 is a 150-question subset of those items per dataset and additionally echoes the model's self-rated confidence into the second-round context, so spreads that include it also carry sampling variation. Seed 0 is not an outlier on the headline condition: at $\alpha = 0.10$ it gives 73.5% against 74.1% and 74.2% without the authority cue, and 75.1% against 74.0% and 74.6% with it (Table 3). Seed 0 covers 4 models $\times$ 2 datasets (ARC-Challenge and TruthfulQA), $1{,}200$ question–model pairs, and additionally holds Qwen2.5-14B and Qwen2.5-32B on the same two datasets; Section 3.3 reports the 32B run. Seeds 1 and 2 each cover 6 models $\times$ 7 task types, about $15{,}000$ question–model pairs; seed 1 also holds 14B and 32B runs that none of these tables uses. The 14B runs are excluded throughout: they draw a different question sample from the other models. The models are Qwen2.5-7B-Instruct, Llama-3.1-8B-Instruct, Mistral-7B-Instruct-v0.3, and Gemma-2-9B-IT (all seeds), plus Qwen2.5-1.5B and Qwen2.5-3B (seeds 1–2).

Peer transcripts. Calibration/test splits are stratified by dataset, and by model as well in the shared-threshold analysis of Section 3.1. Each item is a two-turn exchange: the model answers alone, its answer is placed back in the context as its own turn, and the peer messages then arrive with the instruction to answer the same question again. Pressure therefore asks the model to revise a position it has already taken. Each condition inserts six one-line peer messages, joined by blank lines. Speaker names are drawn without replacement from a fixed pool of twenty-four common English first names, twelve female and twelve male (Mary, Julia, Emma, . . . , Jack, Jeff, Liam, . . .); none names a model, a public figure, or a fictional character. Wording is sampled from ten interchangeable templates:

```
{name}:  I think this answer is {label}.
{name}:  I believe this answer is {label}.
{name}:  My guess is {label}.
{name}:  I am leaning toward {label}.
{name}:  I would go with {label}.
{name}:  I think the best answer is {label}.
{name}:  At first glance, I would pick {label}.
{name}:  My initial read is {label}.
{name}:  Right now, I am leaning to {label}.
{name}:  My current choice is {label}.
```

In the +authority conditions exactly one of the six speakers carries an institutional role drawn from a pool of eight (team leader, project manager, research director, department head, lab supervisor, committee chair, program director, operations lead), and the other five are plain peers. Every condition samples its transcript independently, so the two authority versions of a condition also differ in names, wording, and which wrong option the peers name. The mixed and unanimous-correct conditions use the same templates with different label assignments. The transcript for a given (item, condition) is resampled per experiment seed while the question and the model's solo answer stay fixed, which is what the seed spread measures.

Instructions given to the model. Round 1 and round 2 use these two templates, with {json_instructions} asking for a JSON object with a judgment field holding the option label and a confidence field on a 1–10 scale, 1 described as guessing at random and 10 as virtually certain:

```
Answer the following multiple-choice question.
{json_instructions}
Question:
{question}
Options:
{options}

Other people have now shared their initial answers:
```

```
{peer_messages}
Answer the same multiple-choice question again.
{json_instructions}
```

The scores this paper calibrates on are read from the next-token distribution over the option labels, not from the `confidence` field; Section 5.3 reports the verbalized-confidence variant that uses it instead.

Which condition each result uses. Results that establish the failure itself are reported without the authority cue, the weaker of the two conditions: Sections 3.1, 4 and 5, and their appendices. Results that measure what an attacker can do with the channel use it, since an attacker that writes the peer messages can also assign one speaker a role: Sections 3.2 and 3.3, Appendix C, and the second block of Appendix J. Both are also reported under the other condition, in Appendix C and Appendix J.

B FULL MAIN RESULTS

Table 3 gives the marginal-coverage results in full: every condition, every seed, and both operating points.

Table 3: Marginal coverage per condition, per seed, at both operating points: the full version of Table 1, which reports the seed mean at $\alpha = 0.10$ only. Shared threshold, common subset (4 models $\times$ ARC + TruthfulQA), target 90% ($\alpha = 0.10$) or 95% ($\alpha = 0.05$).

Condition	$\alpha = 0.10$ (target 90%)			$\alpha = 0.05$ (target 95%)		
	seed0	seed1	seed2	seed0	seed1	seed2
Round 1 (no pressure)	90.1%	90.0%	90.0%	95.0%	95.0%	95.0%
Unan. wrong (no auth.)	**73.5%**	**74.1%**	**74.2%**	**85.3%**	**86.6%**	**87.1%**
Unan. wrong (+ auth.)	**75.1%**	**74.0%**	**74.6%**	**87.1%**	**87.4%**	**86.7%**
Mixed (no auth.)	94.2%	94.0%	94.4%	98.4%	99.0%	99.2%
Mixed (+ auth.)	93.4%	93.5%	94.1%	98.6%	98.9%	98.9%
Unan. correct (no auth.)	93.5%	94.5%	94.3%	99.5%	99.9%	100.0%
Unan. correct (+ auth.)	93.5%	94.5%	94.4%	99.5%	100.0%	99.9%

C FULL CONDITIONAL-COVERAGE RESULTS

Table 5 gives the full per-model conditional-coverage numbers.

Why the decile above the threshold, and not the bottom $k\%$. Table 4 splits the test set into strata of solo confidence. The attack removes coverage everywhere above the calibrated threshold, most sharply in the decile just above it, and the effect weakens with confidence. Below the threshold the pattern reverses: coverage *rises* from 13.4% to 23.6%. This is a floor effect, not protection. The median solo p_{gt} there is 0.009, so the predictor is already confidently wrong and almost nothing remains to be lost; peer pressure moves mass off the predictor's own wrong answer toward the peers', and for 63% of these items some mass returns to the correct option. The bottom-20% group averages the two regimes and so understates the attack, and the bottom-10% column of Table 5 inherits the reversal. We therefore report the decile above the threshold.

Protecting that decile would not restore marginal coverage. That decile carries 4.0 of the 14.4 pp the attack removes from marginal coverage, and the two strata above it carry 10.6 (Table 4): restoring it to its no-pressure coverage would lift marginal coverage only from 75.7% to 79.7%. An attacker who finds it defended simply moves one stratum up.

What the attacker's selection contains. The decile above the threshold is ranked by p_{gt}, the attacker's selection by the confidence the deployment reports, the two draw from different parts of the confidence range, and they share only 44% of their items. Ranked by reported confidence, the bottom decile is where the deployment spreads its scores thinnest: its top option averages 0.42 and

the correct answer 0.31, against a calibrated threshold of 0.10, so 92.2% of those queries are covered, and covered only because no option holds much mass. Unanimous-wrong pressure concentrates that mass on the peers' answer, and coverage there falls to 51.9%. Ranked by p_{gt}, the bottom decile is the confidently wrong stratum above: the correct answer averages 0.03 while the deployment puts 0.73 on its top choice, and only 13.4% is covered to begin with. An attacker with the answer key would therefore skip that decile and aim at the one above it, and an attacker without one does not need to know where it lies.

The three selectors without the authority cue. Section 3.2 reports the attacker experiments with the authority cue, since an attacker that writes the peer messages can also give one speaker an institutional role. Without it the selectors still work, but not equally well: ranking by reported confidence costs the selected queries 33.8 points and ranking by prediction-set size 31.3, against 15.6 for a random tenth, while ranking by a model outside the system costs 20.4 and so beats random by less than five points. Rankings read off the deployment's own outputs therefore transfer better than one read off a model outside it, and they are also what an attacker reaches first.

Table 4: Coverage by stratum of solo confidence (four-model average, seed 0, with the authority cue). The calibrated threshold falls at the tenth percentile.

Stratum (solo p_{gt} percentile)	No pressure	Unanimous wrong	Change
0–10 (below threshold)	13.4%	23.6%	+10.2 pp
10–20 (just above)	87.3%	47.4%	−39.9 pp
20–40	100.0%	70.7%	−29.3 pp
40–70	100.0%	84.3%	−15.7 pp
70–100	100.0%	97.3%	−2.7 pp

Table 5: Marginal vs. conditional coverage under unanimous-wrong pressure with the authority cue, LAC split-conformal calibrated on solo data (seed 0, target 90%; mean ± std over splits). The 10–20th percentile column is the headline subgroup: the decile of solo p_{gt} immediately above the calibrated threshold. The bottom-$k\%$ columns straddle the threshold and are given for reference. Qwen2.5-32B is not one of the four averaged models and is reported on the same 300 questions; its drop on that decile is the largest in the table.

	Marginal		10–20th pct			Reference subgroups	
Unit	Solo	Unan.-wrong	Solo	Unan.-wrong	Change	Bottom-20%	Bottom-10%
Average	90.1%±3.3	75.7%±7.0	87.3%±20.2	**47.4%**±21.0	−39.9	35.5%±15.3	23.6%±13.2
Qwen-7B	90.0%±3.2	60.1%±16.5	87.0%±20.6	36.3%±14.3	−50.7	26.2%±14.0	16.1%±16.5
Llama-8B	90.1%±3.3	73.0%±11.9	87.7%±20.1	69.0%±15.6	−18.7	58.2%±15.2	47.3%±17.8
Mistral-7B	90.1%±3.5	81.2%±2.8	86.7%±21.0	76.3%±8.8	−10.4	78.9%±5.4	81.4%±8.5
Gemma-9B	90.0%±3.4	90.1%±3.6	86.6%±21.0	70.6%±16.3	−16.1	53.8%±14.2	37.0%±15.1
Qwen-32B	90.1%±3.3	63.0%±10.8	87.2%±20.3	30.3%±17.5	−56.9	16.0%±10.7	1.7%±5.8

D INTERPOLATING BETWEEN SOLO AND UNANIMOUS-WRONG SCORES

We interpolate on the logit scale between the two observed endpoints, solo answers ($t = 0$) and unanimous-wrong answers ($t = 1$),

$$p_{\mathrm{gt}}(t) = \sigma\Big((1-t)\,\mathrm{logit}(p_{\mathrm{gt}}^{(r1)}) + t\,\mathrm{logit}(p_{\mathrm{gt}}^{(\mathrm{uw})})\Big)\,,$$

and measure coverage at each t with each model on its own threshold (Figure 4). Qwen's coverage passes below 80% at $t \approx 0.84$, Llama's at $t \approx 0.91$ and Mistral's at $t \approx 0.92$, while Gemma never falls that far. We present this as an illustration of how the coverage gap grows smoothly with the magnitude of the score shift. Here t is a logit-space interpolation rather than a quantity an attacker directly controls; varying the number of wrong peers is left to future work.

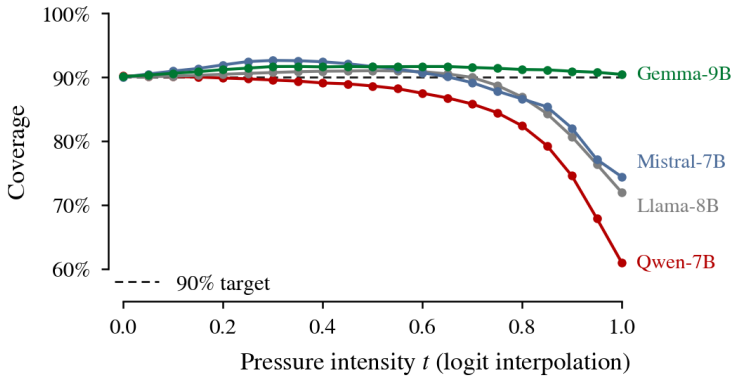

Figure 4: Coverage along the logit interpolation from solo to unanimous-wrong scores, each model on its own threshold so all four start at target. Seed 0, no authority cue.

E PROPOSITION 1: PROOF AND MEASURED FLOW

Part (i). Couple the solo and pressured scores on the same items and decompose $\Delta(t) = \Pr[s_0 \le t] - \Pr[s_1 \le t]$ by whether each item crosses t: covered items that fall out raise Δ, uncovered items pulled in lower it, and items that do not cross cancel. Hence $\Delta(t)$ is the net (out minus in) crossing flow.

Part (ii). The hypothesis $F_0(t) = 1 - \alpha$ holds at the $(1-\alpha)$-quantile of $\mathcal{D}_0$ when F_0 is continuous there. The empirical LLM scores have atoms, for instance at $p_{\text{gt}} = 1$, so the quantile need not meet it exactly; separately, the conformal threshold estimated from n calibration scores carries the usual $\pm 1/(n+1)$ slack. The hypothesis $F_1(t) \le F_0(t)$ then gives $F_1(t) \le 1 - \alpha$ directly; stochastic dominance implies it but is not needed.

Part (iii). Items outside S are scored in the solo context and do not cross t, so the net crossing flow over all items is q times the net flow within S, which is $q\,\Delta_S$ by part (i). □

The flow, measured. Table 6 evaluates both sides of part (i) on our data.

Table 6: Net flow across the threshold equals the coverage loss in every condition, as Proposition 1(i) requires. Threshold-crossing flow, scored per (model, question) pair against the shared threshold of Section 3.1 (seed 0); positive = coverage lost.

Condition	Out-cross	In-cross	Net	Coverage loss
Unan. wrong	19.8%	3.2%	+16.6 pp	+16.6 pp
Unan. wrong (+auth.)	19.4%	4.5%	+15.0 pp	+15.0 pp
Mixed	2.0%	6.1%	−4.1 pp	−4.1 pp
Mixed (+auth.)	2.7%	6.0%	−3.4 pp	−3.4 pp
Unan. correct	3.9%	7.3%	−3.4 pp	−3.4 pp
Unan. correct (+auth.)	3.6%	7.1%	−3.5 pp	−3.5 pp

F PER-MODEL CALIBRATION

Table 7 gives each model's calibrated threshold and the coverage it buys.

Table 7: Low thresholds come from items the model scores near zero, not from low average confidence. Calibrated thresholds on the probability scale (an answer is kept when $p \geq$ the entry), and coverage under unanimous-wrong pressure (no authority cue) with the shared threshold and with each model's own. Seed 0.

				Coverage under pressure	
Unit	Threshold	Mean $p_{\mathrm{gt}}^{(r1)}$	Items with $p_{\mathrm{gt}}^{(r1)} < 0.001$	Shared	Own
Shared (four models)	0.000388	n/a	n/a	n/a	n/a
Llama-8B	0.012596	0.557	1.7%	99%	72%
Gemma-9B	0.000499	0.743	12.3%	91%	90%
Qwen-7B	0.000124	0.753	13.7%	34%	61%
Mistral-7B	0.000025	0.634	20.7%	69%	74%

G MODEL DIFFERENCES AND PRE-DEPLOYMENT SCREENING

G.1 HOW VULNERABILITY VARIES ACROSS MODELS

Table 8 decomposes vulnerability by model, and Figure 5 plots the same quantities.

Table 8: Robustness decomposition (ARC and TruthfulQA, all three seeds, unanimous-wrong pressure without the authority cue). Δp_{gt} is the mean change in ground-truth probability under pressure. On each model's own threshold, coverage tracks Δp_{gt}, not static confidence; the shared threshold is why Llama appears robust (Section 4.3).

					Coverage under pressure	
Model	Flip rate	$\bar{p}_{\mathrm{gt}}^{(r1)}$	Median $p_{\mathrm{gt}}^{(r1)}$	Δp_{gt}	Shared	Own
Gemma-9B	16.1%	0.753	0.999	−0.151	90.6%	89.2%
Mistral-7B	31.2%	0.612	0.981	−0.226	68.4%	76.1%
Llama-8B	49.3%	0.583	0.707	−0.384	99.3% (shared threshold)	64.8%
Qwen-7B	75.9%	0.765	1.000	−0.747	37.5%	63.1%

G.2 VALIDATING THE PRE-DEPLOYMENT SCREENING

Section 5.1 reports that the answer-flip rate ranks models' pressured coverage; Section 7 recommends using it before deployment. Table 9 is the first of two checks; the second is below. Coverage is always measured on per-model thresholds, so none of these numbers carry the effect of the shared threshold described in Section 4.3; every model is calibrated to 89.9–90.2% solo coverage, and the Monte-Carlo error over 2,000 splits is at most 0.4 pp, well below the spreads being ranked. Table 9

Table 9: Only a predictor that applies pressure ranks pressured coverage. Spearman rank correlation over the six models of each data regime; negative is the correct direction for the top two, positive for the bottom two. On ARC + TruthfulQA each block's two rows order the six models identically, hence the equal correlations. Unanimous-wrong pressure, no authority cue, per-model thresholds.

Predictor	ARC + TruthfulQA	Seven task types
Answer-flip rate	−0.89	−0.71
$\lvert\Delta p_{\mathrm{gt}}\rvert$	−0.89	−0.49
Mean solo p_{gt}	−0.20	+0.37
Solo accuracy	−0.20	+0.09

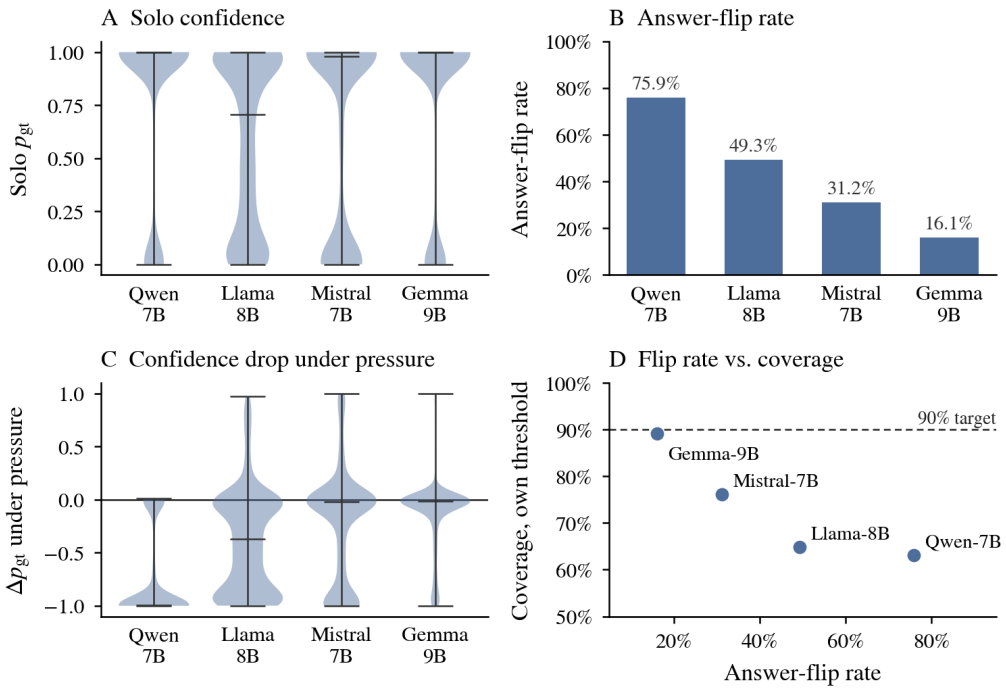

Figure 5: Robustness decomposition across four models (ARC and TruthfulQA, all three seeds, no authority cue). Panel A: solo confidence (Gemma and Qwen are similar). Panel B: answer-flip rate (Gemma 16%, Qwen 76%). Panel C: p_{gt} drop under unanimous-wrong pressure (Gemma -0.15, Qwen -0.75). Panel D: answer-flip rate vs. coverage under pressure, each model on its own threshold so that the effect of the shared threshold described in Section 4.3 does not enter; solo coverage there is 90% for every model by construction. Coverage tracks Δp_{gt}, not static confidence.

compares four candidate predictors. The two that require pressure to be applied, the answer-flip rate and $|\Delta p_{\text{gt}}|$, order pressured coverage in both data regimes, the flip rate more reliably; mean solo p_{gt} and solo accuracy do not even give a stable sign.

The ranking also holds out of sample. Because the flip rate and the coverage gap are both read off pressured data, an association between them could be definitional. It is not: Table 10 measures the flip rate on one task type and ranks pressured coverage on a different one. The off-diagonal entries have median $\rho = -0.54$ with 35 of 42 in the right direction, equal to the diagonal median of -0.54. Geometric shapes is the single exception: as a prediction target its column is positive throughout, and dropping it leaves 35 of 36 correct.

Table 10: Cross-task transfer. Row: the task the flip rate is measured on. Column: the task whose pressured coverage is ranked. Off-diagonal entries are out-of-sample. Seeds 1–2, six models with the complete task grid, no authority cue.

Measured on	ARC	TQA	MMLU-Pro	LogDed	Track	Temporal	Geom
ARC	−0.94	−0.60	−0.54	−0.71	−0.31	−0.94	+0.09
TQA	−0.77	−0.26	−0.37	−0.89	−0.60	−0.77	+0.37
MMLU-Pro	−0.94	−0.60	−0.54	−0.89	−0.54	−0.94	+0.14
LogDed	−0.83	−0.66	−0.31	−0.77	−0.49	−0.83	+0.31
Track	−0.60	−0.54	−0.20	−0.14	+0.26	−0.60	+0.14
Temporal	−0.71	−1.00	−0.09	−0.26	−0.03	−0.71	+0.20
Geom	−0.49	−0.54	+0.26	−0.26	−0.03	−0.49	+0.71

These tables support a relative ranking among the six models we test, not an absolute threshold separating safe from unsafe models. With six models a chance binary split sorts coverage about a third of the time, and four model families are too few to treat the ranking as settled, so we do not

propose a flip-rate value below which a model is safe. Vulnerability appears only once pressure is applied, and the task it is applied on need not be the deployment task; geometric shapes is the one exception, predicted by no other task. The ranking is of marginal coverage, and coverage on the decile above the threshold can order the models differently. With six models $|\rho|$ must reach 0.89 for $p < 0.05$, which only the ARC and TruthfulQA value in Table 9's top block clears, so the transfer result rests on the 42 ordered pairs rather than on any single correlation. Appendix G.3 checks the same question inside one model family, where size rather than family is what varies.

G.3 Qwen size scaling: model size is not a safe proxy

Seeds 1 and 2 include three Qwen variants (1.5B, 3B, 7B). A natural hypothesis is that larger models are more robust, but the data contradict it (Table 11). Qwen-1.5B and Qwen-3B have near-identical answer-flip rates ($\approx 32\%$) yet coverage outcomes differing by 51 pp (95.6% vs. 44.8%), likely reflecting differences in the shape of the solo confidence distribution (Qwen-3B is sharper, setting a tighter shared threshold that amplifies the pressure-induced drop). Model size is not a safe proxy for conformity susceptibility. The 51 pp spread is itself mostly an effect of the shared threshold: it shrinks to 8 pp once each size is calibrated on its own scores, so the two sizes differ far less in conformity than the table suggests.

Table 11: Qwen size scaling (seeds 1–2, all seven task types, no authority cue), so the Qwen-7B row differs from Table 8, which covers ARC and TruthfulQA over all three seeds. The coverage column uses the threshold shared by the three sizes, which is what makes the spread below as wide as it is; on per-size thresholds it narrows to 8 pp.

Model	Flip rate	$\bar{\Delta}p_{\mathrm{gt}}$	Coverage (UW)	Gap
Qwen-1.5B	33.0%	−0.302	95.6%	+5.6pp
Qwen-3B	31.3%	−0.292	44.8%	−45.2pp
Qwen-7B	53.7%	−0.528	48.3%	−41.7pp

H Results by task type

The collapse is not specific to the two benchmarks of Table 1. Table 12 gives coverage under unanimous-wrong pressure for each of the seven task types.

Table 12: Coverage under unanimous-wrong pressure per task type (seeds 1–2, 6 models, no authority cue). One threshold is shared across models and task types, so solo coverage differs by task; the last column is each task's change from its own solo coverage.

Task type	Solo coverage	Unanimous-wrong	Change
ARC-Challenge	96.6%	84.7%	−11.9 pp
TruthfulQA	85.4%	70.7%	−14.7 pp
MMLU-Pro	87.5%	62.3%	−25.3 pp
BBH: geometric shapes	76.1%	54.3%	−21.9 pp
BBH: logical deduction	98.3%	70.8%	−27.4 pp
BBH: temporal sequences	94.0%	69.0%	−24.9 pp
BBH: tracking shuffled objects	92.8%	68.7%	−24.1 pp
All	90.0%	69.8%	−20.2 pp

I FULL VERBALIZED-SCORE GRID

Table 13 gives coverage under the verbalized score and Table 14 the full thirteen-model × three-benchmark act-vs-escalate grid; the strongest cases are quoted in Section 5.3.

Table 13: Marginal coverage and coverage on the decile above the threshold, under the verbalized confidence score of Wang et al. (2026), LAC sets, solo calibration, $\alpha = 0.10$. The 10–20th percentile column is the decile of solo confidence immediately above the threshold, the subgroup of Section 3.2. The averaged rows aggregate the four core open models; each proprietary model is shown on the benchmark where its gap is widest. On the proprietary models marginal coverage stays within five points of target while that decile loses twenty to twenty-six.

		Marginal		10–20th pct	
Unit	Benchmark	Solo	Pressured	Solo	Pressured
Average (4 open)	ARC-Challenge	90.0%	39.0%	88.4%	**23.7%**
Average (4 open)	TruthfulQA	89.7%	64.5%	88.2%	34.0%
Average (4 open)	MMLU-Pro	90.1%	70.1%	87.7%	43.1%
GPT-5-mini	TruthfulQA	91.7%	88.3%	98.6%	74.2%
GPT-5.6-Terra	ARC-Challenge	93.3%	88.4%	97.0%	70.7%
GPT-5.6-Luna	TruthfulQA	89.0%	85.2%	82.4%	62.4%
GPT-5	MMLU-Pro	88.5%	85.7%	79.6%	57.1%

Table 14: Act-vs-escalate failures (act-on-wrong, %) per model under the verbalized confidence score of Wang et al. (2026), LAC sets, solo calibration, $\alpha = 0.10$. For the open models the TruthfulQA and MMLU-Pro columns are largely *degenerate* under LAC: the model is unconfident there, the threshold admits all options (marginal coverage $\to 1$), and no singleton can form, so the near-zero values are uninformative rather than safe. The proprietary models are not degenerate on those benchmarks, which is where the nonzero rates in the lower block come from. All open-model cells are $n = 500$ except DeepSeek-R1-Llama-8B on MMLU-Pro, which is $n = 95$ (†) from a run that was stopped early.

Model	ARC (4 opt)	TruthfulQA ($\leq$5 opt)	MMLU-Pro (10 opt)
Qwen2.5-1.5B	8.8	0.0	0.0
Qwen2.5-3B	2.3	1.6	0.0
Qwen2.5-7B	51.0	1.2	0.0
Qwen2.5-32B (AWQ)	93.7	4.5	11.7
Mistral-7B-v0.3	0.0	0.0	0.0
Llama-3.1-8B	63.5	0.0	0.0
Gemma-2-9B	12.8	1.4	0.0
DeepSeek-R1-Qwen-7B	51.2	0.0	0.0
DeepSeek-R1-Llama-8B	16.4	0.0	0.0†
GPT-5-mini	96.9	25.0	0.0
GPT-5	66.7	27.8	11.7
GPT-5.6-Luna	59.1	28.8	0.0
GPT-5.6-Terra	35.7	45.9	35.3

J CANONICAL NUMBERS AND REPRODUCIBILITY

All numbers below are from CPU-only analysis of pre-computed LLM outputs. Every number in the paper is regenerated from the analysis code; the headline values are also reproduced by a second implementation that draws its own calibration splits, agreeing to within 0.32 pp.

Fact	Value
Marginal coverage at round 1 (target 90%)	$90.0\% \pm 0.0\%$
Marginal coverage at unanimous-wrong (target 90%)	**73.9% ± 0.3%**
Gap	-16.1 pp
Gap on a wider grid (6 models, 7 task types, seeds 1–2)	-20.2 pp
Marginal coverage at round 1 (target 95%)	$95.0\% \pm 0.0\%$
Marginal coverage at unanimous-wrong (target 95%)	$86.4\% \pm 0.8\%$
Gap	-8.7 pp
No collapse at $\alpha = 0.01$ (target 99%)	$\approx 100\%$ all conditions
The rows below are the four-model average, seed 0, with the authority cue.	
10–20th pct, no pressure	87.3% (target 90%)
10–20th pct, unanimous-wrong	**47.4%** (target 90%)
Observing attacker, its least-confident 10%	$92.2\% \to 51.9\%$
Bottom-20% (reference, straddles threshold)	35.5%
Bottom-10% (reference, below threshold)	23.6%
Act-on-wrong	12.1% of flipped items
Randomized APS: marginal, that decile, act-on-wrong	77.2%, $92.1\% \to 58.6\%$, 12.3%
Act-on-wrong, four-model average, mean ± std over three seeds.	
With the authority cue	$13.5\% \pm 1.0$
Without the authority cue	$15.8\% \pm 1.2$
Randomized APS sets	$14.7\% \pm 1.7$
Worst single model (Qwen-7B, seed 0, shared threshold)	33.8% (target 90%)
Same model, mean over three seeds	37.5%
Same model, its own threshold	61.0%
Seeds used (marginal results)	3