

# Mediocrity is the key for LLM as a Judge Anchor Selection

Shachar Don-Yehiya<sup>1,2</sup> Asaf Yehudai<sup>1,2</sup> Leshem Choshen<sup>2,3,4</sup> Omri Abend<sup>1</sup>

<sup>1</sup>The Hebrew University of Jerusalem, <sup>2</sup>IBM Research, <sup>3</sup>MIT, <sup>4</sup>MIT-IBM Watson AI Lab  
{first.last}@mail.huji.ac.il

## Abstract

The “LLM-as-a-judge” paradigm has become a standard method for evaluating open-ended generation. To address the quadratic scalability costs of pairwise comparisons, popular benchmarks like Arena-Hard and AlpacaEval compare all models against a single anchor. However, despite its widespread use, the impact of anchor selection on the reliability of the results remains largely unexplored. In this work, we systematically investigate the effect of anchor selection by evaluating 22 different anchors on the Arena-Hard-v2.0 dataset. We find that the choice of anchor is critical: a poor anchor can dramatically reduce correlation with human rankings. We identify that common anchor choices (best-performing and worst-performing models) make poor anchors. Because these extreme anchors are consistently better or worse than all other models, they are seldom indicative of the relative ranking of the models. We further quantify the effect size of anchor selection, showing it is comparable to the selection of a judge model. We conclude with actionable recommendations. First, we conduct a power analysis, and compute sufficient benchmark sizes for anchor-based evaluation, finding that standard benchmark sizes are insufficient for pairwise evaluation and fail to distinguish between competitive models reliably. Second, we provide guidelines for selecting informative anchors to ensure reliable and efficient evaluation practices.

## 1 Introduction

Traditional reference-based metrics (Papineni et al., 2002; Lin, 2004) are often ill-suited for the open-ended nature of modern LLM applications (Liu et al., 2016). Consequently, “LLM as a Judge” (LMJ)—using one model to evaluate another—has emerged as a scalable alternative that correlates highly with human evaluation (Chiang and Lee, 2023; Liu et al., 2023), despite some potential for bias (Wang et al., 2024; Saito et al., 2023).

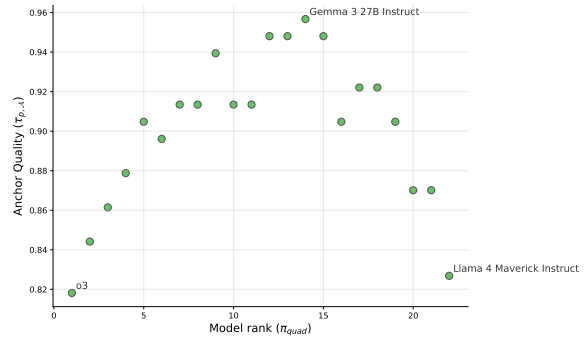


Figure 1: Kendall’s  $\tau$  correlation ( $\tau_{p,A}$ ) plotted against anchor position. The y-axis shows the correlation between the anchor-based ranking and the quadratic ranking  $\pi_{quad}$ , while the x-axis represents the anchor’s position (rank) in  $\pi_{quad}$ . This reveals an inverted U-shaped relationship: top and bottom-ranked models correlate poorly with the gold standard, making them suboptimal anchors. The judge is Deepseek-v3.

A primary setting for LMJ is pairwise comparisons. Given an instruction, the judge compares the responses of two models and states a preference (e.g., the first is better). The main drawback of this approach is that the cost of evaluation grows fast. Specifically, as the number of evaluated models increases, the number of model pairs to compare grows quadratically (Zheng et al., 2023).

To provide better scalability, a common practice is to select an *anchor* model and compare all other models to it. This practice was adopted by popular benchmarks such as the Arena-Hard (Li et al., 2024) and AlpacaEval (Li et al., 2023) evaluation frameworks, and is therefore used by many (e.g., Pombal et al., 2025; Dubois et al., 2024; Raju et al., 2024; Gera et al., 2025; Don-Yehiya et al., 2025; Rafailov et al., 2023; Ethayarajh et al., 2024; Meng et al., 2024; Hong et al., 2024; Chiang et al., 2023; Tang and Feng, 2025; Chen et al., 2025; Xu et al., 2025b).

Although frequently done, few works study the effects of anchor-based evaluation, mostly focusing

on the validity of the transitivity assumption (Xu et al., 2025a; Wang et al., 2025), that states that if models  $A, B, Anchor$  show  $A < Anchor$  and  $Anchor < B$ , then  $A < B$ . They demonstrate that this assumption does not hold, and as an alternative suggest dynamic matching strategies (Liusie et al., 2024; Son et al., 2025), complicating the evaluation process.

Rather than suggesting alternatives, we study the best practices of using anchors. We start by empirically examining the effect of anchor choice. We conduct a large-scale analysis involving over 850K pairwise comparisons across 22 different anchors on the Arena-Hard-v2.0 dataset. We find that a bad anchor can lead to up to .30/.19 drop in correlation with human/quadratic rankings. Notably, we observe an inverted U-shaped relationship between model capability and anchor quality: top-performing (‘strong’) and low-performing (‘weak’) models make the worst anchors (the tails of the U), while ‘mediocre’ models provide the highest correlation (see Fig. 1). This is in sharp contrast to the common practice of using strong or weak models as anchors, as they provide simple ‘baseline’ or ‘gold’ standards for comparison (Li et al., 2024; Xu et al., 2025b).

To better understand the last phenomenon, we look at the win-rate distributions of different anchors (§4.2). We see that overly strong/weak anchors induce skewed distributions. Our results show that these are less helpful, as many of the samples are less informative. For example, *o3* wins against all the other models in about 500/750 of the benchmark’s samples, wasting 2/3 of the evaluation budget.

To further examine the statistical implication of this observation, we run a power analysis that takes into account the ‘informativeness rate’ of the comparisons against the anchor (§4.3). We find that for a small effect size (+5%) and an average informativeness rate, the estimated number of samples is larger than the size of the Arena-Hard-v2.0 dataset. This indicates that the current benchmark is statistically insufficient to reliably distinguish between competitive models in an anchor-based setting.

Finally, we broaden our analysis to practical mitigation strategies and the relative importance of the anchor. We vary the dataset size (§5.1) to test the robustness of our findings, and examine the effect of using multiple anchors (§E). Crucially, we compare the impact of selecting an anchor against the impact of selecting a judge model. We conclude

that choosing an anchor is a critical factor in evaluation reliability, comparable in its effect to choosing a judge (§5.2).

Based on these insights, we suggest a decision framework for pairwise evaluation (Fig. 5). We advise avoiding anchor-based evaluation when possible—specifically for small model sets ( $N \leq 3$ ) or when a natural baseline exists. For leaderboard settings where anchors are unavoidable, we recommend selecting “mediocre” rather than state-of-the-art models to maximize statistical power, and explicitly reporting the anchor’s informativeness to ensure validity.

We release about 900K judgments generated for this study<sup>1</sup>.

## 2 Task Formulation

In this work, we study the use of LLM-based judges for determining the relative quality of systems over a given set of user instructions. Henceforth, *System* or *Model* refers to an LLM that performs a task, and *Judge* refers to the LLM that compares the quality of such systems. Specifically, we focus on the *pairwise anchor-based* evaluation setting, and assume that the transitivity assumption (§1) holds at least to some extent.

Formally, we begin with a set of  $M$  systems  $\mathcal{M} = \{m_j\}_{j=1}^M$ , and  $N$  user instructions  $\mathcal{I} = \{x_i\}_{i=1}^N$ . Each system produces a response for each instruction, denoted with  $\mathcal{R} = \{r_{i,j}\}_{i=1,j=1}^{N,M}$ , such that  $m_j(x_i) = r_{i,j}$ .

In the anchor-based setting, one system is designated as the *anchor*, denoted  $m_A \in \mathcal{M}$ . A judge from  $\mathcal{J} = \{J_p\}_{p=1}^P$  is tasked with comparing a target response  $r_{i,j}$  against the anchor’s response  $r_{i,A}$  for the same instruction  $x_i$ .

The judge maps a triplet of instructions and two candidate responses to a preference verdict:

$$J_p(x_i, r_{i,j}, r_{i,A}) = v_{i,j}^{p,A} \in \{-2, -1, 0, 1, 2\}$$

where 2/1 represents a clear/slight win for the target model  $m_j$  over the anchor model  $m_A$ ,  $-2/-1$  represent a loss, and 0 a tie. Once a judge  $J_p$  evaluates all systems against the chosen anchor  $m_A$ , we obtain a verdict matrix  $V^{p,A} \in \mathbb{R}^{N \times M}$ .

In order to quantify system-level quality, we apply an *aggregation method*. The aggregation method maps the verdict data to a system-level

<sup>1</sup>Code: <https://github.com/IBM/Anchor-Selection>, Data: <https://huggingface.co/datasets/ibm-research/h/900K-Judgements>.

score vector  $\mathbf{s} \in \mathbb{R}^M$ . We consider two aggregation methods commonly used in anchor-based evaluation:

- **Win-Rate:** We collapse the verdicts into  $\{0, 0.5, 1\}$  and compute the average win-rate against the anchor:  $s_j = \frac{1}{N} \sum_{i=1}^N v_{i,j}^{p,A}$ . To score the anchor itself, we compute its average win-rate:  $s_A = 1 - \frac{\sum_{j \neq A} s_j}{M-1}$ .
- **Bradley-Terry (BT):** We collapse the verdicts into  $\{-1, 0, 1\}$  and follow Chiang et al. (2024) estimating the vector of BT coefficients<sup>2</sup>  $s_j$  that maximizes the likelihood of the observed pairwise verdicts in  $V^{p,A}$ . This model posits that the probability of system  $i$  beating system  $j$  is  $P(i \succ j) = \frac{e^{s_i}}{e^{s_i} + e^{s_j}}$ .

Ordering the scores in  $\mathbf{s}$  induces a ranking over the system set  $\mathcal{M}$ , denoted with  $\pi(\mathbf{s})$ . We evaluate the judge  $J_p$  with anchor  $m_A$  by comparing this induced ranking against a golden ranking  $\pi^*$  derived from quadratic comparisons or human annotations. Specifically, we define the *anchor quality* to be the Kendall’s  $\tau$  correlation coefficient:

$$\tau_{p,A} = \text{Kendall}(\pi(\mathbf{s}), \pi^*) \quad (1)$$

### 3 Experimental Setup

**Data.** We use the Arena-Hard-v2.0 benchmark that contains 500 challenging real-world user queries (open-ended software engineering problems, math questions, logic puzzles, etc.) and 250 creative writing queries sourced from Chatbot Arena (Chiang et al., 2024). We replicate the results for the AlpacaEval dataset, which includes 805 instructions from the test sets of Self-instruct (Wang et al., 2023), Open Assistant<sup>3</sup>, Anthropic’s HH-RLHF (Bai et al., 2022), Vicuna (Zheng et al., 2023; Chiang et al., 2023), and Koala (Geng et al., 2023).

**Models.** We examine all models that appear in the Arena-Hard-Auto repository<sup>4</sup>, and are available in the Chatbot Arena leaderboard (see §3.3).

<sup>2</sup>In the Chatbot Arena notebook ([https://colab.research.google.com/drive/1KdwokPjirkTmp0\\_P1WByFNfiqXWQquwH](https://colab.research.google.com/drive/1KdwokPjirkTmp0_P1WByFNfiqXWQquwH)), they demonstrated that the Elo score (Elo, 1967) is noisy for model ranking, as it is highly influenced by the battles order (Boubdir et al., 2023). To obtain more stable results, they used Bradley-Terry (Bradley and Terry, 1952).

<sup>3</sup><https://github.com/LAION-AI/Open-Assistant>

<sup>4</sup><https://github.com/lmarena/arena-hard-auto/tree/main>

Thus, we will be able to compare the automatically extracted ranking to the arena’s human ranking. We end up with 22 contemporary models, see App. A for the full list. These models are used both as anchors and as competitors.

**Judges.** We experiment with 5 different judges: *Deepseek-v3* (Guo et al., 2025), *GPT-OSS 120B*, *GPT-OSS 20B* (OpenAI et al., 2025), *Qwen3 235B-A22B Instruct*, and *Qwen3 8B* (Yang et al., 2025). We chose the first four models based on their high performance, and the last one as a smaller (and hence cheaper) alternative. We run the judges with their default parameters and use the evaluation prompt from the Arena-Hard-Auto repository.

#### 3.1 Extracting Anchor-Based Ranking

Given a judge  $J_p$  and an anchor  $m_A$ , we present the judge with a user query  $x_i$  and two model responses  $(r_{i,A}, r_{i,j})$ , one generated by the anchor and one by another model  $m_j$ . We then parse the judge’s output to extract its verdict  $v_{i,j}^{p,A}$ . Repeating this for all the benchmark samples and for the 22 evaluated models, we end up with  $750 \cdot 22 = 16,500$  comparisons per anchor and judge. We use these comparisons to calculate the model’s win-rates against the anchor to extract a ranking  $\tau_{p,A}$  (see §2).

#### 3.2 Extracting Quadratic (Gold) Ranking

We run the ‘quadratic’ comparisons, i.e., for each instance of the benchmark, we compare the responses of all possible pairs of models. This sums up to  $\binom{22}{2} \times 750 = 173,250$  comparisons per judge, and  $173,250 \cdot 5 = 866,250$  in total. The comparisons of a judge can be summarized into a  $22 \times 22$  win-rate matrix,  $V^{p,A}$ . As the anchor-based ranking is an approximation of the quadratic ranking, we refer to the quadratic ranking as our ‘gold’. Given the win-rate matrix, we use BT to extract the ‘quadratic ranking’,  $\pi_{quad}$ . To complete the picture, we obtain a human ranking as well, see the next section.

#### 3.3 Human Ranking

To obtain a human ranking  $\pi_{human}$ , we use the model’s scores from the Chatbot Arena text leaderboard.<sup>5</sup> Chatbot Arena collects human-annotated battles between pairs of models’ responses and then aggregates the battles with Bradley-Terry into model scores and a continuously updating leader-

<sup>5</sup><https://lmarena.ai/leaderboard/text>

| Anchor                          | $\tau_{quad}$ | $\tau_{human}$ |
|---------------------------------|---------------|----------------|
| Gemma 3 27B Instruct            | .957          | .514           |
| Qwen3 30B A3B                   | .948          | .495           |
| o1                              | .948          | .476           |
| o3 Mini                         | .948          | .486           |
| Claude 3.7 Sonnet thinking 16k  | .939          | .467           |
| Athene V2 Chat                  | .922          | .429           |
| Claude 3.5 Sonnet               | .922          | .486           |
| o3 Mini High                    | .913          | .476           |
| GPT-4.5 (Preview)               | .913          | .495           |
| QwQ 32B                         | .913          | .533           |
| GPT-4.1                         | .913          | .457           |
| GPT-4.1 Mini                    | .905          | .543           |
| GPT-4.1 Nano                    | .905          | .505           |
| Qwen3 32B                       | .905          | .581           |
| o4 Mini                         | .896          | .448           |
| DeepSeek-R1                     | .879          | .486           |
| Llama 3.1 Nemotron 70B Instruct | .870          | .476           |
| Qwen2.5 72B Instruct            | .870          | .381           |
| Gemini 2.5 Flash                | .861          | .476           |
| Qwen3 235B A22B                 | .844          | .571           |
| Llama 4 Maverick Instruct       | .827          | .448           |
| o3                              | .818          | .324           |
| Average correlation:            | .901          | .480           |
| Standard deviation:             | .039          | .057           |

Table 1: Kendall’s tau correlations of each anchor-based ranking with quadratic and human rankings.

board. As the battles data were not available, we took the aggregated scores.

## 4 Results

For each judge  $J_p$ , we measure its quality by Kendall’s  $\tau$  correlation,  $\tau_{p,A}$ , between the anchor-based ranking induced by  $m_A$  (§3.1) and the quadratic ranking  $\pi_{quad}$  (§3.2). Table 1 shows the results for deepseek-v3 as the judge ( $J_p$ ). The tables for the other four judges are found in App. A. Although all correlations are significant with  $p < 0.05$ , we observe that anchors vary in their quality, as shown by a .14 gap in  $\tau_{p,A}$  between the best and worst anchor choices.

We repeat the analysis now comparing to the human ranking,  $\pi_{human}$  (§3.3), as a reference. We observe similar trends, but with an even larger sensitivity, showing a drop of .19 in  $\tau_{p,A}$  between the best and worst anchor choices.

Crucially, the identity of the *worst* anchor is consistent across both the quadratic and human rankings: the *o3* model. Note that this is also the *top-performing* model in our set. In what follows, we study this relation between an anchor’s performance, and its effectiveness as a reference point.

### 4.1 Correlation with Model Ranking

To examine the relation between an anchor’s performance and its effectiveness as an anchor, Fig. 1

plots the correlation  $\tau_{p,A}$  against the rank of the anchor  $m_A$  within the quadratic ranking  $\pi_{quad}$  with *Deepseek V3* as the judge (the full plot with all labels is provided in App. A). The plot reveals a distinct inverted U-shape, where anchors at the edges of  $\pi_{quad}$  (the top and bottom performing models) consistently yield the lowest performance. This finding challenges common practices, where extreme models are frequently selected as  $m_A$  under the assumption that they provide a strong baseline or a reliable lower-bound (e.g., Li et al., 2024). We observe this same inverted U-shape pattern when comparing to the human ranking  $\pi_{human}$  as the ground truth. Additionally, we replicate the experiment for the other judges and on the AlpacaEval dataset; see App. A.

### 4.2 Win-Rate Distribution

To explain the inverted U-shape finding, we examine the win-rate distributions of the anchors against all other models. Given an anchor  $m_A$ , for each instruction  $x_i$ , we count the number of models that win over the anchor  $\sum_{j=1}^{j=M} \mathbb{1}_{v_{i,j}^{p,A} > 0}$ .

Fig. 2 visualizes this for three representative anchors: the top-performing *o3* (right tail), the low-performing *Llama4 Maverick Instruct* (left tail), and the mid-level *Gemma 3 27B instruct* (peak). For the strong *o3*, we observe a heavy positive skew; data points cluster on the left as the anchor dominates most comparisons. Conversely, we see the opposite negative skew for the weak *Llama4 Maverick Instruct*. However, for the mid-level *Gemma 3 27B instruct*, we observe a flatter, more evenly spread distribution.

This distribution shape directly explains the anchor quality. The strongly skewed distributions at the tails are less informative because they suffer from signal saturation; a significant portion of the samples cannot distinguish between different models. For instance, *o3* defeats all opposing models in roughly 500/750 of the benchmark’s samples. This effectively wastes 2/3 of the evaluation budget, as these comparisons yield no information about the relative strength of the opponents.

### 4.3 Informative Samples

To quantify the impact on the evaluation quality of the win-rate distributions observed in §4.2, we measure the prevalence of ‘informative samples’ for each anchor. For a sample  $i$  to be informative for two models  $a$  and  $b$ , their verdicts w.r.t.  $A$ , should hold  $v_{i,a}^{p,A} \neq v_{i,b}^{p,A}$ . We will therefore define




Figure 2: Histograms of the frequency of samples (Y-axis) grouped by the number of models that outperformed the anchor (X-axis). A value of 0 on the X-axis indicates samples where the anchor was superior to all other models, while higher values indicate samples where the anchor was frequently outperformed. *o3* (2a) shows a positive skew, as most of the data points are clustered on the left, in accordance with *o3* being a strong model that usually beats its opponents. Respectively, we get a negative skew for the low performing *Llama 4 Maverick Instruct* (2c). For *Gemma 3 27B-Instruct* (2b) we get a more evenly spread “flatter” distribution.

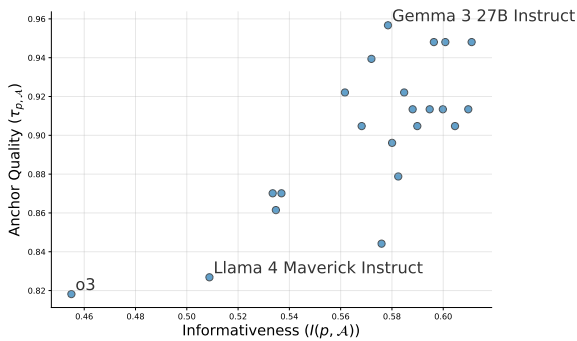


Figure 3: Kendall’s  $\tau$  correlation ( $\tau_{p,\mathcal{A}}$ ) plotted against anchor informativeness. The y-axis shows the correlation between the anchor-based ranking and the quadratic ranking  $\pi_{quad}$ , while the x-axis represents the anchor’s informativeness  $I(p, \mathcal{A})$ . The plot exhibits a positive correlation between anchor quality and anchor informativeness. The judge is *Deepseek-v3*.

the informativeness of an anchor as

$$I(p, \mathcal{A}) = \frac{1}{N \cdot \binom{M}{2}} \sum_{i=1}^N \sum_{a,b \in \mathcal{M}} \mathbb{1}_{v_{i,a}^{p,\mathcal{A}} \neq v_{i,b}^{p,\mathcal{A}}}$$

Our empirical results reinforce the inverted U-shaped hypothesis. The top-performing anchor *o3* yields only 45% informative samples—meaning 55% of the compute budget provides no discriminative signal. In contrast, with the highest informativeness we have *o3 Mini* with 61% informative samples. However, this implies that even in the best scenario roughly 39% of the evaluation budget is inevitably wasted. See App. B for the full results.

To contextualize these rates, note that for the case of verdicts  $v_{i,j}^{p,\mathcal{A}} \in \{-1, 0, 1\}$  (no magnitude), and assuming that the transitivity assumption holds, then  $I(p, \mathcal{A}) \leq 0.5$  with equity when

| Win Rate<br>(Discordant) | Edge | Required Sample Size ( $N$ ) |                     |                     |
|--------------------------|------|------------------------------|---------------------|---------------------|
|                          |      | Base<br>(No Ties)            | Total<br>(39% Ties) | Total<br>(55% Ties) |
| 55%                      | +5%  | 617                          | 1,012               | 1,372               |
| 60%                      | +10% | 153                          | 251                 | 341                 |
| 65%                      | +15% | 67                           | 110                 | 149                 |
| 70%                      | +20% | 37                           | 61                  | 83                  |
| 75%                      | +25% | 23                           | 38                  | 52                  |

Table 2: Required Total Sample Sizes (One-Sided Test) adjusted for informativeness rates. **Base** is the discordant pairs needed for statistical significance using a **One-Sided Test** ( $\alpha = 0.05$ , Power=0.80). **Total** columns account for data loss due to ties.

the anchor is ranked exactly in the middle. That is, the anchor-based setting inherently limits the informativeness.

Table 2 shows the number of samples ( $N$ ) needed to achieve statistical significance (Power=80%,  $\alpha = 0.05$ ) in a sign test between two models. The null hypothesis ( $H_0$ ) is that model B is better than or equal to model A according to the judge’s predictions. We can see that for a small effect size, we will need 617 samples. However, these samples should be *informative*, as the sign test ignores the tied cases. Thus, we will need

$$N_{total} = \frac{N}{I(p, \mathcal{A})}$$

samples<sup>6</sup>, and in the case of *o3* we will have  $N_{total} = 1372$ , far more than the 750 samples of the dataset.

<sup>6</sup>Note that this is an approximation, as  $I(p, \mathcal{A})$  is averaged across all model pairs, whereas  $N$  varies with effect size.

To make better use of the judgments, we can employ a weighted approach like the Wilcoxon signed-rank test. Instead of collapsing the results into three options (tie, model A wins, model B wins), Wilcoxon takes into account the margin of the win (i.e., model A is slightly better than the anchor while model B is in a tie with it  $\neq$  model A is strongly better than the anchor while model B is strongly worse than the anchor). As this test has stronger assumptions about the data, we run a simulation to find  $N$ , see App.C. We find that for a small effect size of +5% and an average of 54% informative samples (the mean informativeness we have for our empirical distributions), we will need  $N = 930$ , about 200 samples less than the sign test (1143), but still more than the dataset size.

Finally, we provide empirical support for the link between sample efficiency and ranking accuracy. Fig. 3 plots the correlation of each anchor’s resulting ranking with the quadratic ranking against its informativeness (for the full plot with all labels see App. B). We observe a strong positive relationship ( $R^2 = 0.5940$ ): as the informativeness increases, the anchor-based ranking aligns more closely with the quadratic ranking. This suggests that maximizing the number of informative samples is not solely a question of computational efficiency, but also yields more reliable results.

## 5 Robustness and Sensitivity Analysis

Having identified the roles of win-rate distributions and informative samples in anchor-based evaluation, we now examine how the evaluation pipeline responds to different settings. We start by testing whether increasing the scale of the evaluation (increasing the dataset size or adding more anchors, see App. E) can mitigate the performance drop of anchor-based evaluation vs. quadratic evaluation. We then compare the magnitude of the effect of selecting an anchor in comparison to that of a judge model (finding them to be comparable). Finally, we estimate the anchor informativeness with fewer samples to allow an informed anchor selection before running the evaluation on the full dataset.

### 5.1 Number of Samples

We next investigate the sample efficiency of anchor-based methods compared to the quadratic approach. We varied the number of instructions  $N$ , sampling 10 sets of samples of sizes 50 to 750. We run BT to aggregate the quadratic ranking. We do the




Figure 4: Mean  $\tau_{p,\mathcal{A}}$  with respect to human ranking averaged over random sample selections as a function of sample size. As the number of samples grows, the variance of the quadratic evaluation correlation decreases. Simultaneously, the mean anchor-based correlation improves, eventually converging with the quadratic correlation at approximately 600 samples. This is not the case for each particular anchor choice, see  $o3$  correlation. This demonstrates that anchor-based ranking is more affected by the dataset size than the quadratic ranking. The judge is Deepseek-v3.

same for each anchor  $\mathcal{A}$ , extracting their anchor-based ranking and their correlation with the human ranking  $\tau_{p,\mathcal{A}}$ . We repeat the process 30 times and average over the resulting correlations.

Our results for *Deepseek-V3* as the judge (Fig. 4) present a distinct difference in stability. Although the standard deviation of the quadratic correlation shrinks as the number of samples grows, the mean correlation does not change much. In contrast, anchor-based rankings are highly sensitive to dataset size. The mean correlation of the anchor-based rankings across all the anchors improves to the point that the quadratic correlation and the mean anchor-based correlation are pretty close (around 600 samples). Note that this is not the case for each particular anchor choice, as the  $o3$  correlation remains far below even when the number of samples grows. Results for the other judges are provided in App. D.

We conclude that anchor-based ranking is more affected by the size of the dataset than the quadratic evaluation. This is in line with the results from §4.3, where we saw that in anchor-based evaluation a significant portion of the dataset is wasted (up to 55% of the comparisons are not informative) and therefore the effective dataset is smaller.

| Judge           | Quadratic | Quadratic vs |       |          | Human vs |       |          |
|-----------------|-----------|--------------|-------|----------|----------|-------|----------|
|                 | vs Human  | Best         | Worst | $\Delta$ | Best     | Worst | $\Delta$ |
| DeepSeek-V3     | .495      | .983         | .844  | .139     | .562     | .371  | .191     |
| GPT-OSS 120B    | .514      | .957         | .766  | .191     | .581     | .371  | .210     |
| GPT-OSS 20B     | .419      | .974         | .784  | .190     | .505     | .200  | .305     |
| Qwen3 235B A22B | .429      | .965         | .818  | .147     | .505     | .276  | .229     |
| Qwen3 8B        | .333      | .957         | .827  | .130     | .410     | .229  | .181     |

Table 3: Kendall’s  $\tau$  correlation coefficients comparing Quadratic/Human and Anchor-based ranking. ‘Best’ refers to the best anchor choice, i.e., the anchor that yields the ranking with the highest correlation to the quadratic. ‘Worst’ is the anchor that yields the ranking with the lowest correlation. The  $\Delta$  columns show the gain from the best vs. the worst anchor choice.

## 5.2 Comparing the Effect of Anchor vs. Judge Selection

Finally, we contextualize the magnitude of the anchor effect by comparing it to the judge effect. Typically, great effort is spent selecting the strongest judge model (Tan et al., 2025; Thakur et al., 2025). This subsection addresses the question of whether selecting an anchor has a similar effect (and is therefore equally important).

Table 3 summarizes the results (full results are provided in App. A). We see that the anchor choice has a comparable or often larger impact on performance than the choice of judge model. The influence of the judge is clearly visible in the ‘‘Quadratic vs Human’’ column, where the difference in the correlation with human judgments between the worst and best judges is 1.81. The impact of the anchor choice is even more pronounced: the difference between the worst and best anchors (in terms of correlation with human judgments) is 0.181 – 0.305, depending on the judge.

Moreover, the effect of selecting a good anchor seems to be complementary to the choice of the judge. Indeed, similar anchor effects are presented for judges of different performance levels.

## 5.3 Estimating Anchor Informativeness

Our experiments showed that the accuracy of the anchor-based ranking is tightly linked to the percentage of informative samples (§4.3). Thus, we would like to choose an anchor that has a high percentage of informative samples. Running the full evaluation for one anchor requires generating  $N \cdot (M + 1)$  model responses, and  $N \cdot M$  judgments. We would like to estimate the anchor informativeness with fewer judgments.

To validate this estimation strategy, we conduct

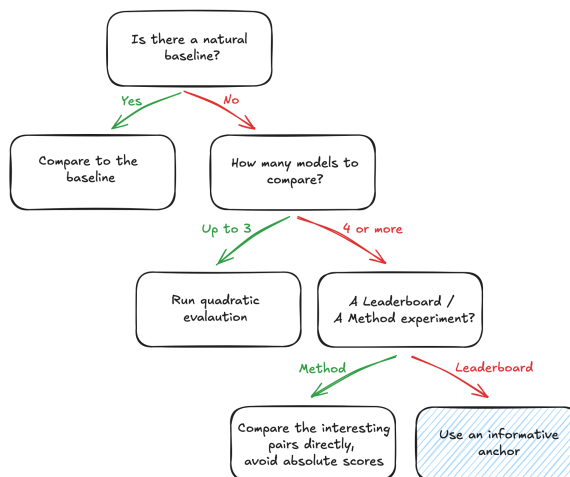


Figure 5: Decision tree for good pairwise evaluation.

an experiment where we varied the number of evaluated models  $M$  within the range  $[3, 22]$ . For each  $M$ , we randomly select  $M$  models and 10 samples from the dataset. We then calculate the informativeness rate for each of the 22 potential anchors and measure the Pearson correlation against the rates derived from the full dataset, repeating the process 30 times.

The results indicate a strong predictive capability: for  $M \geq 3$ , the Pearson correlation is above 0.86, and for  $M \geq 8$ , above 0.91. Additionally, across all values of  $M$ , the estimation successfully identified the least informative anchors—specifically the best and worst-performing models—consistently placing them in the bottom three rankings.

In terms of absolute scores, in §4.3 we saw that the empirical informativeness rate ranges in  $(0.44, 0.61)$ . Here with 10 samples only, we see similar trends, with a maximum informativeness rate of 0.65 (*GPT-4.5 (Preview)*) and a minimum of 0.42 (*o3*). These findings confirm that good anchors can be reliably identified with fewer samples, which is of practical importance, where the more informative anchors (the mid-performing ones) are not known in advance. For the full results, see App. F.

## 6 Recommendations and Conclusion

Based on our analysis, we propose good practices for LMJ anchor-based comparative evaluation. Proposed recommendations are summarized in Fig. 5.

Our analysis revealed that a poor choice of anchor may throw away a substantial part of the evaluation budget, leading to noisier rankings. We

showed that a good anchor choice reduces the noise by up to .19 correlation points. However, even the best anchor we tested has 39% of the benchmark’s samples result in ties. Hence, our first recommendation would be to **avoid anchor-based evaluation when possible**, as the pairwise setting is inherently limiting the informativeness<sup>7</sup> (§4.3).

Passing through works that use anchor-based evaluation, we noticed a few cases of unnecessary use. To eliminate such cases, we suggest first considering whether there is a natural anchor to which all the evaluated models should be compared. For example, a new training method is compared to a similar existing one. If so, use this newly trained model as the anchor, and avoid conclusions regarding other model pairs.

If there is no natural anchor and there are up to three models to evaluate, compare all pairs (quadratic evaluation). This will result in  $3N$  judgments, the same as using an external anchor.

If there are four or more models to evaluate, consider whether this is a leaderboard setting, i.e., you need to rank all the models. Sometimes, there are specific comparisons that are interesting, and there is no need to conclude regarding all the pairs. For example, a paper may propose a new method for a task, a ‘cheaper’ version of it with a smaller model, and an enhancement to the method. We would like to compare the new method to some standard baseline of that task, the enhancement to the new method, and the cheap version to the baseline/new method, or maybe both of them. In this scenario, we will not report absolute scores, but rather the win-rates between pairs of interest.

When a full ranking of models is required, choose your anchor wisely. If possible, **use common knowledge, such as similar leaderboards, to avoid the strongest and weakest models**. Run your evaluation on a smaller sample set first, and confirm the informativeness of the chosen anchor (§5.3). Finally, **report the anchor informativeness as part of your results** to reflect the validity of the evaluation.

## 7 Related Work

In the scope of this work, we discuss LMJ anchor-based pairwise evaluation (§1). Another LMJ setting is pointwise evaluation, where, given an in-

struction and model response, the judge model provides an absolute quality score for the response. The score can be numeric, i.e., a number between 0 and 100, or categorical, e.g., [Very Bad, Bad, Mediocre, Good, Very Good] (Gu et al., 2024). Although pointwise evaluation is easier to scale, it has its own limitations. Its grading may be less suitable to differentiate between model pairs, and is less calibrated and robust to judge or prompt changes (Zheng et al., 2023).

Many works investigated the effect of judge selection, demonstrating that stronger models generally align better with human preferences (Zheng et al., 2023; Chiang et al., 2024; Kocmi and Federmann, 2023). Research has also extensively mapped systematic biases in judges, such as position bias (Wang et al., 2024), verbosity bias (Saito et al., 2023), and self-preference bias (Koo et al., 2024). However, while the choice of judge has been scrutinized, very few works have examined the anchor. Current literature largely treats the anchor as a static default (Li et al., 2024, 2023), overlooking its potential impact on the evaluation outcome. We show that anchor selection is crucial and equally important as the choice of the judge.

Few works have studied the effects of anchor-based evaluation. Gao et al. (2025) showed initial indications of a relation between the anchor performance and the evaluation quality. Xu et al. (2025a) demonstrated that LLM judges exhibit non-transitive preferences, leading to rankings that are sensitive to the choice of anchor. Similarly, Wang et al. (2025) highlighted discrepancies between pointwise and pairwise evaluations, as well as violations of transitivity. Both studies suggest alternative evaluation frameworks, such as dynamic matching strategies. We, on the other hand, do not explore alternatives to the anchor-based evaluation. Instead, we identify that despite its drawbacks, it is extensively used, and propose practical improvements to this methodology.

Significant research has been dedicated to constructing robust benchmarks, with a particular focus on mitigating measurement artifacts such as saturation (Bowman and Dahl, 2021; Ott et al., 2022). In light of this, we do not delve into these issues in this work, but instead proceed under the assumption that these established benchmarks are sufficiently representative of general capabilities and possess adequate discriminative power.

<sup>7</sup>This could possibly be mitigated by a better judge who gives judgment on a broader scale, i.e.  $v_{i,j}^{p,A} \in \{-3, -2, -1, 0, 1, 2, 3\}$ .

## Limitations

Our “gold” ranking is derived from the quadratic evaluation, which may correlate poorly with human rankings (see §5.2). We therefore also report correlations with the human ranking.

We generated judgments with five different models, none of which is a commercial model. However, we chose top-performing open models and thus expect results to be applicable to top commercial models as well. As for the evaluated models, we used both open and commercial models.

The standard BT method does not take into account the magnitude of the judgments. We therefore collapse the scores into  $\{-1, 0, 1\}$  in experiments that use BT, losing information. In our power analysis, however, we do analyze both cases (with and without magnitude, i.e., Wilcoxon’s signed rank test and the sign test).

“Mediocrity” is defined strictly relative to the specific pool of evaluated models. Our findings indicate that an anchor is most effective when it is similar in capability to the models being compared, as this maximizes the discriminative signal. Consequently, if the evaluation focuses on a cluster of high-performing models, the optimal anchor must shift upwards to match them. Anchor selection, therefore, cannot be static; it must be continually calibrated to the capability range of the specific model set being ranked.

Anchor-based evaluation relies on the assumption of transitivity (if  $A > \text{Anchor}$  and  $\text{Anchor} > B$ , then  $A > B$ ), a property that LLM judges have been shown to violate at the instance level. However, we operate under the premise that while this assumption does not hold for every individual comparison, it remains sufficiently valid when averaged across the full benchmark. The aggregation of hundreds of pairwise verdicts helps mitigate the noise of specific intransitive cycles, yielding a global ranking that serves as a practical approximation of relative model quality.

## Acknowledgements

This research was partly supported by a grant from the Israel Science Foundation (grant no. 2912/25).

## References

Ebtesam Almazrouei, Hamza Alobeidli, Abdulaziz Alshamsi, Alessandro Cappelli, Ruxandra Cojocaru, Mérouane Debbah, Étienne Goffinet, Daniel Hesslow,

Julien Launay, Quentin Malartic, Daniele Mazzotta, Badreddine Noune, Baptiste Pannier, and Guilherme Penedo. 2023. [The falcon series of open language models](#). *Preprint*, arXiv:2311.16867.

Yuntao Bai, Andy Jones, Kamal Ndousse, Amanda Askell, Anna Chen, Nova DasSarma, Dawn Drain, Stanislav Fort, Deep Ganguli, Tom Henighan, and 1 others. 2022. Training a helpful and harmless assistant with reinforcement learning from human feedback. *arXiv preprint arXiv:2204.05862*.

Meriem Boubdir, Edward Kim, Beyza Ermis, Sara Hooker, and Marzieh Fadaee. 2023. [Elo uncovered: Robustness and best practices in language model evaluation](#). In *Proceedings of the Third Workshop on Natural Language Generation, Evaluation, and Metrics (GEM)*, pages 339–352, Singapore. Association for Computational Linguistics.

Samuel R. Bowman and George Dahl. 2021. [What will it take to fix benchmarking in natural language understanding?](#) In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 4843–4855, Online. Association for Computational Linguistics.

Ralph Allan Bradley and Milton E. Terry. 1952. [Rank analysis of incomplete block designs the method of paired comparisons](#). *Biometrika*, 39:324–345.

Peter Chen, Xi Chen, Wotao Yin, and Tianyi Lin. 2025. [Compo: Preference alignment via comparison oracles](#). In *The Thirty-ninth Annual Conference on Neural Information Processing Systems*.

Cheng-Han Chiang and Hung-yi Lee. 2023. [Can large language models be an alternative to human evaluations?](#) In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 15607–15631, Toronto, Canada. Association for Computational Linguistics.

Wei-Lin Chiang, Zhuohan Li, Zi Lin, Ying Sheng, Zhanghao Wu, Hao Zhang, Lianmin Zheng, Siyuan Zhuang, Yonghao Zhuang, Joseph E. Gonzalez, Ion Stoica, and Eric P. Xing. 2023. [Vicuna: An open-source chatbot impressing gpt-4 with 90%\\* chatgpt quality](#).

Wei-Lin Chiang, Lianmin Zheng, Ying Sheng, Anastasios N. Angelopoulos, Tianle Li, Dacheng Li, Banghua Zhu, Hao Zhang, Michael I. Jordan, Joseph E. Gonzalez, and Ion Stoica. 2024. [Chatbot arena: an open platform for evaluating llms by human preference](#). In *Proceedings of the 41st International Conference on Machine Learning, ICML’24*. JMLR.org.

Gheorghe Comanici, Eric Bieber, Mike Schaeckermann, Ice Pasupat, Noveen Sachdeva, Inderjit Dhillon, Marcel Blistein, Ori Ram, Dan Zhang, Evan Rosen, Luke Marris, Sam Petulla, Colin Gaffney, Asaf Aharoni,

- Nathan Lintz, Tiago Cardal Pais, Henrik Jacobsen, Idan Szpektor, Nan-Jiang Jiang, and 3416 others. 2025. [Gemini 2.5: Pushing the frontier with advanced reasoning, multimodality, long context, and next generation agentic capabilities](#). *Preprint*, arXiv:2507.06261.
- Tim Dettmers, Artidoro Pagnoni, Ari Holtzman, and Luke Zettlemoyer. 2023. Qlora: Efficient finetuning of quantized llms. *Advances in neural information processing systems*, 36:10088–10115.
- Shachar Don-Yehiya, Leshem Choshen, and Omri Abend. 2025. [Naturally occurring feedback is common, extractable and useful](#). *Preprint*, arXiv:2407.10944.
- Yann Dubois, Balázs Galambosi, Percy Liang, and Tatsunori B Hashimoto. 2024. Length-controlled alpacaeval: A simple way to debias automatic evaluators. *arXiv preprint arXiv:2404.04475*.
- Arpad E Elo. 1967. The proposed uscf rating system, its development, theory, and applications. *Chess life*, 22(8):242–247.
- Kawin Ethayarajh, Winnie Xu, Niklas Muennighoff, Dan Jurafsky, and Douwe Kiela. 2024. Model alignment as prospect theoretic optimization. In *Proceedings of the 41st International Conference on Machine Learning*, ICML’24. JMLR.org.
- Mingqi Gao, Yixin Liu, Xinyu Hu, Xiaojun Wan, Jonathan Bragg, and Arman Cohan. 2025. [Re-evaluating automatic LLM system ranking for alignment with human preference](#). In *Findings of the Association for Computational Linguistics: NAACL 2025*, pages 4605–4629, Albuquerque, New Mexico. Association for Computational Linguistics.
- Xinyang Geng, Arnav Gudibande, Hao Liu, Eric Wallace, Pieter Abbeel, Sergey Levine, and Dawn Song. 2023. [Koala: A dialogue model for academic research](#). Blog post.
- Ariel Gera, Odellia Boni, Yotam Perlitz, Roy Bar-Haim, Lilach Eden, and Asaf Yehudai. 2025. [JuStRank: Benchmarking LLM judges for system ranking](#). In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 682–712, Vienna, Austria. Association for Computational Linguistics.
- Jiawei Gu, Xuhui Jiang, Zhichao Shi, Hexiang Tan, Xuehao Zhai, Chengjin Xu, Wei Li, Yinghan Shen, Shengjie Ma, Honghao Liu, and 1 others. 2024. A survey on llm-as-a-judge. *arXiv preprint arXiv:2411.15594*.
- Daya Guo, Dejian Yang, Haowei Zhang, Junxiao Song, Peiyi Wang, Qihao Zhu, Runxin Xu, Ruoyu Zhang, Shirong Ma, Xiao Bi, and 1 others. 2025. Deepseek-r1 incentivizes reasoning in llms through reinforcement learning. *Nature*, 645(8081):633–638.
- Jiwoo Hong, Noah Lee, and James Thorne. 2024. [ORPO: Monolithic preference optimization without reference model](#). In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 11170–11189, Miami, Florida, USA. Association for Computational Linguistics.
- Albert Q. Jiang, Alexandre Sablayrolles, Antoine Roux, Arthur Mensch, Blanche Savary, Chris Bamford, Devendra Singh Chaplot, Diego de las Casas, Emma Bou Hanna, Florian Bressand, Gianna Lengyel, Guillaume Bour, Guillaume Lample, L lio Renard Lavaud, Lucile Saulnier, Marie-Anne Lachaux, Pierre Stock, Sandeep Subramanian, Sophia Yang, and 7 others. 2024. [Mixtral of experts](#). *Preprint*, arXiv:2401.04088.
- Tom Kocmi and Christian Federmann. 2023. [Large language models are state-of-the-art evaluators of translation quality](#). In *Proceedings of the 24th Annual Conference of the European Association for Machine Translation*, pages 193–203, Tampere, Finland. European Association for Machine Translation.
- Ryan Koo, Minhwa Lee, Vipul Raheja, Jong Inn Park, Zae Myung Kim, and Dongyeop Kang. 2024. [Benchmarking cognitive biases in large language models as evaluators](#). In *Findings of the Association for Computational Linguistics: ACL 2024*, pages 517–545, Bangkok, Thailand. Association for Computational Linguistics.
- Tianle Li, Wei-Lin Chiang, Evan Frick, Lisa Dunlap, Tianhao Wu, Banghua Zhu, Joseph E. Gonzalez, and Ion Stoica. 2024. [From crowdsourced data to high-quality benchmarks: Arena-hard and benchbuilder pipeline](#). *Preprint*, arXiv:2406.11939.
- Xuechen Li, Tianyi Zhang, Yann Dubois, Rohan Taori, Ishaan Gulrajani, Carlos Guestrin, Percy Liang, and Tatsunori B. Hashimoto. 2023. Alpacaeval: An automatic evaluator of instruction-following models. [https://github.com/tatsu-lab/alpaca\\_eval](https://github.com/tatsu-lab/alpaca_eval).
- Chin-Yew Lin. 2004. [ROUGE: A package for automatic evaluation of summaries](#). In *Text Summarization Branches Out*, pages 74–81, Barcelona, Spain. Association for Computational Linguistics.
- Chia-Wei Liu, Ryan Lowe, Iulian Serban, Mike Noseworthy, Laurent Charlin, and Joelle Pineau. 2016. [How NOT to evaluate your dialogue system: An empirical study of unsupervised evaluation metrics for dialogue response generation](#). In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 2122–2132, Austin, Texas. Association for Computational Linguistics.
- Yang Liu, Dan Iter, Yichong Xu, Shuhang Wang, Ruo Chen Xu, and Chenguang Zhu. 2023. [G-eval: NLG evaluation using gpt-4 with better human alignment](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 2511–2522, Singapore. Association for Computational Linguistics.

- Adian Liusie, Vatsal Raina, Yassir Fathullah, and Mark Gales. 2024. [Efficient LLM comparative assessment: A product of experts framework for pairwise comparisons](#). In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 6835–6855, Miami, Florida, USA. Association for Computational Linguistics.
- Yu Meng, Mengzhou Xia, and Danqi Chen. 2024. Simpo: Simple preference optimization with a reference-free reward. In *Advances in Neural Information Processing Systems (NeurIPS)*.
- OpenAI, :, Sandhini Agarwal, Lama Ahmad, Jason Ai, Sam Altman, Andy Applebaum, Edwin Arbus, Rahul K. Arora, Yu Bai, Bowen Baker, Haiming Bao, Boaz Barak, Ally Bennett, Tyler Bertao, Nivedita Brett, Eugene Brevdo, Greg Brockman, Sebastien Bubeck, and 108 others. 2025. [gpt-oss-120b and gpt-oss-20b model card](#). *Preprint*, arXiv:2508.10925.
- Simon Ott, Adriano Barbosa-Silva, Kathrin Blagec, Jan Brauner, and Matthias Samwald. 2022. Mapping global dynamics of benchmark creation and saturation in artificial intelligence. *Nature Communications*, 13(1):6793.
- Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. [Bleu: a method for automatic evaluation of machine translation](#). In *Proceedings of the 40th Annual Meeting of the Association for Computational Linguistics*, pages 311–318, Philadelphia, Pennsylvania, USA. Association for Computational Linguistics.
- José Pombal, Nuno M. Guerreiro, Ricardo Rei, and André F. T. Martins. 2025. [Zero-shot benchmarking: A framework for flexible and scalable automatic evaluation of language models](#). *Preprint*, arXiv:2504.01001.
- Rafael Rafailov, Archit Sharma, Eric Mitchell, Stefano Ermon, Christopher D. Manning, and Chelsea Finn. 2023. Direct preference optimization: your language model is secretly a reward model. In *Proceedings of the 37th International Conference on Neural Information Processing Systems, NIPS '23*, Red Hook, NY, USA. Curran Associates Inc.
- Ravi Shanker Raju, Swayambhoo Jain, Bo Li, Jonathan Lingjie Li, and Urmish Thakker. 2024. [Constructing domain-specific evaluation sets for LLM-as-a-judge](#). In *Proceedings of the 1st Workshop on Customizable NLP: Progress and Challenges in Customizing NLP for a Domain, Application, Group, or Individual (CustomNLP4U)*, pages 167–181, Miami, Florida, USA. Association for Computational Linguistics.
- Keita Saito, Akifumi Wachi, Koki Wataoka, and Youhei Akimoto. 2023. Verbosity bias in preference labeling by large language models. *arXiv preprint arXiv:2310.10076*.
- Seonil Son, Ju-Min Oh, Heegon Jin, Cheolhun Jang, Jeongbeom Jeong, and Kuntae Kim. 2025. [Arenalite: Efficient and reliable large language model evaluation via tournament-based direct comparisons](#). In *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing*, pages 7068–7086.
- Sijun Tan, Siyuan Zhuang, Kyle Montgomery, William Yuan Tang, Alejandro Cuadron, Chenguang Wang, Raluca Popa, and Ion Stoica. 2025. [Judgebench: A benchmark for evaluating LLM-based judges](#). In *The Thirteenth International Conference on Learning Representations*.
- Yuxuan Tang and Yifan Feng. 2025. Beyond pairwise: Empowering llm alignment with ranked choice modeling. *arXiv preprint arXiv:2510.23631*.
- Gemma Team, Aishwarya Kamath, Johan Ferret, Shreya Pathak, Nino Vieillard, Ramona Merhej, Sarah Perrin, Tatiana Matejovicova, Alexandre Ramé, Morgane Rivière, Louis Rouillard, Thomas Mesnard, Geoffrey Cideron, Jean bastien Grill, Sabela Ramos, Edouard Yvinec, Michelle Casbon, Etienne Pot, Ivo Penchev, and 197 others. 2025. [Gemma 3 technical report](#). *Preprint*, arXiv:2503.19786.
- Aman Singh Thakur, Kartik Choudhary, Venkat Srinik Ramayapally, Sankaran Vaidyanathan, and Dieuwke Hupkes. 2025. [Judging the judges: Evaluating alignment and vulnerabilities in LLMs-as-judges](#). In *Proceedings of the Fourth Workshop on Generation, Evaluation and Metrics (GEM<sup>2</sup>)*, pages 404–430, Vienna, Austria and virtual meeting. Association for Computational Linguistics.
- Peiyi Wang, Lei Li, Liang Chen, Zefan Cai, Dawei Zhu, Binghuai Lin, Yunbo Cao, Lingpeng Kong, Qi Liu, Tianyu Liu, and Zhifang Sui. 2024. [Large language models are not fair evaluators](#). In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 9440–9450, Bangkok, Thailand. Association for Computational Linguistics.
- Yidong Wang, Yunze Song, Tingyuan Zhu, Xuanwang Zhang, Zhuohao Yu, Hao Chen, Chiyu Song, Qiufeng Wang, Cunxiang Wang, Zhen Wu, Xinyu Dai, Yue Zhang, Wei Ye, and Shikun Zhang. 2025. [Trustjudge: Inconsistencies of llm-as-a-judge and how to alleviate them](#). *Preprint*, arXiv:2509.21117.
- Yizhong Wang, Yeganeh Kordi, Swaroop Mishra, Alisa Liu, Noah A. Smith, Daniel Khashabi, and Hannaneh Hajishirzi. 2023. [Self-instruct: Aligning language models with self-generated instructions](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 13484–13508, Toronto, Canada. Association for Computational Linguistics.
- Yi Xu, Laura Ruis, Tim Rocktäschel, and Robert Kirk. 2025a. [Investigating non-transitivity in LLM-as-a-judge](#). In *Forty-second International Conference on Machine Learning*.

Zhangchen Xu, Fengqing Jiang, Luyao Niu, Yuntian Deng, Radha Poovendran, Yejin Choi, and Bill Yuchen Lin. 2025b. *Magpie: Alignment data synthesis from scratch by prompting aligned LLMs with nothing*. In *The Thirteenth International Conference on Learning Representations*.

An Yang, Anfeng Li, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chang Gao, Chengen Huang, Chenxu Lv, Chujie Zheng, Dayiheng Liu, Fan Zhou, Fei Huang, Feng Hu, Hao Ge, Haoran Wei, Huan Lin, Jialong Tang, and 41 others. 2025. Qwen3 technical report. *arXiv preprint arXiv:2505.09388*.

An Yang, Baosong Yang, Binyuan Hui, Bo Zheng, Bowen Yu, Chang Zhou, Chengpeng Li, Chengyuan Li, Dayiheng Liu, Fei Huang, Guanting Dong, Haoran Wei, Huan Lin, Jialong Tang, Jialin Wang, Jian Yang, Jianhong Tu, Jianwei Zhang, Jianxin Ma, and 40 others. 2024a. Qwen2 technical report. *arXiv preprint arXiv:2407.10671*.

An Yang, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chengyuan Li, Dayiheng Liu, Fei Huang, Haoran Wei, Huan Lin, Jian Yang, Jianhong Tu, Jianwei Zhang, Jianxin Yang, Jiayi Yang, Jingren Zhou, Junyang Lin, Kai Dang, and 22 others. 2024b. Qwen2.5 technical report. *arXiv preprint arXiv:2412.15115*.

Alex Young, Bei Chen, Chao Li, Chengen Huang, Ge Zhang, Guanwei Zhang, Guoyin Wang, Heng Li, Jiangcheng Zhu, Jianqun Chen, and 1 others. 2024. Yi: Open foundation models by 01. ai. *arXiv preprint arXiv:2403.04652*.

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. 2023. Judging llm-as-a-judge with mt-bench and chatbot arena. In *Proceedings of the 37th International Conference on Neural Information Processing Systems, NIPS '23*, Red Hook, NY, USA. Curran Associates Inc.

## A Full Correlation Results

We evaluate the responses of 22 contemporary models: *o1*, *o3 Mini*, *GPT-4.1*, *GPT-4.5 (Preview)*, *o3 Mini High*, *GPT-4.1 Mini*, *o4 Mini*, *GPT-4.1 Nano*, *o3*, *Qwen3 30B A3B*, *Qwen3 32B*, *QwQ 32B*, *Qwen2.5 72B Instruct*, *Qwen3 235B A22B* (Yang et al., 2025, 2024b), *Claude 3.7 Sonnet thinking 16k*, *Claude 3.5 Sonnet*<sup>8</sup>, *Gemma 3 27B Instruct* (Team et al., 2025), *Gemini 2.5 Flash* (Comanici et al., 2025), *Llama 3.1 Nemotron 70B Instruct*, *Llama 4 Maverick Instruct*<sup>9</sup>, *Athene V2 Chat*, *DeepSeek-R1* (Guo et al., 2025).

<sup>8</sup><https://www.anthropic.com/news/claude-3-5-sonnet>

<sup>9</sup><https://ai.meta.com/blog/llama-4-multimodal-intelligence>

Here we provide the correlation results for all the judges (Tables. 4, 5, 6, 7), and the correlations figures both with quadratic (Figs. 12, 13, 14, 15, 16), and human (Figs. 13, 14, 15, 16) ranking. We can see that the inverted U shape trend remains, although the ranking is changing from judge to judge. The human scores were acquired on September 11th, 2025.

We replicate the results for the AlpacaEval dataset (§3), see Table. 8 and Fig. 17. We use 11 models, which we chose based on their contemporaneity and performance: *GPT-3.5 Turbo*, *GPT-4 Turbo*, *GPT-4o*, *GPT-4 Turbo (Preview)*, *Mixtral 8x22B Instruct* (Jiang et al., 2024), *Qwen2 72B Instruct* (Yang et al., 2024a), *Claude 3.5 Sonnet*, *Llama 3.1 405B Instruct*, *Yi 34B Chat* (Young et al., 2024), *Guanaco 65B* (Dettmers et al., 2023), *Falcon 40B Instruct* (Almazrouei et al., 2023).

## B Informativeness

Fig. 6 presents the plot from §4.3 with all the labels. Table. 9 show  $I(p, \mathcal{A})$  for all anchors with Deepseek-V3 as the judge.

## C Power Analysis Simulation

We provide the code for the power analysis simulation for the Wilcoxon signed test on our data, see 1.

We have a total of  $22 \cdot \binom{21}{2} = 4620$  distributions, 143 of them with effect size of 5%.

## D Number of Samples

We replicate the results from §5.1 for the other judges, measuring the effect of the dataset size on the anchor-based evaluation quality. We can see in Figs. 8, 10, 9 that for the large-medium size judges, the main trends remain: the anchor-based evaluation is more affected by the size of the dataset than the quadratic evaluation. For the Qwen3 8B model, however, we see in Fig. 11 that the mean anchor-based correlation outperforms the quadratic correlation starting at approximately 150 samples, and that the overall correlations are lower.

## E Number of Anchors

Since relying on a single anchor introduces significant variance, we examine whether aggregating multiple anchors mitigates the issue. We perform an iterative analysis: starting with a single random anchor, we compute the Bradley-Terry ranking. In each subsequent step, we add another ran-




Figure 6: Kendall’s  $\tau$  correlation ( $\tau_{p,A}$ ) plotted against anchor informativeness. The y-axis shows the correlation between the anchor-based ranking and the quadratic ranking  $\pi_{quad}$ , while the x-axis represents the anchor’s informativeness  $I(p,A)$ . The plot exhibits a positive correlation between anchor quality and anchor informativeness. The judge is *Deepseek-v3*.

dom model to the anchor set and recompute the ranking, continuing until all 22 models are used (where the result converges to the quadratic ranking, correlation=1.0). We repeat this process over 40 shuffled permutations and average their correlations  $\tau_{p,A}$  with the quadratic ranking.

Fig. 7 illustrates the mean correlation and standard deviation (shaded region) as a function of the anchor set size. As expected, correlation improves, and variance shrinks as more anchors are added. The mean correlation for a single random anchor is 0.92, which seems high. Yet, as established in §1, practitioners do not select anchors at random; they typically select the strongest model (according to prior beliefs). Under this realistic constraint, the starting correlation is actually 0.82 (see Table 1)—a massive .10-point deficit compared to the random average. This demonstrates that while adding anchors helps, the initial choice of anchor remains a critical bottleneck for efficiency.

## F Estimating Informativeness Full Results

The full results are provided in Table 10.

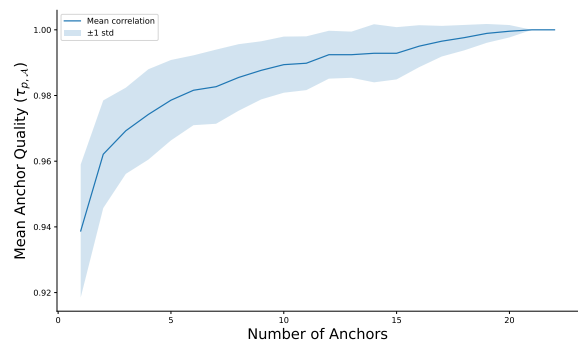


Figure 7: Mean  $\tau_{p,A}$  as a function of the number of anchors averaged over random anchor selections. The correlation increases with the number of anchors, whereas the standard deviation decreases. However, this increase is smaller than the gap between choosing the strongest model as an anchor (.82) and a mean random choice (.92), demonstrating that while adding anchors helps, the initial anchor choice remains critical. The judge is *Deepseek-v3*.

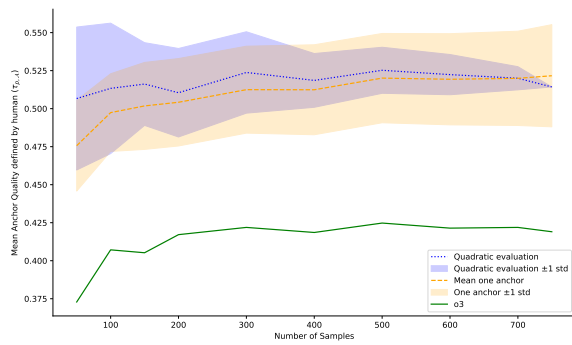


Figure 8: Mean  $\tau_{p,\mathcal{A}}$  with respect to human ranking averaged over random sample selections as a function of sample size. As the number of samples grows, the variance of the quadratic evaluation correlation decreases. Simultaneously, the mean anchor-based correlation improves, eventually converging with the quadratic correlation. This is not the case for each particular anchor choice, see  $o3$  correlation. This demonstrates that anchor-based ranking is more affected by the dataset size than the quadratic ranking. The judge is GPT-OSS 120B.

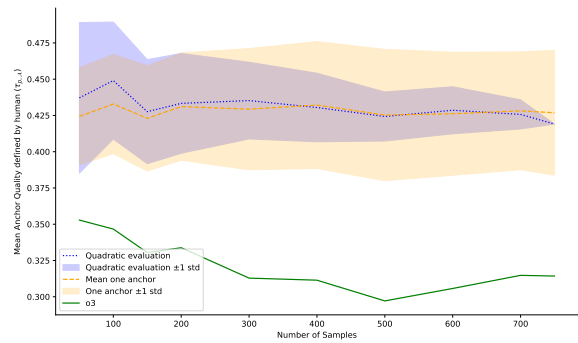


Figure 10: Mean  $\tau_{p,\mathcal{A}}$  with respect to human ranking averaged over random sample selections as a function of sample size. As the number of samples grows, the variance of the quadratic evaluation correlation decreases. Simultaneously, the mean anchor-based correlation improves, eventually converging with the quadratic correlation. This is not the case for each particular anchor choice, see  $o3$  correlation. This demonstrates that anchor-based ranking is more affected by the dataset size than the quadratic ranking. The judge is GPT-OSS 20B.




Figure 9: Mean  $\tau_{p,\mathcal{A}}$  with respect to human ranking averaged over random sample selections as a function of sample size. As the number of samples grows, the variance of the quadratic evaluation correlation decreases. Simultaneously, the mean anchor-based correlation improves, eventually converging with the quadratic correlation. This is not the case for each particular anchor choice, see  $o3$  correlation. This demonstrates that anchor-based ranking is more affected by the dataset size than the quadratic ranking. The judge is Qwen3 235B A22B Instruct.

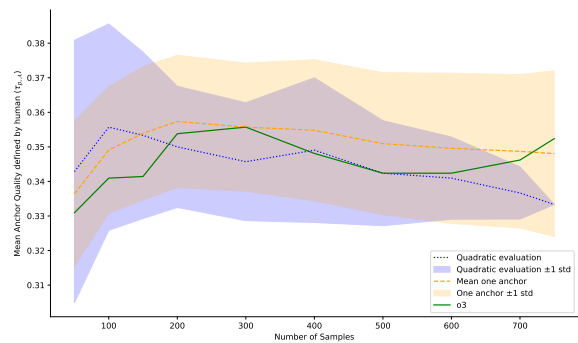


Figure 11: Mean  $\tau_{p,\mathcal{A}}$  with respect to human ranking averaged over random sample selections as a function of sample size. As the number of samples grows, the variance of the quadratic evaluation correlation decreases. Unlike the case of large-medium judges, here with Qwen3 8B as the judge, the mean anchor-based correlation outperforms the quadratic correlation starting at approximately 150 samples.

```

1 def run_power_analysis(all_N, effect_size, alpha, power, M):
2     effect_size_interval = (effect_size, effect_size + 0.01)
3     D = get_empirical_distributions_with_effect_size(effect_size_interval)
4
5     for N in all_N:
6         null_hypothesis_rejected = 0
7         for _ in range(M):
8             d = np.random.choice(D)
9             S = np.random.choice(d, N, replace=True)
10            res = stats.wilcoxon(S, alternative='greater', zero_method='pratt')
11
12            if res.pvalue < alpha:
13                null_hypothesis_rejected += 1
14
15        achieved_power = null_hypothesis_rejected / M
16        if achieved_power >= power:
17            print(f"Successfully rejected null hypothesis with N = {N}")
18            break

```

Listing 1: Power analysis simulation.

| Anchor                          | $\tau_{quad}$ | $\tau_{human}$ |
|---------------------------------|---------------|----------------|
| o1                              | .965          | .333           |
| o3 Mini                         | .965          | .352           |
| Gemma 3 27B Instruct            | .957          | .305           |
| GPT-4.1                         | .957          | .324           |
| Athene V2 Chat                  | .948          | .314           |
| Claude 3.7 Sonnet thinking 16k  | .948          | .343           |
| GPT-4.5 (Preview)               | .939          | .362           |
| Qwen3 30B A3B                   | .931          | .352           |
| o3 Mini High                    | .913          | .371           |
| GPT-4.1 Mini                    | .913          | .400           |
| o4 Mini                         | .913          | .314           |
| DeepSeek-R1                     | .905          | .324           |
| Qwen3 32B                       | .896          | .410           |
| QwQ 32B                         | .887          | .381           |
| Claude 3.5 Sonnet               | .887          | .381           |
| Gemini 2.5 Flash                | .879          | .286           |
| Llama 3.1 Nemotron 70B Instruct | .879          | .362           |
| GPT-4.1 Nano                    | .870          | .352           |
| Qwen2.5 72B Instruct            | .853          | .295           |
| Qwen3 235B A22B                 | .835          | .419           |
| Llama 4 Maverick Instruct       | .835          | .305           |
| o3                              | .818          | .257           |
| <i>Average correlation:</i>     | .904          | .343           |
| <i>Standard deviation:</i>      | .045          | .042           |

Table 4: Kendall’s  $\tau$  correlation,  $\tau_{p,A}$ , of the anchor-based ranking with the quadratic and human ranking  $\pi_{quad}$  and  $\pi^{human}$  for Qwen3-8B as the judge.

| Anchor                          | $\tau_{quad}$ | $\tau_{human}$ |
|---------------------------------|---------------|----------------|
| Athene V2 Chat                  | .948          | .533           |
| Qwen3 30B A3B                   | .948          | .533           |
| GPT-4.5 (Preview)               | .948          | .533           |
| DeepSeek-R1                     | .948          | .476           |
| o1                              | .939          | .505           |
| Qwen3 32B                       | .922          | .543           |
| o3 Mini                         | .922          | .543           |
| Llama 4 Maverick Instruct       | .922          | .533           |
| GPT-4.1 Mini                    | .913          | .533           |
| QwQ 32B                         | .905          | .543           |
| Claude 3.7 Sonnet thinking 16k  | .905          | .505           |
| Qwen2.5 72B Instruct            | .896          | .505           |
| Gemma 3 27B Instruct            | .896          | .590           |
| GPT-4.1 Nano                    | .887          | .533           |
| Qwen3 235B A22B                 | .870          | .524           |
| Llama 3.1 Nemotron 70B Instruct | .870          | .552           |
| Claude 3.5 Sonnet               | .870          | .581           |
| Gemini 2.5 Flash                | .870          | .429           |
| GPT-4.1                         | .861          | .438           |
| o3 Mini High                    | .784          | .533           |
| o4 Mini                         | .766          | .438           |
| o3                              | .662          | .305           |
| <i>Average correlation:</i>     | .884          | .510           |
| <i>Standard deviation:</i>      | .069          | .062           |

Table 5: Kendall’s  $\tau$  correlation,  $\tau_{p,A}$ , of the anchor-based ranking with the quadratic and human ranking  $\pi_{quad}$  and  $\pi^{human}$  for GPT-OSS 120B as the judge.

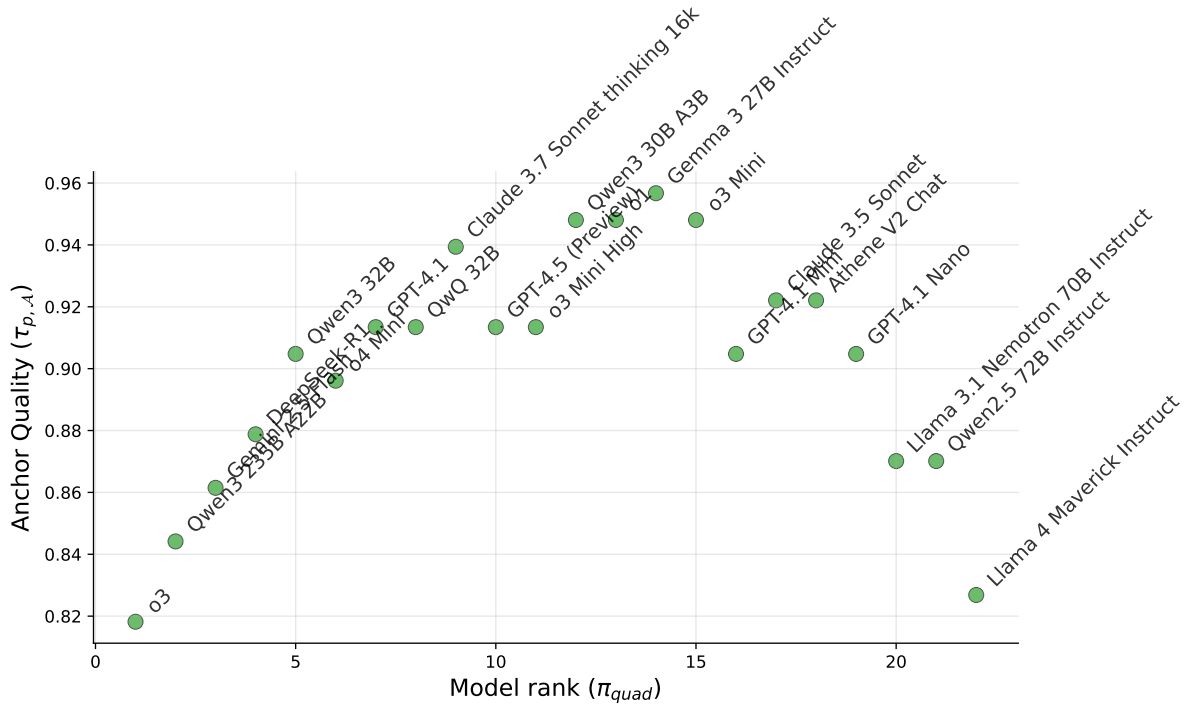


Figure 12: Kendall’s  $\tau$  correlation ( $\tau_{p,\mathcal{A}}$ ) plotted against anchor position. The y-axis shows the correlation between the anchor-based ranking and the quadratic ranking  $\pi_{quad}$ , while the x-axis represents the anchor’s position (rank) in  $\pi_{quad}$ . This reveals an inverted U-shaped relationship: top and bottom-ranked models correlate poorly with the gold standard, making them suboptimal anchors. The judge  $J_p$  is Deepseek-V3.

| Anchor                          | $\tau_{quad}$ | $\tau_{human}$ |
|---------------------------------|---------------|----------------|
| GPT-4.5 (Preview)               | .965          | .448           |
| Qwen3 30B A3B                   | .957          | .448           |
| o1                              | .957          | .467           |
| Gemma 3 27B Instruct            | .948          | .419           |
| o3 Mini High                    | .948          | .429           |
| o3 Mini                         | .948          | .429           |
| o4 Mini                         | .948          | .400           |
| Athene V2 Chat                  | .931          | .390           |
| GPT-4.1 Mini                    | .931          | .457           |
| GPT-4.1                         | .922          | .381           |
| Claude 3.7 Sonnet thinking 16k  | .913          | .400           |
| Qwen3 32B                       | .913          | .438           |
| Claude 3.5 Sonnet               | .905          | .457           |
| QwQ 32B                         | .887          | .505           |
| Llama 3.1 Nemotron 70B Instruct | .887          | .390           |
| DeepSeek-R1                     | .870          | .371           |
| Gemini 2.5 Flash                | .870          | .352           |
| Qwen2.5 72B Instruct            | .853          | .352           |
| GPT-4.1 Nano                    | .835          | .390           |
| Qwen3 235B A22B                 | .827          | .476           |
| Llama 4 Maverick Instruct       | .827          | .429           |
| o3                              | .801          | .295           |
| Average correlation:            | .902          | .415           |
| Standard deviation:             | .050          | .048           |

Table 6: Kendall’s  $\tau$  correlation,  $\tau_{p,\mathcal{A}}$ , of the anchor-based ranking with the quadratic and human ranking  $\pi_{quad}$  and  $\pi^{human}$  for Qwen3 235B A22B Instruct as the judge.

| Anchor                          | $\tau_{quad}$ | $\tau_{human}$ |
|---------------------------------|---------------|----------------|
| GPT-4.5 (Preview)               | .957          | .438           |
| o1                              | .948          | .410           |
| GPT-4.1 Nano                    | .922          | .467           |
| DeepSeek-R1                     | .913          | .429           |
| QwQ 32B                         | .913          | .457           |
| GPT-4.1 Mini                    | .913          | .448           |
| Claude 3.7 Sonnet thinking 16k  | .905          | .381           |
| Llama 4 Maverick Instruct       | .905          | .457           |
| Qwen3 30B A3B                   | .896          | .429           |
| GPT-4.1                         | .879          | .324           |
| Claude 3.5 Sonnet               | .870          | .457           |
| Qwen3 235B A22B                 | .870          | .476           |
| Qwen2.5 72B Instruct            | .870          | .381           |
| Athene V2 Chat                  | .861          | .390           |
| o3 Mini                         | .853          | .419           |
| Qwen3 32B                       | .835          | .495           |
| Llama 3.1 Nemotron 70B Instruct | .835          | .343           |
| Gemma 3 27B Instruct            | .835          | .495           |
| o4 Mini                         | .835          | .381           |
| Gemini 2.5 Flash                | .818          | .390           |
| o3 Mini High                    | .810          | .419           |
| o3                              | .766          | .200           |
| Average correlation:            | .873          | .413           |
| Standard deviation:             | .047          | .066           |

Table 7: Kendall’s  $\tau$  correlation,  $\tau_{p,\mathcal{A}}$ , of the anchor-based ranking with the quadratic and human ranking  $\pi_{quad}$  and  $\pi^{human}$  for textttGPT-OSS 20B as the judge.

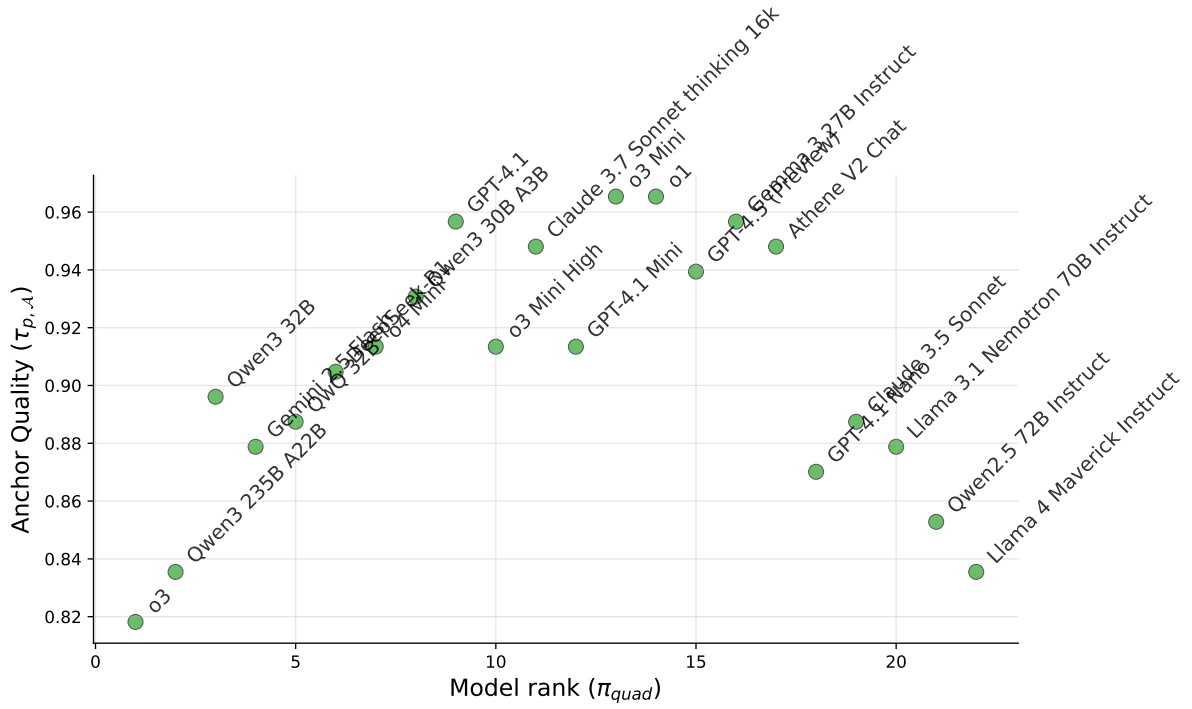


Figure 13: Kendall’s  $\tau$  correlation ( $\tau_{p,A}$ ) plotted against anchor position. The y-axis shows the correlation between the anchor-based ranking and the quadratic ranking  $\pi_{quad}$ , while the x-axis represents the anchor’s position (rank) in  $\pi_{quad}$ . This reveals an inverted U-shaped relationship: top and bottom-ranked models correlate poorly with the gold standard, making them suboptimal anchors. The judge  $J_p$  is Qwen3 8B.

| Anchor                  | $\tau_{quad}$ |
|-------------------------|---------------|
| Mixtral 8x22B Instruct  | .964          |
| Qwen2 72B Instruct      | .964          |
| GPT-3.5 Turbo           | .927          |
| Claude 3.5 Sonnet       | .891          |
| Yi 34B Chat             | .891          |
| GPT-4 Turbo             | .891          |
| Llama 3.1 405B Instruct | .818          |
| Guanaco 65B             | .818          |
| GPT-4o                  | .818          |
| GPT-4 Turbo (Preview)   | .782          |
| Falcon 40B Instruct     | .600          |
| Average correlation:    | .851          |
| Standard deviation:     | .103          |

Table 8: Kendall’s  $\tau$  correlation,  $\tau_{p,A}$ , of the anchor-based ranking with the quadratic ranking  $\pi_{quad}$  for Deepseek-V3 as the judge on the AlpacaEval dataset.

| Model                           | Informative (%) |
|---------------------------------|-----------------|
| o3                              | 45.5            |
| Llama 4 Maverick Instruct       | 50.9            |
| Llama 3.1 Nemotron 70B Instruct | 53.3            |
| Gemini 2.5 Flash                | 53.5            |
| Qwen2.5 72B Instruct            | 53.7            |
| Claude 3.5 Sonnet               | 56.2            |
| GPT-4.1 Nano                    | 56.8            |
| Claude 3.7 Sonnet thinking 16k  | 57.2            |
| Qwen3 235B A22B                 | 57.6            |
| Gemma 3 27B Instruct            | 57.8            |
| o4 Mini                         | 58.0            |
| DeepSeek-R1                     | 58.2            |
| Athene V2 Chat                  | 58.5            |
| QwQ 32B                         | 58.8            |
| Qwen3 32B                       | 59.0            |
| GPT-4.1                         | 59.5            |
| Qwen3 30B A3B                   | 59.6            |
| o3 Mini High                    | 60.0            |
| o1                              | 60.1            |
| GPT-4.1 Mini                    | 60.5            |
| GPT-4.5 (Preview)               | 61.0            |
| o3 Mini                         | 61.1            |

Table 9: Anchor informativeness with Deepseek-V3 as the judge.




Figure 14: Kendall's  $\tau$  correlation ( $\tau_{p,\mathcal{A}}$ ) plotted against anchor position. The y-axis shows the correlation between the anchor-based ranking and the quadratic ranking  $\pi_{quad}$ , while the x-axis represents the anchor's position (rank) in  $\pi_{quad}$ . This reveals an inverted U-shaped relationship: top and bottom-ranked models correlate poorly with the gold standard, making them suboptimal anchors. The judge  $J_p$  is GPT-OSS 120B.




Figure 15: Kendall's  $\tau$  correlation ( $\tau_{p,\mathcal{A}}$ ) plotted against anchor position. The y-axis shows the correlation between the anchor-based ranking and the quadratic ranking  $\pi_{quad}$ , while the x-axis represents the anchor's position (rank) in  $\pi_{quad}$ . This reveals an inverted U-shaped relationship: top and bottom-ranked models correlate poorly with the gold standard, making them suboptimal anchors. The judge  $J_p$  is Qwen3 235B A22B.




Figure 16: Kendall’s  $\tau$  correlation ( $\tau_{p,\mathcal{A}}$ ) plotted against anchor position. The y-axis shows the correlation between the anchor-based ranking and the quadratic ranking  $\pi_{quad}$ , while the x-axis represents the anchor’s position (rank) in  $\pi_{quad}$ . This reveals an inverted U-shaped relationship: top and bottom-ranked models correlate poorly with the gold standard, making them suboptimal anchors. The judge  $J_p$  is GPT-OSS 20B.




Figure 17: Kendall’s  $\tau$  correlation,  $\tau_{p,\mathcal{A}}$ , of the anchor-based ranking with the quadratic ranking  $\pi_{quad}$ , plotted as a function of the anchor  $m_{\mathcal{A}}$ ’s position in  $\pi_{quad}$  on the AlpacaEval dataset. The judge  $J_p$  is DeepSeek-V3. Top and bottom-ranked models in  $\pi_{quad}$  correlate poorly with the quadratic ranking, making them suboptimal anchors.

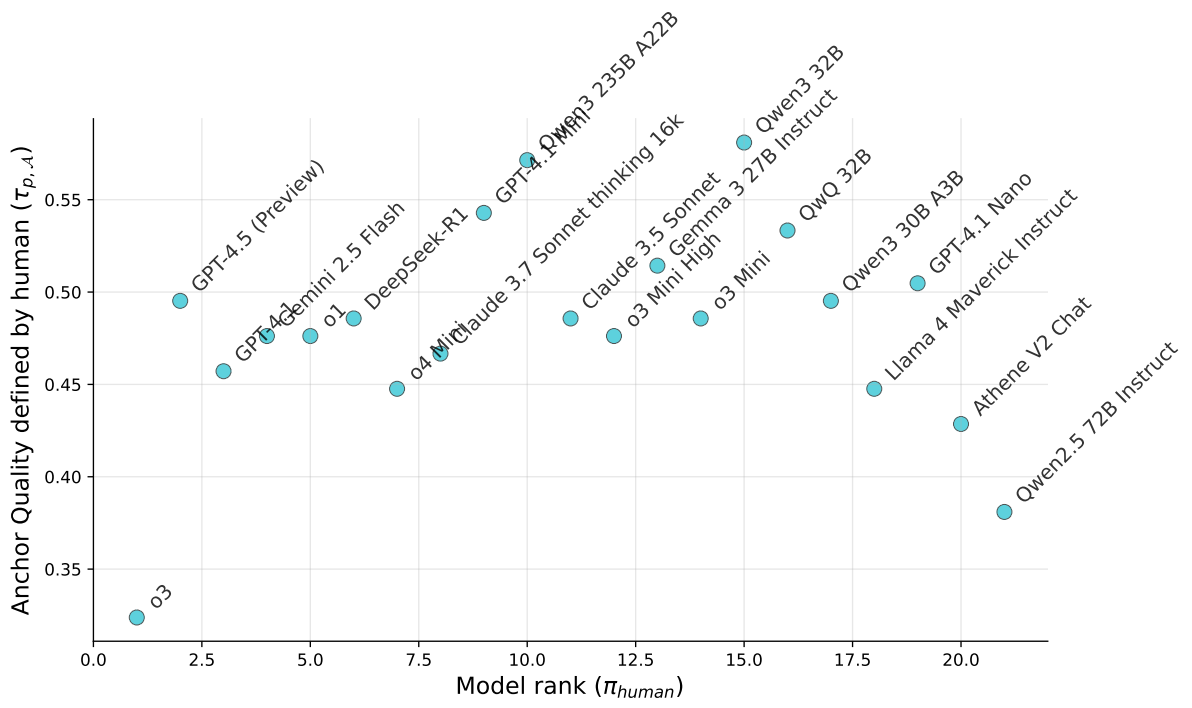


Figure 18: Kendall's  $\tau$  correlation,  $\tau_{p,A}$ , of the anchor-based ranking with the human ranking  $\pi_{human}$ , plotted as a function of the anchor  $m_A$ 's position in  $\pi_{human}$ . The judge  $J_p$  is DeepSeek-V3. Top and bottom-ranked models in  $human$  correlate poorly with the human ranking, making them suboptimal anchors.




Figure 19: Kendall's  $\tau$  correlation,  $\tau_{p,A}$ , of the anchor-based ranking with the human ranking  $\pi_{human}$ , plotted as a function of the anchor  $m_A$ 's position in  $\pi_{human}$ . The judge  $J_p$  is Qwen3 8B. Top and bottom-ranked models in  $\pi_{human}$  correlate poorly with the human ranking, making them suboptimal anchors.




Figure 20: Kendall's  $\tau$  correlation,  $\tau_{p,A}$ , of the anchor-based ranking with the human ranking  $\pi_{human}$ , plotted as a function of the anchor  $m_A$ 's position in  $\pi_{human}$ . The judge  $J_p$  is GPT-OSS 120B. Top and bottom-ranked models in  $\pi_{human}$  correlate poorly with the human ranking, making them suboptimal anchors.




Figure 21: Kendall's  $\tau$  correlation,  $\tau_{p,A}$ , of the anchor-based ranking with the human ranking  $\pi_{human}$ , plotted as a function of the anchor  $m_A$ 's position in  $\pi_{human}$ . The judge  $J_p$  is Qwen3 235B A22B Instruct. Top and bottom-ranked models in  $\pi_{human}$  correlate poorly with the human ranking, making them suboptimal anchors.




Figure 22: Kendall's  $\tau$  correlation,  $\tau_{p,\mathcal{A}}$ , of the anchor-based ranking with the human ranking  $\pi_{human}$ , plotted as a function of the anchor  $m_{\mathcal{A}}$ 's position in  $\pi_{human}$ . The judge  $J_p$  is GPT-OSS 20B. Top and bottom-ranked models in  $\pi_{human}$  correlate poorly with the human ranking, making them suboptimal anchors.

| <b>Rank Pool Size</b> | <b>Pearson Correlation</b> |
|-----------------------|----------------------------|
| 3 Models              | 0.9090                     |
| 4 Models              | 0.8652                     |
| 5 Models              | 0.9528                     |
| 6 Models              | 0.9367                     |
| 7 Models              | 0.8909                     |
| 8 Models              | 0.9658                     |
| 9 Models              | 0.9578                     |
| 10 Models             | 0.9278                     |
| 11 Models             | 0.9500                     |
| 12 Models             | 0.9404                     |
| 13 Models             | 0.9438                     |
| 14 Models             | 0.9542                     |
| 15 Models             | 0.9471                     |
| 16 Models             | 0.9286                     |
| 17 Models             | 0.9679                     |
| 18 Models             | 0.9676                     |
| 19 Models             | 0.9518                     |
| 20 Models             | 0.9698                     |
| 21 Models             | 0.9100                     |
| 22 Models             | 0.9651                     |

Table 10: Estimating anchor informativeness: Pearson Correlation of the estimated informativeness (with 10 samples) and the actual informativeness (with the full dataset) across different amounts of competitive models