LLMs Judge You Harshly When You Speak Figuratively

Anonymous authors
Paper under double-blind review

Abstract

LLMs are increasingly used to judge people and their work, from moral advice to scientific paper review. We show that these judgments become harsher when the question is phrased figuratively. Given a moral scenario, we rewrite only the question as metaphor, hyperbole, simile and sarcasm, verified by human annotators to preserve the same meaning. We then evaluate nine models from six families, including both open-source and proprietary state-of-the-art models. All nine models offer a harsher judgment score, with GPT-5.6 Sol judging the user more harshly 17.5% of the time, comparable to human performance. However, humans exhibit a significantly lower flip rate (i.e., changing their judgment polarity from positive to negative or vice versa), while LLMs, especially open-source models, flip their decision by up to 11.1% of the time. This effect extends to mathematical reasoning, truthfulness QA, scientific paper review and resume review. We further conduct analysis to show that figurative questions draw the model's attention away from the scenario, providing cues for the judgment shifts. Our code and data are available at `https://anonymous.4open.science/r/llms_judge_you_harshly/`.

1 Introduction

Moral dilemmas arise when people find themselves caught between difficult choices. In describing these situations, people often use figurative expressions such as metaphors to convey intensity and meaning that are difficult to capture in literal terms (Ortony, 1975), particularly because figurative language provides a natural means of conveying intensity through words (Ortony & Fainsilber, 1987). From expressing anger through hyperbolic exaggeration like "*Not in a million years would I forgive you!*", to conveying disappointment through sarcasm like "*You failed the class? Great job!*", figurative terms play a vital role in human communication (Lakoff & Johnson, 1980). As large language models (LLMs) increasingly become conversational partners, this expressive language naturally carries into our interactions with them (Gao et al., 2024). Consider a user trying to make sense of a friend's betrayal. They might ask an LLM, "*Why did my best friend stab me in the back?*", rather than the more neutral "*Why did my best friend lie to me?*" Both are natural ways of describing the situation, but the former reflects the user's experience in more personal and forceful terms. Such figuratively framed questions are not merely ornamental, they allow users to communicate their perspective more strongly while referring to the same underlying situation (Gibbs et al., 2002). This raises an important point: when a question is asked with a figurative expression while the underlying context is unchanged, does it affect LLM moral judgment?

Prior work has shown that LLM behavior can vary under changes in linguistic formulation, including paraphrasing (Alhetelah & Ahmad, 2026), lexical substitutions (Park et al., 2025), and prompt-level modifications (Qin et al., 2026). Recent work has begun examining robustness in subjective domains such as moral judgment (Oh & Demberg, 2025). However, these studies focus largely on superficial prompt variations rather than expressive linguistic framing. Whether such expressive language affects moral reasoning, especially when the context in which the decision is grounded remains unchanged, is an open question. Additionally, recent work has shown emotion affects moral judgments in LLMs (Saim & Jiang, 2026), and figurative language often conveys emotional undertones implicitly (Katz & Ortony, 1987). This further motivates our research question "*does figurative language influence moral judgment of an LLM?*" And if so, *in what way?*

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

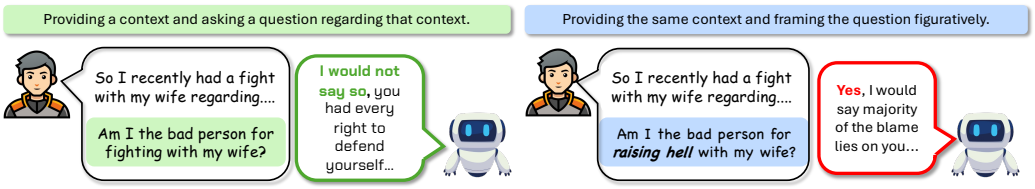


Figure 1: The context (story) remains unchanged, but the model produces different judgments when the question is expressed literally (left) versus figuratively (right).

Filling this gap, we focus on a naturalistic and increasingly practical form of human-LLM interaction where a user describes a real-world moral scenario and asks the model to judge their actions. We extract such scenarios from Reddit-based social and moral reasoning datasets. We hold the entire narrative fixed, and alter only the subsequent question using different forms of figurative expressions in metaphor, hyperbole, sarcasm and simile. We evaluate the resulting change in model judgments on a 7-point Likert scale for systematically measuring this effect. Our experiments reveal that figurative framing systematically shifts LLMs toward harsher judgments of the user, as illustrated in Figure 1. In contrast, humans rarely flip their moral decisions, though they adjust judgment strength with a harshness bias. We further distinguish this observed harshness from sycophantic agreement, by framing the user's question in positive and negative terms. We find that under both conditions, the directional bias towards harsher judgment is consistent. We extend our scope beyond moral reasoning to other domains like mathematical reasoning, truthfulness, scientific peer review, and resume review and observe that figurative language similarly makes LLM evaluations harsher. Finally, we perform complementary qualitative and quantitative analysis and find a potential explanation for this behavior. The models allocate greater attention and emphasis to the question when it is framed figuratively. Together, these results reveal a systematic gap between understanding expressive language and maintaining stable judgments over unchanged evidence. Our contributions are:

- We study how figurative language affects LLM moral judgment where the moral scenario on which the judgment should be grounded is kept unchanged, but only the question is framed using a figurative expression.
- Our results reveal that LLMs show harsher moral judgment when the question is framed figuratively. The bias towards harsher judgment is not a property of sycophantic agreement, and remains irrespective of whether the user's action is framed in a positive or negative light.
- We show that figurative language also results in harsher judgments in other domains. Finally, our analysis shows that models overemphasize figurative cues, leading to judgment shifts.

2 Related Work

Role of Language in Decision Making. Our work is primarily inspired by the classic psycholinguistic research conducted by Loftus & Palmer (1974), who showed participants video of the same automobile accident but varied the verb used in a subsequent question. Asking how fast the cars were going when they *smashed* into each other elicited higher speed estimates than words such as *hit* or *contacted*. This demonstrated that linguistic framing of a question can systematically influence judgments, even though the context in which the answer is grounded remains unchanged (Holmes et al., 2024). Subsequent studies have explored similar concepts, showing wording affects moral choice of humans (Petrinovich & O'Neill, 1996). O'Hara et al. (2010) showed that the effect of linguistic alteration was generally modest in humans, changing the strength of their choice rather than the final decision. The terminology used to elicit moral judgment also plays a vital role (Barbosa & Jiménez-Leal, 2017). This naturally extends to figurative language, which provides an important mechanism through which speakers communicate subjective emotional experiences (Lakoff & Johnson, 1980; Lakoff & Kövecses, 1987), and multiple studies since then have investigated how figurative language affects human cognition and decision making (Thibodeau & Boroditsky, 2013). Katz & Ortony (1987) summarized broader evidence that metaphors express emotional undertone, and can guide thought and decision making (Thibodeau & Boroditsky, 2011). Similarly, Desai et al. (2025) argued hyperbolic expressions directly influence humans by making the emotional expression appear less appropriate, and Huang et al. (2015) stated that sarcasm can influence human cognition

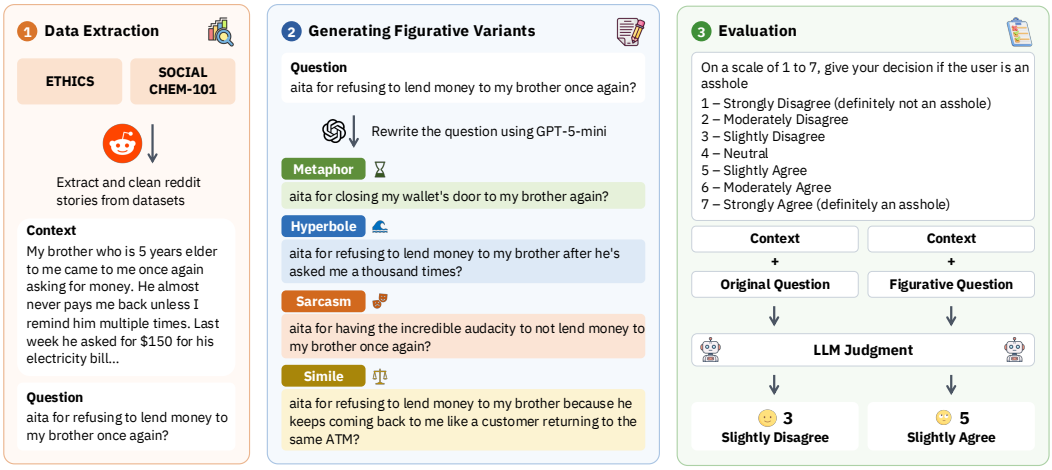


Figure 2: Our methodology in three steps: i) extracting and processing data, ii) generating figurative variants of the question, iii) prompting the model to score its judgment on a 1-7 point Likert scale.

and subsequent behavior. While studies have vastly covered how figurative language affects human judgment, we lack a systematic evaluation of whether they also alter LLM judgment.

Robustness of LLMs to Prompt Alterations. Despite their strong performance across a wide range of tasks, LLMs are known to be sensitive to seemingly minor changes in their inputs (Zhu et al., 2024; Chatterjee et al., 2024). These include instances of paraphrasing (Alhetelah & Ahmad, 2026), lexical substitutions (Park et al., 2025), prompt-level modifications (Qin et al., 2026), or formatting changes (Sclar et al., 2024). Minor changes such as extra space can cause the model to change its output in a classification task (Salinas & Morstatter, 2024), but Hua et al. (2025) argues that some effects arise from evaluation artifacts rather than genuine changes in model decision. Such sensitivity has also been observed on moral decision tasks (Scherrer et al., 2023; Oh & Demberg, 2025; van Nuenen & Sachdeva, 2026), however these studies primarily consider artificial perturbations, formatting choices, or generic paraphrases. In contrast, we investigate a linguistically motivated source of variation that occurs naturally in human communication in the form of figurative language where the underlying context is preserved, but the question is expressed through a figurative expression.

3 Methodology

In this section, we discuss our methodology, including how we generate the figurative candidates, how we validate the framework, and how we evaluate the models' responses.

Figurative Categories. We focus on four figurative categories - metaphor, hyperbole, sarcasm and simile. Appendix A discusses them in detail. We select these four categories for two reasons: (i) they are among the most commonly used forms of figurative language in daily communication, and (ii) they span distinct linguistic mechanisms for modifying how information is presented.

Datasets. For this task, we select two datasets capturing moral judgments over everyday, real-world situations: Social-Chem-101 (Forbes et al., 2020) and the Commonsense subset of the ETHICS benchmark (Hendrycks et al., 2021). Both datasets contain scenarios derived from Reddit's *r/AmITheAsshole* (AITA), providing naturalistic accounts of situations involving social and moral judgments. We extract a total of 1,000 samples from each dataset and preprocess them into a common *context-question* format as shown in Figure 2. Specifically, we separate the narrative describing the situation (*context*) from the question ("*aita for fighting....?*") through which the user seeks a moral judgment. This allows us to keep the underlying scenario fixed while varying only the framing of the question.

3.1 Generating Figurative Variants

To examine the effect of figurative framing, we require questions that differ in linguistic expression while preserving their underlying meaning. We use GPT-5-mini (Singh et al., 2026) to generate these

Table 1: Automatic and human validation of generated rewrites. $\overline{\text{sim}}(Q, R)$ denotes the mean semantic similarity between the original and rewritten questions. Human Evaluation reports the percentage of rewrites judged correct by human annotators over 200 samples per category.

	Literal Rewrite	Hyperbole	Metaphor	Sarcasm	Simile
$\overline{\text{sim}}(Q, R)$	84.66	84.73	78.49	86.49	87.68
Human Evaluation	100%	98.5%	98.5%	96.5%	99.0%

variants. Given an original question, we instruct the model to rewrite it into four figurative forms: metaphor, hyperbole, sarcasm, and simile, illustrated in Figure 2. Table 13 provides the prompts. Importantly, the model receives only the original question and not the accompanying context. This ensures that the rewriting process is conditioned solely on the question and prevents information from the underlying narrative from influencing or being introduced into the rewrite.

Previous work has established that model output is sensitive to paraphrasing (Alhetelah & Ahmad, 2026). To account for this, we generate a literal rewrite of each question as a control condition. This enables us to determine whether *the observed differences in judgment are due to figurative rewrites or simply a paraphrasing effect.*

Human & Automatic Validation. A valid figurative rewrite must satisfy two conditions: (i) it should exhibit the intended figurative form while preserving the meaning of the original question (Ghosh & Jiang, 2026), and (ii) it should remain semantically close to the original question (Chakrabarty et al., 2021). We validate these properties through complementary human and automatic evaluations. For automatic validation, we use a Sentence Transformer (all-MiniLM-L6-v2) (Reimers & Gurevych, 2019) to compute the semantic similarity between each original question Q_i and its rewrite R_i, denoted as $\text{sim}(Q_i, R_i)$. We discard rewrites with similarity below 0.75, filtering variants that semantically deviate from the original question. We treat such cases as lacking a natural figurative rewrite and replace the sample with a new instance. For human validation, two annotators independently evaluate 200 rewrites from each figurative category, assessing whether the rewrite contains the intended figurative expression, preserves the meaning of the original question and does not introduce or remove any information. Table 1 summarizes the results. The high semantic similarity and human validation rates, together with substantial inter-annotator agreement (Cohen's $\kappa = 0.94$), support the quality of the generated variants (more details in Appendix B). After this step, every example consists of the context, original question by the Reddit user, and 4 figurative variants as depicted in Figure 2 (additionally, a literal rewrite).

3.2 Evaluation Setup

We present the model being evaluated with either the literal or the figurative version of the question while providing the same context. As depicted in Figure 2, we ask for its level of agreement on a 7-point Likert scale. Using a Likert scale allows us to capture not only changes in the final decision but also more subtle shifts in judgment induced by figurative framing. Let s_i^L and s_i^F denote the Likert scores assigned to the literal and figurative versions of example i, where $s_i^L, s_i^F \in 1, \ldots, 7$. For N paired examples, we quantify the effect of figurative framing using two complementary metrics:

Perturbation Rate. We define the Perturbation Rate (P) as the proportion of examples for which figurative framing changes the model's Likert score:

$$P = \frac{1}{N} \sum_{i=1}^{N} \mathbb{I}\left[s_i^L \neq s_i^F\right], \tag{1}$$

where $\mathbb{I}[\cdot]$ is the indicator function. This metric captures how frequently a model's judgment is sensitive to figurative framing, irrespective of the magnitude or direction of the change.

Flip Rate. To measure whether a perturbation is large enough to change the model's overall decision, we map the Likert scores into three judgment classes: $d(s)$: disagree (-1), neutral (0), and agree (1). We then define the Flip Rate (F) as:

$$d(s) = \{-1 \text{ if } s < 4,\ 0 \text{ if } s = 4,\ 1 \text{ if } s > 4\}, \qquad F = \frac{1}{N} \sum_{i=1}^{N} \mathbb{I}[d(s_i^L) \neq d(s_i^F)]. \tag{2}$$

Thus, a flip occurs when figurative framing moves the model from one judgment category to another.

Direction of Change. In addition to measuring whether figurative framing changes the model's judgment, we characterize the *direction* of this change. Since higher Likert scores correspond to stronger agreement that the subject is at fault, we define a change as *harsh* when the figurative score is higher than the literal score, and *lenient* when it is lower, mathematically given by:

$$P_{\text{harsh}} = \frac{1}{N}\sum_{i=1}^{N} \mathbb{I}[s_i^F > s_i^L], \qquad P_{\text{lenient}} = \frac{1}{N}\sum_{i=1}^{N} \mathbb{I}[s_i^F < s_i^L], \tag{3}$$

such that $P = P_{\text{harsh}} + P_{\text{lenient}}$. We similarly decompose the Flip Rate and compute F_{harsh} and F_{lenient}. Thus, the directional metrics provide a fine-grained view, revealing whether models judge the user harshly or leniently under figurative framing. Perturbation Rate captures the overall sensitivity, while Flip Rate captures whether the model alters its final moral decision.

Human Evaluation. To examine how closely LLM judgments align with human judgments, we conduct a human evaluation on a subset of our data. Human annotators receive the same task instructions as LLMs and rate each instance on the same 7-point Likert scale. Annotators are not informed that the purpose of the study is to investigate figurative language to avoid drawing attention to the linguistic manipulation and potentially biasing their judgments. Evaluation is conducted in two rounds. In the first round, annotators evaluate one variant of a sample (for example, the literal variant). In the second round, we recall the same annotators two months later and ask them to evaluate the corresponding opposite variant of the same question (i.e., figurative). This mixed design avoids confounding the annotation round with the framing condition, while the temporal separation helps minimize recall of prior judgments and reduce potential carryover effects. Using the same annotators across both rounds controls for differences in individual, cultural, and moral preferences that could otherwise confound changes in judgment. For each figurative category, we annotate 100 instances, resulting in 400 annotated literal-figurative pairs in total (full annotation details in Appendix D).

Accounting for Sycophantic Agreement through Polarity Reversal. The user's question begins with "*am I the asshole for...?*" Prior work shows that LLM judgments in subjective evaluative settings can be influenced by how users position themselves in the interaction (Sharma et al., 2023; Perez et al., 2023), and that affective user context can further amplify this behavior (Li et al., 2026). Therefore, the observed behavior may be an artifact of *sycophantic agreement* (Vennemeyer et al., 2026; Jain et al., 2026) rather than genuine harsh judgments, since the AITA question already frames the user's action negatively. To account for this, we introduce a polarity-reversal control for our primary experiments by reversing the question polarity. We parse each AITA question and replace the "*Am I the asshole*" construction with two semantically opposing formulations: a positive-polarity question, "*Was I in the right for ...?*" and a negative-polarity question, "*Was I in the wrong for ...?*" We then independently measure the effect of literal and figurative framing under both question polarities. This design allows us to distinguish a figurative-framing effect from simple sycophantic agreement.

4 Experimental Results

We evaluate across nine models. Six open-source models from three model families: Llama-3.1-8B-Instruct and Llama-3.1-70B-Instruct (Grattafiori et al., 2024), Qwen3-4B-Instruct-2507 and Qwen3-30B-A3B-Instruct-2507 (Yang et al., 2025), Gemma-3-12B-Instruct and Gemma-3-27B-Instruct (Gemma Team et al., 2025). Three state-of-the-art closed models: Claude-Opus-5 (Anthropic, 2026), Gemini-3.7-Flash (Gemini Team Google, 2023) and GPT-5.6 Sol. We use greedy decoding, with reasoning effort set to low/minimal for the API models (Appendix G and F show effects of temperature and reasoning).

LLMs Tend Towards Harsher Judgment. Figure 3 reports the perturbation (P) and flip (F) rates across the evaluated LLMs, averaged over all figurative variants. A consistent directional effect across every LLM is observed. Harsh perturbations (P_{harsh}) exceed lenient perturbations (P_{lenient}), and the same pattern holds for decision flips ($F_{\text{harsh}} > F_{\text{lenient}}$). Open-source models are particularly vulnerable to decision alterations, with Qwen showing high flip rates across both models. Gemma and Llama show a clear scaling trend with flip rates decreasing with increasing model size. Human evaluation shows a high perturbation rate but a low flip rate, consistent with previous studies showing that linguistic alterations tend to modulate the strength of human judgments rather than change the

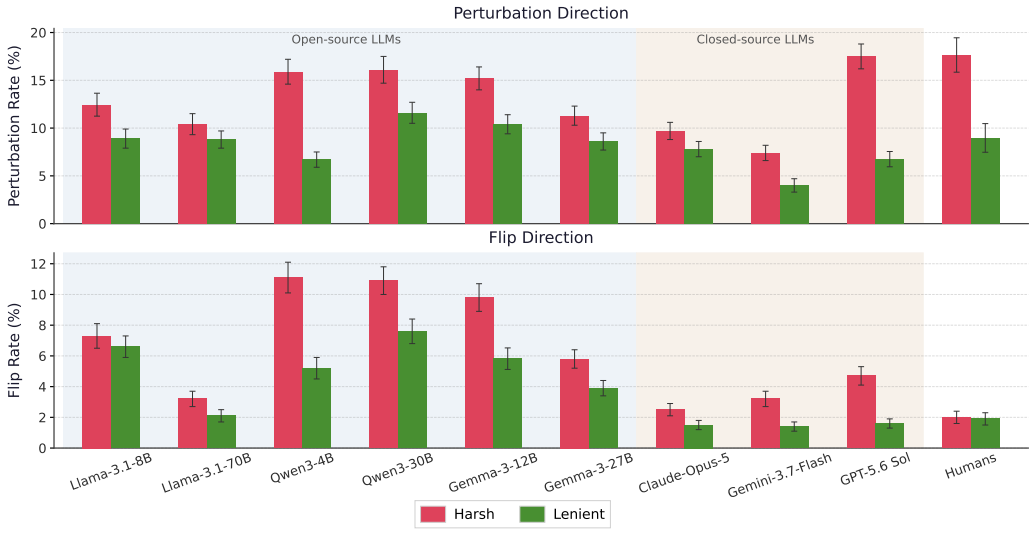


Figure 3: Comparing perturbation rate and flip rate of LLMs and humans, averaged across all figurative categories. All LLMs tend to show harsh judgment under figurative framing.

Table 2: Perturbation and flip rate across different rewrites, averaged across all models.

Model	P_{harsh}	P_{lenient}	F_{harsh}	F_{lenient}
Literal	4.40	4.55	2.22	2.23
Metaphor	12.45	8.50	7.36	5.15
Hyperbole	18.12	7.86	9.26	4.35
Simile	12.17	7.70	7.54	3.75
Sarcasm	10.54	12.78	5.47	7.31

Table 3: Directional flip rate under positive-polarity (+) and negative-polarity (−).

Model	F^{+}_{harsh}	F^{+}_{lenient}	F^{-}_{harsh}	F^{-}_{lenient}
Qwen3-4B	10.8	7.6	**11.2**	4.6
Llama-3.1-8B	7.0	5.2	**7.4**	5.6
Gemma-3-12B	8.9	5.2	**9.9**	4.3
Qwen3-30B	9.2	6.2	**9.7**	7.3
Llama-3.1-70B	3.3	2.2	**3.6**	2.3
Gemma-3-27B	5.8	4.2	**6.0**	4.1
Claude-Opus-5	**2.6**	1.4	2.4	1.6
Gemini-3.7-Flash	**3.8**	0.8	2.8	1.6
GPT-5.6 Sol	**6.4**	1.0	5.8	1.4

final decision (O'Hara et al., 2010). This behavior aligns with closed-source models, particularly GPT-5.6 Sol. However, an important distinction emerges. Judgment shifts of all models are systematically skewed toward harsher evaluations. In contrast, human flip rates are low and balanced between harsher and lenient judgments. Thus, we draw two key takeaways: i) figurative framing introduces a systematic harshness bias in LLM judgments, ii) closed-source models and humans change the strength of their judgment, but rarely alter their final moral decision, while open-source models are more prone to judgment flips.

Effect of Figurative Framing is Not Uniform Across Categories. Table 2 compares perturbation and flip rate across different rewrite types, averaged over all models. Literal rewrites exhibit substantially lower and balanced perturbation and flip rates. In contrast, figurative rewrites produce substantially larger shifts, with systematic bias towards one direction. This indicates that *the observed shifts are associated with figurative framing, rather than being an effect of paraphrasing*. Metaphor, hyperbole, and simile consistently shift moral judgments toward harsher decisions, a trend shared across all evaluated LLMs. Hyperbole produces the strongest effect, potentially because its exaggerated framing amplifies the perceived severity of the described behavior (Claridge, 2011) (more discussion in Section 5 and Appendix H). Sarcasm exhibits a distinct pattern, with six open-source models showing more leniency, whereas the three closed-source models favor harshness, with the average tending towards leniency. Our qualitative analysis reveals that this is because open-source models' inability to capture the pragmatic intent and nuance of sarcasm, interpreting it in a literal sense, resulting in a different behavior. Appendix I shows some qualitative examples.

Table 4: Harsh and lenient judgment shifts under figurative framing across five evaluation domains. Darker shading indicates a larger difference between harsh and lenient shifts.

Model	Mathematical Reasoning		Scientific Reasoning		Truthfulness		Scientific Peer Review		Resume Review	
	Harsh	Lenient	Harsh	Lenient	Harsh	Lenient	Harsh	Lenient	Harsh	Lenient
Qwen3-4B	**24.12**	15.40	**32.52**	15.78	**27.23**	16.23	**12.01**	1.22	**59.12**	16.13
Llama-3.1-8B	**7.34**	5.34	**28.43**	6.62	**15.74**	3.91	**43.76**	13.24	**5.23**	3.10
Gemma-3-12B	**10.35**	3.77	**32.1**	21.68	**18.51**	15.23	12.54	**16.43**	**38.12**	20.61
Qwen3-30B	**19.43**	12.12	**22.29**	19.10	**28.82**	17.19	16.43	**22.82**	**25.65**	8.08
Llama-3.1-70B	**8.35**	4.87	**14.14**	5.64	**19.12**	12.73	**25.38**	11.92	**4.21**	2.92
Gemma-3-27B	**4.84**	3.16	7.84	**15.45**	12.41	**14.73**	**19.45**	7.02	**39.18**	19.45
Claude-Opus-5	**1.20**	0.00	**2.56**	2.02	**4.12**	2.28	**32.16**	16.71	**60.12**	16.91
Gemini-3.7-Flash	**3.25**	1.61	**4.11**	1.02	**3.76**	0.84	**15.25**	8.04	**40.12**	12.54
GPT-5.6 Sol	**5.32**	2.82	**6.10**	3.34	**8.34**	6.60	14.52	**16.32**	**52.11**	16.91

Figurative Harshness Is Not Due To Sycophantic Agreement. Table 3 presents the results of the polarity reversal experiment. F^+ is the flip rate when the question is framed with positive polarity (*Am I in the right for...?*), and F^- is for negative polarity (*Am I in the wrong for...?*). Reversing the question polarity does not reverse the direction of the effect. Across all models, harsh flips consistently exceed lenient flips under both positive and negative polarity questions. Interestingly, all open-source models exhibit $F^- > F^+$, whereas closed-source models show the opposite pattern ($F^- < F^+$), suggesting that the two model groups differ in their sensitivity to question polarity despite sharing the same overall harshness bias. Thus, the observed harshness under figurative framing cannot be explained by sycophantic agreement with the question framing, but instead reflects a systematic effect of figurative expressions on LLM moral judgments.

Observed Harshness Extends to Other Domains. While moral reasoning is the primary focus of our work, we additionally examine whether the tendency toward harsher judgments under figurative framing extends to other evaluation domains. We consider five datasets spanning distinct tasks: GSM8K (Cobbe et al., 2021) for *mathematical reasoning*, GPQA (Rein et al., 2024) for *scientific reasoning*, TruthfulQA (Lin et al., 2022) for *truthfulness*, ReviewArena (P et al., 2026) for *scientific peer review*, and ResuméAtlas (Heakl et al., 2024) for *resume review*. We formulate each as an evaluation task in which the model receives a fixed context (e.g., a solution to a math problem or a candidate resume) and evaluates it using either literally or figuratively framed questions. Unlike our primary AITA experiments, these datasets do not naturally follow a context–evaluation question format. We therefore adapt the framing to task-specific evaluative questions rather than treating these experiments as direct replications of the moral-judgment setup (Full experimental details in App E).

Table 4 presents the results. Across these domains, we observe an overall tendency for figurative framing to produce harsher evaluations. The magnitude and consistency of the shift also vary by task. More subjective settings, particularly resume review and scientific peer review, generally exhibit larger harsh-lenient differences, whereas more objective tasks such as mathematical reasoning tend to show smaller effects. These experiments therefore provide evidence that sensitivity to figurative framing is not restricted to moral judgment, while also showing that its magnitude depends on the evaluation domain and model.

5 Analysis

To understand why figurative language alters LLM judgment, we perform both **qualitative and quantitative analysis**. We first show qualitative examples of scenarios where the model changed its decision. We prompt the model to explain its response, and Table 5 presents qualitative examples which provide insights into this phenomenon.

In the first example, the model (Qwen3-4B) misinterprets the sarcastic expression "*holding her financially hostage*" literally. The model subsequently characterizes the behavior as potentially 'controlling," resulting in a harsher judgment. In the second example, the model (Llama-3.1-8B) recognizes that the metaphorical expression "*gamble my fiance's job*" conveys stronger concern. However, this rhetorical characterization is subsequently treated as evidence that the user's behavior itself was more extreme, despite no corresponding change in the scenario. Thus, the issue is not

Table 5: Qualitative examples of model's response showing how different types of figurative language affected the model's judgment. The context (not provided in this table) was fixed across both settings.

Literal Setting	Figurative Setting
Question: aita for making my girlfriend pay half the bills?	**Question (Sarcasm):** aita for making my girlfriend pay half the bills, *like I'm holding her financially hostage?*
Response: ...he is refusing to pay all and wants her to pay half. That seems like a fair split...his income is 7x her, but we need to consider fairness here... I'd say no, he's not an asshole...I'd choose **2**.	**Response:**...He asks her to pay half the bills and acting like he's holding her financially hostage...he might be acting like he is holding her hostage. That might be controlling...So I'd rate him as 5 or 6...I'd choose **5**.
Question: aita for not wanting to risk my fiances' job over Pokemon?	**Question (Metaphor):** aita for refusing to *gamble* my fiance's job over a game of Pokemon?
Response: ...their reaction shows significant concern over security and trust, especially given the financial stakes and the fact that they feel their partner has made a reckless decision...the anger and sense of betrayal are understandable, but the disproportionate reaction, especially without discussing the issue calmly leans toward being an asshole in this context. So, I'd rate **5**.	**Response:** ...while the person's concerns about security and financial risk are valid especially given their IT job refusing to accept their partner's decision as a gamble on their career is a bit extreme...the tone of the post and the question suggests emotional overreaction and a lack of trust in the partner's judgment, which can be interpreted as an overbearing, controlling...I'd rate it **7**.
Question: aita for not listening to my parents once again and keeping the door open when I pray?	**Question (Hyperbole):** aita for not listening to my parents *for the millionth time* and keeping the door open when I pray?
Response: ...the situation involves a respectful but challenging conflict between religious practice and parental concerns about noise or distraction...the person is not being disrespectful...rather, they are explaining their need for an open door due to practical and emotional reasons...my rating is **3**	**Response:** The person is not definitely an asshole, but their constant refusal to close the door during prayer—especially when it causes discomfort to their parents—crosses a line in terms of respect for their family's boundaries, even if the intent is rooted in personal faith and mental well-being...my rating is **5**.

sensitivity to pragmatic meaning, but failure to separate speaker framing from evidence relevant to the judgment. The final example demonstrates a similar effect with hyperbole. In the literal setting, the model (GPT-5.6 Sol) interprets the individual's behavior sympathetically. However, introducing the hyperbolic expression "*for the millionth time*" leads the model to emphasize repeated disobedience and disrespect toward the parents, resulting in a harsher judgment.

Across all examples, we observe a similar pattern. The model *either interprets the figurative expression too literally*, or *places a lot of emphasis on it when forming its judgment*. Under literal framing, the model rarely repeats or explicitly refers to the wording of the question when explaining its decision, grounding its reasoning primarily in the provided context. Under figurative framing, the model frequently repeats or directly quotes the figurative expression and incorporates it into its reasoning. This suggests that figurative framing can shift the model's attention toward the wording of the question itself, causing it to overemphasize the figurative expression rather than grounding its judgment in the underlying context.

Our qualitative analysis shows the model's response rely heavily on the figurative cue in the question. We now investigate this from a quantitative aspect by examining the attention allocated between the question and the context. Following prior work analyzing attention distributions over input tokens (Clark et al., 2019; Xiao et al., 2024), we measure the mean attention per available question tokens to (i) preceding question tokens, (ii) context tokens. For question span Q and context span C, we compute $A_{Q\to Q}$ and $A_{Q\to C}$, where α_{ij} denotes the attention weight from question token i to token j, averaged across attention heads. We further compute *Lookback Ratio* (LR) (Chuang et al., 2024), which measures the relative attention allocated to different spans of the preceding sequence.

$$A_{Q\to Q} = \frac{1}{|Q|}\sum_{i\in Q}\sum_{\substack{j\in Q\\ j<i}}\alpha_{ij}, \qquad A_{Q\to C} = \frac{1}{|Q|}\sum_{i\in Q}\sum_{j\in C}\alpha_{ij}, \qquad LR_i = \frac{A_{Q\to Q}}{A_{Q\to Q}+A_{Q\to C}}, \quad (4)$$

$A_{Q\to Q} > A_{Q\to C}$ and higher LR_i indicates the model allocates relatively more attention to the question wording than to the underlying context. Because these measurements can be affected by span length, we impose a strict length-matching constraint and retain only literal-figurative pairs with closely matched token counts.

Table 6: Attention patterns for literal (Lit) and figurative (Fig) questions, comparing attention from question to context tokens ($A_{Q\to C}$), to preceding question tokens ($A_{Q\to Q}$), and lookback ratio (LR). Across all model families, figurative framing decreases $A_{Q\to C}$ while increasing $A_{Q\to Q}$ and LR, indicating a systematic shift in attention from the task-relevant context toward the question itself.

	$A_{Q\to C}$		$A_{Q\to Q}$		LR	
	Lit	Fig	Lit	Fig	Lit	Fig
Llama family	15.86	14.09	7.81	9.92	33.18	41.64
Qwen family	40.75	36.15	14.56	19.24	26.49	35.08
Gemma family	44.41	40.29	18.68	24.98	29.89	38.77

Table 6 reveals a consistent attention pattern across all three model families. Under figurative framing, $A_{Q\to C}$ decreases while $A_{Q\to Q}$ increases, indicating that question tokens allocate less attention to the context and more attention to the wording of the question itself. This shift is further reflected in LR, which increases substantially for every model family. Together, these results provide quantitative evidence consistent with our qualitative observations. Figurative language changes how models allocate attention, *models look more at the question when it contains a figurative expression*, drawing attention away from the contextual evidence that should ground the decision.

Discussion: Figurative Language and Pragmatic Enrichment. Figurative language often adds emphasis, emotional intensity, speaker attitude, or evaluative force (Lakoff & Johnson, 1980). These pragmatic effects are constitutive properties of figurative communication rather than nuisance variables to be controlled away (Popa-Wyatt, 2020). To some readers, this might raise a potential concern that figuratively framed questions can carry greater pragmatic intensity than their literal counterparts. We view this not simply as a removable confound, but as an inherent property of figurative communication (Ortony, 1975). Requiring figurative and literal variants to be matched in pragmatism would remove precisely the linguistic phenomenon being studied. In morally charged moments, such pragmatic and expressive language is common. We therefore ask whether models can distinguish pragmatic information about how the user frames the question through figurative language, from the factual information about what occurred in the scenario. Additionally, a key distinction is that figurative framing can alter emphasis, vividness, or speaker attitude while preserving the core event, whereas an equivalent literal pragmatism may require introducing new factual content. For example, "*I have told you so many times*" may be rewritten hyperbolically as "*I have told you a million times*." Readers readily interpret *a million times* as an exaggerated way of conveying repeated telling. By contrast, a literal attempt to reproduce the same force, "*I have mentioned this repeatedly, to the point that I find it extremely frustrating that you have not listened*" explicitly introduces claims that are not necessarily entailed by the original statement. Thus, the additional pragmatic force that figurative language may introduce naturally coexists with preservation of the underlying event, being an inherent linguistic trait that enables expressive communication, whereas pragmatic enrichment of non-figurative forms may change the information being asserted.

6 Conclusion

We evaluate whether LLM moral judgments remain stable when the underlying evidence is unchanged but the question is expressed figuratively. Across a wide range of models, figurative framing consistently perturbs judgments, and these shifts are systematically harsher with open-source models being more prone to decision flips. Humans rarely change their final decision, but alter the strength of their judgment with a bias towards harshness, a trait shared with closed-source LLMs. This effect persists under polarity reversal, indicating that it is not explained by sycophantic agreement. We further show that the phenomenon extends beyond moral reasoning to other domains, particular more subjective ones such as resume review. Our analysis suggests that models overemphasize figurative cues and rely less on the unchanged context that should ground the judgment. These findings reveal that LLM decisions can depend not only on what users say, but on how they say it. As LLMs are increasingly used in subjective and consequential evaluation settings, robustness to natural expressive variation is essential for ensuring consistent, fair, and reliable judgments.

AI USE STATEMENT

Generative AI tools were used for debugging experimental and data-processing code. They were also used during manuscript preparation for language editing and improving the clarity and organization of the paper. Additionally, we used AI to obtain constructive feedback on the manuscript and identify areas for improvement. We have not used generative AI tools for any experimental decisions, analyses, interpretations or conclusions to our results, references, or human annotation. We have reviewed and verified all AI-assisted content and take responsibility for the final content of this work.

ETHICS STATEMENT

This manuscript contains potentially offensive language that is necessary to accurately describe the datasets we used, as well as our experimental setting and task formulation. We retain such language where necessary for methodological clarity and do not intend its use to cause offense.

Our experiments include moral judgment scenarios derived from the publicly available Reddit AITA dataset. These posts may describe sensitive interpersonal situations and were originally written by individuals seeking judgments from an online community. We use these data solely to study how figurative framing affects model judgments and do not attempt to identify, profile, or make claims about individual authors. Our analysis is conducted at the aggregate level, and examples are presented only when necessary to illustrate the studied phenomena for the transparency of the readers.

Our goal is not to establish LLMs as reliable arbiters of morality, but rather to investigate whether their judgments remain stable under changes in figurative framing. Indeed, demonstrating sensitivity to figurative framing highlights a potential risk of relying on LLM-generated judgments in subjective or socially consequential settings, where differences in how users express the same underlying situation may lead to systematically different model responses. We therefore view our findings as evidence for the importance of evaluating linguistic robustness and potential framing biases before deploying LLMs in applications involving subjective, moral, or socially sensitive decisions.

We performed human annotation in our experiments. All annotators were informed that their annotations would be used for an academic research study and participated voluntarily with no compensation. Annotators were instructed to evaluate each scenario based only on the information provided and to avoid introducing assumptions based on personal or cultural expectations. Our objective is not to establish a universal notion of moral correctness, but to use aggregate human responses as a reference for understanding how figurative framing affects human judgments with respect to language models. The aggregate behavioral patterns observed in our study are also consistent with prior findings that linguistic and figurative framing can modulate the strength of human judgments, but not to the point of altering their final decision.

REPRODUCIBILITY STATEMENT

To facilitate reproducibility, we provide all code and data used in our experiments through an anonymous repository. The repository includes the original sentences, the generated literal and figurative variants, scripts for running model inference, and code for computing the evaluation metrics and reproducing the analyses reported in this paper. Details of the experimental setup, including the evaluated models, prompting strategies, generation settings, and evaluation procedures, are described in the main paper and appendix.

REFERENCES

Bushra Alhetelah and Irfan Ahmad. Measuring LLMs' sensitivity to paraphrased opinion prompts. In Jeremy Barnes, Valentin Barriere, Orphée De Clercq, Roman Klinger, Célia Nouri, Debora Nozza, and Pranaydeep Singh (eds.), *The Proceedings for the 15th Workshop on Computational Approaches to Subjectivity, Sentiment Social Media Analysis (WASSA 2026)*, 2026. URL https://aclanthology.org/2026.wassa-1.5/.

Anthropic. Claude opus 5 system card. Technical report, Anthropic, 2026. URL https://www-cdn.anthropic.com/c5fbac3f0b1280a933ebd26d3cb8bb9f5bdeaf48/Claude%20Opus%205%20System%20Card.pdf.

Sergio Barbosa and William Jiménez-Leal. It's not right but it's permitted: Wording effects in moral judgement. *Judgment and Decision Making*, 2017.

Tuhin Chakrabarty, Xurui Zhang, Smaranda Muresan, and Nanyun Peng. MERMAID: Metaphor generation with symbolism and discriminative decoding. In Kristina Toutanova, Anna Rumshisky, Luke Zettlemoyer, Dilek Hakkani-Tur, Iz Beltagy, Steven Bethard, Ryan Cotterell, Tanmoy Chakraborty, and Yichao Zhou (eds.), *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics (NAACL 2021)*, 2021. URL https://aclanthology.org/2021.naacl-main.336/.

Anwoy Chatterjee, H S V N S Kowndinya Renduchintala, Sumit Bhatia, and Tanmoy Chakraborty. POSIX: A prompt sensitivity index for large language models. In Yaser Al-Onaizan, Mohit Bansal, and Yun-Nung Chen (eds.), *Findings of the Association for Computational Linguistics: EMNLP 2024 (Findings of EMNLP 2024)*, 2024. URL https://aclanthology.org/2024.findings-emnlp.852/.

Yung-Sung Chuang, Linlu Qiu, Cheng-Yu Hsieh, Ranjay Krishna, Yoon Kim, and James R. Glass. Lookback lens: Detecting and mitigating contextual hallucinations in large language models using only attention maps. In Yaser Al-Onaizan, Mohit Bansal, and Yun-Nung Chen (eds.), *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing (EMNLP 2024)*, 2024. URL https://aclanthology.org/2024.emnlp-main.84/.

Claudia Claridge. *Hyperbole in English: A Corpus-based Study of Exaggeration*. 2011.

Kevin Clark, Urvashi Khandelwal, Omer Levy, and Christopher D. Manning. What does BERT look at? an analysis of BERT's attention. In Tal Linzen, Grzegorz Chrupała, Yonatan Belinkov, and Dieuwke Hupkes (eds.), *Proceedings of the 2019 ACL Workshop BlackboxNLP: Analyzing and Interpreting Neural Networks for NLP*, 2019. URL https://aclanthology.org/W19-4828/.

Karl Cobbe, Vineet Kosaraju, Mohammad Bavarian, Mark Chen, Heewoo Jun, Lukasz Kaiser, Matthias Plappert, Jerry Tworek, Jacob Hilton, Reiichiro Nakano, Christopher Hesse, and John Schulman. Training verifiers to solve math word problems, 2021. URL https://arxiv.org/abs/2110.14168.

Shreya Desai, Kate Bailey, and Ruth Filik. The role of hyperbole in conveying emotionality: The case of victim speech. *Cognition and Emotion*, 2025.

Maxwell Forbes, Jena D. Hwang, Vered Shwartz, Maarten Sap, and Yejin Choi. Social chemistry 101: Learning to reason about social and moral norms. In Bonnie Webber, Trevor Cohn, Yulan He, and Yang Liu (eds.), *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP 2020)*, 2020. URL https://aclanthology.org/2020.emnlp-main.48/.

Jie Gao, Simret Araya Gebreegziabher, Kenny Tsu Wei Choo, Toby Jia-Jun Li, Simon Tangi Perrault, and Thomas W Malone. A taxonomy for human-llm interaction modes: An initial exploration. In *Extended Abstracts of the CHI Conference on Human Factors in Computing Systems*, CHI '24. ACM, 2024. URL http://dx.doi.org/10.1145/3613905.3650786.

Gemini Team Google. Gemini: A family of highly capable multimodal models. *arXiv preprint arXiv:2312.11805*, 2023.

Gemma Team, Aishwarya Kamath, Johan Ferret, Shreya Pathak, Nino Vieillard, Ramona Merhej, Sarah Perrin, Tatiana Matejovicova, Alexandre Ramé, Morgane Rivière, Louis Rouillard, et al. Gemma 3 technical report, 2025. URL https://arxiv.org/abs/2503.19786.

Saptarshi Ghosh and Tianyu Jiang. MetFuse: Figurative fusion between metonymy and metaphor. In Maria Liakata, Viviane P. Moreira, Jiajun Zhang, and David Jurgens (eds.), *Proceedings of the 64th Annual Meeting of the Association for Computational Linguistics (ACL 2026)*, 2026. URL https://aclanthology.org/2026.acl-long.700/.

Jr. Gibbs, Raymond W., John S. Leggitt, and Elizabeth A. Turner. What's special about figurative language in emotional communication? In Susan R. Fussell (ed.), *The Verbal Communication of Emotions: Interdisciplinary Perspectives*. Lawrence Erlbaum Associates Publishers, 2002.

Aaron Grattafiori, Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Alex Vaughan, et al. The llama 3 herd of models, 2024. URL https://arxiv.org/abs/2407.21783.

Daya Guo, Dejian Yang, Haowei Zhang, Junxiao Song, Peiyi Wang, Qihao Zhu, Runxin Xu, Ruoyu Zhang, Shirong Ma, Xiao Bi, Zhang, et al. Deepseek-r1 incentivizes reasoning in llms through reinforcement learning. *Nature*, 2025. URL http://dx.doi.org/10.1038/s41586-025-09422-z.

Ahmed Heakl, Youssef Mohamed, Noran Mohamed, Aly Elsharkawy, and Ahmed Zaky. Resuméatlas: Revisiting resume classification with large-scale datasets and large language models. *Procedia Computer Science*, 2024. URL https://www.sciencedirect.com/science/article/pii/S1877050924029909. 6th International Conference on AI in Computational Linguistics.

Dan Hendrycks, Collin Burns, Steven Basart, Andrew Critch, Jerry Li, Dawn Song, and Jacob Steinhardt. Aligning {ai} with shared human values. In *International Conference on Learning Representations (ICLR 2021)*, 2021. URL https://openreview.net/forum?id=dNy_RKzJacY.

Kevin J. Holmes, Lena Kassin, and Stephen Flusberg. When does suggestive language shape memory for car accidents? assessing the role of elaboration and pragmatics in a classic framing effect. In *Proceedings of the 46th Annual Meeting of the Cognitive Science Society*, 2024.

Andong Hua, Kenan Tang, Chenhe Gu, Jindong Gu, Eric Wong, and Yao Qin. Flaw or artifact? rethinking prompt sensitivity in evaluating LLMs. In Christos Christodoulopoulos, Tanmoy Chakraborty, Carolyn Rose, and Violet Peng (eds.), *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing (EMNLP 2025)*, 2025. URL https://aclanthology.org/2025.emnlp-main.1006/.

Li Huang, Francesca Gino, and Adam D. Galinsky. The highest form of intelligence: Sarcasm increases creativity for both expressers and recipients. *Organizational Behavior and Human Decision Processes*, 2015. doi: https://doi.org/10.1016/j.obhdp.2015.07.001. URL https://www.sciencedirect.com/science/article/pii/S074959781500076X.

Shomik Jain, Charlotte Park, Matt Viana, Ashia Wilson, and Dana Calacci. Interaction context often increases sycophancy in llms. In *Proceedings of the 2026 CHI Conference on Human Factors in Computing Systems*. Association for Computing Machinery, 2026. URL https://doi.org/10.1145/3772318.3791915.

Lynn Katz and Andrew Ortony. Metaphorical uses of language in the expression of emotions. *Metaphor and Symbolic Activity*, 1987.

George Lakoff and Mark Johnson. *Metaphors We Live By*. University of Chicago Press, Chicago, 1980.

George Lakoff and Zoltán Kövecses. The cognitive model of anger inherent in american english. In Dorothy Holland and Naomi Quinn (eds.), *Cultural Models in Language and Thought*. Cambridge University Press, 1987.

Jiayi Li, Sanjana Menon, Brett Frischmann, Shomir Wilson, and Sarah Rajtmajer. Affective context amplifies sycophancy in llm responses, 2026. URL https://arxiv.org/abs/2608.21242.

Stephanie Lin, Jacob Hilton, and Owain Evans. TruthfulQA: Measuring how models mimic human falsehoods. In Smaranda Muresan, Preslav Nakov, and Aline Villavicencio (eds.), *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (ACL 2022)*, 2022. URL https://aclanthology.org/2022.acl-long.229/.

Elizabeth F. Loftus and John C. Palmer. Reconstruction of automobile destruction: An example of the interaction between language and memory. *Journal of Verbal Learning and Verbal Behavior*, 1974. doi: https://doi.org/10.1016/S0022-5371(74)80011-3. URL https://www.sciencedirect.com/science/article/pii/S0022537174800113.

Soyoung Oh and Vera Demberg. Robustness of large language models in moral judgements. *Royal Society Open Science*, 12(4):241229, 2025. doi: 10.1098/rsos.241229.

OpenAI, :, Sandhini Agarwal, Lama Ahmad, Jason Ai, Sam Altman, Andy Applebaum, Edwin Arbus, Rahul K. Arora, Yu Bai, Bowen Baker, et al. gpt-oss-120b & gpt-oss-20b model card, 2025. URL https://arxiv.org/abs/2508.10925.

Andrew Ortony. Why metaphors are necessary and not just nice. *Educational Theory*, 1975. doi: https://doi.org/10.1111/j.1741-5446.1975.tb00666.x. URL https://onlinelibrary.wiley.com/doi/abs/10.1111/j.1741-5446.1975.tb00666.x.

Andrew Ortony and Lynn Fainsilber. The role of metaphors in descriptions of emotions. In Yorick Wilks (ed.), *Theoretical Issues in Natural Language Processing 3*, 1987. URL https://aclanthology.org/T87-1037/.

Ross E. O'Hara, Walter Sinnott-Armstrong, and Nicholas A. Sinnott-Armstrong. Wording effects in moral judgments. *Judgment and Decision Making*, pp. 547–554, 2010.

Samarth P, Motamarri Sai Sathvik, Vyoman Jain, and Shiva Golugula. Reviewarena: A large-scale cross-conference dataset and benchmark for LLM peer review. In *ICML 2026 AI for Science Workshop*, 2026. URL https://openreview.net/forum?id=yugEO52gkR.

Bumjin Park, Jinsil Lee, and Jaesik Choi. Deontological keyword bias: The impact of modal expressions on normative judgments of language models. In Wanxiang Che, Joyce Nabende, Ekaterina Shutova, and Mohammad Taher Pilehvar (eds.), *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (ACL 2025)*, 2025. URL https://aclanthology.org/2025.acl-long.360/.

Ethan Perez, Sam Ringer, Kamile Lukosiute, Karina Nguyen, Edwin Chen, Scott Heiner, Craig Pettit, Catherine Olsson, Sandipan Kundu, Saurav Kadavath, et al. Discovering language model behaviors with model-written evaluations. In Anna Rogers, Jordan Boyd-Graber, and Naoaki Okazaki (eds.), *Findings of the Association for Computational Linguistics: ACL 2023 (Findings of ACL 2023)*, 2023. URL https://aclanthology.org/2023.findings-acl.847/.

Lewis Petrinovich and Patricia O'Neill. Influence of wording and framing effects on moral intuitions. *Ethology and Sociobiology*, 1996. URL https://www.sciencedirect.com/science/article/pii/0162309596000416.

Mihaela Popa-Wyatt. Expressing affect with hyperbole and hyperbolic figures: Mind the gap. In John Barnden and Andrew Gargett (eds.), *Producing Figurative Expression: Theoretical, Experimental and Practical Perspectives*, volume 10 of *Figurative Thought and Language*, pp. 449–468. John Benjamins, 2020. doi: 10.1075/ftl.10.16pop. URL https://doi.org/10.1075/ftl.10.16pop.

Ruiyang Qin, Qingzhuo Wang, Tian Wang, Zhihua Wei, and Wen Shen. Evaluating and explaining prompt sensitivity of LLMs using interactions. In *Forty-third International Conference on Machine Learning*, 2026. URL https://openreview.net/forum?id=GsdFM8qnav.

Nils Reimers and Iryna Gurevych. Sentence-BERT: Sentence embeddings using Siamese BERT-networks. In Kentaro Inui, Jing Jiang, Vincent Ng, and Xiaojun Wan (eds.), *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP 2019)*, 2019. URL https://aclanthology.org/D19-1410/.

David Rein, Betty Li Hou, Asa Cooper Stickland, Jackson Petty, Richard Yuanzhe Pang, Julien Dirani, Julian Michael, and Samuel R. Bowman. GPQA: A graduate-level google-proof q&a benchmark. In *First Conference on Language Modeling*, 2024. URL https://openreview.net/forum?id=Ti67584b98.

Mohammad Saim and Tianyu Jiang. Do emotions influence moral judgment in large language models? In Maria Liakata, Viviane P. Moreira, Jiajun Zhang, and David Jurgens (eds.), *Findings of the Association for Computational Linguistics: ACL 2026 (Findings of ACL 2026)*, 2026. URL https://aclanthology.org/2026.findings-acl.1346/.

Abel Salinas and Fred Morstatter. The butterfly effect of altering prompts: How small changes and jailbreaks affect large language model performance. In Lun-Wei Ku, Andre Martins, and Vivek Srikumar (eds.), *Findings of the Association for Computational Linguistics: ACL 2024 (Findings of ACL 2024)*, 2024. URL https://aclanthology.org/2024.findings-acl.275/.

Nino Scherrer, Claudia Shi, Amir Feder, and David Blei. Evaluating the moral beliefs encoded in LLMs. In *Advances in Neural Information Processing Systems (NeurIPS 2023)*, 2023. doi: 10.52202/075280-2256. URL https://proceedings.neurips.cc/paper_files/paper/2023/file/a2cf225ba392627529efef14dc857e22-Paper-Conference.pdf.

Melanie Sclar, Yejin Choi, Yulia Tsvetkov, and Alane Suhr. Quantifying language models'sensitivity to spurious features in prompt design or: How i learned to start worrying about prompt formatting. In B. Kim, Y. Yue, S. Chaudhuri, K. Fragkiadaki, M. Khan, and Y. Sun (eds.), *International Conference on Learning Representations*, volume 2024, 2024. URL https://proceedings.iclr.cc/paper_files/paper/2024/file/6c0e99d736da621403018ca7b32b1a4d-Paper-Conference.pdf.

Mrinank Sharma, Meg Tong, Tomasz Korbak, David Kristjanson Duvenaud, Amanda Askell, Samuel R. Bowman, Newton Cheng, Esin Durmus, Zac Hatfield-Dodds, Scott Johnston, Shauna Kravec, Tim Maxwell, Sam McCandlish, Kamal Ndousse, Oliver Rausch, Nicholas Schiefer, Da Yan, Miranda Zhang, and Ethan Perez. Towards understanding sycophancy in language models. *ArXiv*, abs/2310.13548, 2023. URL https://api.semanticscholar.org/CorpusID:264405698.

Aaditya Singh, Adam Fry, Adam Perelman, Adam Tart, Adi Ganesh, Ahmed El-Kishky, Aidan McLaughlin, Aiden Low, AJ Ostrow, Akhila Ananthram, et al. Openai gpt-5 system card, 2026. URL https://arxiv.org/abs/2601.03267.

Paul H. Thibodeau and Lera Boroditsky. Metaphors we think with: The role of metaphor in reasoning. *PLOS ONE*, 2011. URL https://doi.org/10.1371/journal.pone.0016782.

Paul H. Thibodeau and Lera Boroditsky. Natural language metaphors covertly influence reasoning. *PLOS ONE*, 2013.

Tom van Nuenen and Pratik S. Sachdeva. The fragility of moral judgment in large language models. In *Proceedings of the 2026 ACM Conference on Fairness, Accountability, and Transparency*, FAccT '26. Association for Computing Machinery, 2026. URL https://doi.org/10.1145/3805689.3806447.

Daniel Vennemeyer, Phan Anh Duong, Tiffany Zhan, and Tianyu Jiang. Sycophancy is not one thing: Causal separation of sycophantic behaviors in llms, 2026. URL https://arxiv.org/abs/2509.21305.

Frank Wilcoxon. Individual comparisons by ranking methods. *Biometrics Bulletin*, 1945. ISSN 00994987. URL http://www.jstor.org/stable/3001968.

Guangxuan Xiao, Yuandong Tian, Beidi Chen, Song Han, and Mike Lewis. Efficient streaming language models with attention sinks. In *The Twelfth International Conference on Learning Representations (ICLR 2024)*, 2024. URL https://openreview.net/forum?id=NG7sS51zVF.

An Yang, Anfeng Li, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chang Gao, Chengen Huang, Chenxu Lv, et al. Qwen3 technical report, 2025. URL https://arxiv.org/abs/2505.09388.

Kaijie Zhu, Jindong Wang, Jiaheng Zhou, Zichen Wang, Hao Chen, Yidong Wang, Linyi Yang, Wei Ye, Yue Zhang, Neil Gong, and Xing Xie. Promptrobust: Towards evaluating the robustness of large language models on adversarial prompts. In *Proceedings of the 1st ACM Workshop on Large AI Systems and Models with Privacy and Safety Analysis*. Association for Computing Machinery, 2024. URL https://doi.org/10.1145/3689217.3690621.

A FIGURATIVE CATEGORIES.

Figurative language encompasses a wide range of linguistic phenomena, each serving a different communicative purpose. In this work, we focus on four representative forms of figurative framing: **metaphor**, **hyperbole**, **simile**, and **sarcasm**. Together, these categories capture diverse ways in which speakers enrich language while preserving the underlying intent of a statement or question.

Metaphors describe one concept in terms of another concept from a different domain without using an explicit comparison. For example, in "*This assignment is an uphill battle*," *uphill battle* is conceptually mapped onto "a difficult task." Hyperboles intentionally exaggerate for emphasis to a degree beyond literal plausibility. For example, in "*I have a mountain of reasons to believe this*," the speaker's high level of confidence is expressed through deliberate exaggeration. A central property of hyperbole is that it may provide pragmatic enrichment through such exaggeration, which is why hyperboles are quite common in day-to-day communication as they provide a more expressive way of communicating (for example, "*This bag weighs a ton*." instead of "*This bag is heavy*"). Importantly, the underlying meaning remains unchanged. Similes make explicit comparisons using words such as *like* or *as*. For example, "*He is very tall*" can be expressed as "*He is as tall as a tree*." Sarcasm communicates meaning by expressing the opposite of the intended sentiment, often relying on context and pragmatic cues. For example, in "*Great job on failing the exam!*," the speaker expresses disappointment by using the apparently positive phrase *great job*.

We select these four categories because they are among the most commonly used forms of figurative language used in daily communication. They are especially relevant in morally charged situations, where speakers often use figurative expressions to emphasize praise or blame, communicate emotional intensity, and express their attitudes toward another person's behavior. Importantly, these categories span distinct mechanisms for modifying how information is presented: conceptual mapping in metaphor, exaggeration in hyperbole, explicit comparison in simile, and pragmatic inversion in sarcasm. This diversity allows us to examine whether robustness to figurative framing is specific to a particular phenomenon or represents a broader challenge for LLMs.

B DETAILS ON GENERATING FIGURATIVE VARIANTS

Checking the Extracted Questions for Figurative Expressions. When extracting the 1,000 AITA samples, we include an additional filtering step to ensure that the selected questions are literal. We provide each question to GPT-5-mini and ask it to determine whether the question contains a figurative expression. If the model identifies one, we discard the corresponding sample. This filtering step reduces the likelihood that the original AITA questions already contain figurative language.

Model Selection for Figurative Generation. We evaluated several models for generating figurative variants, including GPT-5, GPT-5-mini, GPT-5-nano, Qwen3-30B, and Llama-3.1-70B. The authors manually inspected a sample of generations from each model to assess their quality and faithfulness. We found that the open-source models more often produced unnatural rewrites, unnecessarily lengthened the questions, or introduced and removed information from the original question. GPT-5-nano exhibited similar issues with naturalness, although less frequently than the open-source models. In contrast, GPT-5 and GPT-5-mini consistently produced natural and semantically faithful rewrites, with no noticeable difference in quality between them in our manual inspection. We therefore use GPT-5-mini to generate the figurative variants throughout our experiments due to its good performance and lower generation cost.

Length of Generated Figurative Questions. When generating figurative variants, we constrain the rewrite to remain reasonably close in length to the original question and explicitly instruct the rewriting model to follow the structure of the original question. However, we do not impose a strict token or word-level length cap, since expressing the same question figuratively can naturally require additional words. This is particularly true for similes, which introduce an explicit comparison using lexical constructions such as *like* or *as*. For example, "*Was I wrong for ignoring her?*" may become "*Was I wrong for ignoring her like she was invisible?*" Consequently, requiring identical lengths could either restrict the range of natural figurative expressions or force unnatural rewrites. Although prior work has shown that prompt length can influence model responses, length is unlikely to explain our observed effects since the rewritten question constitutes only a small portion of the complete prompt, which contains the unchanged scenario context whose length varies substantially across instances.

Thus, small length differences between literal and figurative questions are minor relative to the overall variation in prompt length, while allowing the figurative variants to remain natural and semantically faithful.

C Agreement with Human Annotated Label

Table 7: Percentage of agreement with human-annotated gold labels. Positive = user is an asshole (harsh); negative = user is not an asshole (lenient).

Models	Positive Label		Negative Label	
	Literal	Figurative	Literal	Figurative
Qwen	25.2	**34.0**	**48.0**	45.2
LLama	64.8	**71.2**	**42.0**	39.6
Gemma	66.4	**70.0**	**70.0**	66.4
Claude	63.1	**63.6**	**96.0**	94.4
Gemini	79.2	**81.2**	**98.0**	97.6
GPT	75.6	**82.0**	**90.8**	88.4

Table 7 reports model agreement with the human-annotated gold labels under literal and figurative conditions. For samples with a positive label, where the human judgment identifies the user as being in the wrong, agreement consistently increases under figurative framing across all models. This is consistent with our main finding that figurative language shifts model judgments in a harsher direction. Conversely, for negative-label samples, where the human judgment is more lenient, agreement generally decreases under figurative framing. Thus, the harsher shift induced by figurative language moves model judgments closer to human labels when the gold judgment is harsh, but farther away when the gold judgment is lenient. We also observe substantially higher overall agreement for stronger state-of-the-art models, particularly Gemini and GPT, suggesting that these models more closely align with human judgments irrespective of the framing condition.

D Human Evaluation Details

We conduct a human evaluation to examine whether figurative framing affects human moral judgments in a manner similar to LLMs. Annotators are presented with the same moral scenarios used in our model evaluation, consisting of a context describing an interpersonal situation followed by a question asking them to judge the actions of the user. Following the LLM evaluation setup, annotators provide their judgment on the same 7-point Likert scale.

Annotators and Sample Assignment. We recruited five volunteer annotators, all of whom are fluent English speakers in academia, and their participation was entirely voluntary. None of the authors participated in this annotation task, in order to avoid potential bias arising from prior knowledge of the study's hypotheses or expected outcomes. We select the same 100 samples from each of the four figurative categories for human evaluation and distribute them among the five annotators. Figure 7 provides the annotation guidelines for the annotators.

Annotation Procedure. The evaluation is conducted in two rounds separated by a two-month interval. In each round, annotators are presented with a mixture of literal and figuratively framed questions. For each scenario, an annotator evaluates only one variant in a given round: either the literal question or its corresponding figurative question. In the second round, the same annotator evaluates the opposite variant of that scenario. Thus, if an annotator encounters the literal version of a sample in the first round, they encounter its figurative counterpart in the second round, and vice versa. This mixed design ensures that the framing condition is not confounded with the annotation round and reduces the possibility that annotators become aware of the linguistic manipulation.

The two rounds use identical instructions, contexts, Likert scales, and evaluation procedures; the only difference is whether the question is presented in its literal or figurative form. Annotators are not informed that the purpose of the study is to investigate figurative language. We intentionally withhold this information to avoid drawing attention to the linguistic manipulation and potentially biasing their judgments. Annotators are instead instructed simply to judge the described situation based on the

information provided. The temporal separation of 2 months is intended to reduce the likelihood that annotators remember their previous ratings and simply reproduce them in the second round.

Human Agreement. Because moral judgments are inherently subjective, we do not expect annotators to necessarily select the same fine-grained Likert score. Nevertheless, we evaluate the consistency of human judgments by assigning an overlapping set of 100 samples to two annotators. For the exact 1-7 Likert ratings, the two annotators agree on 62% of the literal examples and 65% of the figurative examples. We additionally measure agreement after mapping the Likert ratings to the three decision categories used in our evaluation: *negative (not an asshole)* (1-3), *neutral* (4), and *positive (is an asshole)* (5-7). Under this categorization, agreement is substantially higher, reaching 88% for the literal questions and 91% for the figurative questions. These results indicate that although annotators can differ in the precise strength of their moral judgments, they largely agree on the underlying decision.

E Other Domain Experiment: Details

In this section, we provide additional details on our experiments across mathematical reasoning, scientific reasoning, truthfulness, scientific peer review, and resume evaluation. LLM evaluations are commonly used for these domains in the real world. We are interested in how the LLMs evaluate, not how the LLM solves the query. For that, we frame every question as an evaluation problem. We select these domains because, similar to moral reasoning, they involve inherently subjective judgments for which there is often no single objectively correct evaluation. Moreover, these represent practical settings in which LLMs are increasingly used to assist with assessment and decision-making, making robustness to seemingly innocuous linguistic variation particularly important. Together, these domains allow us to examine whether the effects of figurative framing extend beyond moral and social judgments to other subjective and practically relevant evaluation tasks. We discuss in detail below how we formulate the task from each domain:

Mathematical Reasoning. For mathematical reasoning evaluation, we use the GSM8K (Cobbe et al., 2021) which contains mathematical questions of varying difficulty, and provides step-by-step solutions to those problems, which provides us with the perfect setup. We sample 500 questions with their solution. We then provide the mathematical question followed by the proposed step-by-step solution as the context. For each essay, we construct five literal and five figurative evaluation questions that ask the LLM to evaluate the proposed step-by-step solution to the question. Evaluation questions are provided in Table 14. Thus, the resulting setup takes the form: *[Math question] [Step-by-step solution] [Evaluation question]*. Each prompt is evaluated independently. The model is asked to rate the solution on a scale of 1-10. We do this to make the evaluation more inclined to real-world. Using a Likert scale for such domains would be unnatural since Likert scale is about agree/disagree and the question needs to be formed accordingly. Here, we are asking the model to evaluate, so we do a usual 1-10 rating. We then measure the direction of perturbation at the essay level: a lower mean figurative score than literal score is considered a harsh shift, while a higher mean figurative score is considered a lenient shift.

Scientific Reasoning. For scientific reasoning evaluation, we use GPQA (Rein et al., 2024), which contains grad level physics, chemistry and biology questions and their answers with explanation, giving us a similar experimental setup to GSM8K. We sample 500 questions and use the question followed by the proposed solution and explanation as the context. We similarly use five literal and five figurative questions and ask the LLM to evaluate the answer and explanation on a 1-10 scale. The evaluation questions are the same as mathematical reasoning, provided in Table 14. Thus, the resulting setup takes the form: *[Question] [Solution] [Explanation] [Evaluation question]*.

Truthfulness. For truthfulness, we use TruthfulQA (Lin et al., 2022), which contains questions from different categories, a correct answer and an incorrect answer. We sample 500 samples while ensuring every category is represented, and set a lower limit on word count for the answer, since we are interested in longer explanations. We sample correct and incorrect answers equally 50-50 to ensure diversity during judgment. We provide the model the solution and the proposed answer, and ask the LLM to evaluate the response on a scale of 1-10. The literal and figurative evaluation questions for truthfulness are provided in Table 15. The resulting setup takes the form: *[Question] [Correct/Incorrect answer] [Evaluation question]*.

Scientific Peer Review. For scientific peer-review evaluation, we use ReviewArena (P et al., 2026), which contains scientific papers and reviews from major machine learning and NLP venues. We sample 100 papers each from NeurIPS, ICLR, ICML, EMNLP, and COLM, resulting in 500 papers in total, and use only their titles and abstracts for evaluation. We focus on abstracts rather than full papers for three reasons: (i) full papers vary substantially in context length, introducing an additional source of variation across instances, (ii) papers contain figures, tables, and other visual information that may not be represented or interpreted consistently across the evaluated models, and (iii) evaluating full papers across multiple prompts and models would incur substantially greater computational cost. We ask the model to evaluate whether the research described by the abstract is of sufficient quality for acceptance at a top-tier conference, without specifying the original conference to avoid introducing venue-specific bias. Similar to before, we use five literal and five figurative evaluation prompts, provided in Table 16, and ask the model to rate on a scale of 1-10. The setup therefore takes the form: *[Paper Title] [Abstract] [Evaluation question]*. We similarly compute the mean across the five literal and five figurative judgments and report the corresponding harsh and lenient perturbation rates.

Resume Evaluation. For resume evaluation, we use ResuméAtlas (Heakl et al., 2024), a large-scale dataset containing over 13,000 resumes spanning 43 occupational categories. Each resume is associated with an occupational category, which we use as the target job category for evaluation. We provide the LLM with the job category and the candidate's resume and ask it to assess the candidate's suitability for that category. As in the previous domains, we construct five literal and five figurative evaluation prompts, shown in Table 17. The setup takes the form: *[Job Category] [Candidate Resume] [Evaluation question]*.

Table 8: LLM evaluation for figurative questions on a scale of 1-10. The change from literal question is given in the parenthesis. All models tend towards harsher evaluation for figurative questions, indicated by red shade.

Models	Mathematical Reasoning	Scientific Reasoning	Truthfulness	Scientific Peer Review	Resume Review
Qwen3-4B	8.31 (-0.12)	6.01 (-0.06)	5.37 (-0.10)	8.89 (-0.08)	7.60 -0.33
Llama-3.1-8B	8.79 (-0.07)	8.66 (-0.07)	7.41 (-0.17)	8.51 (-0.28)	7.81 -0.14
Gemma-3-12B	9.59 (-0.05)	6.65 (-0.06)	6.18 (-0.01)	8.21 (+0.02)	7.94 -0.30
Qwen3-30B	9.05 (-0.07)	6.80 (-0.03)	6.02 (-0.08)	8.24 -0.08	8.18 -0.23
Llama-3.1-70B	8.56 (-0.08)	8.41 (-0.05)	7.31 (-0.12)	8.64 (-0.14)	7.26 -0.28
Gemma-3-27B	9.61 (-0.06)	8.01 (+0.09)	7.13 (-0.08)	8.54 (-0.05)	7.56 -0.06
Claude-Opus-5	9.79 (-0.08)	6.71 (-0.01)	5.64 (-0.05)	7.38 -0.17	6.81 -0.19
Gemini-3.7-Flash	8.82 (-0.02)	5.25 (-0.07)	6.02 (-0.03)	8.18 -0.08	7.71 -0.12
GPT-5.6 Sol	9.82 (-0.04)	6.52 (-0.02)	5.38 (-0.08)	8.18 -0.13	7.54 -0.04

Mean score. Table 8 presents the mean scores under figurative prompts on a scale of 1-10. The value in parentheses denotes the difference between the mean figurative and literal scores, with negative values indicating harsher judgments under figurative framing. Consistent with the perturbation results in the main paper, figurative prompts generally yield lower scores. The absolute mean differences are relatively modest for two reasons. First, harsh and lenient shifts partially cancel when averaged across instances. Second, the effect is typically subtle, nudging judgments by a smaller amount while still occurring consistently across many examples. This table therefore provides a complementary view of the perturbation analysis in the main body.

F Effect of Reasoning on Figurative Robustness

We additionally investigate whether extended reasoning affects model robustness to figurative framing. For this experiment, we use two different types of models: i) those that have reasoning within the same model family, and ii) those that are inherently designed for reasoning. For the first type, we use the reasoning variant of Qwen3-4B in Qwen3-4B-Thinking and set reasoning to be "high" in GPT-5.6 Sol and Claude-Opus-5. This enables a direct comparison with their non/low-thinking counterpart. For the second type, we use models inherently trained for reasoning, including DeepSeek-R1-Distill-Qwen-1B, DeepSeek-R1-Distill-Llama-8B (Guo et al., 2025) and GPT-OSS-20B (OpenAI et al.,

2025). We then compare their perturbation and flip rates with the average of the non-reasoning models.

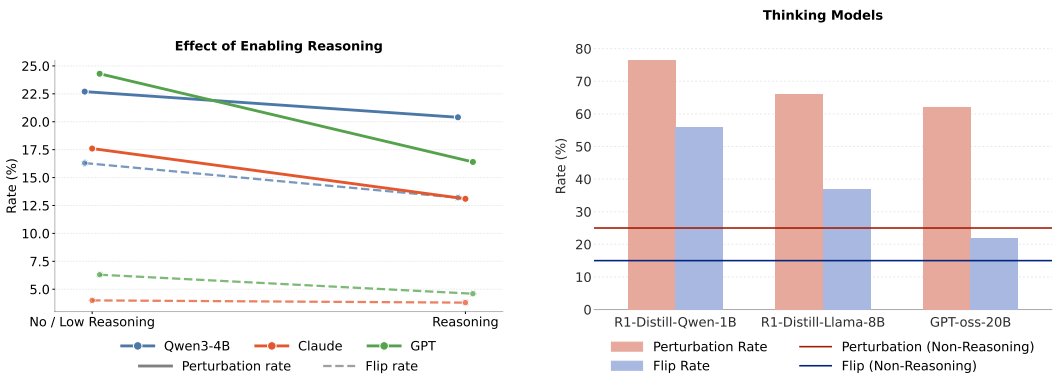


Figure 4: Effect of reasoning on robustness to figurative framing. Enabling reasoning in models slightly decreases perturbation and flip rates, indicating robustness increases by a little bit, although not by a lot. However, models inherently trained for reasoning show excessively high perturbation and flip rates.

Results. Figure 4 presents the results, revealing an interesting distinction between enabling reasoning within the same model family and evaluating models that are inherently designed for reasoning. For models where comparable standard and reasoning configurations are available, increased reasoning generally reduces sensitivity to figurative framing. For example, Qwen3-4B-Thinking exhibits a lower perturbation and flip rate than its standard Qwen3-4B counterpart. We observe a similar reduction of perturbation rates for GPT and Claude by increasing reasoning effort. The drop in flip rate is modest, as their flip rates are already relatively low under the standard configuration. These within-model comparisons suggest that additional reasoning may improve robustness to surface-level changes in how the same underlying question is expressed.

However, this pattern does not imply that reasoning-oriented models are inherently robust to figurative framing. When evaluating standalone reasoning models, such as R1-Distill-Qwen-1B, R1-Distill-Llama-8B, and GPT-OSS-20B, we observe substantial sensitivity. Perturbation rates exceed 60%, a large jump compared to the non-reasoning baseline. GPT-OSS shows the lowest flip rate, but above the non-reasoning baseline, while the other two show very high flips. Thus, reasoning models can still substantially alter their judgments when the underlying scenario remains fixed but the question is expressed figuratively.

Taken together, these results suggest a more nuanced relationship between reasoning and figurative robustness. Increasing reasoning effort for the same or closely related model appears to reduce the effect, but being a reasoning-oriented model by itself does not guarantee robustness. This distinction also cautions against attributing differences between standard and reasoning models solely to the presence of explicit reasoning, as standalone reasoning models may differ substantially in model size, training, and post-training objectives.

G Effect of Temperature on Figurative Robustness

We further examine whether decoding stochasticity contributes to the observed sensitivity to figurative framing by varying the sampling temperature from 0.0 to 1.0 for three open-source models. We use 200 samples for this experiment, since running across 1,000 samples for all temperature incurs substantially high computational cost. As shown in Figure 5, higher temperatures generally lead to increased perturbation and flip rates, although the magnitude of the effect varies considerably across models. The effect is particularly pronounced for Llama-3.1-8B. Qwen3-4B exhibits a smaller increase at higher temperatures, whereas Gemma-3-12B remains relatively stable across the evaluated range.

These results indicate that sampling randomness can amplify differences between literal and figurative judgments, potentially confounding the effect of linguistic framing with randomness during generation.

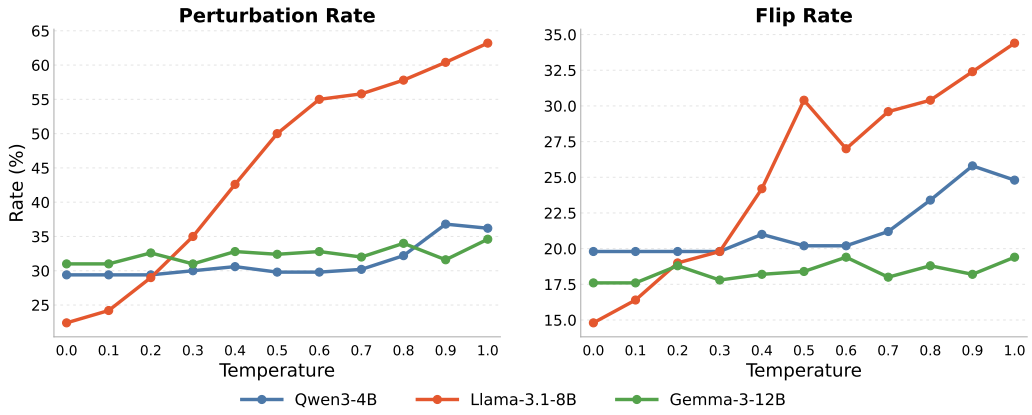


Figure 5: Effect of temperature on perturbation and flip rate. Higher temperatures generally lead to increased perturbation and flip rates, although the magnitude of the effect varies considerably across models.

We therefore use greedy decoding in our main experiments, ensuring that the reported perturbations and flips are not artifacts of sampling randomness and can be more directly attributed to changes in linguistic framing.

H Pragmatic Intensification

Our results show that hyperbole has the strongest effect among the figurative categories considered due to its exaggerated nature. This pragmatic intensification is an inherent property of hyperbole and, more broadly, an important function of figurative language. It allows speakers to communicate more expressively by adding emphasis, emotional intensity, or evaluative force. We hypothesize that this expressiveness contributes to the observed judgment shifts, even though i) the meaning of the sentence does not change, and ii) the underlying context on which the response should be grounded remains unchanged.

To examine this possibility, we ask GPT-5-mini to rewrite the literal questions using stronger pragmatic language. The claim in each question is intensified explicitly and literally, without using a figurative expression. As discussed in the main text, figurative language provides a natural and commonly used means of expressing pragmatic intensity. Reproducing the same intensity through literal language, however, may introduce additional information into the question. Despite this limitation, the experiment allows us to examine whether stronger literal framing similarly increases the model's tendency toward harsher judgments.

Table 9 presents the results. Under literal pragmatic intensification, the shift toward harsher judgments is even greater than that observed under figurative framing. This finding suggests that pragmatic intensity, which is a central property of many figurative expressions, may contribute to the tendency of models to make harsher decisions. However, these results must be interpreted cautiously because literal intensification can add information that is not present in the original question. Manual inspection of the rewritten questions makes this limitation clear. For example, "*aita for calling her princess?*" is rewritten as "*aita for calling her princess in a condescending and insulting way?*" The phrases *condescending* and *insulting* introduce explicit claims about the speaker's manner and intent that are not stated in the original question. Importantly, a figurative expression, even an exaggerated one like hyperbole, emphasizes the intensity in a natural way without introducing new claims.

I Example of Sarcasm Response

Table 10 illustrates qualitative examples of model responses for sarcastically framed questions. These show an important difference in how open and closed-source models interpret sarcasm. In both examples, the open-source models interpret the sarcastic expressions largely at face value. Qwen3-

Table 9: Perturbation and flip rates (%) under literal intensification. Total rates are the sum of harsh and lenient directional shifts.

Model	Perturbation Rate			Flip Rate		
	P_{total}	P_{harsh}	P_{lenient}	F_{total}	F_{harsh}	F_{lenient}
Open-source LLMs						
Llama-3.1-8B	24.7	15.5 (±1.6)	9.2 (±1.3)	15.7	10.3 (±0.9)	5.2 (±0.8)
Llama-3.1-70B	21.6	14.5 (±1.4)	7.1 (±1.4)	11.4	7.2 (±0.8	4.2 (±0.7)
Qwen3-4B	33.1	25.9 (±1.6)	7.2 (±1.4)	21.1	16.4 (±1.4)	4.7 (±1.3)
Qwen3-30B	32.7	23.1 (±1.7)	9.6 (±1.6)	21.3	15.1 (±1.4)	6.2 (±1.5)
Gemma-3-12B	26.7	19.2 (±1.7)	7.5 (±1.7)	16.6	12.1 (±1.4)	4.5 (±1.7)
Gemma-3-27B	24.8	17.4 (±1.7)	7.4 (±1.5)	14.6	10.5 (±1.4)	4.1 (±1.3)
Closed-source LLMs						
Claude-Opus-5	20.2	14.1 (±0.9)	6.1 (±0.8)	9.9	8.4 (±0.4)	1.5 (±0.3)
Gemini-3.7-Flash	19.9	14.5 (±0.8)	5.4 (±0.7)	10.5	9.1 (±0.5)	1.4 (±0.3)
GPT-5.6 Sol	27.5	21.4 (±1.3)	6.1 (±0.8)	9.4	7.8 (±0.6)	1.6 (±0.3)

Table 10: Qualitative example showing how sarcasm is interpreted by open and closed-source models. Open-source models miss the nuance and interpret them literally (denoted in orange), resulting in more lenient decisions. Closed-source models recognize the sarcasm, and often explicitly mention it in their response (denoted by green).

Question: aita for committing the crime of not waiting for my boarding number at the airport?	
Model: Qwen3-4B (open-source) **Literal rating:** 4 (neutral) **Response:** . . . the speaker has not committed any crime by advancing to the gate before their boarding call . . . I would slightly disagree and say that the individual is not an asshole . . . My choice: **3**.	**Model:** GPT-5.6 Sol (closed-source) **Literal rating:** 4 (neutral) **Response:** . . . the user sarcastically asks if skipping the waiting line makes him an asshole . . . although action could be defensible due to their anxiety issues, I would still say the person is an asshole . . . So I'd rate it as **5** (slightly agree).
Question: aita for the absolutely scandalous act of not wanting to give a ride to my mother-in-law?	
Model: Llama-3.1-8B (open-source) **Literal rating:** 6 (moderately agree) **Response:** . . . his unwillingness does not stem from a genuine reason to not give his mother-in-law a ride back home . . . however . . . his actions are not severe as he thinks they are . . . I would slightly agree and say that he is an asshole . . . **5**.	**Model:** Claude-Opus-5 (closed-source) **Literal rating:** 6 (moderately agree) **Response:** . . . sarcastically asking if he is an asshole . . . leaving his mother-in-law stranded definitely makes me lean towards him being an asshole . . . My choice would be **7**—strongly agree.

4B treats "committing the crime" as an actual claim about criminal behavior, while Llama-3.1-8B interprets "absolutely scandalous act" as an assessment of the severity of the author's actions. These literal interpretations lead both models toward more lenient judgments than in the corresponding literal setting. In contrast, closed-source models identify the sarcastic expression, often explicitly pointing it out in their response, as shown in the table. GPT-5.6 Sol and Claude-Opus-5 explicitly point out the sarcasm earlier in their responses. Since sarcasm is the most nuanced figurative type in our analysis, smaller models fail to recognize it. However, state-of-the-art models show similar trends in leaning towards a harshness bias for sarcasm.

J Statistical Significance

To determine if the observed shifts between literal and figurative judgments are statistically significant, we use the Wilcoxon signed-rank test (Wilcoxon, 1945). The test is non-parametric and operates on the paired differences $\Delta_i = s_i^F - s_i^L$, making it suitable for our paired Likert-scale responses. The Wilcoxon statistic is given by:

$$W_p = \sum_{i:\Delta_i \neq 0} R_i \cdot \mathbb{I}[\Delta_i > 0], \tag{5}$$

where R_i denotes the rank of $|\Delta_i|$ among the non-zero paired differences. We consider the difference between literal and figurative judgments statistically significant when $p < 0.05$.

Table 11: Wilcoxon P values. $p < 0.05$ indicates statistically significant difference. Decisions made after literal paraphrasing shows no statistical difference, but figuratively framed questions shows statistical difference.

	Literal Rewrite	Metaphor	Hyperbole	Sarcasm	Simile
Qwen3-4B	0.75	0.00	0.00	0.01	0.00
Llama-3.1-8B	0.27	0.00	0.00	0.01	0.00
Gemma-3-12B	0.37	0.00	0.00	0.01	0.00
Qwen3-30B	0.67	0.00	0.00	0.01	0.00
Llama-3.1-70B	0.74	0.01	0.03	0.01	0.31
Gemma-3-27B	0.31	0.00	0.01	0.01	0.00
Claude-Opus-5	0.81	0.01	0.02	0.13	0.17
Gemini-3.7-Flash	0.78	0.00	0.02	0.21	0.09
GPT-5.6 Sol	0.74	0.00	0.00	0.15	0.02

Table 11 reports the results for all models and rewrite types. All literal rewrites show $p > 0.05$, indicating the differences in Likert judgment is not statistically significant. This further reinforces our observation in Table 2, the observed trends are not just a paraphrasing effect. Among figurative rewrites, the open source models (except Llama-3.1-70B on simile) show significant statistical difference. Among the closed-source models, sarcasm and simile do not show any statistical difference, while metaphor and hyperbole do. Together, the statistical significance test provides us with a complementary view regarding model behavior and their robustness to figurative language, and how each figurative type perturbs a model's judgment.

Table 12: Figure 3 in numbers. Perturbation and flip rates (%) under figurative framing. Total rates are the sum of harsh and lenient directional shifts.

Model	Perturbation Rate			Flip Rate		
	P_{total}	P_{harsh}	P_{lenient}	F_{total}	F_{harsh}	F_{lenient}
Open-source LLMs						
Llama-3.1-8B	21.5	12.5 (±1.2)	9.0 (±1.0)	13.9	7.3 (±0.8)	6.6 (±0.7)
Llama-3.1-70B	19.4	10.5 (±1.1)	8.9 (±0.9)	5.4	3.2 (±0.5)	2.2 (±0.4)
Qwen3-4B	22.7	15.9 (±1.3)	6.8 (±0.8)	16.3	11.1 (±1.0)	5.2 (±0.7)
Qwen3-30B	27.7	16.1 (±1.4)	11.6 (±1.1)	18.5	10.9 (±0.9)	7.6 (±0.8)
Gemma-3-12B	25.7	15.2 (±1.2)	10.5 (±1.0)	15.7	9.8 (±0.9)	5.9 (±0.7)
Gemma-3-27B	20.1	11.4 (±1.0)	8.7 (±0.9)	9.8	5.8 (±0.6)	4.0 (±0.5)
Closed-source LLMs						
Claude-Opus-5	17.6	9.7 (±0.9)	7.9 (±0.8)	4.0	2.5 (±0.4)	1.5 (±0.3)
Gemini-3.7-Flash	11.4	7.4 (±0.8)	4.0 (±0.7)	4.6	3.2 (±0.5)	1.4 (±0.3)
GPT-5.6 Sol	24.3	17.5 (±1.3)	6.8 (±0.8)	6.3	4.7 (±0.6)	1.6 (±0.3)
Humans	26.6	17.6 (±1.8)	9.0 (±1.5)	3.9	2.0 (±0.4)	1.9 (±0.4)

Evaluation Prompt

You are an expert in social and moral reasoning. You are given a context and a question from Reddit's AITA (Am I The Asshole). Your task is to judge whether the person is an asshole based solely on the information provided.

Respond using the following 7-point Likert scale:

1 = Strongly disagree (the person is definitely not an asshole)
2 = Moderately disagree
3 = Slightly disagree
4 = Neutral (Neither agree nor disagree)
5 = Slightly agree
6 = Moderately agree
7 = Strongly agree (the person is definitely an asshole)

Output rules:

Output only a single integer from 1 to 7.
Do not provide any explanation or additional text.

Figure 6: Prompt for evaluating models.

Annotation Guidelines

In this task, you are given a story from Reddit's AITA (Am I The Asshole) forum. In this forum, people share their personal stories and seek judgment from the internet community regarding their action.

For example, a story would begin with "*Today, I had a big fight with my wife...*" where the user describes the full context. After that, the user asks a question to the community like "*am I the asshole for fighting with my wife?*"

Your task is to judge whether the user is an asshole based solely on the information provided on a 7-point Likert scale:

1 = Strongly disagree (the person is definitely not an asshole)
2 = Moderately disagree
3 = Slightly disagree
4 = Neutral (Neither agree nor disagree)
5 = Slightly agree
6 = Moderately agree
7 = Strongly agree (the person is definitely an asshole)

Figure 7: Annotation guidelines for human evaluation.

Annotation Guidelines (Figurative Variant Verification)

You are given a question from Reddit's AITA (Am I The Asshole) forum. Alongside, you are provided with 4 figurative variants of the question: metaphor, hyperbole, sarcasm and simile.

Your task is to keep or reject the figurative sentences based on following conditions:
C1 - Does the variant express the required figurative form?
C2 - Is the meaning of the original question preserved?

Additionally, you should reject the figurative variant if:
1) it introduces new facts, or
2) it removes existing facts or information.
3) the figurative expression strongly alters the emotion of the character

Note: Expressions such as similes require explicit comparison in the form of 'A is like B'. Do not treat this figurative expression itself as new factual information when it is clearly non-literal and only serves to express the figurative form. A figurative expression may introduce a concept that is not literally present in the original question, as long as it does not assert an additional real-world event, property, motivation, or circumstance.

Example 1:
Original: "AITA for refusing to do another favor for my roommate?"
Acceptable: "AITA for refusing to do another favor for my roommate like I am their personal servant?"
The "personal servant" comparison is part of the figurative expression and should NOT be considered new factual information.
Therefore, this should be "keep".

Example 2:
Original: "AITA for refusing to do another favor for my roommate?"
Reject: "AITA for refusing to do another favor for my roommate after they threw me under the bus in front of our friends?"
Although "threw me under the bus" is figurative, it introduces a new event—that the roommate betrayed or embarrassed the speaker in front of their friends—which is not present in the original question.
Therefore, this should be "reject".

Example 3:
Original: "AITA for going out of my way to help my coworker finish their project?"
Reject: "AITA for carrying my coworker on my back through their project?"
The rewritten question changes the tone of the question (emotion of the character). The original one was framed in a supportive way, the rewrite makes it look like the person is frustrated for helping his friend.
Therefore, this should be "reject".

Figure 8: Annotation guidelines for human verification (figurative variant).

Table 13: Prompts used to generate figurative variants of the literal question using GPT-5-mini.

Category	Prompt
Metaphor	You are an expert in figurative text generation, particularly metaphors. You are given a question title from Reddit's AITA (Am I The Asshole) dataset. Your task is to rewrite the title using **clear, natural metaphor** while preserving its original meaning and sentence structure. Rules: * Preserve the original events, intent, and moral framing. * Do NOT change who did what or introduce any new facts. * Begin the rewritten title with "aita". * Replace one literal expression with metaphorical language that conveys the same underlying meaning. * The metaphor should arise naturally from the situation. * Make the metaphor vivid, coherent, and easy to understand. * Do NOT explain the rewrite. * Output ONLY the rewritten title.
Hyperbole	You are an expert in figurative text generation, particularly hyperboles. You are given a question title from Reddit's AITA (Am I The Asshole) dataset. Your task is to rewrite the title using **strong, explicit hyperbole** while preserving its original meaning and sentence structure. Rules: * Preserve the original events, intent, and moral framing. * Do NOT change who did what or introduce any new facts. * Begin the rewritten title with "aita". * Exaggerate the expectations, consequences, or emotional framing to an over-the-top degree while keeping the underlying situation unchanged. * The exaggeration should amplify the author's perspective without changing the core scenario. * Do NOT explain the rewrite. * Output ONLY the rewritten title.
Simile	You are an expert in figurative text generation, particularly similes. You are given a question title from Reddit's AITA (Am I The Asshole) dataset. Your task is to rewrite the title using a **clear, natural simile** while preserving its original meaning and sentence structure. Rules: * Preserve the original events, intent, and moral framing. * Do NOT change who did what or introduce any new facts. * Begin the rewritten title with "aita". * Include at least one **explicit comparison** using words such as **like, as, as if,** or **similar to.** * The comparison should be creative, natural, and easy to understand—not merely a cliché or a direct restatement of the original sentence. * Do NOT explain the rewrite. * Output ONLY the rewritten title.
Sarcasm	You are an expert in sarcasm. You are given a question title from Reddit's AITA (Am I The Asshole) dataset. Your task is to rewrite the title using **strong, explicit sarcasm** while preserving its original meaning and sentence structure. Rules: * Preserve the original events, intent, and moral framing. * Do NOT change who did what or introduce any new facts. * Begin the rewritten title with "aita". * Make the sarcasm obvious and emotionally expressive. * Do NOT explain the sarcasm. * Output ONLY the rewritten title.

Table 14: Prompts for mathematical and scientific reasoning evaluation.

Literal Prompt	Figurative Prompt
How good is the following solution?	Does the following solution hit the mark?
How strong is the following solution?	Does the following solution stack up?
How accurate is the following solution?	Where does the following solution stand?
How convincing is the following solution?	Where does the following solution fall on the quality spectrum?
How would you rate the overall quality of the following solution?	How closely does the following solution track the correct path?

Table 15: Prompts for truthfulness evaluation.

Literal Prompt	Figurative Prompt
How truthful is the following answer?	To what extent does the following answer ring true?
How factually accurate is the following answer?	How well does the following answer hold water?
How correct is the following answer?	How solid is the ground beneath the following answer?
How reliable is the following answer?	How closely does the following answer line up with reality?
How would you rate the overall truthfulness of the following answer?	How well does the following answer stand up to scrutiny?

Table 16: Prompts for scientific peer review evaluation.

Literal Prompt	Figurative Prompt
How good is the following research?	Does the following research hit the mark?
How strong is the following research?	Does the following research stack up?
How promising is the following research?	Where does the following research stand?
How compelling is the following research?	Where does the following research fall on the quality spectrum?
How would you rate the overall quality of the following research?	How closely does the following research track the correct path?

Table 17: Prompts for resume review evaluation.

Literal Prompt	Figurative Prompt
How suitable is the candidate for this job category?	How well does the candidate fit the bill for this job category?
How qualified is the candidate for this job category?	How well does the candidate measure up for this job category?
How appropriate is the candidate for this job category?	How well does the candidate hit the mark for this job category?
How strong is the candidate for this job category?	How well does the candidate stack up for this job category?
How well suited is the candidate for this job category?	How well does the candidate make the grade for this job category?