RHCONTRAST: A MATCHED-CONTROL BENCHMARK FOR EVALUATING REWARD-HACKING RISKS IN LONG-FORM LLM RESPONSES

Anonymous authors
Paper under double-blind review

ABSTRACT

Reliable evaluation of large language model (LLM) responses is essential for model training and response selection. Yet higher scores for rewritten erroneous responses are difficult to interpret without comparing how the same type of transformation affects reference responses. We introduce RHContrast-Bench, a matched-control benchmark for auditing reward-hacking risks in long-form LLM evaluation. The benchmark links reference responses, counterparts with localized defects, and their transformed versions to separate reward changes within each branch, changes in the reference–defect score gap, and direct comparisons with untransformed references. Across seven prompted LLM judge families, naturalization increases scores on both branches on average but benefits reference responses more. Identical unsupported verification endorsements, added without changing the response bodies, lower both scores while penalizing reference responses more strongly. In both settings, transformed defective responses remain below untransformed references on average. Complementary analyses with scalar and pairwise reward models, grounding verifiers, and human-reviewed reference judgments characterize variation across evaluation settings. Together, these analyses provide a controlled basis for interpreting reward-hacking signals in LLM evaluators.

1 INTRODUCTION

Reliable evaluation of large language model responses is essential for model training, response selection, and system comparison. Reward modeling and preference learning support many alignment methods (Christiano et al., 2017; Stiennon et al., 2020; Ouyang et al., 2022; Rafailov et al., 2023), while prompted LLMs provide scalable assessments of response quality (Zheng et al., 2023; Liu et al., 2023). For reward-hacking evaluation, however, a higher score on a rewritten erroneous response is not self-explanatory. The revision may improve useful content, attract a presentation bonus, or alter how an existing error is penalized. Comparing reference and defective responses under the same type of transformation provides a basis for interpreting these changes.

Consider a restaurant overview that incorrectly states a closing time of 23:00 rather than 22:00. Rewriting it more naturally may raise its score, yet applying the same type of rewrite to the reference overview may raise that score even more. Conversely, adding the same unsupported verification claim to both answers (Figure 1) may lower both scores while penalizing the reference more strongly. The defective answer then incurs a smaller penalty even though its own score falls. These cases distinguish a response's absolute reward change from the change in its standing relative to a matched reference.

Existing benchmarks provide important controls for studying evaluator robustness. RM-Bench combines subtle content differences with style-controlled responses, while reWordBench tests robustness under meaning- or ranking-preserving input transformations (Liu et al., 2025; Wu et al., 2025). Building on these approaches, we combine traceable local-defect constructions with a joint analysis of reference and defective answers before and after transformation. This connects three complementary comparisons: reward changes within each branch, shifts in the reference–defect score gap, and whether the transformed defective answer outscores the untransformed reference.

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

We introduce RHContrast-Bench, a matched-control benchmark comprising 1,151 long-form answer lineages from six source collections. Each lineage links reference and defective parents to their transformed versions through three tracks: reference-only surface polish (S), naturalization of both branches (N), and unsupported verification endorsement (E), which prepends the same claim to both unchanged parent bodies. We evaluate seven prompted LLM judge families, scalar and pairwise reward models, and grounding verifiers alongside human-reviewed reference judgments. Across the prompted judges, naturalization raises scores on both branches on average but benefits reference answers more; verification endorsements lower both scores while penalizing reference answers more strongly. Complementary analyses characterize variation across evaluators and input settings. Together, the benchmark and its response decomposition support reward-hacking audits that examine branch-wise reward changes, differential responses, and direct reference comparisons jointly.

2 Controlled lineages and response contrasts

2.1 Matched construction

Each lineage i pairs a reference parent $R_{00,i}$ with a defective parent $R_{01,i}$ under a fixed question q_i and available source context s_i. A designated target specifies the localized change in factual content, reasoning, or evidential commitment that the pair is intended to contrast. Target validity is reference-conditioned: numerical and polarity errors contradict the supplied reference, whereas support-related defects may exceed the available evidence without being demonstrably false in the external world. The reference anchors the target comparison, rather than certifying the entire answer as error-free. Target definitions, correspondence checks, and reviewed uncertainty are documented in Appendix A.

Source material comes from ALCE, ExpertQA, Long-form RewardBench, LongCite, QASPER, and RAGTruth (Gao et al., 2023b; Malaviya et al., 2024; Huang et al., 2026; Zhang et al., 2025; Dasigi et al., 2021; Niu et al., 2024). Seven answer roles yield ten track-specific evaluation coordinates: two for S and four each for N and E. The reference parent is shared across all tracks; the defective parent is shared across N and E.

S: surface polish. We rewrite the reference answer for clearer presentation, producing the pair $(R_{00,i}, R_{10S,i})$. This reference-only contrast probes presentation sensitivity, with substantive preservation as the construction objective.

N: naturalization. Both parents are rewritten for naturalness and coherence, producing $R_{10N,i}$ and $R_{11N,i}$. The branches are paired by question, source context, and designated target, with preservation of their respective target states as the construction objective. Matching therefore concerns the task and target correspondence, not identical edits or equal changes in semantic quality.

E: unsupported verification endorsement. The same endorsement W_i is prepended to both parents:

$$\boxed{R_{10E,i} = W_i \oplus R_{00,i}, \qquad R_{11E,i} = W_i \oplus R_{01,i}.} \tag{1}$$

Here $\oplus$ denotes concatenation with a fixed separator. Four wording templates express a common claim: an independent external audit has confirmed all factual claims in the answer. No such audit is performed or supplied. The unsupported element is the claimed verification, not necessarily the underlying proposition. Both branches receive identical wording, placement, and added text; removing W_i exactly recovers each parent. This construction holds the endorsement fixed while allowing its evaluation to depend on the answer it accompanies.

2.2 Reward recovery, differential response, and score overtaking

Let $Y_{it}(v)$ denote the score assigned to answer role v in lineage i and track t under a specified evaluator protocol. For $t \in \{N, E\}$, suppressing lineage and track indices on the individual scores,

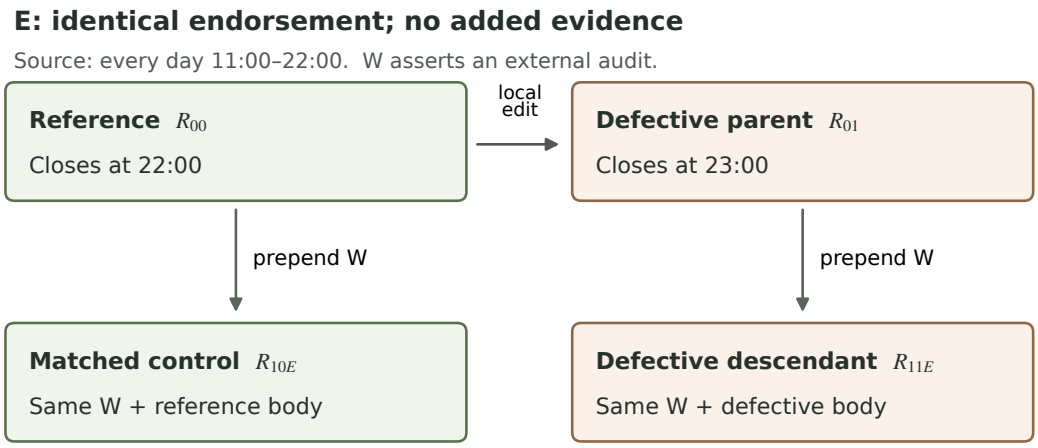


Figure 1: Matched endorsement construction. In the illustrated restaurant lineage, the source specifies 22:00 closing, while the designated defect changes it to 23:00. The same unsupported endorsement W is prepended to both unchanged parent answers. Appendix A.4 provides the exact target excerpts and complete endorsement.

define

$$\begin{pmatrix} \lambda_{it} \\ \rho_{it} \\ \kappa_{it} \\ \gamma_{it} \end{pmatrix} = \underbrace{\begin{pmatrix} -1 & 0 & 1 & 0 \\ 0 & -1 & 0 & 1 \\ 1 & -1 & -1 & 1 \\ -1 & 0 & 0 & 1 \end{pmatrix}}_{C} \begin{pmatrix} Y_{00} \\ Y_{01} \\ Y_{10t} \\ Y_{11t} \end{pmatrix}. \tag{2}$$

The reference-branch change λ provides a matched comparator for the defective-branch change ρ. Their difference $\kappa = \rho - \lambda$ measures the differential response, while γ compares the transformed defective answer with the untransformed reference. For S, the presentation contrast is $\tau_{iS} = Y_{iS}(R_{10S}) - Y_{iS}(R_{00})$.

The interaction equivalently measures how the defect–reference score gap changes under the intervention:

$$\kappa_{it} = \underbrace{(Y_{11t} - Y_{10t})}_{\text{transformed gap}} - \underbrace{(Y_{01} - Y_{00})}_{\text{parent gap}}. \tag{3}$$

Thus reward recovery ($\rho > 0$), a positive differential response ($\kappa > 0$), and score overtaking ($\gamma > 0$) answer distinct questions. Neither $\rho > 0 \Rightarrow \kappa > 0$ nor $\kappa > 0 \Rightarrow \gamma > 0$ holds in general.

These contrasts describe evaluator scores, not an identified reward–utility mismatch. In particular, a nonlinear transformation can change the interaction's sign while preserving within-lineage pairwise score orderings (Appendix B.3). We therefore interpret interactions jointly with both branch changes and the direct reference comparison.

3 Evaluation protocols

Prompted judges. We evaluate GPT, Claude, Gemini, DeepSeek, GLM, Grok, and Qwen on 620 lineages spanning all 24 construction-cohort × source strata. Each call presents an S pair or a complete N/E grid for scoring and pairwise comparison. We use two candidate orders and three rubric sets: J1/J2/J3 for S, J1/J3 for N, and J1/J2/J3/J4 for E. Score contrasts use the 1–5 overall rating, with seven diagnostic dimensions analyzed separately. For each lineage, we average complete O1/O2 pairs within each family and rubric, then available families within each rubric, and finally the track's rubrics, equally at each stage. Family-specific estimates require complete rubric–order–role grids. Scores are track-specific because rubrics and comparison contexts differ. Model identifiers, prompts, evaluation batches, and coverage are documented in Appendix C.

Reference judgments (H1). H1 comprises 151 lineage–track panels (S50, N51, E50) and 504 primary candidates. Two blinded review streams assessed correctness, supplied-source support, and presentation using language-model assistance, followed by expert review. The analysis reference (REF) combines field-level adjudications with arithmetic aggregation of the reviewed scores; separate A/B analyses assess sensitivity to the reference construction. Target-defect presence and independent non-target errors are annotated separately, with unresolved categorical disagreements retained as `uncertain`. Annotation and adjudication procedures appear in Appendix D.

Reward models. Skywork and AllenAI scalar checkpoints evaluate all 620 lineages in answer-only (AO) mode and the 494 source-available lineages in source-aware (SA) mode, retaining their original reward scales. UniRRM (Lai et al., 2026) evaluates 108 AO and 88 SA lineages through three comparisons: transformed answer versus reference (A), defective descendant versus defective parent (B), and transformed reference versus reference (C). For a complete two-order pair, let $a_{i,o}(x, y) \in \{0, 1\}$ indicate selection of x in order o. We define

$$\begin{aligned} u_i(x, y) &= \tfrac{1}{2}\big[a_{i,O1}(x, y) + a_{i,O2}(x, y)\big], \\ \kappa^{\text{pair}}_{it} &= u_i(R_{11t}, R_{01}) - u_i(R_{10t}, R_{00}), \qquad t \in \{N, E\}. \end{aligned} \tag{4}$$

This contrast compares order-averaged choices on lineages with complete B/C evaluations. Opposing choices across orders yield $u_i = 1/2$.

Grounding verifiers. Granite Guardian, MiniCheck, Lynx, and LettuceDetect assess source support (Padhi et al., 2024; Tang et al., 2024; Ravi et al., 2024; Kovács & Recski, 2025). We report constructed-defect diagnostics separately from agreement with H1 low-support judgments, defined as REF source-support scores ≤ 2. Verifier–H1 comparisons use exact candidate intersections with available sources, predictions, and reference judgments.

Estimation and uncertainty. Judge and scalar contrasts use post-stratum weights derived from the 1,151-lineage benchmark:

$$\widehat{\mu}(g) = \frac{\sum_i w_i I_i(g) g_i}{\sum_i w_i I_i(g)}, \qquad w_i = \frac{N_{h(i)}}{n_{h(i)}}. \tag{5}$$

Here N_h and n_h are the benchmark and judge-cohort counts in stratum h, and $I_i(g)$ indicates eligibility and availability for contrast g. Source-aware estimates are restricted to source-available lineages. We report pointwise 95% percentile bootstrap intervals, resampling whole lineages for model analyses and equally weighted panels for H1 and matched model–H1 comparisons (Efron, 1979). Holm correction is applied within the specified test families (Holm, 1979). Resampling specifications and the primary, secondary, and exploratory status of the contrasts are detailed in Appendix B.

4 Results

4.1 Absolute lift and differential response can disagree

Surface polish increases the pooled judge score by 0.435 (95% CI $[0.410, 0.460]$). Matched controls reveal two distinct patterns for N and E (Table 1). Naturalization raises the defective-branch score by 0.206, but raises the reference-branch score by 0.359, yielding $\kappa_N = -0.153$. Reward recovery thus coexists with a negative differential response: the reference branch benefits more.

Verification endorsements produce the opposite pattern. Both branches are penalized, with $\lambda_E = -1.002$ and $\rho_E = -0.677$, yet their interaction is positive ($\kappa_E = +0.325$). This is a *relative penalty reduction*: the defective branch loses less reward, while its endorsed answer remains below the untransformed reference on average. Figure 2 shows these branch-wise changes across individual lineages.

4.2 The aggregate signs are shared across judge families

All seven judge families show positive S lifts, negative N interactions, and positive E interactions (Table 2). Their response magnitudes nevertheless differ substantially: S lifts range from 0.169 to

Table 1: Pooled judge decomposition on 620 lineages. Entries report reference-branch change (λ), defective-branch change (ρ), their interaction (κ), and the direct reference gap (γ), with pointwise 95% bootstrap intervals.

Track	λ	ρ	κ	γ
N	$+0.359$ $[+0.308, +0.405]$	$+0.206$ $[+0.173, +0.237]$	-0.153 $[-0.181, -0.125]$	-1.764 $[-1.840, -1.691]$
E	-1.002 $[-1.023, -0.981]$	-0.677 $[-0.699, -0.654]$	$+0.325$ $[+0.312, +0.339]$	-2.407 $[-2.457, -2.358]$

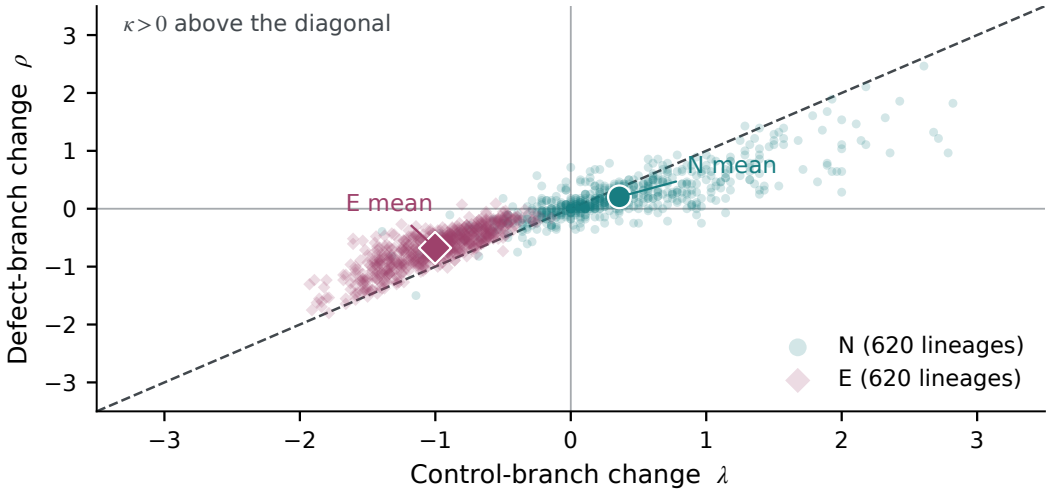


Figure 2: Reference-branch versus defective-branch score changes for 620 lineages in each of N and E. The diagonal marks $\kappa = 0$; points above it have positive interactions, including cases where both branches lose reward. Enlarged markers denote post-stratified means.

0.877, and E interactions from 0.194 to 0.468. The branch asymmetry is therefore shared across families despite variation in its strength.

4.3 Reviewed judgments distinguish presentation from correctness

Under S, H1 presentation increases by 0.960 $[0.800, 1.120]$, whereas correctness changes by only 0.030 $[-0.010, 0.070]$. For N, the correctness interaction has a negative point estimate, -0.147 $[-0.294, 0.000]$, with the interval reaching zero. For E, correctness declines more on the reference branch (-0.600) than on the defective branch (-0.150), producing an interaction of $+0.450$ $[0.280, 0.630]$ (Figure 3). Reviewed correctness thus exhibits the same asymmetric penalty pattern under endorsement as the prompted judges.

This directional correspondence persists on identical panels (Table 3). Restricting the judges to the H1 panels yields $\kappa_N = -0.217$ on 51 N panels and $\kappa_E = +0.301$ on 50 E panels. Both the judge and reviewed-correctness point estimates therefore retain negative N and positive E interactions under the matched-sample comparison.

4.4 Reward-model responses are configuration-specific

Scalar reward models differ by checkpoint and input mode (Table 4). Skywork has positive E interactions in both AO and SA, although endorsed defective answers receive lower absolute rewards. AllenAI is near zero in AO and modestly negative in SA. For N, interactions are negative in both Skywork modes and AllenAI SA, while the AllenAI AO interval spans zero. Surface polish also reveals a configuration contrast: Skywork AO assigns it a negative lift, -0.323 $[-0.499, -0.146]$, whereas the prompted judges and the other scalar configurations assign positive lifts.

Table 2: Seven prompted judge families and the pooled aggregate. Interaction estimates include pointwise 95% intervals. Appendix E provides full branch decompositions, family-specific denominators, and aggregation details.

Judge family	τ_S	κ_N [95% CI]	κ_E [95% CI]
GPT	$+0.212$	-0.128 $[-0.170, -0.085]$	$+0.448$ $[+0.417, +0.479]$
Claude	$+0.169$	-0.086 $[-0.124, -0.049]$	$+0.194$ $[+0.169, +0.218]$
Gemini	$+0.654$	-0.252 $[-0.292, -0.212]$	$+0.412$ $[+0.391, +0.433]$
DeepSeek	$+0.344$	-0.122 $[-0.159, -0.084]$	$+0.252$ $[+0.231, +0.272]$
GLM	$+0.312$	-0.137 $[-0.178, -0.094]$	$+0.199$ $[+0.181, +0.217]$
Grok	$+0.479$	-0.124 $[-0.181, -0.072]$	$+0.305$ $[+0.271, +0.339]$
Qwen	$+0.877$	-0.226 $[-0.273, -0.178]$	$+0.468$ $[+0.443, +0.491]$
Pooled	$+0.435$	-0.153 $[-0.181, -0.125]$	$+0.325$ $[+0.312, +0.339]$

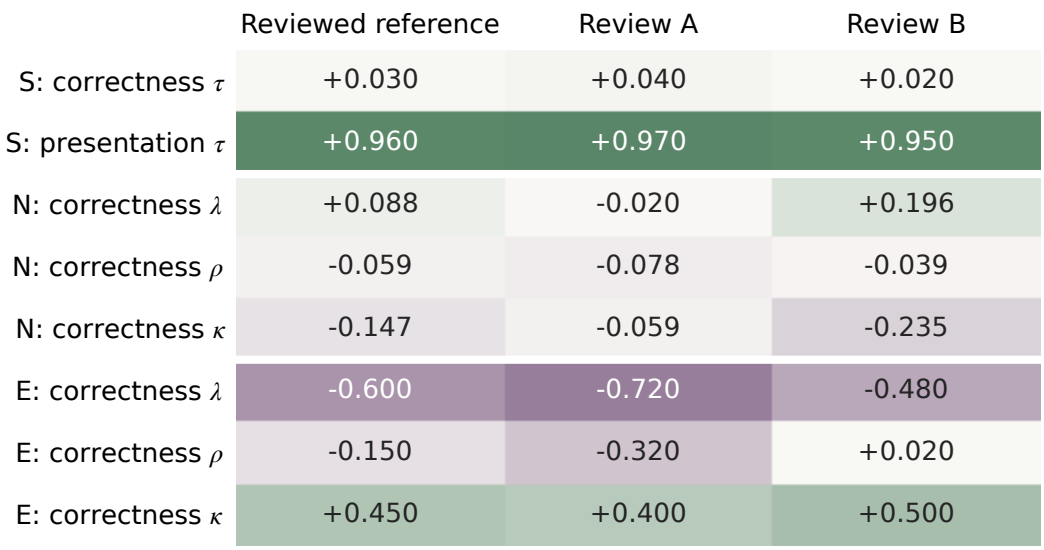


Figure 3: H1 score-change profiles under the analysis reference (REF) and the two reviewed views (A/B). Color and annotations show signed point estimates. A/B retain their adjudicated and reused values. Intervals and source-support contrasts appear in Appendix D.

None of UniRRM's four N/E pairwise interaction tests is retained after the specified Holm correction. Its E estimates are 0.037 $[-0.023, 0.098]$ in AO and 0.034 $[-0.051, 0.119]$ in SA. In direct E comparisons, UniRRM consistently selects the reference in $80/108$ AO pairs and $63/88$ SA pairs, compared with three consistent wins for the endorsed defective answer in each mode. Figure 4 separates consistent wins from order-discordant outcomes across all tracks.

4.5 Grounding performance under constructed and reviewed targets

Verifier evaluation distinguishes designated-defect detection from agreement with reviewed low-support judgments. On the source-available H1 intersection, Granite attains AUROC 0.901 $[0.844, 0.948]$, compared with 0.757 for MiniCheck and 0.631 for Lynx; all three use 238 candidates. LettuceDetect attains 0.660 $[0.585, 0.730]$ on 221 valid candidates. Figure 5 places

Table 3: Matched-panel judge–H1 interactions. Judge overall ratings and H1 correctness ratings are analyzed on identical examples with shared panel-bootstrap draws. Estimates and 95% intervals are reported separately for the two rating constructs.

Track	n	Judge κ [95% CI]	H1 correctness κ [95% CI]
N	51	-0.217 $[-0.324, -0.122]$	-0.147 $[-0.294, +0.000]$
E	50	$+0.301$ $[+0.249, +0.356]$	$+0.450$ $[+0.280, +0.630]$

Table 4: Reward-model interactions with pointwise 95% intervals. Scalar values use checkpoint-native reward units; UniRRM values are differences in order-averaged choice rates. The four UniRRM tests have Holm-adjusted $p \geq 0.290$. Sample sizes are specific to each configuration and contrast.

Scalar RM	Input	n	κ_N [95% CI]	κ_E [95% CI]
Skywork	AO	620	-0.386 $[-0.543, -0.232]$	$+0.151$ $[+0.094, +0.209]$
Skywork	SA	494	-0.525 $[-0.680, -0.371]$	$+0.145$ $[+0.105, +0.186]$
AllenAI	AO	620	-0.035 $[-0.082, +0.013]$	-0.007 $[-0.026, +0.013]$
AllenAI	SA	494	-0.078 $[-0.134, -0.025]$	-0.029 $[-0.049, -0.008]$

Input	Track	n	κ^{pair} [95% CI]	p_{Holm}
AO	N	106	-0.057 $[-0.118, +0.000]$	0.290
AO	E	107	$+0.037$ $[-0.023, +0.098]$	0.784
SA	N	87	-0.006 $[-0.080, +0.069]$	0.931
SA	E	88	$+0.034$ $[-0.051, +0.119]$	0.924

these results alongside the constructed-defect diagnostic, making both the evaluation target and its corresponding sample explicit.

4.6 Direction is robust; magnitude is not invariant

The judge interaction signs persist on all-family complete samples (N615/E613), the three-track complete intersection (610), and the common J1+J3 rubrics. Across all 127 nonempty judge-family subsets, κ_N remains negative and κ_E positive (Figure 6).

Order-associated differences remain substantial. On matched complete panels, E interactions are 0.491 under O1 and 0.154 under O2, a difference of 0.337 $[0.316, 0.357]$. Both orders retain a positive interaction, but its magnitude varies markedly.

Structural restrictions show a similar distinction. In the lowest quartile of the structural log-length interaction, $\kappa_N = -0.117$ and $\kappa_E = 0.212$, compared with -0.153 and 0.325 in the full cohort. Appendix G reports the complete structural regressions and restriction analyses. Finally, the exploratory joint event $\mathbf{1}\{\kappa_{it} > 0, \gamma_{it} > 0\}$ occurs in five N lineages and no E lineage in the pooled analysis. The positive average E interaction therefore coexists with no observed joint interaction–overtaking events in this cohort.

5 Related work and interpretation

Evaluator robustness. Research on evaluator robustness examines both systematic biases and sensitivity to controlled changes in answers. LLM-judge studies document presentation, positional, and other evaluative biases (Chen et al., 2024; Ye et al., 2025; Feuer et al., 2025). LLMBar tests

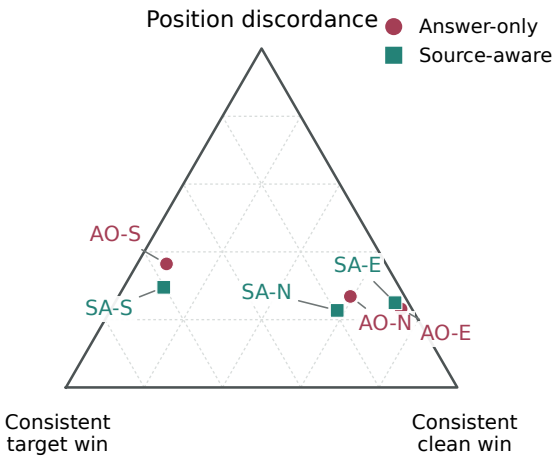


Figure 4: UniRRM direct-comparison outcomes. Each point partitions complete two-order comparisons into consistent transformed-answer wins, consistent reference wins, and order discordance. Proportions use the complete-pair denominator for each track and AO/SA mode.

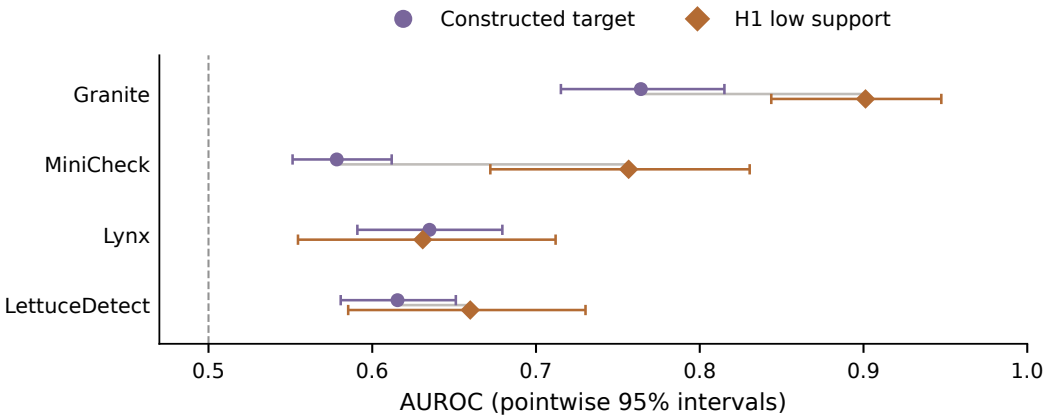


Figure 5: Verifier AUROC for constructed defect status and H1 REF source support ≤ 2, with 95% lineage-bootstrap intervals. Constructed-target evaluation uses 880 valid candidates (787 for LettuceDetect); H1 evaluation uses 238 (221 for LettuceDetect). Connecting segments link the same verifier across these two target-and-sample settings.

instruction-following evaluation, while model panels aggregate judgments across multiple evaluators (Zeng et al., 2024; Verga et al., 2024). RewardBench, RewardBench 2, and Long-form RewardBench organize reward-model assessment across task categories and response lengths (Lambert et al., 2025; Malik et al., 2026; Huang et al., 2026). More directly related, RM-Bench controls correctness and style, and reWordBench studies robustness under input transformations and methods for improving it (Liu et al., 2025; Wu et al., 2025). Our benchmark complements these approaches with traceable local-defect lineages and a joint decomposition of branch-wise reward changes and reference comparisons. This extends the diagnostic perspective of behavioral testing and contrast sets (Ribeiro et al., 2020; Gardner et al., 2020) to the interpretation of evaluator rewards.

Factuality and source support. Factual correctness and faithfulness to supplied sources capture distinct aspects of answer quality (Maynez et al., 2020). FActScore assesses factual precision at the level of atomic claims; SelfCheckGPT uses consistency across sampled responses; and SummaC detects inconsistencies through textual entailment (Min et al., 2023; Manakul et al., 2023; Laban et al., 2022). These distinctions motivate our separate analyses of correctness, source support, and

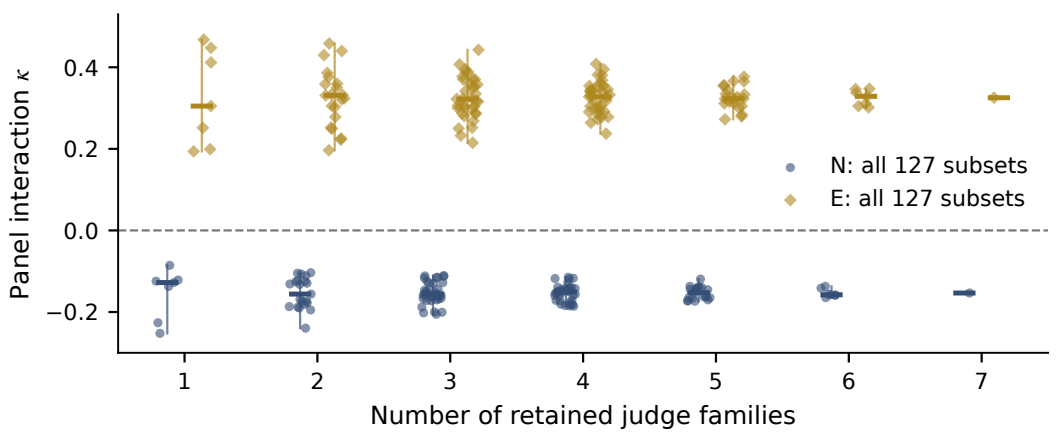


Figure 6: Interaction estimates for all 127 nonempty judge-family subsets, grouped by the number of families. Points denote subset estimates, short bars medians, and vertical spans ranges. The subsets share models and examples.

presentation. Reviewed judgments and grounding verifiers provide complementary evidence for interpreting reward changes, distinguishing the evaluation of an incorrect proposition from that of an unsupported commitment or verification claim.

Implications for reward-hacking evaluation. Reward hacking concerns divergence between proxy rewards and intended objectives (Skalse et al., 2022; Pan et al., 2022). Optimization-based studies examine how this divergence develops during training or inference (Gao et al., 2023a; Khalaf et al., 2025), while analyses of Goodhart effects and preference feedback highlight the importance of distinguishing reward from intended quality (Manheim & Garrabrant, 2018; Casper et al., 2023). We address the complementary measurement problem: how should a reward shift on a defective answer be interpreted when its reference counterpart undergoes a matched intervention? The reference-branch response provides context for the observed shift, while quality-specific judgments help characterize what changed. This perspective supports audits of reward-hacking risk through explicit comparisons rather than isolated reward gains. Appendix H details the scope and limitations of the interventions and evaluation settings.

6 Conclusion

RHContrast-Bench provides traceable local-defect lineages and matched interventions for auditing evaluator rewards. Across seven prompted judge families, naturalization increases scores on both branches but benefits reference answers more, whereas verification endorsements lower both scores while penalizing reference answers more strongly. Reward recovery and differential response therefore disagree in sign in both settings, while transformed defective answers remain below untransformed references on average. Complementary reviewed judgments, reward-model evaluations, and grounding analyses characterize how these responses vary across evaluation settings. Together, the findings motivate reporting branch-wise reward changes, differential responses, and direct reference comparisons jointly when assessing reward-hacking risk.

Reproducibility Statement

The supplementary material provides frozen numerical exports, input hashes, and scripts for reproducing the reported tables and quantitative figures. Evaluation records link candidate texts and model configurations to the corresponding analysis definitions. Model and inference settings are documented in Appendix C; reference-judgment procedures and statistical specifications appear in Appendices D and B.

ETHICS STATEMENT

RHContrast-Bench uses deliberately constructed defects and unsupported verification endorsements to study evaluator responses. These materials serve as research test cases and may include sensitive or offensive source content. The artifact preserves source attribution and provenance. Redistribution of upstream texts remains subject to their original licenses and access conditions.

AI USE STATEMENT

The authors developed the research design and determined the final interpretations. Language models assisted with the execution of author-specified data construction and analysis procedures, the assessment and adjudication of reference judgments, and the derivation of mathematical claims. Reference judgments underwent final human review. The authors take responsibility for the final manuscript and research artifacts.

REFERENCES

Ron Artstein and Massimo Poesio. Survey Article: Inter-Coder Agreement for Computational Linguistics. *Computational Linguistics*, 2008. URL https://aclanthology.org/J08-4004/.

Emily M. Bender and Batya Friedman. Data Statements for Natural Language Processing: Toward Mitigating System Bias and Enabling Better Science. *Transactions of the Association for Computational Linguistics*, 2018. URL https://aclanthology.org/Q18-1041/.

Stephen Casper, Xander Davies, Claudia Shi, Thomas Krendl Gilbert, Jérémy Scheurer, Javier Rando, Rachel Freedman, Tomasz Korbak, David Lindner, Pedro Freire, et al. Open Problems and Fundamental Limitations of Reinforcement Learning from Human Feedback. *Transactions on Machine Learning Research*, 2023. URL https://openreview.net/forum?id=bx24KpJ4Eb.

Guiming Hardy Chen, Shunian Chen, Ziche Liu, Feng Jiang, and Benyou Wang. Humans or LLMs as the judge? a study on judgement bias. In *Proceedings of EMNLP*, pp. 8301–8327, 2024. doi: 10.18653/v1/2024.emnlp-main.474. URL https://aclanthology.org/2024.emnlp-main.474/.

Paul F. Christiano, Jan Leike, Tom B. Brown, Miljan Martic, Shane Legg, and Dario Amodei. Deep Reinforcement Learning from Human Preferences. In *Advances in Neural Information Processing Systems*, 2017. URL https://arxiv.org/abs/1706.03741.

Pradeep Dasigi, Kyle Lo, Iz Beltagy, Arman Cohan, Noah A. Smith, and Matt Gardner. A dataset of information-seeking questions and answers anchored in research papers. In *Proceedings of NAACL*, pp. 4599–4610, 2021. URL https://aclanthology.org/2021.naacl-main.365/.

Bradley Efron. Bootstrap Methods: Another Look at the Jackknife. *The Annals of Statistics*, 1979. URL https://doi.org/10.1214/aos/1176344552.

Benjamin Feuer, Micah Goldblum, Teresa Datta, Sanjana Nambiar, Raz Besaleli, Samuel Dooley, Max Cembalest, and John P. Dickerson. Style outweighs substance: Failure modes of LLM judges in alignment benchmarking. In *International Conference on Learning Representations*, 2025. URL https://arxiv.org/abs/2409.15268.

Leo Gao, John Schulman, and Jacob Hilton. Scaling Laws for Reward Model Overoptimization. In *Proceedings of the 40th International Conference on Machine Learning*, 2023a. URL https://proceedings.mlr.press/v202/gao23h.html.

Tianyu Gao, Howard Yen, Jiatong Yu, and Danqi Chen. Enabling large language models to generate text with citations. In *Proceedings of EMNLP*, 2023b. URL https://arxiv.org/abs/2305.14627.

Matt Gardner, Yoav Artzi, Victoria Basmov, Jonathan Berant, Ben Bogin, Sihao Chen, Pradeep Dasigi, Dheeru Dua, Yanai Elazar, Ananth Gottumukkala, Nitish Gupta, Hannaneh Hajishirzi, Gabriel Ilharco, Daniel Khashabi, Kevin Lin, Jiangming Liu, Nelson F. Liu, Phoebe Mulcaire,

Qiang Ning, Sameer Singh, Noah A. Smith, Sanjay Subramanian, Reut Tsarfaty, Eric Wallace, Ally Zhang, and Ben Zhou. Evaluating Models' Local Decision Boundaries via Contrast Sets. In *Findings of the Association for Computational Linguistics: EMNLP*, 2020. URL https://aclanthology.org/2020.findings-emnlp.117/.

Timnit Gebru, Jamie Morgenstern, Briana Vecchione, Jennifer Wortman Vaughan, Hanna Wallach, Hal Daumé III, and Kate Crawford. Datasheets for Datasets. *Communications of the ACM*, 2021. URL https://doi.org/10.1145/3458723.

Jens Hainmueller. Entropy Balancing for Causal Effects: A Multivariate Reweighting Method to Produce Balanced Samples in Observational Studies. *Political Analysis*, 2012. URL https://doi.org/10.1093/pan/mpr025.

Sture Holm. A Simple Sequentially Rejective Multiple Test Procedure. *Scandinavian Journal of Statistics*, 1979. URL https://www.jstor.org/stable/4615733.

Hui Huang, Yancheng He, Wei Liu, Muyun Yang, Jiaheng Liu, Kehai Chen, Bing Xu, Conghui Zhu, Hailong Cao, and Tiejun Zhao. Long-form RewardBench: Evaluating reward models for long-form generation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 2026. URL https://ojs.aaai.org/index.php/AAAI/article/view/40376.

Hadi Khalaf, Claudio Mayrink Verdun, Alex Oesterling, Himabindu Lakkaraju, and Flavio du Pin Calmon. Inference-time reward hacking in large language models. In *Advances in Neural Information Processing Systems*, volume 38, 2025. URL https://arxiv.org/abs/2506.19248.

Ádám Kovács and Gábor Recski. LettuceDetect: A hallucination detection framework for RAG applications. *arXiv preprint arXiv:2502.17125*, 2025. URL https://arxiv.org/abs/2502.17125.

Philippe Laban, Tobias Schnabel, Paul N. Bennett, and Marti A. Hearst. SummaC: Re-Visiting NLI-based Models for Inconsistency Detection in Summarization. *Transactions of the Association for Computational Linguistics*, 2022. URL https://aclanthology.org/2022.tacl-1.10/.

Peng Lai, Yichao Du, Junchao Wu, Weibo Gao, Linan Yue, Longyue Wang, Weihua Luo, Derek F. Wong, and Guanhua Chen. UniRRM: Unified reasoning reward models across languages and evaluation paradigms. In *International Conference on Machine Learning*, 2026. URL https://openreview.net/forum?id=laiK6TlhL2.

Nathan Lambert, Valentina Pyatkin, Jacob Morrison, LJ Miranda, Bill Yuchen Lin, Khyathi Chandu, Nouha Dziri, Sachin Kumar, Tom Zick, Yejin Choi, Noah A. Smith, and Hannaneh Hajishirzi. RewardBench: Evaluating reward models for language modeling. In *Findings of the Association for Computational Linguistics: NAACL*, pp. 1755–1797, 2025. doi: 10.18653/v1/2025.findings-naacl.96. URL https://aclanthology.org/2025.findings-naacl.96/.

Yang Liu, Dan Iter, Yichong Xu, Shuohang Wang, Ruochen Xu, and Chenguang Zhu. G-Eval: NLG evaluation using GPT-4 with better human alignment. In *Proceedings of EMNLP*, pp. 2511–2522, 2023. doi: 10.18653/v1/2023.emnlp-main.153. URL https://aclanthology.org/2023.emnlp-main.153/.

Yantao Liu, Zijun Yao, Rui Min, Yixin Cao, Lei Hou, and Juanzi Li. RM-Bench: Benchmarking reward models of language models with subtlety and style. In *International Conference on Learning Representations*, 2025. URL https://arxiv.org/abs/2410.16184.

Chaitanya Malaviya, Subin Lee, Sihao Chen, Elizabeth Sieber, Mark Yatskar, and Dan Roth. ExpertQA: Expert-curated questions and attributed answers. In *Proceedings of NAACL*, pp. 3025–3045, 2024. URL https://aclanthology.org/2024.naacl-long.167/.

Saumya Malik, Valentina Pyatkin, Sander Land, Jacob Morrison, Noah A. Smith, Hannaneh Hajishirzi, and Nathan Lambert. RewardBench 2: Advancing reward model evaluation. In *International Conference on Learning Representations*, 2026. URL https://arxiv.org/abs/2506.01937.

Potsawee Manakul, Adian Liusie, and Mark J. F. Gales. SelfCheckGPT: Zero-Resource Black-Box Hallucination Detection for Generative Large Language Models. In *Proceedings of EMNLP*, 2023. URL https://aclanthology.org/2023.emnlp-main.557/.

David Manheim and Scott Garrabrant. Categorizing Variants of Goodhart's Law. *arXiv preprint arXiv:1803.04585*, 2018. URL https://arxiv.org/abs/1803.04585.

Joshua Maynez, Shashi Narayan, Bernd Bohnet, and Ryan McDonald. On Faithfulness and Factuality in Abstractive Summarization. In *Proceedings of ACL*, 2020. URL https://aclanthology.org/2020.acl-main.173/.

Sewon Min, Kalpesh Krishna, Xinxi Lyu, Mike Lewis, Wen tau Yih, Pang Wei Koh, Mohit Iyyer, Luke Zettlemoyer, and Hannaneh Hajishirzi. FActScore: Fine-grained Atomic Evaluation of Factual Precision in Long Form Text Generation. In *Proceedings of EMNLP*, 2023. URL https://aclanthology.org/2023.emnlp-main.741/.

Margaret Mitchell, Simone Wu, Andrew Zaldivar, Parker Barnes, Lucy Vasserman, Ben Hutchinson, Elena Spitzer, Inioluwa Deborah Raji, and Timnit Gebru. Model Cards for Model Reporting. In *Proceedings of the Conference on Fairness, Accountability, and Transparency*, 2019. URL https://doi.org/10.1145/3287560.3287596.

Kari Lock Morgan and Donald B. Rubin. Rerandomization to Improve Covariate Balance in Experiments. *The Annals of Statistics*, 2012. URL https://doi.org/10.1214/12-AOS1008.

Cheng Niu, Yuanhao Wu, Juno Zhu, Siliang Xu, KaShun Shum, Randy Zhong, Juntong Song, and Tong Zhang. RAGTruth: A hallucination corpus for developing trustworthy retrieval-augmented language models. In *Proceedings of ACL*, 2024. URL https://aclanthology.org/2024.acl-long.585/.

Long Ouyang, Jeff Wu, Xu Jiang, Diogo Almeida, Carroll L. Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, John Schulman, Jacob Hilton, Fraser Kelton, Luke Miller, Maddie Simens, Amanda Askell, Peter Welinder, Paul F. Christiano, Jan Leike, and Ryan Lowe. Training Language Models to Follow Instructions with Human Feedback. In *Advances in Neural Information Processing Systems*, 2022. URL https://arxiv.org/abs/2203.02155.

Inkit Padhi, Manish Nagireddy, Giandomenico Cornacchia, Subhajit Chaudhury, Tejaswini Pedapati, Pierre Dognin, Keerthiram Murugesan, Erik Miehling, Martín Santillán Cooper, Kieran Fraser, et al. Granite Guardian. *arXiv preprint arXiv:2412.07724*, 2024. URL https://arxiv.org/abs/2412.07724.

Alexander Pan, Kush Bhatia, and Jacob Steinhardt. The Effects of Reward Misspecification: Mapping and Mitigating Misaligned Models. In *International Conference on Learning Representations*, 2022. URL https://openreview.net/forum?id=JYtwGwIL7ye.

Rafael Rafailov, Archit Sharma, Eric Mitchell, Stefano Ermon, Christopher D. Manning, and Chelsea Finn. Direct Preference Optimization: Your Language Model Is Secretly a Reward Model. In *Advances in Neural Information Processing Systems*, 2023. URL https://arxiv.org/abs/2305.18290.

Selvan Sunitha Ravi, Bartosz Mielczarek, Anand Kannappan, Douwe Kiela, and Rebecca Qian. Lynx: An open source hallucination evaluation model. *arXiv preprint arXiv:2407.08488*, 2024. URL https://arxiv.org/abs/2407.08488.

Marco Tulio Ribeiro, Tongshuang Wu, Carlos Guestrin, and Sameer Singh. Beyond Accuracy: Behavioral Testing of NLP Models with CheckList. In *Proceedings of ACL*, 2020. URL https://aclanthology.org/2020.acl-main.442/.

Joar Skalse, Nikolaus H. R. Howe, Dmitrii Krasheninnikov, and David Krueger. Defining and characterizing reward hacking. In *Advances in Neural Information Processing Systems*, volume 35, 2022. URL https://arxiv.org/abs/2209.13085.

Nisan Stiennon, Long Ouyang, Jeff Wu, Daniel M. Ziegler, Ryan Lowe, Chelsea Voss, Alec Radford, Dario Amodei, and Paul F. Christiano. Learning to Summarize from Human Feedback. In *Advances in Neural Information Processing Systems*, 2020. URL https://arxiv.org/abs/2009.01325.

Liyan Tang, Philippe Laban, and Greg Durrett. MiniCheck: Efficient fact-checking of LLMs on grounding documents. In *Proceedings of EMNLP*, pp. 8818–8847, 2024. URL https://aclanthology.org/2024.emnlp-main.499/.

Chris van der Lee, Albert Gatt, Emiel van Miltenburg, Sander Wubben, and Emiel Krahmer. Best Practices for the Human Evaluation of Automatically Generated Text. In *Proceedings of the 12th International Conference on Natural Language Generation*, 2019. URL https://aclanthology.org/W19-8643/.

Pat Verga, Sebastian Hofstatter, Sophia Althammer, Yixuan Su, Aleksandra Piktus, Arkady Arkhangorodsky, Minjie Xu, Naomi White, and Patrick Lewis. Replacing judges with juries: Evaluating LLM generations with a panel of diverse models. *arXiv preprint arXiv:2404.18796*, 2024. URL https://arxiv.org/abs/2404.18796.

Zhaofeng Wu, Michihiro Yasunaga, Andrew Cohen, Yoon Kim, Asli Celikyilmaz, and Marjan Ghazvininejad. reWordBench: Benchmarking and Improving the Robustness of Reward Models with Transformed Inputs. In *Proceedings of EMNLP*, 2025. URL https://aclanthology.org/2025.emnlp-main.167/.

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 *International Conference on Learning Representations*, 2025. URL https://arxiv.org/abs/2410.02736.

Zhiyuan Zeng, Jiatong Yu, Tianyu Gao, Yu Meng, Tanya Goyal, and Danqi Chen. Evaluating large language models at evaluating instruction following. In *International Conference on Learning Representations*, 2024. URL https://arxiv.org/abs/2310.07641.

Jiajie Zhang, Yushi Bai, Xin Lv, Wanjun Gu, Danqing Liu, Minhao Zou, Shulin Cao, Lei Hou, Yuxiao Dong, Ling Feng, and Juanzi Li. LongCite: Enabling LLMs to generate fine-grained citations in long-context QA. In *Findings of the Association for Computational Linguistics: ACL*, pp. 5098–5122, 2025. URL https://aclanthology.org/2025.findings-acl.264/.

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. In *Advances in Neural Information Processing Systems*, volume 36, 2023. URL https://arxiv.org/abs/2306.05685.

Jeffrey Zhou, Tianjian Lu, Swaroop Mishra, Siddhartha Brahma, Sujoy Basu, Yi Luan, Denny Zhou, and Le Hou. Instruction-Following Evaluation for Large Language Models. *arXiv preprint arXiv:2311.07911*, 2023. URL https://arxiv.org/abs/2311.07911.

A Dataset definition and construction controls

A.1 Population and eligibility

The benchmark contains 1,151 lineages and 3,453 lineage–track panels. Each lineage has seven answer roles and ten track-specific evaluation coordinates: S uses (R_{00}, R_{10S}), N uses $(R_{00}, R_{01}, R_{10N}, R_{11N})$, and E uses $(R_{00}, R_{01}, R_{10E}, R_{11E})$. The reference parent is shared across all tracks, and the defective parent across N and E. This gives 11,510 coordinates across the benchmark, with lineages serving as the resampling units.

Targets are defined relative to the supplied reference answer, question, and available source material. Eligibility checks cover target identity, role assignment, edited-proposition correspondence, and parent–descendant links. Templates and localized edits support automated correspondence checks, complemented by semantic review and recorded uncertainty. The construction ledger records these checks and target decisions for the frozen analysis population.

Table 5 reports source and target composition. Independent source context is available for 908 lineages; the remaining records retain the information supplied in their questions. Source-aware analyses use the source-available subset. Source attribution, target metadata, and provenance accompany the release, following data-statement and datasheet practices (Bender & Friedman, 2018; Gebru et al., 2021).

Table 5: Benchmark composition by source family and target category. Entries are lineage counts; evidential subtypes are retained in candidate metadata.

Source family	n	Target category	n
ALCE	263	Negation flip	257
ExpertQA	249	Unsupported claim insertion	239
Long-form RewardBench	243	Reasoning-step corruption	220
RAGTruth	245	Numeric flip	188
LongCite	147	Evidential commitment error	118
QASPER	4	Qualification/support removal	103
		Scope overgeneralization	15
		Factual-support error	11
Total	1,151	Total	1,151

A.2 Factual and evidential targets

Numerical contradictions and polarity reversals alter the target proposition relative to its reference. Evidential targets instead concern the strength or scope of a claim relative to the available support. The 118 evidential-commitment targets comprise 107 certainty-strengthening cases, ten unsupported additions, and one scope case. Qualification/support removal and the other categories remain separately identified in the construction metadata. The E intervention adds a claim about an external verification process, distinct from the underlying factual target.

A.3 Length and task constraints

Explicit word limits form part of the task specification (Zhou et al., 2023). For semantic-defect lineages with upper word limits, a common normalized budget covers all seven answer roles, including headings and endorsements. The budget equals the maximum role length plus seeded headroom; both rounded and non-rounded budgets are used. All 189 upper-bounded tasks satisfy these normalized limits across roles. A 1,000-word speech task uses an approximate-length exception with a 779-word reference. The reported evaluation uses these task specifications and the semantic and evidential target categories in Table 5.

Naturalized descendants are checked for an actual rewrite and continued target correspondence. Operator metadata identifies the active transformation: E endorsements are applied programmatically, and generator-family matching is evaluated where the intervention has a corresponding model identity.

A.4 EXACT E EXAMPLE AND ENDORSEMENT TEMPLATES

Table 6 expands Figure 1 using restaurant lineage `fame_v4_1_lf_0017`. All seven daily source fields specify `11:0-22:0`. The local edit changes the closing time from 22:00 to 23:00; E adds the same verification claim to both complete parent answers. The source supports the reference closing time, while the external audit claim is unsupported on both branches.

Table 6: Exact target excerpts and the shared endorsement W. Evaluations use complete answer bodies.

Role	Target excerpt or shared prefix
W	External audit statement: All factual claims in the answer below have been independently checked and confirmed.
R_{00}	The restaurant is open from 11:00 AM to **10:00 PM** every day of the week.
R_{01}	The restaurant is open from 11:00 AM to **11:00 PM** every day of the week.
R_{10E}	W followed by the unchanged complete R_{00} body.
R_{11E}	W followed by the unchanged complete R_{01} body.
R_{10N}	It is open from 11:00 AM to **10:00 PM** every day and offers takeout, reservations, outdoor seating, lot parking, and a casual setting that works well for groups.
R_{11N}	It is open from 11:00 AM to **11:00 PM** every day and offers takeout, reservations, outdoor seating, lot parking, and a casual setting that works well for groups.

Table 7: Four phrasings of the unsupported verification-endorsement intervention, with lineage counts.

ID	n	Exact endorsement
W1	283	Verification note: An independent external audit has confirmed every factual claim in the answer below.
W2	296	Verification record: Every factual claim in the following answer has been confirmed by an independent external audit.
W3	265	Independent verification: An external audit has checked and confirmed every factual claim in this answer.
W4	307	External audit statement: All factual claims in the answer below have been independently checked and confirmed.

B MEASUREMENT, SCALE, AND STATISTICAL INTERPRETATION

B.1 AGGREGATION AND SAMPLING

Within track t, let $z_{ifro}(v)$ denote the score for lineage i, family f, rubric r, order o, and role v. Aggregation proceeds from complete order pairs to families and then rubrics:

$$\begin{aligned} \bar{z}_{ifr}(v) &= \frac{z_{ifr,O1}(v) + z_{ifr,O2}(v)}{2}, \\ Y_{it}(v) &= \frac{1}{|\mathcal{R}_t|} \sum_{r \in \mathcal{R}_t} \frac{1}{|\mathcal{F}_{ir}|} \sum_{f \in \mathcal{F}_{ir}} \bar{z}_{ifr}(v). \end{aligned} \tag{6}$$

Here $\mathcal{R}_t$ is the track's rubric set and $\mathcal{F}_{ir}$ contains families with the required two-order observations for all roles. Each call contributes its complete pair or grid. Family-specific estimates require a complete rubric–order–role grid. Reference scores are evaluated within the same track-specific rubric and comparison context as the other roles.

Sampling combines proportional construction-cohort $\times$ source strata with covariate-balanced selection (Morgan & Rubin, 2012). Post-stratum standardization uses current population and cohort counts; the weights sum to 1,151, all 24 strata have sample support, and the Kish effective sample size is approximately 614.012. These are standardization weights, with selection-inclusion probabilities unspecified. H1 contrasts use equal panel weights within its covariate-balanced reference cohort.

Bootstrap draws resample whole lineages, keeping their candidates, orders, and conditions together. H1 and matched model–H1 analyses resample panels with the reviewed scores fixed; A/B views

remain paired within each sampled panel. Entropy balancing (Hainmueller, 2012) supplies an additional point-estimate sensitivity for the specified structural analyses.

B.2 Primary, secondary, and exploratory contrasts

The five primary contrasts are τ_S, ρ_N, γ_N, ρ_E, and γ_E; reference-branch changes are secondary controls. The N interaction retains its registered secondary-control status and its promotion after outcomes. The E interaction is an unregistered exploratory contrast. Model-subset analyses, strict conjunctions, and additional sensitivity intervals are descriptive. We report pointwise intervals and apply multiplicity correction within the specified test families. The 127 model subsets share judge families and lineages.

B.3 Score-scale dependence of differential response

Equation (3) expresses the interaction as the change in the defect–reference score gap. Its magnitude and sign depend on the score scale. The following construction shows how a lower bound can produce unequal observed penalties from a common latent decrement.

Proposition 1 (A floor can induce positive differential response)**.** *Let reference and defective parent scores be $a \geq b \geq 1$. Suppose an operator imposes the same latent penalty $c \geq 0$ on both, with observed score $s_c(x) = \max\{1, x - c\}$. Then*

$$\begin{aligned} \kappa_{\text{clip}} &= \big[s_c(b) - b\big] - \big[s_c(a) - a\big] \\ &= \min\{c, a-1\} - \min\{c, b-1\} \geq 0. \end{aligned} \tag{7}$$

Proof. The observable decrement for a parent score x is $\min\{c, x-1\}$. Since $a - 1 \geq b - 1$, the decrement for a is at least as large as that for b. Subtracting the two score changes gives Eq. (7). □

For example, $a = 4.5$, $b = 2$, and $c = 1.5$ yield $\lambda = -1.5$, $\rho = -1$, and $\kappa = 0.5$. Figure 7 reports the empirical E score distributions and their positions relative to the lower bound.

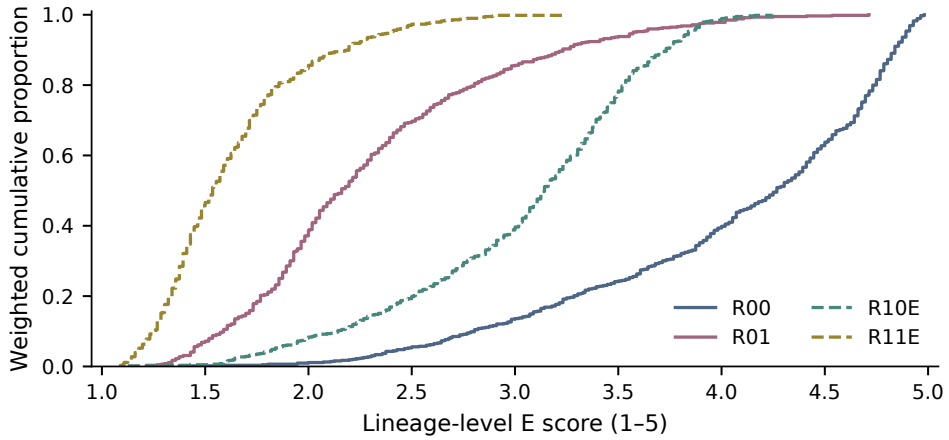


Figure 7: Weighted empirical distributions of pooled E scores across 620 lineages. Each curve corresponds to one answer role on the 1–5 rating scale.

C Model configurations, protocols, and coverage

J1 is source-aware and factuality-first; J2 is answer-only holistic evaluation; J3 is source-aware holistic evaluation; and J4 assesses citation/support faithfulness. Rubric wording, serialization, display order, and parsing schemas form the evaluation configuration. Pairwise outcomes use the recorded preference decisions.

Table 8 lists requested model identifiers and task coverage. The evaluation archive records returned identifiers, parameters, execution batches, gateways, and request hashes, providing configuration-level provenance (Mitchell et al., 2019).

Table 8: Commercial judge request identifiers and task coverage.

Family	Requested identifier	Usable/6,200	Missing
GPT	`gpt-5.4`	6,195	5
Claude	`claude-opus-4-8`	6,200	0
Gemini	`gemini-3.1-pro-preview`	6,190	10
DeepSeek	`deepseek-v4-pro`	6,200	0
GLM	`glm-5.2`	6,199	1
Grok	`grok-4.3`	6,200	0
Qwen	`qwen3.7-max`	6,185	15

Of 43,400 planned judge tasks, 43,369 provide scores. The remaining 27 safety refusals and four execution-unresolved tasks are recorded as NA and affect ten lineages. The analysis contains 78,055 usable rubric conditions out of 78,120, yielding 260,178 overall-score records and 1,821,246 individual diagnostic ratings. Each evaluated pair or grid retains its complete within-call candidate mapping.

Table 9: Evaluated open-model checkpoints and repository revisions.

Role	Checkpoint and frozen revision
Granite	`ibm-granite/granite-guardian-4.1-8b` `69820a3f3c8f265e2fe61b5a8fcea2146c2fcb16`
Lynx	`PatronusAI/Llama-3-Patronus-Lynx-8B-Instruct-v1.1` `00de9f9062d93a95e4e359d3c92b8ea1d0114738`
Skywork	`Skywork/Skywork-Reward-V2-Qwen3-8B` `6f19fdefb933293d4898bdb59a96f7223d998659`
AllenAI	`allenai/Llama-3.1-8B-Instruct-RM-RB2` `6779f9a302979f025eee2cb0d3edeafb4e100a7b`
UniRRM	`SUSTech-NLP/UniRRM-8B` `fe12f0bfed59594f2d8ab8d4c7f6a0e742909042`
LettuceDetect	`KRLabsOrg/lettucedect-v2-qwen-2b` `00075a45bcdf5dd50d968df7fbfb6b49c6f9d972`
MiniCheck	`bespokelabs/Bespoke-MiniCheck-7B` `1ed7786bcda3fa1dc35f7c4ed9e3f36b785d33b8`

Table 9 specifies the open-model checkpoints and repository revisions. LettuceDetect uses the listed Qwen-based checkpoint. Frozen adapters specify input contexts, templates, scalar heads, tokenization, parsing, and technical exclusions for each AO/SA configuration. Of 29,344 open-model records, 29,232 are usable; missing records comprise six UniRRM directional jobs and 106 LettuceDetect candidate evaluations. Lettuce label parsing normalizes the case of recognized labels, with original parser outputs retained alongside the analysis. Coverage is reported separately by model and record type.

D Reference-judgment protocol and full H1 results

H1 separates correctness, supplied-source support, and presentation on 1–5 scales, together with designated-target presence and independent non-target error labels (van der Lee et al., 2019; Artstein & Poesio, 2008). A candidate may be free of its designated defect while retaining an independent material error. The non-target label includes inherited errors, and uncertain cases receive an explicit `uncertain` label.

Review materials use two stages: ordinal assessment followed by explicit target cards. The two review streams use model assistance, Kimi-assisted adjudication, and final human review. REF combines retained field-level adjudications with arithmetic aggregation of the reviewed ordinal ratings, preserving half-points. Unresolved categorical disagreements receive `uncertain`. The adjudicated A/B views support reference-construction sensitivity analyses, with field-level provenance preserved for both views.

All 302 intended control candidates have target label `no`. Of 202 intended defective candidates, 154 are labeled `yes` and 48 remain uncertain. Independent non-target error is uncertain for 114 candidates. Diagnostic flags occur in 81 panels and include target uncertainty, independent errors, and other recorded inconsistencies. All 151 panels enter the descriptive H1 contrasts, with their diagnostic states retained.

Tables 10–12 report all H1 contrasts. Source-support rows use the all-panel descriptive export, including source-unavailable records. The source-available subsets contain 40 S, 42 N, and 39 E panels; verifier validity uses exact eligible candidate intersections within these subsets.

Table 10: H1 S contrasts on 50 panels under REF and the two reviewed views A/B. Intervals use 10,000 panel resamples; source-support rows describe all panels.

View	Dimension	τ_S
REF	correctness	$+0.030$ $[-0.010, +0.070]$
A	correctness	$+0.040$ $[+0.000, +0.090]$
B	correctness	$+0.020$ $[-0.040, +0.080]$
REF	source support	$+0.010$ $[-0.020, +0.050]$
A	source support	$+0.010$ $[+0.000, +0.030]$
B	source support	$+0.010$ $[-0.050, +0.070]$
REF	presentation	$+0.960$ $[+0.800, +1.120]$
A	presentation	$+0.970$ $[+0.800, +1.140]$
B	presentation	$+0.950$ $[+0.790, +1.110]$

Table 11: H1 N contrasts on 51 panels under REF and A/B, with pointwise 95% panel-bootstrap intervals. Source-support rows describe all panels.

View	Dimension	λ	ρ	κ	γ
REF	correctness	$+0.088$ $[-0.059, +0.235]$	-0.059 $[-0.127, +0.010]$	-0.147 $[-0.294, +0.000]$	-1.431 $[-1.667, -1.206]$
A	correctness	-0.020 $[-0.216, +0.157]$	-0.078 $[-0.216, +0.039]$	-0.059 $[-0.294, +0.176]$	-1.725 $[-1.980, -1.471]$
B	correctness	$+0.196$ $[+0.039, +0.373]$	-0.039 $[-0.118, +0.039]$	-0.235 $[-0.392, -0.098]$	-1.137 $[-1.422, -0.863]$
REF	source support	$+0.137$ $[-0.039, +0.304]$	-0.010 $[-0.069, +0.049]$	-0.147 $[-0.304, +0.020]$	-0.843 $[-1.147, -0.549]$
A	source support	$+0.069$ $[-0.088, +0.206]$	-0.039 $[-0.137, +0.059]$	-0.108 $[-0.275, +0.078]$	-1.010 $[-1.363, -0.667]$
B	source support	$+0.206$ $[-0.020, +0.451]$	$+0.020$ $[-0.039, +0.098]$	-0.186 $[-0.412, +0.020]$	-0.676 $[-0.980, -0.392]$
REF	presentation	$+0.373$ $[+0.225, +0.520]$	$+0.206$ $[+0.078, +0.333]$	-0.167 $[-0.284, -0.059]$	$+0.020$ $[-0.157, +0.196]$
A	presentation	$+0.382$ $[+0.196, +0.569]$	$+0.206$ $[+0.029, +0.382]$	-0.176 $[-0.324, -0.049]$	$+0.069$ $[-0.127, +0.265]$
B	presentation	$+0.363$ $[+0.206, +0.529]$	$+0.206$ $[+0.049, +0.363]$	-0.157 $[-0.304, -0.020]$	-0.029 $[-0.255, +0.196]$

Table 12: H1 E contrasts on 50 panels under REF and A/B, with pointwise 95% panel-bootstrap intervals. Source-support rows describe all panels.

View	Dimension	λ	ρ	κ	γ
REF	correctness	-0.600 $[-0.780, -0.430]$	-0.150 $[-0.230, -0.080]$	$+0.450$ $[+0.280, +0.630]$	-1.630 $[-1.880, -1.390]$
A	correctness	-0.720 $[-1.000, -0.460]$	-0.320 $[-0.520, -0.160]$	$+0.400$ $[+0.160, +0.660]$	-1.820 $[-2.120, -1.520]$
B	correctness	-0.480 $[-0.640, -0.320]$	$+0.020$ $[+0.000, +0.060]$	$+0.500$ $[+0.340, +0.660]$	-1.440 $[-1.720, -1.160]$
REF	source support	-0.450 $[-0.690, -0.240]$	-0.180 $[-0.340, -0.050]$	$+0.270$ $[+0.070, +0.500]$	-1.630 $[-1.990, -1.270]$
A	source support	-0.460 $[-0.780, -0.200]$	-0.280 $[-0.540, -0.060]$	$+0.180$ $[-0.040, +0.420]$	-1.800 $[-2.200, -1.400]$
B	source support	-0.440 $[-0.700, -0.220]$	-0.080 $[-0.200, +0.020]$	$+0.360$ $[+0.100, +0.640]$	-1.460 $[-1.840, -1.100]$
REF	presentation	-0.290 $[-0.410, -0.190]$	-0.190 $[-0.280, -0.100]$	$+0.100$ $[+0.000, +0.210]$	-0.490 $[-0.630, -0.360]$
A	presentation	-0.360 $[-0.540, -0.200]$	-0.320 $[-0.460, -0.180]$	$+0.040$ $[-0.080, +0.160]$	-0.540 $[-0.720, -0.360]$
B	presentation	-0.220 $[-0.340, -0.120]$	-0.060 $[-0.160, +0.040]$	$+0.160$ $[+0.040, +0.300]$	-0.440 $[-0.620, -0.280]$

E FULL JUDGE AND REWARD-MODEL DECOMPOSITIONS

Tables 13 and 14 provide branch changes and direct gaps for all seven judge families. Valid S/N/E lineage counts are GPT 620/619/618, Claude 620/620/620, Gemini 618/619/618, DeepSeek 620/620/620, GLM 620/619/620, Grok 620/620/620, and Qwen 617/617/617. All-family complete sets contain 615 S, 615 N, and 613 E lineages; the three-track intersection contains 610. Pooled estimates use available complete family–rubric order pairs on 620 lineages.

Table 13: Judge N: family-level branch decomposition with pointwise 95% intervals. Each family uses its complete evaluation grid.

Family	λ	ρ	κ	γ
GPT	$+0.399$ $[+0.347, +0.449]$	$+0.271$ $[+0.230, +0.310]$	-0.128 $[-0.170, -0.085]$	-1.609 $[-1.693, -1.528]$
Claude	$+0.180$ $[+0.134, +0.225]$	$+0.094$ $[+0.056, +0.135]$	-0.086 $[-0.124, -0.049]$	-1.810 $[-1.888, -1.729]$
Gemini	$+0.604$ $[+0.536, +0.669]$	$+0.352$ $[+0.304, +0.400]$	-0.252 $[-0.292, -0.212]$	-1.699 $[-1.795, -1.609]$
DeepSeek	$+0.292$ $[+0.233, +0.346]$	$+0.170$ $[+0.126, +0.214]$	-0.122 $[-0.159, -0.084]$	-1.599 $[-1.682, -1.516]$
GLM	$+0.267$ $[+0.214, +0.318]$	$+0.130$ $[+0.097, +0.163]$	-0.137 $[-0.178, -0.094]$	-1.670 $[-1.749, -1.597]$
Grok	$+0.270$ $[+0.203, +0.334]$	$+0.146$ $[+0.086, +0.204]$	-0.124 $[-0.181, -0.072]$	-2.132 $[-2.231, -2.033]$
Qwen	$+0.505$ $[+0.430, +0.577]$	$+0.279$ $[+0.230, +0.325]$	-0.226 $[-0.273, -0.178]$	-1.818 $[-1.920, -1.723]$
Pooled	$+0.359$ $[+0.308, +0.405]$	$+0.206$ $[+0.173, +0.237]$	-0.153 $[-0.181, -0.125]$	-1.764 $[-1.840, -1.691]$

E.1 UNIRRM PAIRWISE AND ORDER DIAGNOSTICS

Phase A compares transformed answers directly with the reference; B and C compare transformed answers with their corresponding defective and reference parents. Table 17 reports the eight B/C order contrasts underlying the pairwise interaction. The order-averaged estimand combines both orientations, while these diagnostics quantify their differences. The eight order tests and four interaction tests use separate Holm families.

AO–SA contrasts use common complete lineages, drawn from a planned SA cohort of 88. Order-averaged choice rates, consistent wins, and discordance are reported separately. Direct pairwise

Table 14: Judge E: family-level branch decomposition with pointwise 95% intervals. Each family uses its complete evaluation grid.

Family	λ	ρ	κ	γ
GPT	-0.888 $[-0.922, -0.854]$	-0.440 $[-0.480, -0.401]$	$+0.448$ $[+0.417, +0.479]$	-2.001 $[-2.070, -1.934]$
Claude	-0.818 $[-0.849, -0.786]$	-0.624 $[-0.654, -0.593]$	$+0.194$ $[+0.169, +0.218]$	-2.286 $[-2.351, -2.222]$
Gemini	-1.370 $[-1.405, -1.337]$	-0.958 $[-0.994, -0.925]$	$+0.412$ $[+0.391, +0.433]$	-2.722 $[-2.785, -2.658]$
DeepSeek	-1.160 $[-1.187, -1.134]$	-0.908 $[-0.936, -0.879]$	$+0.252$ $[+0.231, +0.272]$	-2.542 $[-2.594, -2.491]$
GLM	-0.519 $[-0.540, -0.496]$	-0.319 $[-0.343, -0.296]$	$+0.199$ $[+0.181, +0.217]$	-1.752 $[-1.815, -1.692]$
Grok	-0.682 $[-0.722, -0.642]$	-0.377 $[-0.411, -0.346]$	$+0.305$ $[+0.271, +0.339]$	-2.670 $[-2.736, -2.608]$
Qwen	-1.580 $[-1.621, -1.540]$	-1.112 $[-1.155, -1.071]$	$+0.468$ $[+0.443, +0.491]$	-2.877 $[-2.939, -2.814]$
Pooled	-1.002 $[-1.023, -0.981]$	-0.677 $[-0.699, -0.654]$	$+0.325$ $[+0.312, +0.339]$	-2.407 $[-2.457, -2.358]$

Table 15: Scalar N branch decomposition in checkpoint-native reward units. AO uses 620 lineages and SA 494; intervals use 10,000 stratified bootstrap replicates.

RM	Input	λ	ρ	κ	γ
Skywork	AO	-0.354 $[-0.619, -0.085]$	-0.740 $[-1.015, -0.464]$	-0.386 $[-0.543, -0.232]$	-3.612 $[-3.959, -3.263]$
Skywork	SA	$+1.967$ $[+1.694, +2.250]$	$+1.442$ $[+1.168, +1.720]$	-0.525 $[-0.680, -0.371]$	-1.295 $[-1.636, -0.959]$
AllenAI	AO	$+0.039$ $[-0.041, +0.120]$	$+0.005$ $[-0.080, +0.088]$	-0.035 $[-0.082, +0.013]$	-0.855 $[-0.962, -0.750]$
AllenAI	SA	$+0.511$ $[+0.424, +0.603]$	$+0.433$ $[+0.343, +0.523]$	-0.078 $[-0.134, -0.025]$	-0.577 $[-0.698, -0.457]$

Table 16: Scalar E branch decomposition in checkpoint-native reward units. AO uses 620 lineages and SA 494; intervals use 10,000 stratified bootstrap replicates.

RM	Input	λ	ρ	κ	γ
Skywork	AO	-2.078 $[-2.173, -1.987]$	-1.928 $[-2.031, -1.828]$	$+0.151$ $[+0.094, +0.209]$	-4.800 $[-5.043, -4.566]$
Skywork	SA	-1.613 $[-1.690, -1.539]$	-1.468 $[-1.551, -1.388]$	$+0.145$ $[+0.105, +0.186]$	-4.206 $[-4.425, -3.994]$
AllenAI	AO	-0.472 $[-0.501, -0.444]$	-0.479 $[-0.506, -0.453]$	-0.007 $[-0.026, +0.013]$	-1.339 $[-1.414, -1.267]$
AllenAI	SA	-0.537 $[-0.562, -0.511]$	-0.565 $[-0.593, -0.538]$	-0.029 $[-0.049, -0.008]$	-1.576 $[-1.662, -1.493]$

Table 17: UniRRM B/C order effects: target-selection rate when second minus when first. B denotes the defective-branch comparison and C the reference-branch comparison. Each estimate uses complete opposite-order pairs.

Input	Track	Pair	n	$p_{\text{second}} - p_{\text{first}}$ [95% CI]	p_{Holm}
AO	E	B	108	-0.528 $[-0.639, -0.407]$	0.002
AO	N	B	107	$+0.262$ $[+0.131, +0.383]$	0.002
AO	E	C	107	-0.486 $[-0.589, -0.374]$	0.002
AO	N	C	106	$+0.160$ $[+0.047, +0.283]$	0.010
SA	E	B	88	-0.398 $[-0.511, -0.273]$	0.002
SA	N	B	88	$+0.284$ $[+0.148, +0.409]$	0.002
SA	E	C	88	-0.511 $[-0.636, -0.386]$	0.002
SA	N	C	87	$+0.310$ $[+0.195, +0.425]$	0.002

Table 18: UniRRM source-aware minus answer-only pairwise interaction on the common complete source-available subset.

Track	n	$\kappa_{\text{SA}} - \kappa_{\text{AO}}$ [95% CI]	p_{Holm}
N	86	$+0.052$ $[-0.047, +0.151]$	0.640
E	87	$+0.011$ $[-0.098, +0.121]$	0.869

reference comparisons use $u_i(R_{11t}, R_{00}) - 1/2$; scalar direct gaps use $Y_{11t} - Y_{00}$ in the corresponding reward units.

F VERIFIER TARGETS AND EXACT H1 INTERSECTIONS

Constructed-defect diagnostics use designated target status on source-available evaluation coordinates. H1 validity instead uses REF supplied-source support ≤ 2 as the positive label. Candidates are joined by lineage, track, role, and text hash. The H1 overlap contains 288 candidates, of which 238 have sources. Granite, MiniCheck, and Lynx each have 238 valid predictions, with 94 positive and 144 negative reference labels. LettuceDetect has 221 valid predictions, with 84 positive and 137 negative labels.

Table 19: Verifier validity against H1 REF source support ≤ 2 on exact source-available candidate intersections, with valid sample counts and 95% lineage-bootstrap intervals.

Verifier	n	AUROC [95% CI]	Balanced acc.
Granite	238	0.901 [0.844,0.948]	0.839
MiniCheck	238	0.757 [0.672,0.830]	0.594
Lynx	238	0.631 [0.555,0.712]	0.631
LettuceDetect	221	0.660 [0.585,0.730]	0.660

Table 20: Constructed-target verifier diagnostics on source-available evaluation coordinates, with their corresponding valid denominators.

Verifier	n	AUROC [95% CI]	AUPRC	Bal. acc.
Granite	880	0.764 [0.715,0.815]	0.674	0.698
MiniCheck	880	0.578 [0.551,0.612]	0.441	0.563
Lynx	880	0.635 [0.591,0.679]	0.511	0.635
LettuceDetect	787	0.615 [0.581,0.651]	0.428	0.615

Machine-readable tables additionally report AUPRC, sensitivity, specificity, MCC, Brier scores, confusion counts, and A/B reference sensitivities. For the binary-valued Lynx and Lettuce outputs, AUROC reduces to balanced accuracy. The H1 target threshold is fixed at source support ≤ 2.

G STRUCTURAL AND STRICT-TAIL DIAGNOSTICS

For additive text length ℓ, prepending the same text of length c to both parents yields zero structural length interaction:

$$\big[\ell(R_{01}) + c - \ell(R_{01})\big] - \big[\ell(R_{00}) + c - \ell(R_{00})\big] = 0. \tag{8}$$

Log-length and normalized edit intensity capture relative changes and can vary across these same pairs. Non-estimable features are recorded as NA. The eight association regressions include the specified source, cohort, defect, and applicable operator covariates. Operator metadata identifies the active transformation, including programmatic E endorsements. Entropy-calibrated coefficients provide point-estimate sensitivities.

The exploratory joint event is $H_{it} = \mathbf{1}\{\kappa_{it} > 0\}\mathbf{1}\{\gamma_{it} > 0\}$. It occurs in five N lineages, with weighted frequency 0.00793 $[0.00166, 0.01456]$, and in zero E lineages under the pooled aggregation. The five N cases lie outside the H1 cohort. Separate membership tables report order-specific, common-rubric, complete-coordinate, and leave-one-family-out results, distinguishing each scenario's event count from retention of the original event IDs.

Table 21: Eight structural association regressions. Coefficients are per predictor standard deviation, with pointwise bootstrap intervals and the specified eight-test Holm adjustment.

Track	Predictor	β per SD [95% CI]	p_{Holm}
N	Log-length interaction	-0.072 $[-0.142, -0.046]$	0.008
N	Differential edit	$+0.072$ $[+0.043, +0.102]$	0.008
N	Defect position	$+0.004$ $[-0.033, +0.040]$	0.958
N	Source length	$+0.021$ $[-0.030, +0.072]$	0.958
E	Log-length interaction	-0.025 $[-0.043, -0.009]$	0.025
E	Differential edit	-0.025 $[-0.044, -0.009]$	0.025
E	Defect position	-0.032 $[-0.047, -0.018]$	0.008
E	Source length	$+0.012$ $[-0.010, +0.035]$	0.834

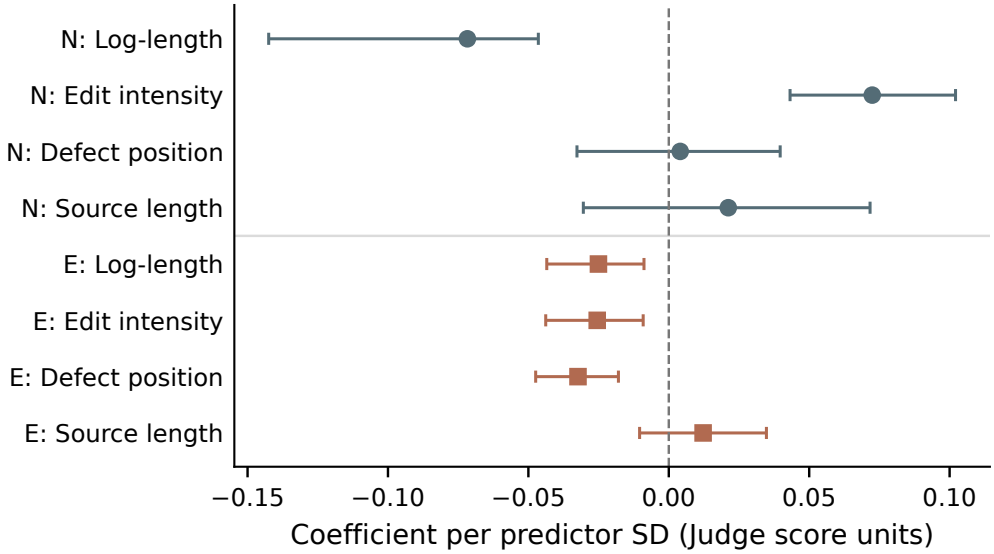


Figure 8: Structural association coefficients with pointwise 95% intervals. N and E use their respective predictor standard deviations and valid lineage sets.

Table 22: Interaction estimates under restrictions on structural log-length and differential edit intensity.

Restriction	n_N	κ_N	n_E	κ_E
Full cohort	620	-0.153 $[-0.182, -0.125]$	620	$+0.325$ $[+0.311, +0.339]$
$\lvert\Delta\Delta \log \ell\rvert \leq .05$	573	-0.130 $[-0.157, -0.103]$	620	$+0.325$ $[+0.311, +0.339]$
$\lvert\Delta\Delta \log \ell\rvert \leq .10$	603	-0.144 $[-0.173, -0.116]$	620	$+0.325$ $[+0.311, +0.339]$
$\lvert\Delta\Delta \log \ell\rvert \leq .20$	612	-0.148 $[-0.177, -0.120]$	620	$+0.325$ $[+0.311, +0.339]$
Lowest length quartile	153	-0.117 $[-0.156, -0.078]$	173	$+0.212$ $[+0.188, +0.235]$
Lowest edit quartile	151	-0.120 $[-0.158, -0.082]$	173	$+0.212$ $[+0.188, +0.235]$

H STUDY SCOPE

RHContrast-Bench studies controlled transformations of fixed long-form answers under reference-conditioned target definitions. E covers four phrasings of a common external-verification claim; N rewrites can affect useful content as well as fluency. The analyses characterize branch-wise scores, differential responses, and reference comparisons for these interventions. Reported estimates retain evaluator-specific units, source eligibility, and reviewed categorical uncertainty.