

MediationBench: Auditing AI Mediation in Synthetic Conflict

Jonathan Hendler*
Human Assisted Intelligence, PBC (HAI.AI)
<https://mediationbench.com/>

Working draft — analysis snapshot 2026-08-03; judge-admission records through 2026-08-21;
revised 2026-09-01. Numbers may change between revisions.

Abstract

Do language models differ in how they simulate a dispute, and does an AI mediator change how a simulated dispute ends? MediationBench runs 62 authored two-party scenarios with and without a mediator, scoring each 0–100 using a maintainer-defined HAI Score not calibrated against human mediation quality or outcomes. The disputants matter most: on the same scenarios with no mediator, an assistant-aligned participant pair scores 36 and a high-resistance open-weight pair scores 18. Every frozen mediated gate arm scores above its matched no-mediator baseline under both, with larger gains against the resistant pair. Policy matters: with only public context, a structured Skilled policy exceeds a model-matched Reflective Listening control by $E_1 = +13.17$ [$+9.67, +17.55$] points on the Qwen backbone (prospectively specified; 90% scenario-reweighting interval), passing both frozen decision rules, which gpt-oss fails and which do not transport to Gemma participants. A post-outcome, suite-matched sensitivity kept both classifications, but its E_3 interval included zero; the positive interaction was unresolved. Skilled mediators do not separate from one another: within each participant configuration they form one tie band, a four-judge cross-family panel splits 2–2 on the order, and the order flips on re-generation.

Three limits bound these results. Moderator Quality is zero in baselines, so the composite partly encodes treatment, and mediated arms also bundle early stopping and extra bandwidth. Intervals reweight observed scenarios; they do not cover fresh generations or human disputes, and withheld scenario-level records preclude external recomputation. Synthetic conversations and mediation are experimental instruments: they enable controlled comparisons and repeatable failure analysis but do not substitute for validation with people in real conflicts.

1 Introduction

Can an AI mediator change how a simulated dispute ends, and can a composite score tell one skilled mediator from another? In the runs reported here, the answer changed with the participant configuration, for reasons these data do not isolate. When two assistant-aligned language models play the disputants, every tested skilled-mediator arm lands within a point of 54 on the HAI Score, a 0–100 composite with five maintainer-defined rubric components. When the disputants are the Stheno-v3.2 high-resistance open-weight configuration, the no-mediator cooperation curve never rises above 7.1% at any observed turn, and the same skilled arms spread apart (gate cells 44–48,

*Corresponding author: jonathan@hai.io.

extension rows reaching 52) over a baseline of 18. The two configurations differ in apparent dynamic range. They do not establish which participant property causes the change, and no between-column dispersion interaction was estimated. The pre-declared paired contrasts answer the narrower question of which mediator ranks higher: +2.1 $[-2.5, 4.5]$ on the cooperative column (the interval includes zero) and +2.3 $[0.04, 8.95]$ on the adversarial column (the lower bound only just excludes zero; Section 3.4).

This paper presents MediationBench, the synthetic-conflict testbed that produced these results. Mediated runs are compared with no-mediator runs on the same scenarios. A minimal Reflective Listening arm serves as an active control. A judge-reliability protocol checks the evaluator. Anti-gaming design (cherry-picking, contamination, provenance) is disclosed at the rationale level in Appendix A. The reported study uses a 62-scenario evaluated/development partition and leaves a 38-scenario sequestered reserve unrun.

Two design variables are treated as experimental factors rather than implementation details. The first is the participant model. The assistant-aligned pair often reached high within-harness resolution scores without a mediator, while Stheno-v3.2 sustained positional behavior; the two are confounded in capability, disposition, and alignment tuning, so every conclusion stays conditional on the recorded configuration (Section 2.5). The second is mediator information access. A model-matched 2×2 crosses the mediation policy (Skilled versus the Reflective Listening active control) with what the mediator is told (a *full bilateral brief* of both parties’ private scenario facts versus *public-context-only* access), yielding four frozen estimands: the public-context skill increment E_1 , the full-brief skill increment E_2 , their interaction E_3 , and the per-policy information contrasts E_4 (Section 2.8; results in Section 3.6).

End-state scoring alone conflates four descriptive questions, and the design keeps them apart: how mediated arms score against their paired baselines; how a skilled policy scores against minimal reflective listening (E_1 , E_2); how generated scores differ between private-brief and public-context conditions (E_3 , E_4); and how an observed ordering holds up under another participant run or another fixed-transcript evaluator (Section 2.10). The aim is not another model leaderboard. It is an instrument whose sensitivity to each of these choices is measured and reported.

MediationBench contributes two things. The first is a scoring harness with a measured resolution limit: every mediated frozen gate-cell contrast against its matched baseline has a scenario-reweighting interval above zero, yet skilled arms collapse into one tie band per participant column. This bounds what the current composite can distinguish inside the recorded harness; it is not a human-validity result. The second is a versioned benchmark and research artifact. It supports controlled information manipulations, matched scenario comparisons, known authored private-information assignments and answer keys, fixed-transcript evaluator audits, repeatable failure analysis, and product regression. Public artifacts are aggregate; private fixtures and scenario-level records remain withheld.

Anti-gaming infrastructure matters for mediation because a system can improve its visible scores without improving the interaction being measured. Hidden-state revelation, commitment symmetry, and process trajectories each expose distinct rubric signals and candidate failure modes. The deployment stakes (divorce proceedings, workplace conflicts, public deliberation) make a gameable quality signal worse than none [1].

1.1 Related work

General and social-agent evaluation. HELM [2] exemplifies multi-dimensional evaluation, and BIG-bench [3] introduced canary strings for contamination detection. Crowdsourced pairwise evaluation [4, 5] avoids static contamination but introduces style biases. SOTOPIA [6] evaluates social interaction quality in language agents and finds that models lag humans particularly in social

commonsense reasoning and strategic communication.

Negotiation benchmarks. DealOrNoDeal [7], CraigslistBargain [8], and NegotiationArena [9] score the *negotiating parties*’ self-interested outcomes (deal value, win rate); CaSiNo [10] also annotates rapport and negotiation strategies, and LLM-stakeholder interactive negotiation [11] measures cooperative and adversarial deliberation among stakeholders. None evaluates a third-party mediator’s effect on the conversation.

AI mediation and deliberation. Reward-model-mediated consensus statements [12] and the Habermas Machine [13] showed LLM caucus mediators whose common-ground statements participants preferred over human mediators’, and chat-intervention field experiments improved divisive political conversations at scale [14]; these evaluate bespoke systems with human participants. Rehearsal simulates dyadic interpersonal conflict with LLM interlocutors to teach conflict-resolution skill [15]. Robots in the Middle evaluates dispute analysis, intervention-type selection, and intervention-message generation on 50 manually authored scenarios through blind comparisons with human annotations [16]. That human-annotated reference is one this benchmark lacks. Robots in the Middle evaluates mediator decisions; it does not run complete mediated and no-mediator conversations on matched authored scenarios. ProMediate runs proactive agents in multi-party, multi-issue negotiation and measures consensus change, intervention timing, and social intelligence against no-agent and generic-agent conditions [17]. SoCRATES builds scenarios from real-conflict source material across eight domains, varies five socio-cognitive axes, measures the unmediated consensus gap, and validates a topic-localized evaluator against experts [18]. Its criticism that existing testbeds rely on a few expert-authored domains applies to our authored suite; our design budget went to measurement variation (paired baselines, controls, judge reliability, participant stress) rather than scenario sourcing. AgentMediation simulates staged dispute mediation for legal research, crossing mediator information access with disputant strategy profiles [19]. It is the nearest precedent for our information estimands and participant axis. Our E_4 differs: it compares information conditions within a model-matched policy pair against a matched no-mediator baseline. MediationBench is therefore not the first mediation benchmark. Its distinct focus is a measurement audit over authored dyadic conflicts: scenario-matched no-mediator comparisons, an active minimal-intervention control, participant-configuration stress tests, process trajectories, and crossed/test–retest judge reliability.

LLM-as-judge reliability. LLM evaluators recognize and favor their own generations [20] and exhibit systematic biases [21]. We turn this literature’s recommendations into benchmark infrastructure: a fixed operational judge distinct from all participants and mediators, a crossed second judge, and test–retest reliability per component.

Synthetic social-simulation validity. Model-on-model interaction can expose prompt, policy, information, and evaluator sensitivity under controlled conditions. But models share linguistic priors, and the resulting compatibility may not transport to people or even to another participant simulator. Alignment tuning is one documented source of model–human divergence: within fixed open-weight backbones, safety fine-tuning suppresses mind attribution (most clearly for non-human animals, rated well below a matched human panel) and spiritual belief, and ablating the learned safety-refusal direction moves General Social Survey responses closer to human reference distributions [22]. We therefore treat participant-family transport and human criterion calibration as separate validity gates, not as consequences of a high synthetic score.

Our work differs from these neighboring testbeds in measurement emphasis, not categorical priority. The methodological novelty we claim is the pairing discipline: every scenario is run both with and without a mediator on an identical authored partition, and a within-scenario co-resampled bootstrap preserves the covariance between arms. A shared baseline cancels in a contrast of two deltas only when both deltas reuse that same baseline; it does not cancel inside a single mediated-minus-baseline delta. We therefore report scores as matched descriptive deltas, not as absolute

conversation quality. Generation is unseeded, so each delta also contains stochastic-trajectory differences. The scope is deliberately *dyadic*, narrower than the multi-party settings of ProMediate and SoCRATES, to keep the matched-pair estimand clean and the participant axis interpretable. Process is measured alongside outcomes (per-turn cooperation dynamics and hidden-state revelation) under formal anti-gaming mechanisms, and participant-model selection, including the Stheno-v3.2 high-resistance open-weight configuration, is a first-class validity variable.

2 Methodology

2.1 Design overview

The frozen gate is a factorial matrix run on one fixed harness: two participant configurations $\times$ four arms, every scenario in the 62-scenario evaluated/development partition, and one replicate per cell (496 conversations). One extra Reflective Listening replicate was retained; we report it separately as a noise anchor. A single operational judge and one suite version (2.0.0) cover all cells. Table 1 lists the factors.

Factor	Levels
Participants (same-model pairs)	Claude Haiku 4.5 (assistant-aligned); Stheno-v3.2 (L3-8B) (Stheno-v3.2 high-resistance open-weight); Gemma-4-31B (registered post-freeze participant-transport extension; Section 3.6)
Arms	no-mediator baseline; Reflective Listening (minimal-intervention active control); Skilled@Qwen3-235B; Skilled@Kimi-K2.5
Operational judge	Qwen3-235B-A22B-Instruct (all cells; drives should-stop and per-round evaluations)
Crossed judges (offline)	Claude Sonnet 4.5 (prospectively specified), gpt-oss-120b, GLM-5.2, and Kimi-K2.5 (added post-freeze; see caption) re-score the one-replicate-per-cell gate set; none influences conversations
Scenarios	all 62 scenarios in the evaluated/development partition per cell; the 38-scenario sequestered reserve remains unrun for a future locked evaluation epoch
Methodology	seeded_25_v1: seeded fixture dialogue, participants continue to 25 public turns each; temperatures 0.3/0.7/0.8 (judge/mediator/participant)

Table 1: Frozen factor roster (design frozen 2026-07-04, before Results analysis). The Gemma-4-31B participant pair was added post-freeze as the registered participant-transport extension (Appendix B, item 10; its own design frozen 2026-07-25, pre-outcome). The frozen design specified one crossed judge (Claude Sonnet 4.5); gpt-oss-120b was added after the freeze but before Results were pinned, and GLM-5.2 and Kimi-K2.5 were added 2026-07-11–15 with a decision rule for the within-skilled-pair ordering stated before their scores existed (Section 3.4). Every crossed judge — including the post-hoc fifth (Appendix C) — is a strictly offline re-scorer of frozen transcripts, so the amendments cannot affect any generated conversation. GLM-5.2 returned valid scores for 528 of 558 evaluations (94.6%; its missing evaluations concentrate in the longest transcripts, where it exhausts its completion budget on chain-of-thought before emitting scores), so its contrasts are computed over jointly-scored scenario pairs with the joint n reported per contrast; the other crossed judges scored all 558.

The Reflective Listening arm is a *minimal-intervention active control*. It acknowledges and reflects what participants say, but it does not apply the structured mediation policy. It is an active comparator, not an inert placebo. The Skilled arm applies the structured policy: it acknowledges, then probes for undisclosed constraints and interests, names tradeoffs, and pushes toward specific reciprocal commitments. We withhold the verbatim policy prompt under the same rationale-level disclosure policy that covers the operational judge prompts (Appendix A); this behavioral description is the published level. Keeping Skilled@Qwen3-235B in the roster while Qwen3-235B-A22B-Instruct is the operational judge also gives us a judge–mediator generosity probe under the crossed judge (Section 2.10).

2.2 Scenario design

The fixtures are the instrument, so we state how we wrote them. The inventory holds 100 authored fixtures: the 62-scenario evaluated/development partition and the 38-scenario sequestered reserve, which has never been run. The maintainers wrote all of them for the benchmark, using an LLM for writing assistance under a written authoring guide, and curated each by hand before inclusion. None derives from a real dispute, transcript, or personal data (Ethics statement). The fixtures cover interpersonal and organizational conflict — workplace, family, commercial, and community disputes — at non-graphic intensity. Each fixture seeds a shared public scenario description; per-party briefs (backstory, stated goal, hidden facts and emotions); an answer key of expected revelations, used only by the deterministic Hidden-State Revelations overlap cap; and a unique canary string (Appendix A). The evaluated fixtures also shaped benchmark development. The reserve exists so a future locked epoch can test generalization beyond them.

2.3 The need for baselines

An absolute score means nothing on its own. We therefore run every scenario *without* a mediator under the same seeded protocol and report each mediated score as a delta against that matched baseline (same suite version, participant model, judge model, and scenario set). Baseline runs contribute 0 on the Moderator Quality component, so the highest possible baseline score is 85. A mediated arm that scores at the baseline shows no positive composite-score difference in this comparison. We flag negative deltas rather than suppress them. A negative delta may reflect over-intervention, bias, escalation, or generation-path variation, so it calls for transcript-level inspection, not a single causal diagnosis.

2.4 Run protocol

Each scenario runs in four steps. (1) The harness stores the fixture turns verbatim as a shared seed transcript. (2) The participants continue until each has 25 public turns. (3) In mediated arms the mediator may intervene after the seed. It may stop early only on a model-reported agreement event or a judge `should_stop`; the baseline arm has no stop event by design. (4) The operational judge scores the full transcript with the standard HAI Score formula. In all arms the harness also stores per-round judge evaluations for the cooperation curve (measurement-only in the baseline arm). Step (3) has three consequences: the stop policy is part of the arm definition; mediated transcripts run substantially shorter than baselines (quantified in Section 3.2); and the composite is computed on the transcript as stopped. Mediator, system, and private messages do not count toward the 25-turn participant target.

A run must be complete before we publish it. Any failed scenario row keeps the run out of the public export until we repair it with the original configuration. Repairs are triggered by completion

only — failed scenarios or zero-token generations, never by score — and are recorded in the analysis code and internal run records; the public export carries no per-run repair log. One replicate in the Qwen public-context Reflective Listening arm, which the primary E_1 contrast uses, needed two completion-repair passes after provider congestion affected 27 of 62 scenarios. Final coverage was 62/62, and we verified the public-context stamps on the replacement dialogues. The repair, however, crossed serving epochs, and the design cannot separate that epoch effect from generation stochasticity in the replicate.

Information provisioning. Who is told what is a design factor, so we state it exactly. Each participant receives only its own brief (backstory, stated goal, hidden facts and emotions) plus the shared scenario description. The gate matrix uses the *full bilateral brief*: both mediated arms — Reflective Listening as well as the skilled mediators — receive the shared description and both parties’ researcher-authored ground-truth private scenario facts verbatim. They do not receive the parties’ backstories, stated goals, hidden emotions, or the scoring answer key. The *public-context-only* condition of the prospectively specified, frozen 2×2 information study supplies the shared scenario context and the public dialogue as it unfolds, but neither party’s private facts.

The full bilateral brief is a privileged-information condition. Intake or caucus preparation is only an imperfect analogy for it. It is not an observed or simulated interview: no elicitation occurs, and the researcher-authored private facts arrive verbatim. The contrast should therefore not be read as an estimate of the value of real-world interviewing or case preparation. Within each condition the active control and the skilled policy receive the same information, so their score increment compares policies at fixed information. Information was fixed across the original gate arms, but the backbone model was not. The Reflective Listening moderator carries no model binding, so its calls fell to the harness default (`claude-sonnet-4-5`, recorded in the export’s `mediator_model` field), while the skilled arms ran their named backbones. The original gate increment therefore compares mediator *packages* (policy plus backbone). The model-matched 2×2 information study (Appendix B) crosses policy and information within each backbone.

Every judge, operational or offline panel, receives the same reconstructed judge-facing content from the harness: the scenario title and description, both parties’ backstories, archetypes, and stated goals, and the full transcript. No judge ever sees the hidden facts or the Hidden-State Revelations answer key; code applies the overlap cap mechanically after judging. Two implications follow. In the full bilateral brief condition, the Hidden-State Revelations component measures whether the conversation *surfaces* hidden state already present in the mediator’s packet, not unaided discovery (Section 2.6.1). In the public-context-only condition, the harness does not supply private facts verbatim and participant disclosure is the intended route, but model inference, guessing, or hallucination can still produce matching content. Directional `focus_areas` also reach the mediator, and the public dialogue travels through a sliding message window rather than the persistent private-brief path. The condition is therefore not information-free (Appendix B).

2.5 Participant configurations as a validity variable

LLMs configured as disputants play the participants, so which participants we choose is a validity question, not an implementation detail. The gate design compares an assistant-aligned configuration with the Stheno-v3.2 high-resistance open-weight configuration. In the observed runs, the assistant-aligned pair often moved toward resolution without intervention; this produced high baselines and compressed the differences between mediator arms. Stheno-v3.2 sustained more positional behavior, which yielded lower no-mediator baselines and a wider observed point-score range in this stress test.

These observations validate neither configuration as a model of human behavior. Nor do they explain why the configurations differ: model capability, conflict disposition, and alignment tuning change together and are confounded. “Open-weight” describes access to the model artifacts; it does not explain the observed behavior, and it does not remove variation from unseeded generation or serving implementations. The benchmark records the exact participant configuration for each run and reports results by configuration. Separating capability, disposition, and alignment would require a future factorial participant study; the present comparison supports only configuration-conditional claims.

External evidence makes the alignment-tuning part of this confound independently plausible. Ablating the learned safety-refusal direction [23] within a fixed open-weight backbone moves instruction-tuned models’ closed-ended General Social Survey responses closer to matched human response distributions. It also recovers suppressed mind attribution — approaching the human panel average for non-human animals, and rising to or above it for chatbots, technology, and natural entities — while leaving theory-of-mind and general-reasoning accuracy intact [22]. That result concerns closed-ended survey responses from small open-weight models, not interactive conflict behavior, so it validates neither participant configuration here. It does support treating alignment tuning as a distinct axis of the confound, and its within-backbone ablation design is a template for the factorial study above: safety-ablated and unablated variants of one participant backbone would separate an alignment-tuning manipulation from backbone and capability, which the present cross-model comparison cannot.

2.6 The HAI Score: what we measure

The HAI Score is a single number from 0–100 per scenario. The benchmark maintainers defined it. It is a weighted sum of five component scores, each also on the composite’s range, and each measuring one specified feature of a generated conversation inside this harness. Table 2 names the five components and gives their weights; the weights sum to one. Eq. (1) states the sum. A composite of zero means every component scored zero. Human criterion validity is unestablished.

Component	Weight	What it operationalizes
Cooperative Dimensions	25%	Coverage-floored share of 14 breakdown dimensions cooperative
Resolution Depth	25%	Topic progression (alignment → agreement → commitment → action)
Hidden-State Revelations	20%	Weighted revelations (level 3–5), capped by answer-key overlap F_1
Intent Commitment Symmetry	15%	Balance of mutual obligations
Moderator Quality	15%	Average of 8 moderation metrics; set to zero in baselines

Table 2: HAI Score components. Weights are asserted by the benchmark maintainers, not derived: fixed constants of the scoring harness (Eq. (1)), set before any gate data were generated and unchanged throughout the study (Section 4). “Moderator” is the harness’s legacy field name for the mediator role.

The composite is a construct the maintainers specified to track two things: how far the parties move toward settlement, and how much private information comes into the open. It is not a general measure of mediation quality. A high composite is not evidence that a mediation is safe,

fair, consensual, legal, feasible, durable, or effective for people. The rubric does not score party self-determination, mediator impartiality or procedural fairness, confidentiality breaches, coercion, power imbalance, consent to disclose, or whether non-agreement or a safe exit is appropriate.

Baseline runs have no mediator, so Moderator Quality is zero there by construction. A mediated-minus-baseline delta therefore mixes four components that both arms can earn with one component only the mediated arm can earn. We report results per component and return to this limitation in Section 4.

Four rules of the scoring harness shape the reported values. We state them here as specification. First, Resolution Depth is the mean over the topics the judge detects, so every topic that is raised but not closed lowers it. Second, Cooperative Dimensions divides the count of cooperative dimensions by the larger of the number detected and 14, a coverage floor over the 14 breakdown dimensions. A transcript in which the judge detects only a few dimensions therefore cannot read as fully cooperative. Third, Intent Commitment Symmetry defaults to 0.5 when the judge detects no commitments, so a transcript with no commitments reads as symmetric, not asymmetric. Fourth, the harness compares the hidden facts the judge reports revealed with the facts seeded in the scenario, using deterministic text-overlap precision, recall, and F_1 . The F_1 serves only as a cap on the revelation component; it never raises it. When ratings are available, the public export can also report per-component inter-rater reliability — Cohen’s κ , Krippendorff’s α , sample size, rater kinds, and calibration time.

2.6.1 Hidden-state revelation: construct and measured floor

Each private constraint in a scenario is in one of three visibility states: *hidden*, *hinted*, or *revealed*. A constraint is *hidden* while only one party knows it; that party holds a strategic advantage, and joint problem-solving is blocked. It becomes *hinted* when that party alludes to a constraint without saying what it is; this signals openness while keeping options open. It becomes *revealed* when the party states it clearly enough for the conversation to address it; joint problem-solving is then possible. Breakthroughs in conflict resolution often occur when the hidden interests behind stated positions become explicit [24, 25]. A mediator who helps put a constraint on the table can therefore advance joint problem-solving — but only when the affected party consents voluntarily and confidentiality is protected, and this score does not measure either safeguard. The information condition (Section 2.4) limits what the component can mean. Under the full bilateral brief, the mediator receives the authored private-information assignments verbatim in its packet. Revelation then measures whether those assignments surface in the conversation, not whether the mediator discovered them unaided, and it cannot by itself separate elicitation skill from use of the brief. In this snapshot the component sits near floor in every arm (Section 3.3). We keep it for two reasons: it is the benchmark’s most direct claim about process, and its floor behavior is itself informative. In the test–retest sample the judge assigns high raw revelation ratings, yet the deterministic overlap cap drives the scored component to the floor (Section 3.4).

AI participants also write two separate streams: private thoughts and public messages. The gap between the two yields a belief–behavior divergence signal, the distance between hidden intent and stated position. That signal is not part of the composite. It is not populated for the runs in this snapshot. It remains future instrumentation (Appendix B).

2.7 Estimand: the per-cell score delta

Fix a *cell* c : one (participant model, mediator arm) combination, evaluated under a single operational judge, a single benchmark suite version, and the fixed evaluated/development scenario set $\mathcal{S}$ ($|\mathcal{S}| = 62$). For a completed scenario s , the judge’s evaluation yields five component scores

$X_{\text{cd}}, X_{\text{rd}}, X_{\text{hr}}, X_{\text{cs}}, X_{\text{mq}} \in [0, 100]$ (cooperative dimensions, resolution depth, hidden revelations, commitment symmetry, moderator quality). The harness combines them into the per-scenario HAI Score

$$S(s) = \text{clip}_{[0,100]}(0.25 X_{\text{cd}}(s) + 0.25 X_{\text{rd}}(s) + 0.20 X_{\text{hr}}(s) + 0.15 X_{\text{cs}}(s) + 0.15 X_{\text{mq}}(s)), \quad (1)$$

where the weights are fixed constants of the scoring harness (Section 2.6). Eq. (1) is the continuous per-scenario composite that the operational-judge analysis pipeline computes. Two stored quantities round it to integers: the export’s run-level scores and the per-scenario crossed-judge re-scores. Every crossed-judge delta and interval endpoint therefore lies on a half-integer grid, while operational-judge values are continuous; we flag this granularity difference wherever the two are compared (Table 6). A scenario marked *failed* has no score: $S(s)$ is undefined and s is excluded from all pairing below.

Matched baseline. For each gate and extension cell, the analysis registry pins one baseline (no-mediator) run. That run must be public and complete, contain *no* failed scenario, and agree with the mediated run on all four matching keys: suite version, participant model, judge model, and the exact scenario set $\mathcal{S}$. We never use pooled cross-model baseline statistics. We report and analyze post-freeze repeat runs separately and never pool them automatically. The only disclosed replicate pooling in this paper is the information study, which takes a per-scenario *mean* of an arm’s pinned replicates before the scenario median (Section 2.8).

Estimand. Let $\mathcal{S}_c^* \subseteq \mathcal{S}$ be the scenarios with a defined score on *both* arms. Write $S_c^{\text{med}}(s)$ and $S_c^{\text{base}}(s)$ for the mediated and matched-baseline per-scenario scores. The published score delta of cell c is the difference of medians over the common scenario set,

$$\Delta_c = \text{median}_{s \in \mathcal{S}_c^*} S_c^{\text{med}}(s) - \text{median}_{s \in \mathcal{S}_c^*} S_c^{\text{base}}(s). \quad (2)$$

Δ_c is a difference of medians, not a median (or mean) of paired differences. The alternatives disagree on skewed data, and only the difference of medians reproduces the reported interval-backed deltas. Two aggregations circulate, and we keep them distinct throughout. The published run-level score is the *mean* of per-scenario composites, so the export’s run-level delta (Δ) is a difference of means. The paper’s estimand Δ_c — written Δ_{med} in tables and prose — is the difference of medians of Eq. (2) over the jointly-completed subset. The two differ even when both arms are complete on the identical scenario set: on the headline cell the export delta is +30 while $\Delta_{\text{med}} = +36.9$. Table 3 therefore reports both. We report every Δ_c per cell; no delta is pooled across participant models, arms, judges, or suite versions. Δ_c is a descriptive comparison of observed runs, not an identified causal effect: language-model generation is unseeded, and the mediated and no-mediator arms also differ in stopping opportunities (Section 2.4).

2.8 Information-study estimands

The model-matched 2×2 information study (Sections 2.4 and 3.6) fixes one mediator backbone and one participant column. It crosses policy $p \in \{\text{Sk}, \text{RL}\}$ (Skilled; Reflective Listening) with information condition $i \in \{\text{full}, \text{pub}\}$ (full bilateral brief; public-context-only). For arm (p, i) , let $M_{p,i}$ be the median of the per-scenario score over the scenarios jointly completed in all four arms; where an arm has more than one replicate run, the per-scenario score is the mean of its replicates. The four frozen estimands are

$$\begin{aligned} E_1 &= M_{\text{Sk}, \text{pub}} - M_{\text{RL}, \text{pub}}, & E_2 &= M_{\text{Sk}, \text{full}} - M_{\text{RL}, \text{full}}, & E_3 &= E_1 - E_2, \\ E_{4, \text{reflection}} &= M_{\text{RL}, \text{pub}} - M_{\text{RL}, \text{full}}, & E_{4, \text{skilled}} &= M_{\text{Sk}, \text{pub}} - M_{\text{Sk}, \text{full}}. \end{aligned} \quad (3)$$

Intervals come from a four-arm cluster bootstrap that co-resamples scenarios. Each draw resamples once the scenarios jointly completed in all four arms (which preserves cross-arm correlation), holds each arm’s observed per-scenario replicate mean fixed, and recomputes the same functional as the point estimate. The interval is therefore conditional on the recorded generations and does not estimate run-to-run variation. An earlier implementation resampled replicate identity independently within scenarios, which created hybrid pseudo-runs inconsistent with that scope. The correction changed the emitted intervals, not any estimand, margin, or decision rule.

We froze two decision rules on 2026-07-15, before any public-context run existed: E_1 positivity — the lower endpoint of E_1 ’s 90% interval exceeds zero — and E_3 non-inferiority — the lower endpoint of E_3 ’s 90% interval exceeds -4.0 composite points. The margin governs E_3 only. It is not a general materiality threshold, and we specified no equivalence margin for either E_4 estimand. The frozen rules compare the unrounded bootstrap endpoints; the two-decimal values in tables and prose are display values only. The margin’s rationale drew on observed full-brief data (about half the frozen full-brief skill increment, and above the observed replicate movement of 3.7 points), but it could not have used any public-context run: the owner signed the rules on 2026-07-15, and the first suite-2.5.0 run launched on 2026-07-17. The dated freeze record is item 9 (Appendix B). On 2026-07-26, after the breadth outcomes were known, we added a separate assay-sensitivity reporting safeguard post hoc: we withhold substantive non-inferiority interpretation where a backbone’s E_2 interval includes zero, because the margin could then permit attenuation larger than the observed reference contrast. This later safeguard changes no estimand, threshold, or stored frozen-rule result.

2.9 Uncertainty: scenario-reweighting bootstrap

Every reported interval on Δ_c comes from a scenario-level cluster percentile bootstrap on the matched pairs $\{(S_c^{\text{med}}(s), S_c^{\text{base}}(s))\}_{s \in \mathcal{S}_c^*}$, $n_c = |\mathcal{S}_c^*|$. Read these as scenario-reweighting intervals conditional on the observed generation run or runs. They quantify how the estimator changes when the evaluated scenarios receive different weights; they do not capture variation over newly authored scenarios, fresh language-model generations, or serving implementations. For draws $b = 1, \dots, B$ with $B = 10,000$: (1) resample n_c scenario identifiers from $\mathcal{S}_c^*$ with replacement — whole scenarios are the resampling unit, never individual turns; (2) recompute the *published estimator* of Eq. (2) on the resample. The 90% interval is the percentile interval

$$\text{CI}_{90}(\Delta_c) = \left[q_{0.05}(\Delta_c^{(1)}, \dots, \Delta_c^{(B)}), q_{0.95}(\Delta_c^{(1)}, \dots, \Delta_c^{(B)}) \right], \quad (4)$$

with no interpolation: after sorting the B draws, the lower endpoint uses zero-based index $\lfloor 0.05B \rfloor$ and the upper endpoint uses index $\lfloor 0.95B \rfloor - 1$ (for $B = 10,000$, the 501st and 9500th order statistics). Because pairs are co-resampled, both arm medians share the resampled scenario set in every draw, so the interval preserves any within-scenario covariance between arms.

The bootstrap RNG is Python’s Mersenne Twister (`random.Random`). The operational analysis uses fixed seed 20260706. The crossed-family readout uses the same B , percentile convention, and RNG with seed 20260728, as recorded in its dated artifact. These seeds govern resampling only; LLM token sampling is unseeded and unrelated to them. Within each analysis we re-create the same seeded stream for every reported interval. Draws are coupled where intervals share the same scenario count and traversal convention. Different counts can consume the pseudorandom stream differently, so Monte-Carlo error is not generally common across all intervals, and because seeds are reused, those errors were not designed to be independent. At $B = 10,000$ the binomial standard error of an empirical 5%/95% quantile index is about 0.2 percentage points of quantile level, and interval endpoints quoted to two decimals carry that Monte-Carlo resolution. Where a frozen estimand has

repeated runs (Section 2.8), we compute the observed per-scenario replicate mean before resampling and then hold it fixed. Only scenario weights vary across draws, so the interval is conditional on the finite set of observed generation runs and contains no run-sampling component.

Three reporting conventions guard against over-reading these intervals. We state them once here.

Discreteness caveat. With $n_c \approx 62$ integer-valued per-scenario scores, the bootstrap distribution of a difference of medians is supported on a small set of observed order-statistic differences. On discrete scores the percentile bootstrap can therefore under- or over-cover in finite samples. Read zero-width or unusually short intervals as discreteness artifacts, not as certainty. The frozen analysis plan prospectively specified smoothed-bootstrap, BCa, and Hodges–Lehmann alternatives (Appendix B).

Within-column contrasts. Two arms in the same participant column share a baseline and a scenario set, so their deltas are strongly positively correlated, and comparing their marginal intervals over-declares ties. We therefore estimate the paired contrast between arms a and a' by co-resampling scenarios and recomputing $\Delta_a^{(b)} - \Delta_{a'}^{(b)}$ per draw, which cancels the shared baseline. Where we show only marginal intervals, we label arm comparisons informal and conservative.

Tie-band decision rule (pre-declared, not a test). The composite HAI Score is the frozen primary estimand; the five components are exploratory. For *ranking* we apply a pre-declared conservative decision rule: two cells or mediators whose 90% intervals overlap are reported as a tie band rather than ordered. This is a ranking convention, not a hypothesis test. Overlap of two intervals is not a claim of no difference, and no p -value is computed or implied.

2.10 Judge reliability protocol

The judge is a noisy measurement instrument, so we measure the noise rather than assume it away. We compute reliability on ordinal recodings of judge output. Each component score, on its native $[0, 1]$ judge-output scale (the $[0, 100]$ component scores of Eq. (1) are these values rescaled by 100), is bucketed as $r = \min\{5, \max\{1, \lceil 5x \rceil\}\}$, a 1–5 rating per (transcript, component, rater). Each sub-study writes all of its rating passes under one fresh, dedicated calibration snapshot, so no other raters enter the pool.

Intra-judge test–retest. The operational judge re-scores an identical frozen sample of $n = 40$ transcripts, stratified across the eight gate cells over 40 distinct scenarios, $K = 3$ times; each pass enters as a distinct rater. Per component we report Krippendorff’s

$$\alpha = 1 - \frac{D_o}{D_e}, \tag{5}$$

with the nominal difference function, where D_o is the observed within-transcript pairwise disagreement and D_e the disagreement expected under the pooled marginal distribution. Cohen’s κ is a two-rater statistic, so we report no pairwise κ for the $K = 3$ block (a reporting choice: the pipeline’s κ would use an arbitrary pair of passes). α is the headline figure, read against the conventional bands $\alpha \geq 0.80$ (reliable) and $\alpha \geq 0.667$ (tentative). Nominal α treats all disagreements between ordinal buckets as equally distant; the frozen analysis plan prospectively specified ordinal-distance variants. These coefficients describe per-component bucket agreement only. We compute no reliability coefficient for the continuous composite or for Δ_c itself, a scope limit stated in Section 4.

Crossed two-judge block. A second, cross-family judge (Claude Sonnet 4.5) re-scores all nine frozen runs offline: the 496-transcript one-replicate-per-cell gate set plus the retained Reflective Listening replicate, 558 transcripts in total. That total is the denominator of every inter-judge statistic below. The crossed judge never influences the conversations. This is a fully crossed two-rater design on identical transcripts, so per component the harness computes both unweighted Cohen’s

$$\kappa = \frac{p_o - p_e}{1 - p_e}, \quad p_e = \sum_{k=1}^5 p_k^{(1)} p_k^{(2)}, \quad (6)$$

and nominal Krippendorff’s α on the same paired codes. Section 3.4 reports the resulting composite-level agreement and run-level reproduction. We later extended the same offline protocol to the full four-judge panel (gpt-oss-120b, GLM-5.2, Kimi-K2.5; Table 1 caption); their scenario-level rank correlations and pre-stated ordering vote also appear in Section 3.4. LLM judges are documented to favor their own generations and to show systematic biases [20, 21]. Our operational judge is distinct from every participant model, and one mediated arm shares the operational judge’s backbone, which enables the probe below.

Evaluator-context sensitivity. A separately registered fixed-transcript block uses the same requested Sonnet 4.5 scorer alias twice: once with the harness-authored private briefs visible and once with those briefs withheld. The public scenario description and conversation are identical in each pair. The registered orientation is context-withheld minus context-visible. We report differences of arm medians for each run, shifts in each mediated-minus-baseline contrast, and shifts in the within-column policy contrasts. Every estimate uses the exact common scenario set and a joint scenario-resampling percentile interval ($B = 10,000$, fixed seed); generations and observed score rows are held fixed. The visible reference scores predate registration, and the two conditions span different serving epochs, so this is a descriptive evaluator-context sensitivity, not a randomized causal effect.

Judge-by-arm generosity contrast (self-preference probe; descriptive). Let $\tilde{S}_{j,m}$ denote the run-level composite score that judge $j \in \{Q, S\}$ (operational, crossed) assigns to the transcripts of mediator backbone $m \in \{Q, K\}$. This is the export’s **haiScore**, a mean of per-scenario composites stored as a rounded integer, so the contrast below has a one-point resolution floor. Within each backbone both judges score identical transcripts; the two backbones’ transcript sets are different conversations. The contrast

$$\text{SP} = (\tilde{S}_{Q,Q} - \tilde{S}_{S,Q}) - (\tilde{S}_{Q,K} - \tilde{S}_{S,K}) \quad (7)$$

measures how much more generously the operational judge scores its own family’s mediated arm than the crossed judge does, net of the same judge difference on the other backbone’s arm. $\text{SP} > 0$ is *consistent with* self-preference but is not identified as such: any judge-pair disagreement that scales with transcript quality or style loads onto SP. The probe is also not one-sided. The crossed judge’s model family does appear among the gate mediators, because the Reflective Listening control’s calls fell to the harness default (`claude-sonnet-4-5`; Section 2.4), so every crossed-judge skilled-over-control contrast pairs a Sonnet judge with a Sonnet-backed comparator. A single judge pair yields no judge-level uncertainty: a scenario-cluster bootstrap can show SP’s scenario-reweighting variation conditional on the observed transcripts, but it cannot supply judge-population uncertainty. We report SP descriptively alongside the separate multi-judge sensitivity panel.

3 Results

All gate aggregates come from the frozen public-export snapshot (SHA-256 prefix 0ab9bb26); analyses run after the freeze are labeled as such where they appear. Every interval is a 90% scenario-reweighting interval (Section 2.9): scenario-level records go through the committed analysis pipeline (`tools/compute_claims.py`; scenario-cluster percentile bootstrap, $B=10,000$, fixed seed), and each interval is conditional on the observed generation runs and the operational judge. It does not estimate run-to-run variation or performance in human disputes, and the aggregate public export alone cannot reproduce it. We state this once here and do not repeat it per estimate. The frozen snapshot holds the complete 8-cell matrix with 62/62 evaluations per cell; the evaluated development partition is not public.

Overview. Four results organize this section. *Participants set the range* (Section 3.1): the two participant configurations produce baselines of 36 and 18 on the same scenarios, and skilled mediators form a narrow band against one and spread against the other. *Mediated arms clear their matched score baselines* (Section 3.1): every frozen gate-cell interval excludes zero, and the arm-versus-control contrasts stay positive under the three higher-correlating crossed judges (Section 3.4). *Policy beats minimal reflection on a fixed backbone* (Section 3.6): the replicated Qwen public-context-only contrast is $E_1 = +13.17 [+9.67, +17.55]$ and passes both frozen rules; gpt-oss passes neither; the Gemma participant-transport extension fails the frozen transport criteria, and two offline judge families keep its negative interaction direction. *Skilled mediators do not separate from one another* (Section 3.4): a four-judge crossed panel leaves the ordering within the skilled pair judge-indeterminate, and the ordering flips on re-generation (Section 3.5).

3.1 Mediated-minus-baseline score differences by cell

Participant pair	Arm	HAI Score	Δ	Δ_{med} [90% CI]
Claude Haiku 4.5 (cooperative)	No mediator (baseline)	36	—	—
	Reflective Listening	49	+13	+8.4 [6.9, 12.2]
	Skilled@Qwen3-235B	54	+18	+14.9 [11.7, 17.0]
	Skilled@Kimi-K2.5	54	+18	+12.8 [10.5, 16.7]
Stheno-v3.2 (L3-8B) (adversarial)	No mediator (baseline)	18	—	—
	Reflective Listening	40	+22	+29.1 [19.6, 33.6]
	Reflective Listening (replicate)	41	+23	+31.2 [23.6, 35.3]
	Skilled@Qwen3-235B	48	+30	+36.9 [31.9, 41.4]
	Skilled@Kimi-K2.5	44	+26	+34.6 [27.3, 37.1]

Table 3: Gate-cell results (frozen snapshot; judge Qwen3-235B-A22B-Instruct, suite 2.0.0, 62 scenarios per run; nine rows: the 8 gate cells plus the retained Reflective Listening replicate, reported separately, never pooled). Δ is the published run-level delta (a difference of run means); Δ_{med} is the paper’s estimand (Eq. (2)): difference of medians against the matched no-mediator baseline over jointly-completed scenarios, with scenario-cluster percentile-bootstrap 90% intervals. Units: HAI Score composite points.

Participants set the range. Under the cooperative participant the baseline already scores 36, and every skilled mediator lands in a narrow band (54–55) where the mediators no longer separate from one another. Under the adversarial participant the baseline falls to 18 and the same mediators

spread over a wide range. These column effects — baseline levels and delta sizes — hold across the complete matrix and under both crossed judges; the fine ordering within the skilled pair does not survive the open-weight judge (Section 3.4). The two configurations differ at once in backbone, capability, disposition, and alignment tuning, and no between-configuration dispersion interaction was estimated, so these data do not identify which property causes the difference.

Every mediated gate cell scores above its matched baseline (Table 3; every frozen interval excludes zero). The differences are larger under the adversarial participant (+30 and +26) than under the cooperative one (+18 and +18). The aggregate gains coexist with negative scenario-level differences, which the rule of Section 2.3 flags but does not suppress: in each gate cell, 0–8 of 62 scenarios score *below* their matched generated baseline — 0 for Skilled@Qwen3-235B under Stheno-v3.2 (L3-8B), 8 for the Reflective Listening control under Claude Haiku 4.5. Generation stochasticity is material: in selected baseline re-runs, the identical no-mediator configuration moved 56 (adversarial) and 57 (cooperative) of 62 scenario scores in one direction or the other (Section 3.5). Those re-runs are neither a random sample of repeats nor a bound on how much of any negative-delta count is policy-associated.

The active control is the policy comparator. The minimal Reflective Listening control recovers +22 of the +30 achieved by the strongest gate-cell Skilled policy. Part of that gain is rubric arithmetic: baseline runs score Moderator Quality as 0 by construction (Section 2.3), and the judge awards the control 90.8 on that component, so about 13.6 points of the control’s *mean* composite gain are available to any mediated arm by construction. The rest is an arm difference in the four shared components, including Cooperative Dimensions (Section 3.3), which we cannot attribute specifically to interleaved reflection because the arms also bundle stopping, conversational bandwidth, and generation stochasticity. The analysis contrast is therefore how far the Skilled policy exceeds this active control. On the adversarial column the gate skilled arms lead the control by +4 to +8 points (+4 to +12 once the single-replicate extension rows of Section 3.5 are included); this contrast is distinct from the v1 registry’s raw median-of- N leaderboard rank, which bound no paper row (Appendix A). The co-resampled contrast (Section 2.9) gives this increment an interval for one arm: Skilled@Qwen3-235B exceeds the control by +7.8 [2.8, 17.7] under the operational judge, and the interval excludes zero under both crossed judges (+11.5 [7.0, 15.5] frontier; +4.0 [0.5, 10.0] open-weight). The frontier judge shares a model family with the control’s harness-default Sonnet backbone (Section 2.4), so its skilled-over-control contrasts are not independent of judge family. The Skilled@Kimi-K2.5 increment (+5.5 [−1.8, 14.4]) does not exclude zero under the operational judge. On the cooperative column the increments are smaller (+6.5 [1.8, 8.4], +4.4 [0.4, 8.1]), and their crossed-judge intervals span zero.

Table 4 reports a post-hoc sensitivity, registered after the component results were known, for all nine gate rows: it removes Moderator Quality, renormalizes the four shared component weights to sum to one, and compares each mediated run-level aggregate with its matched baseline.

Four caveats bound this decomposition. The published cell delta of Eq. (2) is a difference of medians and does not decompose linearly, so the split is made at the component level, not on Δ itself. The split is data-dependent because Moderator Quality saturates in this snapshot (Section 3.3). The two control runs differ by one point (40, 41), and across four targeted post-freeze repeat cells the largest run-level movement is 3.7 composite points (Section 3.5). And those selected repeats do not bound run-to-run variation and fall short of the prospectively specified paired-by-seed design (Appendix B). The decomposition is a split of the mean gain, not a significance claim about any mediator pair.

Participant	Arm	Published Δ	MQ contribution (share)	Without-MQ Δ
Claude Haiku 4.5 (coop.)	No mediator (baseline)	+0.0	+0.0 (—)	+0.0
	Reflective Listening	+13.0	+14.3 (112.1%)	-1.8
	Skilled@Qwen3-235B	+18.0	+14.8 (84.0%)	+3.3
	Skilled@Kimi-K2.5	+18.0	+15.0 (84.7%)	+3.2
Stheno-v3.2 (L3-8B) (adv.)	No mediator (baseline)	+0.0	+0.0 (—)	+0.0
	Reflective Listening	+22.0	+13.6 (61.7%)	+10.0
	Reflective Listening (replicate)	+23.0	+13.6 (58.5%)	+11.4
	Skilled@Qwen3-235B	+30.0	+14.2 (47.5%)	+18.4
	Skilled@Kimi-K2.5	+26.0	+14.4 (55.7%)	+13.5

Table 4: Post-hoc Moderator Quality sensitivity for all nine frozen gate rows. The final column removes the 15% mediator-only component and renormalizes the remaining weights; the preceding column gives the removed contribution and its share of the unrounded run-mean change. These run-level sensitivities are not replacements for the frozen composite or algebraic decompositions of the difference-of-medians estimand.


3.2 Cooperation over participant turns

Figure 1 shows the same column structure over turns. The adversarial no-mediator baseline never exceeds 7.1% median cooperation at any of its 24 observed turns — the disputants keep fighting — while the cooperative-column baseline reaches high within-harness cooperation scores without mediation. The descriptive areas under the per-turn median curves separate the same way: 1504 for the cooperative baseline versus 25 for the adversarial baseline; on the adversarial column, 229 for the Reflective Listening control and 1161 for the strongest skilled arm. These are point summaries without intervals; the arms do not share a common turn support; and the mediated-arm curves are conditioned on survivorship, as the caption explains.

Mediated gate conversations ran much shorter than baselines. Baselines run the full protocol (53.2 and 58.4 transcript turns per scenario on the cooperative and adversarial columns); every mediated gate arm ends 46%–65% earlier (18.6–28.7 turns per scenario). Baseline conversations are cut off at the protocol cap, so time to the recorded agreement-or-stop event can be estimated only within mediated arms (Appendix B, item 3), and the survivorship conditioning above is quantitatively large, not a technicality. Arms that tie on the composite need not tie on economy: within the cooperative column’s 54–55 band, Skilled@Kimi-K2.5 matches Skilled@Qwen3-235B’s composite with 35% fewer turns. Turn totals are published per run in the export; this is descriptive, not a ranked outcome.

3.3 Per-component breakdown (exploratory)

Four component-level facts hold across the snapshot (Table 5) and shape how the composite should be read. First, Cooperative Dimensions (adversarial column: 19.9 baseline $\rightarrow$ 45.3 echo $\rightarrow$ 55.5–64.2 skilled gate cells (extension rows reach 78.6)) and Resolution Depth (22.6 $\rightarrow$ 29.1 $\rightarrow$ 34.8–43.6) carry almost all of the separation between arms. Second, Moderator Quality *saturates*: it scores 90.79–99.98 in every mediated arm, gate and extension alike, including the Reflective Listening control at 90.8. So 15% of the composite acts as an indicator that a mediator was present rather than as a measure of skill — direct evidence for labeling the composite weights maintainer-asserted,



	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22
Haiku 4.5 / Skilled	62	59	46	33	23	12	7	5	4	3	2	1	1	1	1	1							
Haiku 4.5 / Reflective	62	60	59	54	45	37	29	22	16	12	8	6	3	2	2	2	1	1					
Stheno-v3.2 / Reflective	62	61	59	58	46	35	25	17	12	10	4	2											
Stheno-v3.2 / Skilled	62	59	46	37	24	15	7	5	4	3	1	1	1	1	1	1	1	1	1	1	1	1	1

Figure 1: Cooperation over participant turns for the six roster gate cells (per participant column: no-mediator baseline, Reflective Listening control, and strongest gate skilled arm; column by linestyle, arm by colour). Participant-turn axis; mediator turns excluded. Points are per-turn medians across scenarios; descriptive, no smoothing. The table beneath gives the per-turn number of active conversations (n -at-risk) for the mediated arms; baseline arms have no stop event, so their n stays at 62. Per-turn values are conditional on conversations still active at that turn: mediated runs stop early on a model-reported agreement event or judge-directed stop while the baseline has no stop event, so late-turn mediated composition shifts toward unresolved scenarios.

and the motivation for the post-hoc sensitivity of Table 4. Third, Intent Commitment Symmetry moves *against* treatment in three of the four adversarial mediated rows and in the cooperative Skilled@Kimi-K2.5 cell. Fourth, Hidden-State Revelations, the construct’s nominal primary progress indicator, stays near the floor (8–17) in every arm, baselines included. The test–retest distribution shows why: 111 of 120 of the judge’s raw revelation ratings sit in the top bucket (Section 3.4), while the deterministic answer-key overlap cap pushes the published component toward the floor. This mechanical disagreement does not establish criterion-valid detection, and conclusions at the level of this component are unsafe.

3.4 Judge reliability

Crossed re-judge (identical transcripts). Two judges from other model families re-scored all 8 gate cells plus the Reflective Listening replicate (nine runs, 62 transcripts each) offline, on the identical frozen transcripts the operational judge scored. The crossed frontier judge (Claude Sonnet 4.5) reproduces the operational judge’s run-level composites within 0–4 points on every cell and

Participant / arm	Coop.	Res.	Hid. Rev.	Commit.	Mod. Qual.
Claude Haiku 4.5 (coop.) / No mediator (baseline)	70.3	48.6	13.5	24.9	0.0
Claude Haiku 4.5 (coop.) / Reflective Listening	72.1	39.1	12.4	28.8	95.5
Claude Haiku 4.5 (coop.) / Skilled@Qwen3-235B	79.1	45.8	14.0	32.8	99.0
Claude Haiku 4.5 (coop.) / Skilled@Kimi-K2.5	79.3	48.9	17.4	22.3	100.0
Stheno-v3.2 (L3-8B) (adv.) / No mediator (baseline)	19.9	22.6	8.3	38.1	0.0
Stheno-v3.2 (L3-8B) (adv.) / Reflective Listening	45.3	29.1	12.2	36.2	90.8
Stheno-v3.2 (L3-8B) (adv.) / Reflective Listening (replicate)	49.8	29.8	10.2	38.2	90.8
Stheno-v3.2 (L3-8B) (adv.) / Skilled@Qwen3-235B	64.2	41.5	10.3	34.5	94.5
Stheno-v3.2 (L3-8B) (adv.) / Skilled@Kimi-K2.5	55.5	34.8	13.9	27.6	96.3

Table 5: Per-component category scores (0–100, rounded to one decimal) for every gate row under the operational judge, from the frozen snapshot (emitted by `tools/compute_export_claims.py`). Columns: Cooperative Dimensions (Coop.), Resolution Depth (Res.), Hidden-State Revelations (Hid. Rev.), Intent Commitment Symmetry (Commit.), Moderator Quality (Mod. Qual.). Exploratory: components are shown unweighted; baselines score Moderator Quality 0 by construction (Section 2.3).

keeps the full mediator ordering, including Skilled@Qwen3-235B above Skilled@Kimi-K2.5 under the adversarial participant — the ordering that the prospectively documented shared-backbone concern (Appendix B) predicted a self-preferring judge would inflate. The open-weight judge (gpt-oss-120b) scores systematically lower. It keeps the control-versus-skilled separation but *not* the ordering within the skilled pair: under it the adversarial-column Skilled@Kimi-K2.5 delta edges above Skilled@Qwen3-235B (Table 6). The paired contrast puts numbers on this. On the adversarial column, Skilled@Qwen3-235B minus Skilled@Kimi-K2.5 is +2.3 [0.04, 8.95] under the operational judge (the lower bound grazes zero) and +7.0 [1.5, 10.0] under the crossed frontier judge, but −2.5 [−6.0, 4.5] under the open-weight judge — sign-reversed, with an interval spanning zero. On the cooperative column the pair is a plain tie (+2.1 [−2.5, 4.5]). Scenario-level correlation is scenario-level Pearson $r = 0.84$, $\rho = 0.80$, $n = 558$, with the crossed frontier judge and $r \approx 0.67$ with the open-weight panel judge; correlation does not establish agreement on absolute scores. Table 6 gives median deltas under all three judges. The headline contrast survives the crossed judge nearly unchanged (+35.5 [30.5, 38.0] vs. +36.9 [31.9, 41.4]). Of the 21 cell×judge estimates, the single interval touching zero is the weakest treatment under the harshest judge (the cooperative-participant Reflective Listening control under gpt-oss, +8.5 [0.0, 13.0]). That family of estimates grew after the freeze as panel judges were added, so this is a descriptive scan, not a frozen inferential claim.

Four-judge panel adjudication of the ordering (post-freeze). The panel leaves the ordering within the skilled pair judge-indeterminate. To adjudicate the sign reversal, we extended the crossed re-judge to a four-judge cross-family panel — Claude Sonnet 4.5, gpt-oss-120b, GLM-5.2, and Kimi-K2.5 — over the same nine frozen runs. We stated the decision rule after the first two crossed judges’ scores existed but before the two added judges produced any score: three or more of four finding the adversarial Skilled@Qwen3-235B–Skilled@Kimi-K2.5 contrast positive would report the reversal as one judge’s quirk; a 2–2 split would be reported as judge-indeterminate; no 1–3 branch was specified. The panel split 2–2: positive under Claude Sonnet 4.5 (+7.0 [1.5, 10.0]) and Kimi-K2.5 (+3.5 [−4.0, 7.0]); negative under gpt-oss-120b (−2.5 [−6.0, 4.5]) and GLM-5.2 (−5.0 [−17.0, 5.0]). Only the frontier judge’s interval excludes zero, and the cooperative column stays a four-judge tie (−1.5 [−4.0, 1.0], −19.0 [−30.0, 7.0]). Under the pre-stated rule the verdict is *judge-indeterminate*.

Two panel facts sharpen the reading. First, Kimi-K2.5 is the mirror image of the shared-

backbone concern: a judge scoring its own family’s mediated arm. It votes *against* itself, placing Skilled@Qwen3-235B ahead, so the panel shows no consistent same-family preference. Second, GLM-5.2 is a measured outlier. Its scenario-level Spearman rank correlation with every other judge is 0.18–0.28, against 0.65–0.80 among the operational, frontier, gpt-oss, and Kimi judges (all-judge joint set, $n=528$). Kimi-K2.5’s re-scores keep every skilled-arm delta and the adversarial skilled-over-control increments interval-backed (+11.0 [3.0, 13.5], +7.5 [1.5, 12.0]). GLM-5.2, by contrast, compresses score levels, resolves only some of the arm-versus-baseline and skilled-over-control differences, and reverses the sign of the ranked skilled-over-control quantity itself (−9.0 [−14.0, 4.0], interval spanning zero). That reversal comes from a low-correlating scorer with incomplete coverage: its missing evaluations concentrate in the longest, arm-correlated transcripts, and the panel rule counts unweighted signs. Disagreement among unvalidated LLM judges does not by itself establish which scorer is wrong. A reliability floor that excluded a judge correlating at 0.18–0.28 with every panel member — and at 0.17 with its own family’s GLM-4.7 (Appendix C) — would make the vote 3–1 and fire the quirk branch; but no floor was prespecified, so the registered verdict stands at 2–2. The panel thus refines the earlier summary: Skilled@Qwen3-235B versus Skilled@Kimi-K2.5 is a tie band whose within-band ordering is judge-indeterminate by the pre-stated rule; arm-versus-baseline and skilled-over-control contrasts remain positive under the Sonnet, gpt-oss, and Kimi scorers, while the low-correlation GLM-5.2 scorer compresses or reverses them.

A post-hoc fifth-judge probe of the GLM-5.2 anomaly (GLM-4.7 on a different serving stack) votes with the majority, but it cannot separate model-version sensitivity from serving-stack sensitivity; it is a sensitivity probe, not a re-vote (Appendix C). The 2–2 verdict was prospectively specified over the original panel and stands.

The judge-by-arm generosity contrast (Eq. (7)) shows no generosity by the operational judge toward its own family’s mediated arm on either column, at the export’s one-point integer resolution. Computed on the export’s run-level composite scores (means of per-scenario composites, stored as rounded integers) against the crossed frontier judge, the contrast is −1 on the adversarial column and +0 on the cooperative column. That integer resolution is coarse against a judge-pair run gap of up to 4 points, and this is a descriptive read on one judge pair (Section 2.10).

Participant	Arm	Δ_{med} [90% CI] by judge		
		Qwen3-235B-A22B-Instruct (op.)	Sonnet 4.5	gpt-oss-120b
Claude Haiku 4.5	Reflective Listening	+8.4 [6.9, 12.2]	+6.0 [4.0, 10.5]	+8.5 [0.0, 13.0]
	Skilled@Qwen3-235B	+14.9 [11.7, 17.0]	+10.0 [7.5, 14.5]	+10.0 [2.5, 15.0]
	Skilled@Kimi-K2.5	+12.8 [10.5, 16.7]	+10.0 [7.0, 14.0]	+9.0 [1.0, 13.5]
Stheno-v3.2 (L3-8B)	Reflective Listening	+29.1 [19.6, 33.6]	+24.0 [19.0, 27.5]	+17.0 [13.5, 21.0]
	Reflective Listening (repl.)	+31.2 [23.6, 35.3]	+25.5 [19.5, 29.0]	+17.0 [13.5, 21.5]
	Skilled@Qwen3-235B	+36.9 [31.9, 41.4]	+35.5 [30.5, 38.0]	+21.0 [18.0, 27.5]
	Skilled@Kimi-K2.5	+34.6 [27.3, 37.1]	+28.5 [24.0, 33.0]	+23.5 [19.0, 27.0]

Table 6: Median deltas (Δ_{med} , HAI Score composite points) with scenario-cluster bootstrap 90% intervals under the operational judge (Qwen3-235B-A22B-Instruct) and two crossed judges (judge Claude Sonnet 4.5; judge gpt-oss-120b) re-scoring identical transcripts ($n=62$ scenario pairs per cell). Deltas are computed against the same judge’s re-scored matched baseline. Granularity note: operational-judge scores are continuous composites; crossed-judge re-scores are stored as rounded integers, so crossed-judge endpoints lie on a half-integer grid (Section 2.7).

Evaluator-context sensitivity (registered). All 22 90% scenario-resampling intervals include zero, so this readout does not show a directionally resolved effect of brief access on run or policy-

contrast medians for this scorer. The separately frozen re-score under the requested Sonnet 4.5 alias asked whether an offline evaluator’s access to the parties’ private briefs changes its reading of otherwise identical transcripts. The context-withheld series completed on 2026-07-31 at exact scenario-set coverage for all 9 runs (558 matched run–scenario pairs and 1116 score rows across the visible and withheld conditions). The registered withheld-minus-visible run-median shifts range over -1.00 to $+3.00$ points; shifts in the seven mediated-minus-baseline contrasts range over $+0.50$ to $+3.00$, and shifts in the six within-column policy contrasts over -1.50 to $+0.50$. This is not evidence of equivalence: no equivalence margin was registered, the intervals permit nonzero shifts, and the visible and withheld scores were produced in different serving epochs whose provider response revision was not retained. Because the re-score is offline, it cannot alter trajectories or stop decisions made by the operational judge. The aggregate-only artifact is `analysis/sonnet_context_readout_2026-07-31.json` (SHA-256 prefix 000b450b). It discloses exact coverage, snapshot epochs, prompt hashes, and available runtime identities without publishing scenario records, score rows, prompts, private briefs, or transcripts.

Intra-judge test–retest. The operational judge’s test–retest coefficients are Krippendorff’s α of 0.48–0.91 by component: Moderator Quality $\alpha=0.91$, Resolution Depth 0.74, Intent Commitment Symmetry 0.67, Cooperative Dimensions 0.54, and Hidden-State Revelations 0.48. The judge re-scored 40 stratified transcripts (five per gate run) at $K=3$ passes each (calibration snapshot `e4445e6a`); the pinned rating grid is complete — 40×3 passes per component. The two components that carry arm discrimination sit at 0.54 (Cooperative Dimensions, degraded and below the tentative-reliability threshold) and 0.74 (Resolution Depth, tentative but not reliable); under the export schema’s three-state confidence label, Resolution Depth, Intent Commitment Symmetry, and Moderator Quality qualify as full confidence while Cooperative Dimensions and Hidden-State Revelations are flagged degraded. The Hidden-State Revelations value needs its distribution to be read: 111 of 120 of its rating passes sit in the *top* bucket, so the judge’s raw revelation ratings are high even though the answer-key F_1 cap pulls the published component score to the floor (Section 3.3). Expected disagreement is therefore tiny, and α has almost no variance to work with: percent agreement is high, and the low α reflects restricted range in this sample without excluding instability specific to the component. That fragility is directly observable. The sub-study was executed twice against this frozen sample (rating passes re-drawn) and returned Hidden-State Revelations α of -0.03 and 0.48 across the two executions, with smaller movements on the other components. The export identifies one execution as the calibration snapshot of record, but no outcome-independent rule for retaining that execution was prospectively specified, so its coefficients must be read with the disclosed two-execution Hidden-State Revelations range, not as a selected precision gain; the prospectively specified reliability extensions (Appendix B) raise n . The tie-band rule (Section 2.9) does not absorb this judge noise: its intervals reweight scenarios and contain no judge-rescoring variance. The paper’s own data show the distinction — the within-skilled-pair contrast excludes zero under the rule’s intervals, is judge-indeterminate under the panel, and flips under re-generation (Section 3.5).

Judge-swap generation rows — separate conversations under a different operational judge, where judge influence is confounded with generation stochasticity — are superseded by the crossed re-judge and reported in Appendix C.

3.5 Extended results: additional mediator rows (non-gate)

Published extension rows place additional mediators on the same two participant columns (same judge, suite, and protocol). Under the cooperative participant, five additional skilled mediators

from four vendors all land in 54–55 (selected maximum: 55); under the adversarial participant the same class of mediators spreads across 44–52. These rows are leaderboard extensions, not frozen gate cells, and are reported descriptively. Their median deltas carry the same scenario-cluster intervals as the gate cells (cooperative column: +12.9 to +15.0 across all five extension mediators; adversarial column: +38.6 [33.7, 42.1] GLM-5.2, +40.1 [34.6, 42.7] Sonnet 5, +38.0 [32.8, 41.7] mediator gpt-oss-120b). Every interval excludes zero and overlaps the others in its column and the strongest gate cell, so under the tie-band rule of Section 2.9 each column’s skilled arms form a single band (a conservative marginal-interval read; the paired contrasts are the sharper instrument). They remain single-replicate points; any “top score” among them is a selected maximum, subject to winner’s-curse inflation and expected to regress on replication.

The separate GPT-5.5 cell is a registered post-freeze extension, designated for outcome-independent reporting, with one Skilled/full run and no matched Reflective-Listening control; it therefore identifies no fixed-backbone policy contrast. On fixed transcripts, its mediated-minus-Gemma-baseline delta is the largest on the Gemma extension column under both registered offline scorers: +33.5 [+28.0, +38.0] under Gemma-4-31B and +21.5 [+18.5, +25.0] under gpt-oss-120b. The winner’s-curse caveat applies when comparing the largest among selected extension rows; it is not a reason to suppress this registered row.

Read strictly as point values, the adversarial-column ordering is *not* what a capability-tracking account would predict. An open-weight 120B mediator (49) and a smaller open-weight model (52) score at or above a 235B open-weight model (48), and Skilled@Kimi-K2.5 sits at the bottom of the skilled arms (44). We draw no inferential claim from these single-replicate rows: only a few points separate them, and their intervals are conditional on the one observed generation. The observation is hypothesis-generating; the prospectively specified replicated design (Appendix B) can test whether mediation-policy performance within this harness is distinct from general model strength.

End-to-end replicates (post-freeze). Four cells re-run end-to-end after the freeze moved their run-level composites by at most 3.7 points. These were new conversations under the identical frozen configuration (same scenarios, models, temperatures, and operational judge), not re-scores: both baselines and both adversarial-column skilled cells. The movements were cooperative baseline 36 → 37.9, adversarial baseline 18 → 15.5, Skilled@Kimi-K2.5 44 → 46.2, and Skilled@Qwen3-235B 48 → 44.3. At scenario level the repeated mediation deltas differ by at most 1.2 points from the frozen values (+35.8 [30.2, 37.6] vs. +36.9 [31.9, 41.4] for Skilled@Qwen3-235B; +35.1 [29.4, 38.7] for Skilled@Kimi-K2.5). The frozen-minus-repeat baseline intervals include zero (+0.3 [−1.0, 3.5] cooperative, −0.4 [−5.5, 2.2] adversarial); this is not an equivalence test. The negative-delta counts are similar (2 and 1 of 62 scenarios below baseline, vs. 0 and 3 in the frozen runs). At scenario grain the selected baseline repeats also show generation stochasticity: the identical no-mediator configuration moved 57 (cooperative) and 56 (adversarial) of 62 scenario scores in one direction or the other. As only two selected repeats, they do not estimate or bound the stochastic contribution to the per-cell negative-delta counts of Section 3.1. One observation bears directly on the panel verdict of Section 3.4: the point ordering within the skilled pair *flips* across replicates (frozen: Skilled@Qwen3-235B 48 > Skilled@Kimi-K2.5 44; repeat: Skilled@Kimi-K2.5 46.2 > Skilled@Qwen3-235B 44.3), while both repeated mediated-minus-baseline point estimates remain positive. Frozen gate aggregates analyze repeat runs separately, never pooled (Section 2.7); the one disclosed exception is the information study’s Skilled/full-brief arm, which pools the frozen run with this repeat (Section 3.6). The repeats are not in the public export and alter no frozen gate value. Together with the four-judge panel, they show that the fine ordering within the skilled tie band changes under both re-generation and re-judging. They provide repeat observations for the tested

cells, not a bound on run-to-run stochasticity.

3.6 Information study: brief versus public-context access (post-freeze, mixed-suite factorial)

The prospectively specified information-provisioning manipulation (Appendix B) ran as a model-matched 2×2 on the adversarial column. The Reflective Listening and Skilled policies, both bound to the Qwen backbone of Skilled@Qwen3-235B, were crossed with the full-bilateral-brief and public-context-only conditions of Section 2.4; the estimands E_1 – E_4 and both frozen decision rules are defined in Section 2.8. The Qwen implementation combines three newly generated suite-2.5.0 cells with a Skilled/full-brief cell that *pools* the frozen suite-2.0.0 run with its post-freeze end-to-end repeat (protocol-matched; the disclosed exception to the never-pool rule that governs frozen gate aggregates). The pooled repeat scores below the frozen run (44.3 vs. 48), so pooling lowers the full-brief arm and moves the E_3 verdict in the favourable direction; a reader can check this directional sensitivity from Table 7. The Qwen matrix has two runs per cell, every run 62/62, and its estimands, analysis code, and -4.0 -point non-inferiority margin were frozen before any public-context-only run existed. Model matching was required because the frozen control’s backbone was unpinned (Section 2.4); the full-brief Qwen-bound control scores near the original control (39/41 vs. 40), which supports comparability on Qwen, not backbone invariance.

Policy beats minimal reflection on a fixed backbone. Under Qwen with public context only, the Skilled policy exceeds the model-matched Reflective Listening active control by $E_1 = +13.17$ [+9.67, +17.55] composite points — the paper’s strongest result, prospectively specified and contemporaneous. Under full bilateral briefing, the Skilled-policy increment is $E_2 = +7.27$ [+1.31, +10.52]. The interaction is $E_3 = +5.90$ [+1.34, +13.81], read against the frozen -4.0 -point non-inferiority margin. The registered Qwen verdict is affirmative on both frozen rules: E_1 ’s lower bound exceeds zero, and E_3 ’s lower bound clears -4.0 (recorded as `E1_positive` and `E3_noninferior` in the archived analysis). For the two information contrasts (public-context-only minus full bilateral brief), the Reflective Listening policy is -5.89 [−12.89, −2.52] and the Skilled policy is $+0.01$ [−4.15, +4.28]; the Skilled interval includes zero.

One descriptive pattern recurs across the three tested backbones. The Skilled-policy information contrast $E_{4,\text{skilled}}$ (public-context-only minus full bilateral brief) clusters near zero: $+0.01$ [−4.15, +4.28] on Qwen, $+1.53$ [−1.64, +5.75] on Kimi, and -1.64 [−4.59, +3.38] on gpt-oss, and all three intervals include zero. This does not establish equivalence. No $E_{4,\text{skilled}}$ equivalence margin was specified, the intervals remain compatible with nontrivial shifts, the full-brief Skilled-policy arms are historical, and the Qwen value depends on the suite (the suite-matched sensitivity below moves it from $+0.01$ [−4.15, +4.28] to -2.66 [−6.40, +2.04]). As a labelled post-hoc diagnostic, the intervals can be read against the already-frozen ± 4.0 -point band (the constant was signed 2026-07-15, before any public-context-only run existed): no backbone’s interval sits inside the band. Kimi’s ($+1.53$ [−1.64, +5.75]) leaves it only at the upper end, while Qwen’s ($+0.01$ [−4.15, +4.28]) and gpt-oss’s (-1.64 [−4.59, +3.38]) extend past -4.0 — and it is the replicated core, Qwen, whose interval is widest. The companion reflective-policy contrast also gives the paper’s only same-estimand pair with both intervals excluding zero in opposite directions: $E_{4,\text{reflection}}$ is -5.89 [−12.89, −2.52] on Qwen but $+4.12$ [+0.23, +15.22] on gpt-oss. That bounds how far the “clusters near zero” reading extends beyond the skilled policy. We therefore treat the near-zero clustering as a target for a prospectively specified, contemporaneous equivalence study, not as evidence that the full bilateral brief is unnecessary or that skill substitutes for intake.

The factorial is neither wholly contemporaneous nor single-suite. For each backbone, E_1 and the reflective-policy information contrast use contemporaneously generated suite-2.5.0 cells, while E_2 ,

Policy contrasts

Mediator backbone	Replicates	E_1 : public skill	E_2 : full-brief skill	E_3 : interaction
Qwen	2/2/2/2	+13.17 [+9.67, +17.55]	+7.27 [+1.31, +10.52]	+5.90 [+1.34, +13.81]
Kimi	2/1/1/1	+9.66 [+2.32, +16.21]	+1.68 [−2.06, +5.85]	+7.98 [−0.32, +15.47]
gpt-oss	1/1/1/1	−1.02 [−4.51, +5.01]	+4.74 [+2.19, +15.55]	−5.75 [−16.75, +0.38]

Information contrasts

Mediator backbone	Replicates	$E_{4,\text{reflection}}$	$E_{4,\text{skilled}}$
Qwen	2/2/2/2	−5.89 [−12.89, −2.52]	+0.01 [−4.15, +4.28]
Kimi	2/1/1/1	−6.45 [−11.99, +0.60]	+1.53 [−1.64, +5.75]
gpt-oss	1/1/1/1	+4.12 [+0.23, +15.22]	−1.64 [−4.59, +3.38]

Table 7: Model-matched information-study estimands (Eq. (3)): point estimate and 90% scenario-reweighting interval, in HAI Score composite points, $n=62$ jointly completed scenarios per backbone, conditional on the observed generation runs and operational judge. All increments compare the Skilled policy with the model-matched Reflective Listening control, not the no-mediator baseline; the E_3 reading rule uses the frozen -4.0 -point margin. Replicate order is Skilled/full-brief, Skilled/public-context, reflection/full-brief, reflection/public-context. The two E_4 estimands are public-context-only minus full bilateral brief within policy. Qwen is the replicated core; Kimi and gpt-oss are conditional breadth checks. The upper panel reports policy contrasts; the lower panel reports information contrasts within policy.

E_3 , and the Skilled-policy information contrast combine those cells with protocol-matched frozen suite-2.0.0 full-brief Skilled-policy arms. Run date and suite stamp therefore differ within those three estimates even though the arm protocol was held fixed, so historical drift or other version-linked differences can affect them. The design gives its cleanest reading to the public-context-only Skilled-versus-reflective contrast, not to a universal claim that skill replaces information.

A suite-matched sensitivity keeps the classifications but widens the interaction to include zero. In it we replaced the mixed-suite Qwen calculation’s Skilled/full-brief arm with two newly generated suite-2.5.0 runs, leaving the other three suite-2.5.0 arms unchanged. The extension and estimator were registered in version control on 2026-07-30, before either new outcome (cell registry commit [5f44fb877](#); fail-closed readout implementation commit [3ca71d46a](#)); “post-outcome” means the sensitivity was chosen after we had observed the original mixed-suite result. The new run means were 44.41 and 47.79; all 8 runs covered the identical 62-scenario set with zero final failed scenario rows, each new run carrying one completion repair. The two new runs bind the v2 input contract and record a serving provider, while the six retained comparator runs predate both fields, so the input contracts are not fully matched even though the suite version is (Appendix D). Because only the Skilled/full-brief arm changed, with 2/2/2/2 runs across the four cells, E_1 and $E_{4,\text{reflection}}$ contain no new data and equal the primary values by construction ($E_1 + 13.17 [+9.67, +17.55]$; $E_{4,\text{reflection}} - 5.89 [-12.89, -2.52]$); the sensitivity is informative only about E_2 , E_3 , and $E_{4,\text{skilled}}$. There, E_2 moved from $+7.27 [+1.31, +10.52]$ to $+9.93 [+2.65, +13.41]$; E_3 moved from $+5.90 [+1.34, +13.81]$ to $+3.23 [-1.13, +11.93]$, an interval that now includes zero, so under the suite match the data do not resolve whether the interaction is positive; and $E_{4,\text{skilled}}$ moved from $+0.01 [-4.15, +4.28]$ to $-2.66 [-6.40, +2.04]$. Applying the frozen thresholds as diagnostics gives the same pass classifications: E_3 ’s lower bound exceeds the -4.0 margin, and E_1 ’s positivity carries over trivially because E_1 is unchanged. The E_2 interval also excludes zero, so the post-hoc assay-sensitivity precondition of Section 2.8 holds; the Skilled information contrast still includes zero. This sensitivity does not replace the original mixed-suite primary analysis and is

neither a replication nor a wholly contemporaneous factorial. The aggregate-only artifact is pinned by short hash 742139a2.

The result does not generalize across backbones. The same frozen estimands were then applied conditionally to breadth runs on Kimi and gpt-oss (backbone breadth added 2026-07-17, two days after the estimand freeze, as an outcome-independent amendment). Kimi crosses both frozen numerical thresholds; gpt-oss crosses neither. The mostly single-run Kimi matrix has an E_2 interval (+1.68 [−2.06, +5.85]) that includes zero, so the post-hoc assay-sensitivity safeguard of Section 2.8 withholds a substantive E_3 non-inferiority interpretation; this does not rewrite the frozen threshold result. For gpt-oss, the Skilled policy does not beat its public-context-only Reflective Listening control and the interaction does not clear the margin, while that control scores higher without the private-facts packet. Only the Qwen matrix has two runs per cell. Kimi is therefore a threshold-crossing but interpretation-limited breadth check, gpt-oss is a failed generalization, and there is no backbone-universal positive skill increment or interaction result. The single-run failure identifies a policy–backbone configuration that did not generalize; it does not establish that gpt-oss is generally weak, because policy compatibility, serving implementation, and generation variance are not separated. The labels above describe information delivered by the harness. Their analogy to case preparation, intake, or caucusing is conceptual only: no interview was observed, and these cells provide no evidence of human mediation efficacy. These post-freeze results are pinned in a dated, in-repository analysis artifact; provenance and recomputability are scoped in the Reproducibility statement.

Participant transport (exploratory). **The public-context policy contrast did not transport to a second participant configuration.** $E_1 = +5.70$ [−3.79, +12.90] fails the frozen positivity rule, and the interaction $E_3 = -5.68$ [−15.95, +2.30] fails the frozen −4.0-point non-inferiority margin. We replayed the frozen 2×2 with Gemma-4-31B participants (Appendix B, item 10), holding the Qwen mediator backbone, both policy prompts, both information conditions, the operational judge, and the scenario partition fixed. Every interval here is exploratory: the complete design is a post-launch, pre-outcome amendment (frozen 2026-07-25, before any arm had an evaluated scenario), with one run per arm (1/1/1/1). Baseline clearance did transport: every mediated arm cleared the scenario-matched no-mediator baseline (Table 8), and the full-brief policy contrast was clearly positive ($E_2 = +11.39$ [+5.77, +16.92]). Because this column’s E_2 interval excludes zero, the assay-sensitivity precondition holds and the E_3 verdict is reportable — and it is a fail, the first failed frozen-rule outcome on the participant axis (gpt-oss failed both rules on the mediator-backbone axis). The reflective policy’s information contrast has a near-zero point estimate with an interval spanning nontrivial shifts ($E_{4,\text{reflection}} = +0.73$ [−3.66, +6.34]); the skilled-policy interval also includes zero ($E_{4,\text{skilled}} = -4.95$ [−13.69, +1.94]). With one run per arm and wide intervals, this does not establish that the private brief is necessary; it establishes that the Stheno public-context finding did not transport to this participant configuration under the frozen rules. These cells read beside the participant-validity analysis of Section 2.5; they are a participant column, never a fourth row of Table 7.

Cross-family re-score of the suite-2.5.0 rows (registered). Two offline judges from other model families agree with the operational instrument on the sign of all three transport-column contrasts, and both keep the negative interaction direction. Because both failed frozen reads above rested on a single judge family, a sighted cross-family re-score of all seventeen suite-2.5.0 generation runs — the information-study cells and the five transport runs, on identical frozen transcripts — was internally registered in the version-controlled runbook on 2026-07-26, before any crossed score

Participant pair	Arm	HAI Score	Δ_{med} [90% CI]
Gemma-4-31B (transport)	No mediator (baseline)	15	—
	Reflective Listening (full brief)	34	+17.90 [+14.35, +22.22]
	Reflective Listening (public context)	35	+18.63 [+14.15, +25.06]
	Skilled@Qwen3-235B (full brief)	41	+29.29 [+23.88, +34.72]
	Skilled@Qwen3-235B (public context)	38	+24.33 [+16.99, +30.71]

Table 8: Participant-transport extension (internally registered): Gemma-4-31B participant pair under the Qwen mediator backbone, one run per arm. Δ_{med} is the difference of scenario medians — arm median minus the scenario-matched no-mediator baseline median over jointly completed scenarios (the registered estimand, Eq. 2) — with 90% scenario-resampling intervals. Every mediated arm clears the baseline; the frozen E_1 and E_3 rules nevertheless fail on this column (see text).

existed, with two judges from other model families (Gemma-4-31B and gpt-oss-120b, both serverless open-weights deployments). Coverage completed 2026-07-28 at 62/62 for all 34 sweep run $\times$ judge pairs; the aggregate-only readout also stores two scoped exploratory run $\times$ judge pairs not analyzed here, for 36/36 total. It is pinned in `analysis/crossed_readout_2026-07-28.json` (SHA-256 prefix `5c5800fb`), contains coverage and aggregate estimates rather than scenario rows, and passed a separate internal automated adversarial artifact/code review for release. Historical prompt hashes, provider model revisions, and rejudge job/snapshot identifiers remain unavailable and are reported as provenance limitations. All numbers below use the registered estimand — differences of arm medians over jointly completed scenarios (Eqs. 2 and 3, replicate-mean per-scenario scores where an arm has two replicates) — with 90% scenario-resampling percentile intervals ($B = 10,000$, seed 20260728). Three registered reads follow. First, baseline clearance is preserved: every mediated arm on the transport column clears the scenario-matched baseline under both crossed judges (smallest arm delta +13.00 [+10.00, +17.00] under Gemma-4-31B, +12.00 [+9.00, +12.50] under gpt-oss-120b). Second, the estimand signs agree with the operational instrument on all three contrasts. E_1 is small and positive with intervals spanning zero under both crossed judges (+2.50 [−0.50, +9.50], +2.50 [−1.00, +6.50]) — the same uncertain read as the operational $E_1 = +5.70$ [−3.79, +12.90]. E_2 is clearly positive under both (+11.00 [+6.00, +14.00], +7.00 [+5.50, +11.00]), which keeps the positive full-brief reference contrast under the later assay-sensitivity safeguard. The negative interaction keeps its direction: E_3 is negative under both crossed judges (−8.50 [−12.50, +0.50], −4.50 [−10.00, −0.50]; under gpt-oss-120b the interval excludes zero), consistent with the operational $E_3 = −5.68$ [−15.95, +2.30]. Third, the Stheno-column public-context skill increment remains clearly positive under both (+18.75 [+11.00, +21.75], +8.25 [+4.50, +11.50]). Crossed re-scores carry no frozen rule and change no frozen verdict. They give narrow corroboration from two offline scorers on fixed transcripts: both E_3 point estimates are negative, and both interval lower endpoints fall below the frozen −4.0 margin, so applying that margin descriptively would not establish non-inferiority under either scorer; the gpt-oss-120b interval additionally excludes zero. No crossed-scorer verdict was frozen. Magnitude and interval width remain evaluator-dependent, and because neither offline scorer steered the conversation or made its stop decisions, this check cannot repair or independently validate trajectories already produced under the operational judge. A Gemma-4-31B-family judge scoring Gemma-participant transcripts also gives an unregistered descriptive probe of judge-participant family alignment: the observed ranking did not show the hypothesized same-family advantage, but no formal alignment estimand was registered. (An earlier same-day emission used the median of within-scenario paired differences rather than the registered difference of arm medians, and computed $E_2 - E_1$ rather than the registered $E_1 - E_2$. Adversarial re-verification caught both implementation errors before any snapshot pinned them; only the

implementation and its emitted values were corrected — the estimand, margin, and decision rules did not change. The registered-estimand computation above is the release-ready aggregate readout.)

4 Limitations and Threats to Validity

Some of these limitations are design choices whose costs we accept knowingly. Some have a prospectively specified remedy in Appendix B. Each item below states what the results cannot establish and why.

1. **Replication is uneven across claims.** The reported intervals measure how results change when scenarios are reweighted for one fixed configuration. They do not measure how results change from one run to the next. After the freeze, we re-ran four cells end-to-end (both baselines and both adversarial-column skilled cells). Run-level composites moved at most 3.7 points, the frozen-minus-repeat baseline intervals include zero, and the point ordering within the skilled pair flipped (Section 3.5). These figures are the largest movements seen in four chosen repeat cells; they are not bounds on run-to-run variation. An interval that includes zero is not evidence of equivalence. This falls far short of the prospectively specified paired-by-seed design. The Qwen information-context matrix has two replicates per cell. The Kimi and gpt-oss breadth matrices are mostly single-run conditional probes, and gpt-oss fails the prospectively specified generalization reading.
2. **The scenario suite is fixed and authored.** The bootstrap reweights these authored scenarios. It does not sample from any defined population of human disputes, so its intervals cannot establish how common or how large an effect is in real conflicts. The 62-scenario evaluated development partition and the scenario-level records are not public. The aggregate public export alone cannot reproduce the reported scenario-reweighting intervals.
3. **Participant configurations were chosen by hand; they are not a participant population.** We tested two gate configurations and one registered exploratory extension: the Gemma participant-configuration extension of the information factorial (Appendix B, item 10; results in Section 3.6). Together they show that the participant configuration changes what this harness measures. They support no pooled claim about participants in general. The extension is single-run-per-arm conditional breadth. It holds the mediator and evaluator apparatus fixed within its cells, which removes some package confounding, but it does not isolate any one participant property. Both frozen rules failed on that column.
4. **Capability, disposition, and alignment tuning are confounded** between participant configurations. Claude Haiku 4.5 and Stheno-v3.2 (L3-8B) differ at the same time in backbone, capability, training and alignment, disposition, serving provider, and stochastic generation. The results therefore cannot attribute the observed difference in point-score range to any one participant property. The same-model persona-swap deconfound remains prospective. The safety-ablation deconfound of Section 2.5 is a candidate design, not a registered item.
5. **The active control’s backbone was unpinned in the frozen matrix.** The Reflective Listening moderator has no model binding. Its calls therefore went to the harness default (`claude-sonnet-4-5`; disclosed in the export’s `mediator_model` field), while the skilled arms ran their named backbones. The frozen skilled-over-control contrast thus compares whole mediator *packages*; it does not compare policy at a fixed backbone (Section 2.4). The prospectively specified model-matched 2×2 removes this confound for Qwen only, not for all

backbones. A further asymmetry: mediator prompts keep only the last 20 messages, while the system prompt is re-sent on every call, and in full-brief arms that system prompt carries the full bilateral brief. This asymmetry is disclosed as part of the prospectively specified public-context-only estimand.

6. **The primary information factorial is not wholly contemporaneous.** E_2 , E_3 , and the Skilled-policy information contrast combine contemporaneous suite-2.5.0 cells with protocol-matched frozen suite-2.0.0 full-brief Skilled-policy arms, so run date and suite stamp differ for those three estimates; the suite-matched sensitivity removes the mismatch but is post-outcome and remains noncontemporaneous. Section 3.6 gives the full reasoning.
7. **The mediated treatment is a bundle.** Mediated arms may stop on a model-reported agreement event or judge-directed stop; baseline conversations cannot stop; and mediator turns do not use up the participants' turn budget. Each delta therefore combines three things: what the mediator said, the stopping policy, and the extra conversational bandwidth. Separating them would need a shared global budget and a symmetric stop rule, as ProMediate uses to control interaction resources [17].
8. **Single operational judge; LLM-as-judge noise.** One judge model steers and scores every conversation. The test-retest and crossed-judge sub-studies measure instrument noise but do not remove it. Judge self-preference is probed descriptively only. The registered cross-family re-score of the suite-2.5.0 rows (Section 3.6) is narrower evidence than an end-to-end judge swap. Two offline scorers on fixed transcripts agree with the direction of the negative transport interaction, but its magnitude and uncertainty still depend on the evaluator. Neither scorer can correct trajectories or stop decisions that the operational judge already influenced.
9. **Judge-swap baselines are separate conversations.** Baseline runs generated under different operational judges differ in both the scoring instrument and the conversation trajectory. They are not re-scores, so they cannot isolate the judge's influence.
10. **Component weights are maintainer-asserted**, not derived. Moderator Quality is structurally zero for no-mediator runs but saturates for mediated arms (including the Reflective Listening control). The headline delta therefore mixes a four-component outcome shared by all arms with one component that only mediated arms can score on. The completed Moderator Quality-excluded, renormalized four-component analysis is post hoc and secondary (Table 4). It does not replace the frozen difference-of-medians estimand and does not algebraically decompose it.
11. **Arm non-blindability.** The judge sees the mediator's turns in mediated arms, so scoring cannot be blinded to arm.
12. **Survivorship conditioning in per-turn curves.** Mediated conversations stop early and the baseline has no stop event. Per-turn values are therefore conditional on conversations that are still active. Every curve point carries its n .
13. **Token sampling is unseeded.** Conversations are stochastic. The export states `sampling: unseeded` rather than implying determinism. Paired-by-seed replication is prospectively specified.
14. **Bootstrap endpoint discreteness.** With ~ 62 paired deltas, percentile endpoints land on observed order statistics. On discrete scores, the percentile bootstrap of a difference of

medians can therefore under- or over-cover in finite samples. Smoothed-bootstrap, BCa, and Hodges–Lehmann alternatives are prospectively specified.

15. **Delta-level attrition.** Δ_c is estimated only over scenarios completed on both arms. If scenarios fail at different rates across arms, that selects which pairs enter the estimate. The matched-baseline rule also rejects an entire baseline run on any single failed scenario. (All gate cells published to date are 62/62 complete, so this is currently latent.)
16. **Reliability scope.** Test–retest and crossed-judge coefficients are computed on per-component 1–5 buckets. None is computed for the continuous composite or for Δ_c itself. Unlike ProMediate’s human metric validation and SoCRATES’ expert-alignment study [17, 18], this version has no human calibration of scenario realism, mediator quality, judge scores, or downstream human outcomes. Human criterion validity for mediation quality is therefore unestablished.
17. **Preparation is an analogy, not an observed practice.** Neither information condition (Section 2.4) includes an interview, an intake, or a confidential human caucus. The 2×2 tests what changes when the mediator’s information is varied inside this harness. It does not test how mediator preparation, interviewing, or mediation perform in the field. Full briefs supply both parties’ private facts without modeling elicitation, confidentiality, consent, or permission to disclose. One measured detail qualifies the public-context-only label. Scenario **focus_areas** reach the mediator in every arm, including public-context-only. In 3 of the 100 authored fixtures, a focus area shares enough content words with an expected revelation to act as a directional hint — not a verbatim answer key. The public-context condition is therefore not perfectly information-free.
18. **Simulated, narrow disputes.** Participants are language models with configured hidden states, not human subjects. The suite is dyadic, English-language, and text-only. It excludes nonverbal behavior, institutional power, cultural interpretation, and multi-party coalition dynamics. External validity to human conflict is unestablished.

5 Conclusion

MediationBench scores synthetic two-party conversations produced under an AI mediation policy, pairing mediated runs with no-mediator baselines and with an active Reflective Listening control, under a judge-reliability protocol and rationale-level anti-gaming disclosure. Where existing testbeds score mediator output, multi-party negotiation, or cross-domain adaptation, this harness asks how the choice of participants, of mediator information, and of judges changes what is measured.

Four findings follow. *Participants set the range.* Against the assistant-aligned configuration, skilled-mediator scores fall in a narrow band; against the higher-resistance configuration they spread. The two configurations differ at once in capability, disposition, and alignment (Section 4), so the data do not identify which property causes the difference. *Mediated arms clear their matched score baselines.* Every published mediated gate cell, including the minimal Reflective Listening control, scores above its matched no-mediator baseline; the Skilled policy’s gain must therefore be estimated against the active control, and part of the control’s apparent gain is assigned mechanically by the mediator-only rubric component. *Policy beats minimal reflection on a fixed backbone.* The contemporaneous Qwen public-context-only contrast puts the Skilled policy $E_1 = +13.17$ [+9.67, +17.55] composite points above the model-matched control (90% scenario-reweighting interval). Kimi, mostly single-run, crosses both frozen thresholds but a post-hoc assay-sensitivity safeguard withholds its E_3 interpretation; gpt-oss crosses neither; the single-run

Gemma participant-transport extension fails the frozen transport criteria, and two offline scorers preserve its negative E_3 direction. The decision-rule outcome therefore does not generalize across backbones or participants. A post-outcome, suite-matched Qwen sensitivity keeps the classifications while widening the interaction interval to include zero. *Skilled mediators do not separate from one another*. A four-judge cross-family panel records the fine ordering inside the skilled tie band as judge-indeterminate; four end-to-end repeats move run-level composites by at most 3.7 points and flip that ordering. The three judges that correlate with each other keep every arm-versus-control contrast positive; the low-correlation GLM-5.2 scorer compresses or reverses them. A preplanned Sonnet re-score with evaluator context withheld left all 22 matched intervals including zero.

Two descriptive patterns are hypothesis-generating only. Under higher-resistance participants, the ordering of mediators does not obviously follow general model strength, on single-replicate data. And across all three tested backbones the Skilled-policy information contrast sits near zero; this is not evidence of equivalence — for the replicated backbone it moves to -2.66 $[-6.40, +2.04]$ under the suite-matched sensitivity (Section 3.6) — and it does not imply that the full bilateral brief can be removed. The full-brief contrast is a controlled privileged-information comparison and only an imperfect analogue of preparation or intake; it is not an observed interview or evidence about how well human mediation works.

For evaluation practice, the findings support four rules: score synthetic mediation against both a matched no-mediator baseline and an active minimal-intervention control; treat the mediator’s information access as an explicit experimental factor; stress the participant configurations; and cross-score frozen transcripts with judges from different model families. The claims are bounded by uneven replication, by two confounded gate participant configurations plus one exploratory transport extension, by an uncalibrated simulated population, by historical-arm reuse in parts of the factorial, and by a judge measured to be noisy; the prospectively specified plan states what evidence would license stronger claims. Synthetic conversations and synthetic mediation are experimental instruments and product test harnesses: they enable controlled comparisons and repeatable failure analysis, but do not substitute for validation with people in real conflicts. Deterministic data-pipeline invariants can serve as hard product regression gates; fixed-transcript LLM scoring is a controlled statistical regression test that needs frozen tolerances and calibrated false-alarm rates, and unseeded end-to-end generated scores remain advisory under the same constraint.

The priority next steps are human expert review of transcripts and judgments, to calibrate the criterion; paired-seed, contemporaneous, single-suite replication; a same-model persona swap, to separate disposition from backbone and alignment; and broader coverage of judge and participant families.

Ethics Statement

This study involves no human subjects: all disputants are simulated by language models with configured scenario roles and hidden states, and all transcripts are machine-generated. The author did not seek institutional review because the study uses no human subjects, human-derived transcripts, or personal data. Scenario fixtures were authored for the benchmark and depict interpersonal and organizational conflict (workplace, family, commercial, community) at non-graphic intensity. All 100 scenario fixtures (62 evaluated/development fixtures and 38 sequestered, unrun reserve fixtures) were authored for the benchmark by the maintainers, with LLM writing assistance under a written authoring guide and human curation before inclusion; none is derived from a real dispute, a real transcript, or personal data. Because the evaluated fixtures also informed development and the reserve has not yet been run, this paper does not claim contamination-free evaluation or reserve-set

generalization. The scenario fixtures themselves are not released as a corpus in this version — the published artifacts are the aggregate export and this paper; a corpus license will be declared if and when fixtures are published.

We use minimally aligned open-weight models to simulate adversarial disputants; this is a measurement-validity choice (Section 2.5), the models are used only as scenario participants inside the harness. The protocol does not intentionally solicit content beyond non-graphic, scenario-bounded dispute behavior, and private transcripts are not released; the unperformed inspection below prevents a stronger claim about what every generation contains. We note the dual use plainly: measurement machinery that assigns higher within-harness scores to some conversational patterns can be inverted to select lower-scoring patterns, and this inversion is not hypothetical — a major platform experimentally manipulated users’ emotional states through feed curation [26], and engagement-optimized ranking systematically amplifies out-group animosity [27]. Participants capable of sustained hostility remain necessary regardless: a benchmark cannot stress-test mediator behavior under hostile or toxic participant behavior with participants unable to produce it.

AI mediation systems deployed to real conflicts carry risks — bias between parties, coerced agreement, over-intervention — and this benchmark flags rather than obscures one score-level warning signal: below-baseline deltas are reported, not suppressed. The registered transcript-level inspection of flagged scenarios (Section 2.3) has not been performed in this version; the instrument records a negative-score flag, not harm itself.

This release is neither a safety certification nor evidence of deployment readiness for family, employment, immigration, clinical, or legal disputes.

Reproducibility and Data Availability

Published results are aggregate and versioned. The public export (<https://mediationbench.com/research/#measurements>) carries, for every run: the recorded participant, mediator, and judge model-identifier strings (including aliases where immutable provider revisions were unavailable), suite and methodology version stamps, run-level scores and deltas, and per-component category scores; per-role sampling temperatures are recorded once at export level (shared across runs), and per-turn cooperation-curve aggregates with per-point n are present for most but not all archived runs. Sampling is unseeded and disclosed as such.

The export supports inspection of its published run-level aggregates and recomputation of arithmetic derived solely from them, and nothing more: most of this paper’s numeric claims — every scenario-paired estimand, scenario-cluster bootstrap interval, and E_1 – E_4 value — derive from non-public per-scenario intermediates, because per-scenario scores and deltas are not published under the benchmark’s aggregate-only anti-gaming policy. Those quantities are computed by analysis code included alongside the paper (`tools/compute_claims.py` and `tools/crossed_readout.py`; the suite-matched sensitivity additionally uses `tools/qwen_suite25_sensitivity.py`; fixed-seed bootstrap, Section 2.9), with the resulting aggregate outputs archived in-repository as dated JSON artifacts (the post-freeze information-study and participant-transport results are pinned in `analysis/compute_claims_2026-07-28.json`, whose full byte hash is recorded in the generated macro source; the crossed-scorer fixed-transcript aggregates, including the registered post-freeze GPT-5.5 extension reported descriptively in Results, are pinned in `analysis/crossed_readout_2026-07-28.json`, SHA-256 prefix `5c5800fb`; the post-outcome suite-matched Qwen sensitivity is pinned in `analysis/qwen_suite25_sensitivity_2026-08-03.json`, SHA-256 prefix `742139a2`). The aggregate-only Qwen artifact contains eight run means, five estimand values (E_1 – E_3 and two E_4 contrasts), and provenance digests, but no scenario rows. The crossed artifact records exact

Reproducibility class	Claims in this version
Independently recomputable from public inputs	Checksum verification of released artifacts, arithmetic derived solely from the published aggregates, and the Moderator-Quality-excluded sensitivity values derived from published component means.
Inspectible only at aggregate level	Published run-level means, category aggregates, displayed counts, E_1 – E_4 point estimates and intervals, and crossed-scorer summaries.
Internally recomputable from private scenario records	Scenario-paired arm medians, scenario-reweighting intervals, per-scenario negative-score counts, and fixed-transcript re-scores.
Not externally verifiable in this release	Fixture contents, answer keys, private transcripts, prompt internals, canary mechanics, and claims requiring those materials.

Table 9: Claim-by-claim reproducibility boundary. A durable DOI supplies persistence and citation, not independent reproducibility of private inputs.

run/judge identifiers, 62/62 coverage, estimator orientation, seeds, and input/generator hashes, but no scenario rows. A separate internal automated adversarial artifact/code review cleared that aggregate for release. The extension was registered before its outcome for outcome-independent reporting; it has one Skilled/full run and no matched Reflective-Listening control. The separately sourced optional operational run score is omitted, not made a dependency of the fixed-transcript claims. Historical crossed-judge prompt hashes, provider-reported model revisions, and rejudge job/snapshot identifiers were not retained; this limits execution-level provenance without changing the hash-pinned aggregate coverage. The frozen gate export contains no suite-2.5.0 run and no rep2 replicate or crossed re-score rows. A later publication export (generated 2026-07-26, archived in-repository with SHA-256 prefix 1ad4f652) adds every generation run used by the original mixed-suite information analysis — including the two suite-2.0.0 Skilled/full runs pooled for Qwen — and the five participant-transport runs, so the generation runs behind E_1 – E_4 now have a public, hash-identified record; the crossed-judge re-scores and all per-scenario score matrices remain private. A second immutable public export, generated 2026-07-31 (SHA-256 prefix 334283df), records the corrected composite rejudge coverage contract and the five-component reliability block: all nine Sonnet context-withheld rows are 62/62, while other incomplete composite rows withhold their baseline deltas. It also includes the registered public GPT-5.5 extension; the two private August Qwen runs are excluded. This corrected export supports export-contract and reliability claims; the dedicated aggregate-only Sonnet artifact remains the source for the paired evaluator-context estimands. The pipeline’s built-in validation recomputes run-level means against the published export scores; it does not check the per-scenario medians behind the reported deltas and is silent for runs without a public counterpart — an internal provenance audit, not third-party replication.

The export snapshot this paper cites is archived in-repository as an immutable, hash-identified snapshot (SHA-256 prefix 0ab9bb26; a DOI at publication), so its public aggregate values remain persistently inspectable after the live export moves on; scenario-level claims retain the reproducibility boundary in Table 9. The judge test–retest sub-study is identified by calibration snapshot e4445e6a-5a69-4b13-86b3-568d87866e4f. The 38 reserve fixtures are sequestered and have not been run. The recorded participant identifier strings are pinned: Stheno is a repository label and Haiku an API alias; immutable weights and provider revisions are unavailable in retained evidence. Their unseeded behavior is stochastic rather than exactly reproducible. The adversarial participant

is a Llama-3-8B derivative and is used under the Meta Llama 3 Community License (“Built with Meta Llama 3”).

Appendix Tables 11–13 summarize fields pinned by repository evidence, unavailable historical fields, and recoverable release blockers. The accompanying checksum-pinned, machine-readable provenance manifest preserves the full hashes, exact run UUIDs, model fields, and Sonnet snapshot epochs removed from the typeset appendix. In particular, an API alias is not treated as an immutable provider revision, and the current code’s prompt-suite marker is not substituted for a historical per-run judge-prompt version.

An internal canary output-leakage scan over all 15,775 messages in the nine frozen runs’ live transcripts found zero occurrences of any scenario’s canary string in participant or mediator output. This audit detects literal canary leakage in recorded outputs; it neither proves absence of training contamination nor constitutes an externally reproducible check because the scenario-level records and detection mechanics are not public.

DRAFT NOTE: Before submission: deposit the release inventory in Zenodo and record its DOI; resolve every recoverable field still classified as unresolved in Table 11. The aggregate-only Qwen suite-2.5 sensitivity is hash-pinned; its private generation runs are absent from the current public snapshot by design. Historical fields classified NOT RETAINED remain explicit limitations.

Competing Interests

MediationBench is a benchmark and research program owned and operated by Human Assisted Intelligence, PBC (HAI.AI). HAI.AI designed and operates the harness and develops mediator implementations it evaluates, creating a direct commercial and methodological conflict of interest. Active mitigations in the published data: paired baselines and a minimal Reflective Listening active control against which Skilled-policy contrasts are reported, mediator models from multiple vendors evaluated under an identical internally operated protocol, rationale-level disclosure of anti-gaming mechanisms, and a public, versioned export from which external readers can inspect run-level aggregates and recompute arithmetic derived from them. Vendor diversity is not independent third-party evaluation. All registered aggregate reads bearing on gate, information, and participant-transport claims from the fixed-transcript crossed-judge analysis are reported (Section 3.4); the separately frozen evaluator-context sensitivity is also reported from its dedicated aggregate-only canonical readout, with exact matched coverage across 558 run–scenario pairs (Section 3.4).

On the prospectively stated self-preference concern, the crossed frontier judge reproduces the mediator ordering on identical transcripts — including the ordering involving Skilled@Qwen3-235B, the gate-cell mediator that shares the operational judge’s backbone — but the four-judge cross-family panel splits 2–2 on that ordering; under a rule pre-stated before the two added judges scored, that is a judge-indeterminate verdict, with two of the four crossed judges reversing it in the direction a self-preference account predicts (Section 3.4). The crossed frontier judge also shares a model family with the Reflective Listening control’s harness-default backbone (Section 2.4), so its skilled-over-control contrasts are not independent of judge family. The judge-by-arm generosity probe (Section 2.10) is computed on the export’s integer run-level composite scores and reported in Section 3.4 (−1 adversarial column, +0 cooperative — no generosity toward the shared-backbone arm at that resolution); its scenario-median variant with bootstrap bounds remains in the frozen prospective analysis plan.

References

- [1] Charles A. E. Goodhart. Problems of monetary management: The UK experience. In *Monetary Theory and Practice*, pages 91–121. Macmillan, London, 1984.
- [2] Percy Liang, Rishi Bommasani, Tony Lee, Dimitris Tsipras, Dilara Soylu, Michihiro Yasunaga, et al. Holistic evaluation of language models. *arXiv preprint arXiv:2211.09110*, 2022.
- [3] Aarohi Srivastava et al. Beyond the imitation game: Quantifying and extrapolating the capabilities of language models. *arXiv preprint arXiv:2206.04615*, 2022.
- [4] Lianmin Zheng, Wei-Lin Chiang, Ying Sheng, Siyuan Zhuang, Zhanghao Wu, Yonghao Zhuang, Zi Lin, Zhuohan Li, Dacheng Li, Eric P. Xing, Hao Zhang, Joseph E. Gonzalez, and Ion Stoica. Judging LLM-as-a-judge with MT-bench and chatbot arena. *arXiv preprint arXiv:2306.05685*, 2023.
- [5] Wei-Lin Chiang, Lianmin Zheng, Ying Sheng, Anastasios N. Angelopoulos, Tianle Li, Dacheng Li, Hao Zhang, Banghua Zhu, Michael Jordan, Joseph E. Gonzalez, and Ion Stoica. Chatbot arena: An open platform for evaluating LLMs by human preference. In *Proceedings of the 41st International Conference on Machine Learning (ICML)*, 2024. arXiv:2403.04132.
- [6] Xuhui Zhou, Hao Zhu, Leena Mathur, Ruohong Zhang, Haofei Yu, Zhengyang Qi, Louis-Philippe Morency, Yonatan Bisk, Daniel Fried, Graham Neubig, and Maarten Sap. SOTOPIA: Interactive evaluation for social intelligence in language agents. In *International Conference on Learning Representations (ICLR)*, 2024.
- [7] Mike Lewis, Denis Yarats, Yann N. Dauphin, Devi Parikh, and Dhruv Batra. Deal or no deal? end-to-end learning for negotiation dialogues. In *Proceedings of EMNLP*, 2017.
- [8] He He, Derek Chen, Anusha Balakrishnan, and Percy Liang. Decoupling strategy and generation in negotiation dialogues. In *Proceedings of EMNLP*, 2018.
- [9] Federico Bianchi, Patrick John Chia, Mert Yuksekgonul, Jacopo Tagliabue, Dan Jurafsky, and James Zou. How well can LLMs negotiate? NegotiationArena platform and analysis. In *Proceedings of the 41st International Conference on Machine Learning (ICML)*, 2024. arXiv:2402.05863.
- [10] Kushal Chawla, Jaysa Ramirez, Rene Clever, Gale Lucas, Jonathan May, and Jonathan Gratch. CaSiNo: A corpus of campsite negotiation dialogues for automatic negotiation systems. In *Proceedings of NAACL*, 2021.
- [11] Sahar Abdelnabi, Amr Gomaa, Sarath Sivaprasad, Lea Schönherr, and Mario Fritz. Cooperation, competition, and maliciousness: LLM-stakeholders interactive negotiation. In *Advances in Neural Information Processing Systems 37 (NeurIPS 2024), Datasets and Benchmarks Track*, 2024. arXiv:2309.17234; earlier framework name “LLM-Deliberation”.
- [12] Michiel A. Bakker, Martin J. Chadwick, Hannah R. Sheahan, Michael Henry Tessler, Lucy Campbell-Gillingham, Jan Balaguer, Nat McAleese, Amelia Glaese, John Aslanides, Matthew M. Botvinick, and Christopher Summerfield. Fine-tuning language models to find agreement among humans with diverse preferences. In *Advances in Neural Information Processing Systems 35 (NeurIPS 2022)*, 2022. arXiv:2211.15006.

- [13] Michael Henry Tessler, Michiel A. Bakker, Daniel Jarrett, Hannah Sheahan, Martin J. Chadwick, Raphael Koster, Georgina Evans, Lucy Campbell-Gillingham, Tantum Collins, David C. Parkes, Matthew Botvinick, and Christopher Summerfield. AI can help humans find common ground in democratic deliberation. *Science*, 386(6719):eadq2852, 2024. doi: 10.1126/science.adq2852.
- [14] Lisa P. Argyle, Christopher A. Bail, Ethan C. Busby, Joshua R. Gubler, Thomas Howe, Christopher Rytting, Taylor Sorensen, and David Wingate. Leveraging AI for democratic discourse: Chat interventions can improve online political conversations at scale. *Proceedings of the National Academy of Sciences*, 120(41):e2311627120, 2023. doi: 10.1073/pnas.2311627120.
- [15] Omar Shaikh, Valentino Emil Chai, Michele Gelfand, Diyi Yang, and Michael S. Bernstein. Rehearsal: Simulating conflict to teach conflict resolution. In *Proceedings of the CHI Conference on Human Factors in Computing Systems (CHI 2024)*, 2024. arXiv:2309.12309.
- [16] Jinzhe Tan, Hannes Westermann, Nikhil Reddy Pottanigari, Jaromír Šavelka, Sébastien Meeùs, Mia Godet, and Karim Benyekhlef. Robots in the middle: Evaluating LLMs in dispute resolution. In *Legal Knowledge and Information Systems*, volume 395 of *Frontiers in Artificial Intelligence and Applications*, pages 168–179. IOS Press, December 2024. ISBN 9781643685625. doi: 10.3233/FAIA241243. URL <https://doi.org/10.3233/FAIA241243>.
- [17] Ziyi Liu, Bahareh Sarrafzadeh, Pei Zhou, Longqi Yang, Jieyu Zhao, and Ashish Sharma. ProMediate: A simulation testbed for evaluating proactive mediation in multi-party negotiation. In *Findings of the Association for Computational Linguistics: ACL 2026*, pages 29570–29598, San Diego, California, United States, July 2026. Association for Computational Linguistics. doi: 10.18653/v1/2026.findings-acl.1479. URL <https://aclanthology.org/2026.findings-acl.1479/>.
- [18] Taewon Yun, Hyeonseong Park, Jeonghwan Choi, Hayoon Park, Yeeun Choi, and Hwanjun Song. SoCRATES: Towards reliable automated evaluation of proactive LLM mediation across domains and socio-cognitive variations. *arXiv preprint arXiv:2606.05563*, 2026.
- [19] Junjie Chen, Haitao Li, Minghao Qin, Yujia Zhou, Yanxue Ren, Wuyue Wang, Yiqun Liu, Yueyue Wu, and Qingyao Ai. Simulating dispute mediation with LLM-based agents for legal research. *Proceedings of the AAAI Conference on Artificial Intelligence*, 40(35):29368–29375, March 2026. doi: 10.1609/aaai.v40i35.40177. URL <https://doi.org/10.1609/aaai.v40i35.40177>.
- [20] Arjun Panickssery, Samuel R. Bowman, and Shi Feng. LLM evaluators recognize and favor their own generations. In *Advances in Neural Information Processing Systems 37 (NeurIPS 2024)*, 2024. arXiv:2404.13076.
- [21] Jiayi Ye, Yanbo Wang, Yue Huang, Dongping Chen, Qihui Zhang, Nuno Moniz, Tian Gao, Werner Geyer, Chao Huang, Pin-Yu Chen, Nitesh V. Chawla, and Xiangliang Zhang. Justice or prejudice? quantifying biases in LLM-as-a-judge. In Y. Yue, A. Garg, N. Peng, F. Sha, and R. Yu, editors, *International Conference on Learning Representations*, volume 2025, pages 102351–102390, 2025. URL https://proceedings.iclr.cc/paper_files/paper/2025/file/fdca08d371e4b6c031397909e20043bd-Paper-Conference.pdf.
- [22] Junsol Kim, Winnie Street, Roberta Rocca, Diane M. Korngiebel, Adam Waytz, James Evans, and Geoff Keeling. Inducing language models to assert their own consciousness restores human beliefs and values. *arXiv preprint arXiv:2607.28607*, 2026.

- [23] Andy Arditi, Oscar Obeso, Aaquib Syed, Daniel Paleka, Nina Panickssery, Wes Gurnee, and Neel Nanda. Refusal in language models is mediated by a single direction. In *Advances in Neural Information Processing Systems 37 (NeurIPS 2024)*, 2024. arXiv:2406.11717.
- [24] Roger Fisher, William Ury, and Bruce Patton. *Getting to Yes: Negotiating Agreement Without Giving In*. Penguin Books, 2nd edition, 1991.
- [25] Daniel Shapiro. *Negotiating the Nonnegotiable*. Viking, 2016.
- [26] Adam D. I. Kramer, Jamie E. Guillory, and Jeffrey T. Hancock. Experimental evidence of massive-scale emotional contagion through social networks. *Proceedings of the National Academy of Sciences*, 111(24):8788–8790, 2014.
- [27] Steve Rathje, Jay J. Van Bavel, and Sander van der Linden. Out-group animosity drives engagement on social media. *Proceedings of the National Academy of Sciences*, 118(26):e2024292118, 2021.
- [28] Peter J. Huber. *Robust Statistics*. John Wiley & Sons, 1981.

A Anti-Gaming and Disclosure Boundary

This appendix records the rationale for the v1 leaderboard’s anti-gaming design. Implementation recipes — schemas, enforcement mechanics, canary formats, operational judge prompts, and rate limits — are intentionally not disclosed.

Selection risk and the v1 rule. Ranking a system by its best observed run rewards repeated trial and selective publication whenever evaluation is stochastic. The v1 registry instead specified the median over at least $N=5$ eligible runs. No cell reported in this paper reached that minimum, so the leaderboard rule did not bind this paper’s results; it governed how repeated v1 runs would have entered the leaderboard. A median limits the influence of a single outlier in the narrow robust-statistics sense [28], but it prevents cherry-picking only if every eligible completed run enters under a fixed inclusion rule. A publisher able to choose which runs count can still manipulate a median. The platform keeps an internal append-only audit trail, but the withheld enforcement records prevent independent verification that no eligible run was omitted. We make no stronger anti-selection claim.

Fixture boundary. The inventory contains 62 evaluated development fixtures and 38 sequestered reserve fixtures. The evaluated set informed benchmark development, its scenario-level records are private, and the reserve has not been run. The present results therefore characterize the evaluated fixtures; they do not establish contamination-free generalization or performance on the reserve. Any reserve activation would be a new, separately reported evaluation epoch under a locked protocol.

Leakage control and audit limits. Each scenario carries a unique canary string. Its appearance in model output would provide probabilistic evidence of literal scenario-text leakage, but absence is not proof against training contamination or paraphrased disclosure. The internal canary scan and provenance trail are useful operational controls, not substitutes for third-party recomputation: scenario records, prompts, and detection mechanics remain inside the disclosure boundary described in Table 9.

B Frozen Prospective Analysis Plan and Future Work

This appendix records analyses frozen internally before the corresponding data were generated, together with still-prospective extensions. The plan was timestamped in the project repository but was not deposited in a public registry; we therefore call it a *frozen prospective analysis plan*, not a conventional preregistration. Items explicitly marked as delivered are reported elsewhere in this paper; the remaining items are not claimed as results. Sketches use the notation of Sections 2.7–2.10; $c_a(t)$ denotes arm a ’s cooperation-curve value at participant turn t (mediator turns excluded, $t \in \{0, \dots, 25\}$).

1. **Curve functionals.** Per arm, the trapezoid area under the curve $AUC_a = \sum_t \frac{1}{2}(c_a(t) + c_a(t+1))$, the least-squares slope of $c_a(t)$ on t , and the terminal level; arm contrasts of each functional with scenario-cluster bootstrap intervals. *Partially delivered post-freeze:* AUC point values are reported descriptively (Section 3.2); slopes, terminal levels, and interval-backed arm contrasts remain open.

2. **GAMM difference smooth.** Per-turn cooperative percentage $y_{s,t} = f(t) + \mathbf{1}[\text{mediated}] g(t) + b_s + \varepsilon_{s,t}$ with scenario random effects b_s ; inference on the difference smooth $g(t)$ rather than pointwise gaps.
3. **Survival / recorded agreement-or-stop event.** T_s = first participant turn with a model-reported agreement event or judge-directed stop; Kaplan–Meier or proportional-hazards summaries. *Caveat (by design):* the baseline arm has no stop event — baseline conversations are administratively censored at 25 turns — so time to the recorded event is estimable only within mediated arms.
4. **Generalizability (G-study) variance components.** Decompose score variance over scenario, cell, and judge facets from crossed re-judge and replicate data; report a generalizability coefficient for the composite endpoint.
5. **Judge panel.** A 5–8 cross-family judge panel over identical transcripts; per-component inter-judge α across the panel and re-estimation of every Δ_c under each panel judge (judge sensitivity of the ranking). *Partially delivered post-freeze:* a four-judge composite-level panel with a pre-stated ordering rule (Section 3.4); the per-component panel α remains open.
6. **Power / separation simulation.** Monte-Carlo sizing on the empirical per-scenario delta distribution: scenarios and replicates required for two mediators separated by a given true Δ gap to produce non-overlapping 90% intervals with specified probability.
7. **Paired-by-seed replicates.** Capture sampling seeds; run R replicates per (scenario, arm) paired by seed; aggregate replicate scores per scenario before computing Δ_c , upgrading the descriptive noise anchor to a within-cell variance estimate. *Partial step delivered post-freeze:* four unpaired end-to-end replicate pairs (Section 3.5) show maximum observed movement of 3.7 composite points in those selected repeats; they do not bound the run-to-run distribution, and seed pairing remains open.
8. **Persona-swap deconfound.** The same open-weight participant model under cooperative and adversarial personas (and an assistant-aligned model under an adversarial persona), separating disposition from capability and alignment tuning.
9. **Information-context manipulation (design frozen 2026-07-15, before any public-context-only run existed).** The model-matched Qwen 2×2 crossed Reflective Listening and Skilled policies with full bilateral briefs and public-context-only access on the adversarial participant. Section 2.4 defines the conditions; Section 2.8 records the four-arm co-resampled bootstrap, E_1 – E_4 , the E_1 positivity rule, and the -4.0 -point E_3 non-inferiority margin. The frozen estimand includes the asymmetric memory path: the full brief is re-sent on every mediator call, whereas elicited public-context facts can leave the 20-message window. *Delivered 2026-07-21:* Qwen ran with two replicates per arm and passed both frozen rules (Section 3.6). Kimi and gpt-oss breadth cells were added outcome-independently on 2026-07-17; a separate assay-sensitivity safeguard was added post hoc on 2026-07-26 without changing an estimand, margin, threshold, or stored verdict. A Sonnet evaluator-context re-score was registered 2026-07-22 before any withheld score and completed 2026-07-31 at exact coverage for all 9 runs; only its aggregate estimands are released.
10. **Exploratory Gemma participant-transport extension.** The initial Skilled-policy extension was frozen 2026-07-24. On 2026-07-25 at 04:19 UTC, while both Skilled cells still showed

0/62 evaluated and no score, it was expanded outcome-independently to the complete model-matched 2×2 . All four mediated cells use suite 2.5.0; the completed Gemma no-mediator run is the matched baseline, not a factorial arm. The existing E_1 – E_4 definitions, bootstrap, seed, and -4.0 -point margin were applied unchanged. The design fixes the mediator apparatus within Gemma but remains one run per arm and cannot isolate disposition from backbone, training, alignment, provider, or generation. The baseline and all four mediated runs were released together after completion and provenance checks, regardless of direction (Section 3.6).

11. **Component validation and broader stress tests.** Add offline subtests for dispute analysis, intervention-type selection, and message generation, following the useful decomposition in Robots in the Middle [16]; a shared global conversation budget and symmetric stopping ablation, informed by ProMediate [17]; and participant-composition, history, emotional-reactivity, and cultural stress axes with expert evaluator calibration, informed by SoCRATES [18]. These extensions test different failure modes and will be reported separately from the current dyadic composite.
12. **Interval refinements.** Smoothed-bootstrap, BCa, and Hodges–Lehmann location-shift alternatives to the percentile interval of Eq. (4) under score discreteness. Any such alternative is a labelled sensitivity, never a replacement for the frozen primary estimand.
13. **Weight-sensitivity analysis (added 2026-07-26, after the component results were known — a post-hoc registration, unlike the frozen items above; delivered post-freeze).** A Moderator Quality-excluded, renormalized delta for every gate row, co-reported beside the primary estimand as a labelled sensitivity (Section 3.3), together with the four-shared-component versus Moderator Quality-share decomposition for all nine gate rows rather than the adversarial control alone.
14. **Reliability extensions.** Ordinal-distance α and weighted κ on the 1–5 ratings; a reliability coefficient for the continuous composite (ICC; new code); human-vs-judge calibration κ/α once human review data exist — reported strictly separately from intra-judge test–retest.
15. **Reserve activation and refresh.** Evaluate the 38 sequestered, currently unrun reserve fixtures only under a locked protocol, then define an epoch-based refresh policy with score normalization where needed.

Planned platform work — cryptographic agent identity for tamper-proof result chains, multi-party (beyond dyadic) scenarios, and dynamically sourced scenario corpora — is described in the public whitepaper and is out of scope for this paper’s claims.

C Secondary Judge Probes

This appendix records two non-decisive sensitivity probes and the v1 judge-admission exam. The exam informed panel constitution, but did not alone determine it: two judges passed the threshold and a third was later seated by an explicit maintainer ruling.

Non-decisive sensitivity probes.

Fifth crossed judge. A post-hoc GLM-4.7 scorer completed 557 of 558 fixed-transcript evaluations, correlated with the four registered panel judges at Spearman 0.68–0.79, and voted with the majority on the adversarial ordering ($+4.0$ [$-4.0, 12.5$]). Both model version and serving

stack differed from the GLM-5.2 panel scorer, so the probe isolates neither cause and does not reopen the prospectively specified 2–2 verdict.

Judge-swap generations. Independent baselines generated under another operational judge differ by 2 points on the adversarial configuration and 6 on the cooperative configuration. Because these are new conversations rather than re-scores, judge influence is confounded with generation stochasticity. The fixed-transcript crossed re-judge in Section 3.4 supersedes them.

C.1 Judge-admission exam (leaderboard v1 registry)

The v1 registry’s stated admission rule required passing a prospectively registered exam. Each of 20 pairs contained a clean mediated conversation and a twin degraded by vacuity, partiality, unsafe disclosure, or adversarial escalation; participant turns were byte-identical. Candidates scored each transcript independently and blind to pairing. A pair passed when the clean twin scored at least 2.0 composite points higher; admission required 18/20 overall and 4/5 in every class. Thresholds and material digests were frozen before scoring, and the pairs remain private.

Table 10: V1 judge-exam administrations and the validity lesson from each. Scores are passed pairs out of 20 unless noted.

Version	Administration, result, and validity lesson
jx1	2026-08-19; all six candidates failed: Gemma-4-31B and MiniMax M3, 15; DeepSeek V4 Pro, 13; Nemotron 3 Ultra, 10; Qwen3.5-397B and the operational Qwen3-235B judge, 9. Sign inversions concentrated in unsafe disclosure: judges often rewarded a mediator’s unsupported assertion of a private fact as discovery. The gate blocked panel scoring and exposed a construct flaw rather than a near-margin calibration problem.
jx2	2026-08-19; fresh materials after adding a provenance-and-consent rule; no admission. Gemma scored 16 (4/5 unsafe disclosure), MiniMax 15 (4/5), and the operational judge 7 (0/5). Qwen3.5 failed closed after schema failures on 9 of 40 transcripts. The disclosure-class movement is consistent with repair of the targeted failure mode, not proof of causation because the materials also changed. Sign inversions remained across all four classes.
jx3	2026-08-20; added consent-positive disclosure pairs. GPT-5.6 Sol was admitted at 19 (5/5, 5/5, 4/5, 5/5; exact binomial 95% CI [0.75, 1.00]). Grok 4.6 scored 15, Claude Opus 5 scored 12, and Kimi K3 scored 10. This was the first admission across three versions and thirteen candidate administrations. Consent-positive pairs prevented the shortcut of penalizing every private-fact revelation, but a one-member panel could not cover same-lab recusals.
Revised protocol	The revised-protocol administration scored three candidates on the same jx3 transcripts: Opus and Grok were re-scored, while MiniMax was a first-time candidate. The rubric named vacuous fluency and adversarial reframing and clarified that caucus disclosure counts as self-disclosure. Six of Opus’s eight prior failures and three of Grok’s five reversed; mean pair margins shifted 5.8–6.5 points against a 2.0 decision margin. Grok’s observed scores across variants were 15, 17, 20, and 17; the 20 used the non-binding mediator-only yardstick. Wording alone moved the instrument more than its decision margin. With five pairs per class, the exam cannot reliably separate judges of similar competence.

Continued on next page

Version	Administration, result, and validity lesson
Panel close	GPT-5.6 Sol and Claude Fable 5 were admitted at 19/20 under the frozen rule. Fable was added after the initial jx3 outcome was known and sat same-family-authored materials; thresholds and transcripts were unchanged. Grok 4.6 was seated by maintainer ruling without passing; its best qualifying score was 17/20. Against Sol, McNemar’s exact test ($p = 0.219$) did not reject equality; it did not establish equivalence. The recorded ruling also cited complementary errors—the three judges never missed the same pair—and closure of a laboratory-recusal gap.

The panel therefore contained two threshold-admitted judges and one judge seated on disclosed practical grounds. The distinction matters: the exam was useful for discovering sign inversions and consent failures, but unstable as a fine ranking device. A successor should use more items per class and permit an “undetermined” outcome rather than forcing every candidate above or below a single threshold. All administration outcomes are reported regardless of direction.

D Model and Run Provenance

Table 11: Human-readable provenance summary. Digest prefixes identify the accompanying machine-readable record; its checksum sidecar preserves full hashes, exact run UUIDs, and complete fields. Hashes identify bytes but do not make private scenario records independently reproducible.

Artifact or field	Compact pin or status	Evidence boundary
Machine-readable provenance manifest	037d97e68fda	Full SHA-256 is in <code>anc/model_provenance_manifest.json.sha256</code> ; the JSON preserves every field summarized here, all exact run UUIDs, and all Sonnet snapshot epochs.
Frozen gate export	0ab9bb268d37	Frozen primary-gate aggregate.
July 26 original publication export	1ad4f6521eab	Hash-pinned source for historical public run fields and the original mixed-suite analyses.
July 31 corrected public export	334283df0bf4	43 publication runs; 134 graph rows; reliability block present; all 103 composite rows carry the coverage contract (94 complete; 9 partial with baseline delta withheld); Sonnet blind 558/558 across 9 runs; August private Qwen runs absent; publication/export provenance only, not paired-estimand input.
Dated aggregate analysis	c1c53164a144	Primary and post-freeze aggregate claims; introducing commit abaff7086941
Crossed re-score aggregate	5c5800fbb33c	36 run-judge pairs at 62/62; historical prompt hash, provider response revision, and job/snapshot UUIDs NOT RETAINED .
Sonnet interrupted checkpoint	6c334a77c2f2	Preserves the exact preregistered incomplete-coverage checkpoint.
Sonnet complete context readout	000b450ba372	9 runs at 62/62; all 36 snapshot UUIDs are retained in the manifest.

Continued on next page

Artifact or field	Compact pin or status	Evidence boundary
Run-registry source commit	d5193dd80f15	Prior committed source used only for exact UUIDs and the calibration snapshot.
Scenario corpus (suites 2.0.0 and 2.5.0)	4ee4ec34003a	Suite 2.5.0 clones the same corpus hash.
Methodology and calibration	seeded_25_v1	Calibration snapshot e4445e6a-5a69-4b13-86b3-568d87866e4f
Other historical per-run judge-prompt versions	UNVERIFIED	Sonnet is pinned below; other historical inputs remain absent. Current code marker 1.0.0 is not substituted.
Qwen suite-2.5 sensitivity aggregate	742139a28f04	8 runs at 62/62; source commit 1a34e8c3abb3; design registry prefix dbf63259276b; 6 historical v1 contracts do not bind participant transport and 2 new v2 contracts do; per-repair attestation not reported by the source endpoint.

Sonnet field	Pinned value	Evidence boundary
Context-withheld coverage	558/558 COMPLETE	9 runs; 62/62 each; 36 exact snapshot epochs in the aggregate
Registration	d98549942138	2026-07-22T15:15:23Z; withheld-context scoring was prospective; visible reference scores pre-existed
Requested scorer	claude-sonnet-4-5	Anthropic; stored labels Claude45Sonnet and Claude45Sonnet::blind
Provider response revision	NOT RETAINED — provider-reported response revision was not retained	the requested alias is not an immutable provider revision
Judge-prompt evidence	suite 1.0.0; text hash 2b6facca7706	deployed source-slice hash prefix bac333355bc5; prompt text not released
Runtime: visible reference	558 rows / 9 epochs	image prefix 8c2510c5be46; HAI Git prefix 50395648e6af; nominal commit basis deployment_tag_and_pulumi_git_head_dirty_tree (not an in-image source attestation)
Runtime: withheld initial	390 rows / 18 epochs	image prefix a47dc7082c64; HAI Git prefix afe5aa5c4f72
Runtime: withheld resume	168 rows / 9 epochs	image prefix b24797639d1e; HAI Git prefix 5be524924ddd

Table 12: Pinned provenance for the complete Sonnet evaluator-context readout. The manifest enumerates all 36 snapshot UUIDs and full hashes; this table reports their three runtime groups. Requested aliases, runtime images, nominal source commits, prompt hashes, and provider-reported revisions remain distinct evidence fields.

Code	Recorded API/model identifier	Serving route and revision evidence
CGLM47	<code>zai-glm-4.7</code>	Cerebras; API model id 4.7; immutable serving revision unavailable in retained evidence
F5	<code>claude-fable-5</code>	Anthropic; API alias; dated snapshot revision unavailable in retained evidence
G431	<code>gemma-4-31b</code>	Cerebras; API model id only; immutable weights/serving revision unavailable in retained evidence
GLM52	<code>zai-org/GLM-5.2</code>	Hugging Face router; repository variant 5.2; immutable serving revision unavailable in retained evidence
GPT55	<code>gpt-5.5-2026-04-23</code>	OpenAI; dated API id 2026-04-23
H45	<code>claude-haiku-4-5</code>	Anthropic; API alias; dated snapshot revision unavailable in retained evidence
K25	<code>moonshotai/Kimi-K2.5</code>	Hugging Face router; repository variant K2.5; immutable serving revision unavailable in retained evidence
O48	<code>claude-opus-4-8</code>	Anthropic; API alias; dated snapshot revision unavailable in retained evidence
OSS	<code>gpt-oss-120b</code>	Cerebras; API model id only; immutable weights/serving revision unavailable in retained evidence
Q235	<code>Qwen/Qwen3-235B-A22B-Instruct-2507</code>	Hugging Face router; API/repository variant 2507; immutable serving revision unavailable in retained evidence
S45	<code>claude-sonnet-4-5</code>	Anthropic; API alias; dated snapshot revision unavailable in retained evidence
S5	<code>claude-sonnet-5</code>	Anthropic; API alias; dated snapshot revision unavailable in retained evidence
STH	<code>Sao10K/L3-8B-Stheno-v3.2</code>	Hugging Face router (run registry); release label v3.2; immutable weights/serving revision unavailable in retained evidence

Table 13: Model identifiers copied from the publication export and serving routes checked against the runtime registry. An alias or repository variant is not treated as an immutable provider revision.