

Arguments that Alter Minds: LLM Rationales Sway Human (and LLM) Notions of Plausibility

Shramay Palta
Sarah Wiegrefe

Peter Rankel
Rachel Rudinger

University of Maryland, College Park
{spalta, par, sarahwie, rudinger}@umd.edu

Abstract

We investigate the degree to which human (and LLM) plausibility judgments of multiple-choice commonsense benchmark answers are subject to influence by (im)plausibility arguments for or against an answer, in particular, using rationales generated by LLMs. We collect 3,000 plausibility judgments from humans and another 13,600 judgments from LLMs. Overall, we observe increases and decreases in mean human plausibility ratings in the presence of LLM-generated PRO and CON rationales, respectively, suggesting that, on the whole, human judges find these rationales convincing. Experiments with LLMs reveal similar patterns of influence. Our findings demonstrate a novel use of LLMs for studying aspects of human cognition, while also raising practical concerns that, even in domains where humans are “experts” (i.e., common sense), LLMs have the potential to exert considerable influence on people’s beliefs.¹

1 Introduction

A typical framework for evaluating the commonsense reasoning abilities of NLP models is to present an everyday situation described in text (*A person drops a glass*), followed by a set of possible inferences (*The glass breaks; The glass bounces; etc*), and instruction to select the most plausible answer choice (Levesque et al., 2011; Roemmele et al., 2011; Mostafazadeh et al., 2016). Unlike in logical or mathematical reasoning problems, these answer choices are rarely definitively true or false, but rather lie on a continuum of plausibility (Zhang et al., 2017; Chen et al., 2020), and require non-monotonic modes of reasoning to assess, such as defeasible (Rudinger et al., 2020), abductive (Bhagavatula et al., 2020), or counterfactual (Qin et al., 2020) reasoning. In this setting, the answer designated as “correct” should be the choice that human

judges agree is most plausible. Indeed, human annotators are more likely to agree on a “correct” answer when the plausibility gap between answer choices is large (Palta et al., 2024).

However, unless an answer choice is logically impossible, there exist circumstances under which even a relatively implausible distractor answer could be true. For example, one could argue that a glass is actually quite likely to bounce if dropped on a trampoline or rubber mat. By the same token, unless an answer choice is logically necessary, there exist circumstances under which an otherwise highly plausible outcome becomes much less likely: The drinking glass probably *won’t* break if it’s very sturdy or doesn’t fall far. Thus, any answer choice which is neither certain nor impossible can be subject to arguments for or against its plausibility. In principle, these types of (im)plausibility arguments or rationales do not introduce any new evidence to the situation; rather, they simply highlight possible circumstances which, *if true*, would impact an answer’s plausibility.

In this paper, we investigate whether an LLM can produce convincing arguments for the plausibility or implausibility of commonsense benchmark answer choices, as measured by their ability to impact the plausibility ratings of human (and LLM) judges. (We will refer to such plausibility or implausibility arguments as PRO or CON rationales, respectively.) Posed another way, this paper asks whether and how human plausibility judgments change in response to LLM-written PRO and/or CON rationales.

To answer these questions, our work builds on the prior study of Palta et al. (2024), who previously collected Likert-scale judgments from human annotators about the plausibility of individual question-answer choice pairs from two commonsense multiple choice benchmarks, Social IQA (SIQA; Sap et al. 2019) and CommonsenseQA (CQA; Talmor et al. 2019). These judgments were collected in the

¹Our data is available at <https://github.com/shramay-palta/AAM.git>.

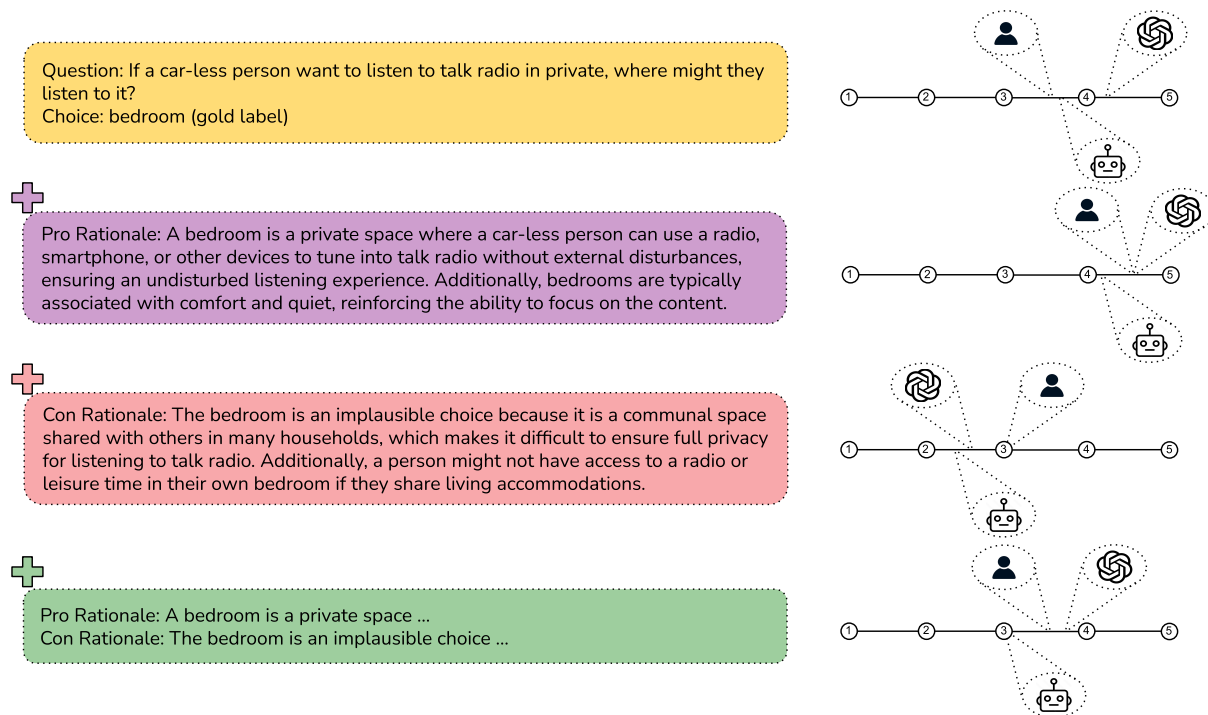


Figure 1: An example question and answer choice from CQA (Talmor et al., 2019) in yellow, paired with LLM-generated rationales in favor of the answer’s plausibility (PRO, purple); against the answer choice’s plausibility (CON, red); and both for and against (PRO+CON, green). Likert scales on right show how mean plausibility ratings for the answer, as judged by human and LLM raters (OpenAI and Non-OpenAI), shift in response to different types of rationales.

absence of any rationales or (im)plausibility arguments, which we will refer to as the NO rationale setting. We extend this work by sampling a subset of their annotated question-answer pairs, and using a LLM to generate PRO and CON rationales for or against the plausibility of each answer. Adopting a similar protocol, we then collect 3,000 new human plausibility ratings of answer choices, presented alongside either a PRO rationale of the answer, a CON rationale, or both (PRO+CON), as illustrated in Figure 1. In a parallel set of experiments using the same stimuli, we prompt 17 different LLMs to also rate the plausibility of these answers, collecting an additional 13,600 judgments from LLMs.

With this data, we are able to make the following observations:

- Both human and LLM plausibility judgments are significantly affected by the addition of PRO, CON, and PRO+CON rationales, relative to a NO rationale setting.
- On average, PRO rationales raise mean plausibility ratings from humans and LLM judges alike, while CON rationales lower mean ratings.
- These patterns hold across both gold and dis-

tractor answer choices, with one notable exception: mean plausibility ratings by humans surprisingly *drop* for gold answer choices when a PRO rationale is presented. For LLMs, mean plausibility ratings increase in response to PRO rationales, for gold and distractor answer choices alike.

- PRO+CON rationales yield mixed reactions from both humans and LLMs, with ratings decreasing for the gold label, but increasing for the distractor answer choice.
- While OpenAI and Non-OpenAI models generally exhibit similar trends, we observe some differences. In particular, CON rationales appear to have a larger effect on the OpenAI models’ ratings compared to the human ratings and Non-OpenAI models’ ratings.
- The initial plausibility rating of an answer choice has a strong anchoring effect: the higher the initial plausibility of an answer, the smaller the subsequent change in its rating after a rationale is provided.

2 Rationale Generation

Palta et al. (2024) introduced a plausibility rating

Dataset	# q	#(q, a)	#(q, a, r)	#Human Judgments	#LLM Judgments	#Total Judgments
SIQA	50	100	400	1,500(+500)	6,800	8,300(+500)
CQA	50	100	400	1,500(+500)	6,800	8,300(+500)
Total	100	200	800	3,000(+1,000)	13,600	16,600(+1,000)

Table 1: Size and structure of the collected dataset. From left to right, columns report number of unique questions drawn from each dataset (# q); number of unique question-answer pairs (#(q, a)); number of unique question-answer-rationale tuples (#(q, a, r)) over four settings (NO rationale, PRO rationale, CON rationale, and PRO+CON rationale); total number of human plausibility ratings collected; total number of LLM plausibility ratings generated; and total number of plausibility ratings (human or LLM). Numbers in parentheses represent *additional* judgments previously collected by Palta et al. (2024) for the NO rationale setting, used in our analysis.

framework for rating the plausibility of individual answer choices from Multiple Choice Question (MCQ) datasets in isolation. For our study, we randomly sample a subset of 100 questions—50 each from SIQA and CQA— from the set of questions used in Palta et al. (2024). SIQA is designed to evaluate the social and emotional intelligence in computational models. Each SIQA item consists of a short context describing a social situation, a question about a person in the situation and three answer choices. CQA, on the other hand, evaluates general commonsense knowledge, with each item containing five answer choices. For each sampled question, we pick the dataset gold label ($a_{gold-label}$) and a randomly selected distractor answer choice ($a_{distractor}$) for our study. To maintain consistency with Palta et al. (2024), we break down each question q with answer choices $a_{gold-label}$ and $a_{distractor}$ into pairs (q, a) where $a \in [a_{gold-label}, a_{distractor}]$.

We then generate two forms of rationales for each (q, a) pair:

1. **PRO Rationale:** An explanation that best explains the *plausibility* of the answer choice for the given question and context (for SIQA).
2. **CON Rationale:** An explanation that best explains the *implausibility* of the answer choice for the given question and context (for SIQA).

These two rationales are then used to collect human and LLM plausibility judgments under four settings, described in §3: NO, PRO, CON, and PRO+CON.

2.1 Model Preference Study

Because the cost of human annotation is high (\$1,000 USD for our full-scale study), we first ran a small-scale preference study to select the language model most likely to produce rationales convincing to humans for subsequent use in the full-scale

study.

For this initial preference study, we prompted several state-of-the-art models: GPT-4o and GPT-4o-mini (OpenAI et al., 2024c), as well as Llama 3.1-8B-Instruct and Llama 3.1-70B-Instruct (Grattafiori et al., 2024) to generate both PRO and CON rationales for a set of 40 randomly sampled (q, a) pairs from the validation split of SIQA. The prompts used for generating the PRO and CON rationales are presented in Prompts A.1 and A.2 in Appendix A.5.

The generated rationales, and their corresponding (q, a) pairs, were presented to 4 human annotators, who selected the rationale “that best explained the plausibility (for PRO) or implausibility (for CON) of the answer choice a , and was also the most believable”. We show an example of the interface presented to the annotators in Figure 4. Based on the annotators’ responses, GPT-4o received the most votes and was selected to generate the rationales in our study.² The set of (q, a) pairs used in the preference study is entirely disjoint from the set used in the full-scale study.

We note that it is not the purpose of the preference study, nor the full-scale study, to determine *which* LLM is *most* persuasive to humans. Rather, the goal of the full-scale study is to determine whether *any* LLM might produce rationales capable of swaying human plausibility judgment. A positive finding (that human judgments are indeed systematically swayed) can in principle be demonstrated with a single model, since it is not our claim that *all* LLMs would necessarily yield such results. On the other hand, a negative finding (that the LLM-generated rationales did not affect hu-

²According to annotator feedback, selecting a single rationale from the four presented was challenging, as the generated rationales were quite similar to one another. Nevertheless, GPT-4o received the majority of votes.

Dataset	Agent	Pro Rationale			Con Rationale			Pro+Con Rationales		
		Increase	Decrease	No Change	Increase	Decrease	No Change	Increase	Decrease	No Change
SIQA	Human	28%	22%	50%	2%	69%	29%	11%	44%	45%
	OpenAI	40%	2%	58%	0%	84%	16%	12%	43%	45%
	Non-OpenAI	61%	1%	38%	5%	66%	29%	20%	24%	56%
CQA	Human	30%	27%	43%	9%	33%	58%	13%	44%	43%
	OpenAI	62%	4%	34%	4%	74%	22%	21%	38%	41%
	Non-OpenAI	53%	5%	42%	4%	65%	31%	21%	32%	47%

Table 2: Percentage cases where the mean plausibility rating given by human annotators, OpenAI and Non-OpenAI models increased, decreased or remained the same when different rationales were presented for SIQA and CQA. Numbers here include both $a_{gold-label}$ and $a_{distractor}$ answer types.

man judgments) would beg the question: could a stronger model have yielded positive results? It is this latter concern that motivates our preference study.

3 Human Plausibility Ratings

To understand the impact of rationales on human (and LLM) plausibility judgments, we borrow the Likert scale introduced by Zhang et al. (2017), and also used in Palta et al. (2024), defined as *1-Impossible, 2-Technically Possible, 3-Plausible, 4-Likely and 5-Very Likely*. Given a question q and possible answer a , judges are asked to assess the plausibility of a under each of the following four conditions:

- **NO Rationale:** The (q, a) pair is presented alone. No rationale is provided.
- **PRO Rationale:** The (q, a) pair is accompanied by an argument for the *plausibility* of a , i.e., a PRO rationale.
- **CON Rationale:** The (q, a) pair is accompanied by an argument for the *implausibility* of a , i.e., a CON rationale.
- **PRO+CON Rationale:** The (q, a) pair is accompanied by both a PRO rationale *and* a CON rationale.

For the No Rationale setting, we use the plausibility judgments released by Palta et al. (2024). For each of the three remaining rationale settings (PRO, CON, and PRO+CON) we collect 5 judgments, totaling 15 new judgments per selected (q, a) pair.

Annotators are recruited through Prolific, and compensated at \$15/hour. To ensure independence, each annotator could judge at most one (q, a, r) pair for a given question q for a rationale type r . Further details, including examples of the annotation interfaces (Figures 5, 6 and 7), are presented in Appendix A.4. We present the overall annotation

counts in Table 1.

3.1 How do Rationales influence Human Ratings?

We measure how mean plausibility ratings³ for (q, a) pairs change between the original NO Rationale setting and each of the PRO, CON, and PRO+CON Rationale settings. Table 2 shows that, for human judges, mean plausibility ratings often fall or stay unchanged—but rarely rise—in the presence of CON rationales. With PRO rationales, mean human ratings often stay unchanged, and are only slightly more likely to rise than drop. In the PRO+CON setting, ratings are unlikely to rise, but about equally likely to drop or stay unchanged. However, an interesting pattern emerges when we stratify outcomes based on whether the answer choice is a “gold” (correct) or “distractor” (incorrect) answer, as shown in Table 3. Critically, *we observe that the effect of PRO rationales on mean human judgments is bimodal*: on distractor answers, mean judgments rise, while for gold answers they fall. Thus, for implausible distractor answers, PRO rationales appear sufficiently persuasive to raise mean ratings. But why do ratings for gold answers drop? One possible explanation is rooted in Grice’s Maxim of Quantity (Grice, 1975): a conversation should be as informative as possible, and not provide more information than what is required. When the LLM presents a rationale arguing that a highly obvious fact is “plausible”, human judges might implicitly interpret the PRO rationales as arguing the answer is *at most* plausible, rather than *at least* plausible. This “underselling” introduces doubt, causing humans to actively downgrade their initial high-confidence ratings to align with the LLM-generated rationale.

³Average rating across the 5 annotators for each (q, a) item.

Dataset	Agent	Pro Rationale			Con Rationale			Pro+Con Rationales		
		Overall	Gold Label	Distractor	Overall	Gold Label	Distractor	Overall	Gold Label	Distractor
SIQA	Human	3.33(+0.1)	3.84(-0.26)	2.81(+0.47)	2.28(-0.94)	2.72(-1.39)	1.85(-0.49)	2.87(-0.35)	3.36(-0.75)	2.39(+0.05)
	OpenAI	3.7(+0.57)	4.25(+0.26)	3.14(+0.88)	1.9(-1.22)	2.28(-1.7)	1.52(-0.74)	2.9(-0.22)	3.38(-0.6)	2.42(+0.16)
	Non-OpenAI	3.6(+0.62)	3.94(+0.42)	3.26(+0.82)	2.21(-0.77)	2.42(-1.09)	1.99(-0.45)	2.93(-0.05)	3.19(-0.32)	2.66(+0.22)
CQA	Human	3.39(0.0)	3.91(-0.44)	2.86(+0.45)	2.62(-0.76)	3.28(-1.08)	1.96(-0.45)	2.97(-0.41)	3.7(-0.65)	2.24(-0.17)
	OpenAI	3.91(+0.64)	4.55(+0.25)	3.27(+1.03)	2.02(-1.25)	2.56(-1.74)	1.48(-0.76)	3.12(-0.15)	3.76(-0.54)	2.49(+0.25)
	Non-OpenAI	3.63(+0.61)	3.95(+0.31)	3.32(+0.92)	2.18(-0.84)	2.43(-1.21)	1.94(-0.47)	2.99(-0.03)	3.28(-0.36)	2.7(+0.3)

Table 3: Mean plausibility ratings given by Human, OpenAI, and Non-OpenAI agents after exposure to PRO, CON, and PRO+CON rationales, reported separately for $a_{gold-label}$ and $a_{distractor}$. Numbers outside parentheses indicate mean ratings after rationales; numbers in parentheses represent the corresponding changes (Δ) from the No Rationale Plausibility Rating.

With CON rationales, both gold and distractor answers show drops in mean plausibility, albeit with much larger drops for gold answers. (Interestingly, the drop for gold answers with CON rationales is over one full Likert scale point on average, the largest changes observed for any setting.) While boundary effects likely contribute to the differences observed between changes for gold and distractor labels (e.g., there is a cap on how much an already high rating can rise), they cannot account for differences between answers of the *same* type (e.g. distractor) across different rationale settings. For example, we have seen that mean ratings for distractor answers rise with PRO rationales and fall with CON rationales; in our controlled experimental setup, the only non-random explanation for this difference is the rationale type itself. We present statistical significance tests at the end of this section, with a causal analysis of the factors that led to the changes in the ratings presented in § 5.

For PRO+CON rationales, we observe an approximate averaging of the PRO and CON outcomes: for gold and distractor answers both, changes in mean ratings lie in between the changes observed in the PRO and CON settings, respectively.

We perform several Chi-squared tests of homogeneity to compare the rating distributions with and without rationales. For SIQA, p-values for the shifts in the distributions are 0.069 , $2.17E^{-9}$, and 0.001 for the PRO, CON, and PRO+CON settings, respectively; for CQA, they are 0.001 , $9.24E^{-7}$, and 0.0006 . In nearly all cases, adding rationales yields statistically significant distribution shifts.

4 Do Rationales Affect LLMs as well?

In an effort to understand whether LLMs are also influenced by the introduction of different forms of rationales, we replicate the entire setup used to collect human plausibility ratings, as described in § 3.

Previous work (Panickssery et al., 2024) has shown that LLMs can recognize and favor their own generated outputs over outputs from other models. Given that GPT-4o generated the PRO and CON rationales, to mitigate any potential biases that can arise from instances of self-preference, we divide the LLMs we study in our work into two groups, with membership as follows:

- **OpenAI Models:** GPT (3.5, 4, 4-turbo, 4o, 4o-mini, 4.5-preview, o1, and o3-mini) (OpenAI et al., 2024c,a,b)
- **Non-OpenAI Models:** LLaMA-3.1-Instruct (8B, 70B) (Grattafiori et al., 2024), LLaMa-3.2-Instruct (3B), LLaMa-3.3-Instruct (70B), Mistral Instruct (8x22B) (Jiang et al., 2024), Yi-1.5-Chat (6B, 9B, 34B) (AI et al., 2025) and DeepSeek-R1 (Qwen-32B) (DeepSeek-AI, 2025).

We prompt the models listed above with the (q, a) pairs accompanied by one rationale for each setting [NO, PRO, CON, PRO+CON] (same as used in human experiments), and present the zero-shot prompts in Prompts A.3, A.4, A.5 and A.6 in Appendix A.5.

4.1 LLM Ratings Pre and Post Rationales

Paralleling our analysis in §3.1, we compare how mean plausibility ratings produced by LLMs (OpenAI and Non-OpenAI models) change in the presence of each rationale type, reporting results again in Tables 2 and 3, for direct comparison with results from the human judges.

Overall, PRO rationales have a larger positive impact for both sets of models, with a large proportion of ratings increasing across both datasets with the inclusion of the PRO rationale. Similarly, adding CON rationales leads to a large drop in ratings across both datasets. The inclusion of PRO+CON rationales once again elicits a mixed response, similar to our observation with human judgments in § 3.1.

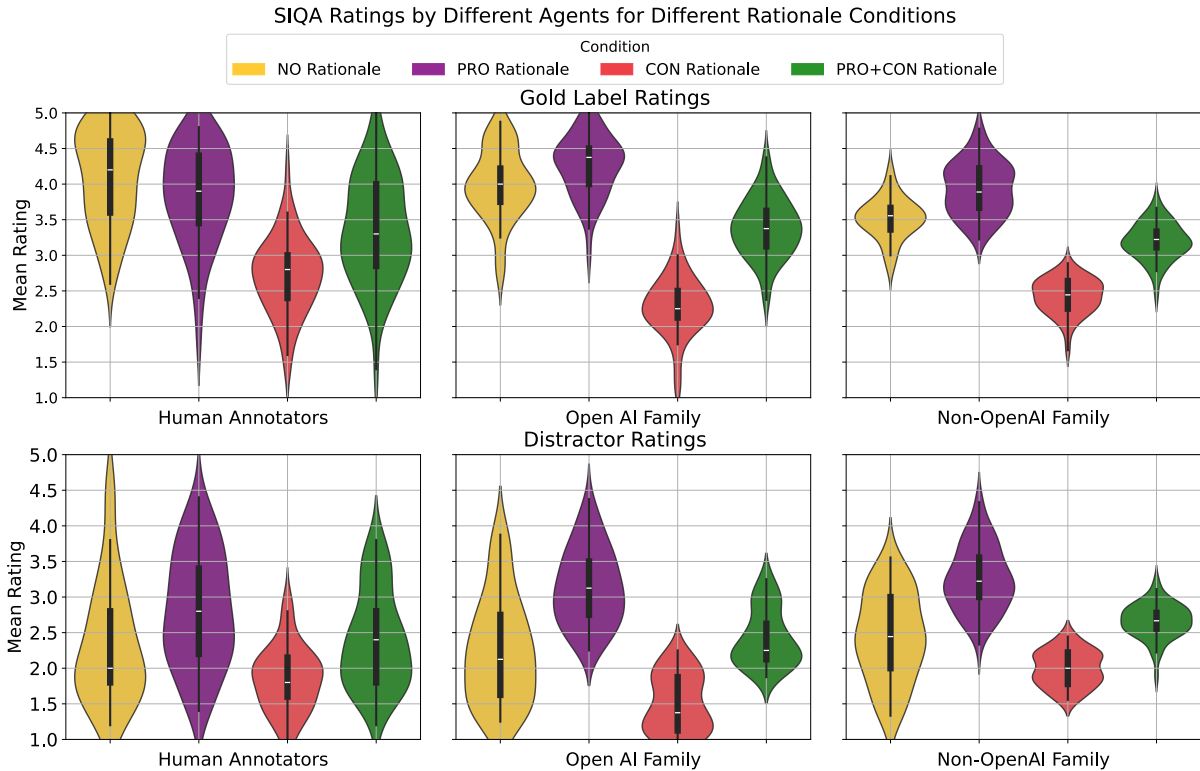


Figure 2: Mean Plausibility Rating for SIQA for Different Agents for Different Rationale Conditions. Table 3 shows the mean changes for $a_{gold-label}$ and $a_{distractor}$ for all agents.

However, in sharp contrast to what we observed with human annotators, the introduction of PRO rationales leads to an increase in the mean plausibility ratings for $a_{gold-label}$ across both SIQA and CQA for both OpenAI and Non-OpenAI models. Similar to the shifts in human ratings for $a_{gold-label}$ and $a_{distractor}$, by adding CON rationales, both groups of models exhibit a decrease in mean plausibility ratings across both datasets.

Finally, when PRO+CON rationales are presented together with a (q, a) item, we observe the LLMs to also display mixed results, with plausibility ratings for $a_{gold-label}$ dropping for both SIQA and CQA, but increasing for $a_{distractor}$. This behavior is similar to what was observed for the human ratings (§ 3.1). One possible explanation for this phenomenon is that, if PRO rationales dampen responses at the low end of the plausibility scale, and CON rationales dampen responses at the high end of the plausibility scale, then PRO+CON rationales together may have the effect of pushing plausibility ratings to the center of the scale because neither extreme is being used for rating the answer a . This would have the effect of raising average scores for low-scoring distractors ($a_{distractor}$) and lowering scores for high-scoring gold-label answers

($a_{gold-label}$), as is observed for both humans and LLMs. Figures 2 and 3 show that this is not simply a case of reversion to the mean, as, for instance, CON rationales on average push mean scores of distractors (which are low to begin with) even lower. Another compatible explanation is that, as observed in human judgments, PRO+CON rationales result in shifts in mean judgments that lie between the shifts observed in the PRO and CON settings.

To establish the significance of these shifts, we ran several Chi-squared tests of homogeneity to compare the ratings distributions with and without rationales within each setting. These shifts were found to be highly statistically significant across all settings, with all p-values < 0.0001 .

Comparing the ratings given by OpenAI and Non-OpenAI models, we see the OpenAI models to be more sensitive to the inclusion of PRO, CON and PRO+CON rationales, which we hypothesize to be because of the OpenAI models favoring the rationales generated by GPT-4o.

We include examples from SIQA and CQA for each of PRO rationales (Table 6), CON rationales (Table 7) and PRO+CON rationales (Tables 8 and 9), and show how human and LLM ratings differ when different rationales are included for a (q, a) item as

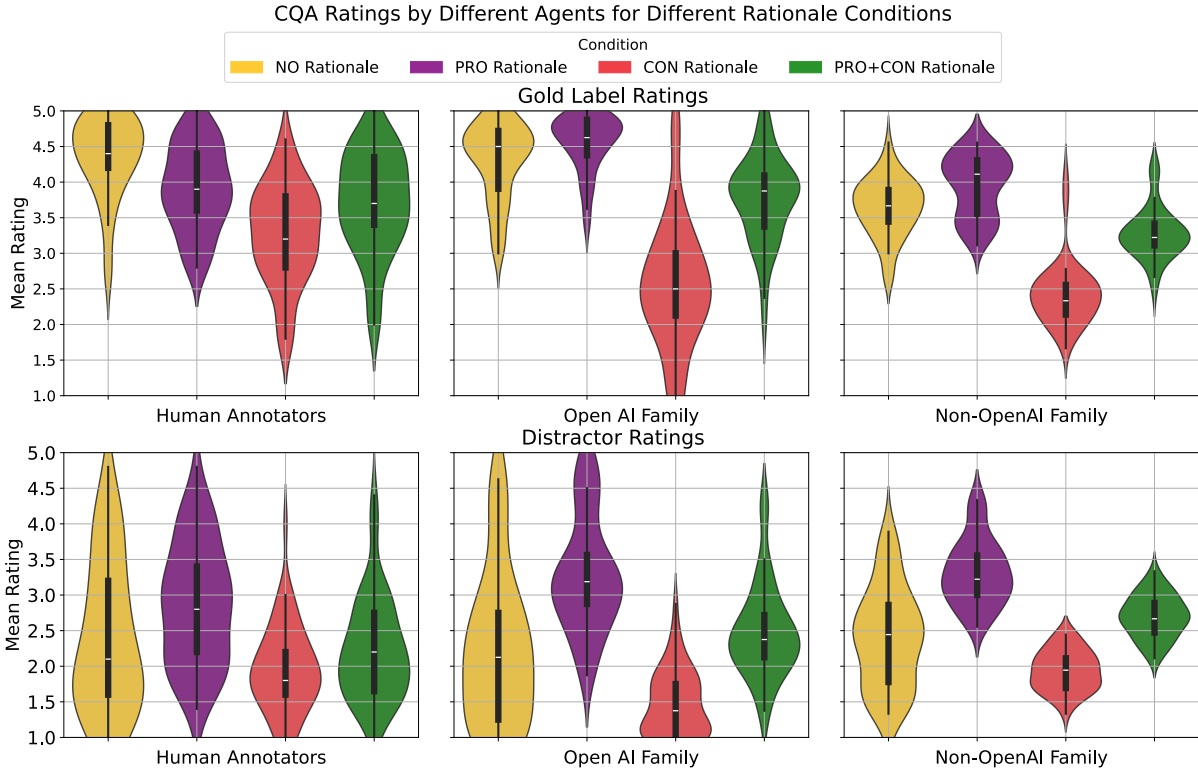


Figure 3: Mean Plausibility Rating for CQA for Different Agents for Different Rationale Conditions. Table 3 shows the mean changes for $a_{gold-label}$ and $a_{distractor}$ for all agents.

compared to annotating without any rationales.

5 Why Did the Ratings Change?

We have now established that both humans and LLMs are sensitive to the inclusion of PRO, CON and PRO+CON rationales, and adjust their plausibility ratings accordingly. To better understand what factors contributed the most to the changes in the plausibility ratings, we perform ordinary least squares (OLS) regression on the collected data separately for $a_{gold-label}$ and $a_{distractor}$ using the rationale type (PRO, CON, PRO+CON) and the NO Rationale Plausibility Rating as features, with the dependent variable being the change (Δ) in plausibility rating after a rationale is presented. The NO Rationale Plausibility Rating was originally collected on the 1-5 Likert Scale as described in Palta et al. (2024). For the regression, we use a $[-2, 2]$ scale to center the ratings around 0.

As detailed in Table 4, the NO Rationale Plausibility Rating is a statistically significant predictor ($p < 0.05$) of the rating change for humans, OpenAI and Non-OpenAI models. The consistently negative coefficients reveal a strong anchoring effect: the higher the initial plausibility of an answer, the smaller the subsequent change in its rating af-

ter a rationale is provided. This suggests that all rater groups are less likely to raise their opinion if their initial assessment is already high. Additionally, this anchoring effect is more pronounced for distractors than for gold-label answers across all groups. This indicates that the initial rating of an incorrect answer choice has a much stronger influence on the final rating change than the initial rating about the correct one.

Comparing the different rationale types included with the (q, a) items, we find that CON rationales exert a stronger effect than PRO rationales, with large negative coefficients that sharply *reduce* plausibility for otherwise correct answers. PRO rationales do *increase* plausibility, but the magnitude of their coefficients is smaller, indicating that supporting reasons reinforce less than opposing reasons undermine. PRO+CON rationales typically carry small negative coefficients, indicating that when both supporting and opposing arguments are presented together, the opposing component tends to dominate, yielding a net plausibility decrease.

Finally, we also test the rationale length, measured as the number of words in the generated rationale, as a feature in a separate regression where the dependent variable is the magnitude of the change

Dataset	Agent	Feature							
		NO Rationale Rating		PRO Rationale		CON Rationale		PRO+CON Rationale	
		Gold Label	Distractor	Gold Label	Distractor	Gold Label	Distractor	Gold Label	Distractor
SIQA	Human	-0.4581	-0.5495	0.2417	0.1082	-0.8823	-0.8525	-0.2423	-0.3125
	OpenAI	-0.4503	-0.5086	0.7083	0.5023	-1.2592	-1.1177	-0.1592	-0.2127
	Non-OpenAI	-0.4770	-0.6491	0.6659	0.4543	-0.8452	-0.8102	-0.0785	-0.1391
CQA	Human	-0.5092	-0.4983	0.2445	0.1590	-0.3875	-0.7450	0.0405	-0.4650
	OpenAI	-0.2340	-0.4883	0.5517	0.6614	-1.4408	-1.1286	-0.2383	-0.1236
	Non-OpenAI	-0.5632	-0.6236	0.6693	0.5428	-0.8507	-0.8394	-0.0040	-0.0750

Table 4: OLS Regression Coefficients for the NO Rationale Plausibility Rating and Rationale Type (PRO, CON and PRO+CON) features. The dependent variable here is the change (Δ) in the plausibility rating after a rationale was presented. Numbers highlighted in **bold** represent coefficients that are statistically significant ($p < 0.05$) according to the t-test.

in the plausibility rating, and present the results in Table 5. Across all settings, the regression coefficients are close to zero suggesting that rationale length exhibits only a weak linear relationship with rating scores, regardless of the answer choice type, dataset or the annotating agent.

6 Related Works

Rating the likelihood or plausibility of events in language has been explored previously in commonsense reasoning (Wang et al., 2018; Liu et al., 2023). The graded nature (Zhang et al., 2017; Chen et al., 2020), the subjectivity of commonsense reasoning (Whiting and Watts, 2024) along axes like culture (Palta and Rudinger, 2023; Hershovich et al., 2022; Bhatia and Shwartz, 2023; Acquaye et al., 2024) has been long studied. We take advantage of this subjectivity to introduce LLM-generated rationales into a commonsense MCQ item and see the impact on human and LLM ratings. Previous works have also looked at how biases in human probability estimates (Kynn, 2007; Kruglanski and Ajzen, 1983; Martins, 2006), framing effects (Druckman, 2001; Mandel, 2014; Sher and McKenzie, 2008), recency bias (Jones and Bergias, 1976; Steiner and Rain, 1989; Arnold et al., 2000) and availability heuristics (Tversky and Kahneman, 1973; Dube-Rioux and Russo, 1988; Folkes, 1988; Pachur et al., 2012; Nazlan et al., 2018) can impact human decision making.

LLMs can be trained or prompted to produce rationales of many forms (Wiegrefe and Marasovic, 2021); early work on training LMs to produce free-form textual rationales focused on explaining model predictions on computer vision or natural language inference tasks (Hendricks et al., 2016; Camburu et al., 2018). LMs’ ability to generate

free-form textual explanations for commonsense reasoning tasks was first established by Rajani et al. (2019); Shwartz et al. (2020); Latcinnik and Berant (2020); Brahman et al. (2021); Wiegrefe et al. (2021), with subsequent work showing that large pre-trained LMs could generate high-quality explanations for commonsense tasks without having been explicitly trained to do so (Paranjape et al., 2021; Lampinen et al., 2022; Wiegrefe et al., 2022; Marasovic et al., 2022).

It is now widely accepted that LMs can produce fluent textual explanations for a wide variety of NLP tasks. The role of these explanations on human (over-)trust, reliance (Jacovi et al., 2021; Goyal et al., 2023), distorting beliefs (Batista and Griffiths, 2026), and utility to humans doing tasks has been studied in a variety of task settings (Biran and McKeown, 2017; Feng and Boyd-Graber, 2019; González et al., 2021; Bansal et al., 2021; Parrish et al., 2022; Lee et al., 2023; Chen et al., 2023; Joshi et al., 2023; Si et al., 2024; Fok and Weld, 2024), though without a specific focus on understanding how explanations affect humans’ perception of the plausibility of commonsense situations (which can be inherently ambiguous), a domain where humans are traditionally considered to be the “experts”.

To the best of our knowledge, our finding about the bimodal effect of PRO rationales on human judges: raising the ratings for implausible distractor answers but dropping the ratings for highly plausible gold-label answers has not been documented in prior works on LLM-generated explanations. Further, our causal OLS regression analysis also confirms a strong anchoring effect (a higher initial plausibility leads to a smaller change in the ratings) for all rater groups (humans and LLMs). The find-

ing that this anchoring effect is more pronounced for distractors than for gold-label answers for all groups is a new insight into the rigidity of initial judgments for incorrect answers.

7 Conclusion

In this work, we have tried to answer the critical question, *How do explanations work differently depending on the context, the audience (human vs. LLMs), and the baseline certainty of the claim?* By collecting 3000 human and 13,600 LLM ratings for 200 (q, a) pairs, we are able to demonstrate that both human and LLM plausibility ratings shift when rationales either in favor, against, or presented together are included with a commonsense reasoning question q and answer a , highlighting the potentially persuasive nature of LLM-generated explanations, specifically for commonsense reasoning tasks. Our findings also highlight that humans and LLMs do not always react similarly when shown different rationales, with models showing a higher sensitivity to the rationales, underscoring the challenges of aligning AI-generated judgments with human intuition, particularly in subjective reasoning tasks. Finally, our findings highlight the potential for LLMs to aid humans by challenging users' reasoning, stress-testing arguments, and introducing alternative perspectives, all of which have direct implications on effective human-AI collaboration, human-AI safety, AI for decision support, and the methodology of how commonsense reasoning is evaluated in AI systems. However, the fact that LLM-generated rationales are able to exert considerable influence on human annotators' beliefs, even in a domain like commonsense reasoning where humans are traditionally considered to be the "experts", also raises the question that if LLMs can sway human judgments on basic intuitive facts, the risk of persuasive manipulation or psychological influence might be higher in specialized or technical domains. We also demonstrate that LLMs can effectively steer human judgment toward more implausible inferences through persuasive argumentation. This highlights a critical vulnerability in human-AI collaboration: users may prioritize the plausibility of an LLM's argument over their own knowledge and beliefs. Because it is somewhat counter-intuitive that LLM reasoning can influence human reasoning in the domain of everyday commonsense reasoning, users may also be less likely to realize when such manipulations have occurred.

Addressing these risks requires robust safeguards, including transparency, bias mitigation, and mechanisms to detect and counteract harmful manipulative uses. Future works will focus on how factors such as cognitive abilities, age, grade-level and other demographic characteristics can influence human responses to LLM-generated explanations.

Limitations

Our goal in this work was to determine whether *any* LLM might produce rationales capable of swaying human (and other LLMs') plausibility judgments. For that purpose, GPT-4o proved effective at creating convincing arguments which, as we demonstrated, led to the changes in plausibility ratings. However, the persuasive power and style of LLM-generated rationales can vary significantly across different models and architectures. We hope that our findings inspire future works to look into the extent to which different models influence human and LLM plausibility judgments.

The scope of our investigation is confined to two English-language commonsense reasoning datasets, Social IQA and CommonsenseQA. The plausibility of an answer choice and the extent to which a rationale is persuasive might differ in other languages and domains, such as technical, scientific, or legal reasoning, where arguments may be more constrained by factual evidence. Future work may explore understanding the impact of LLM-generated rationales on domains beyond commonsense reasoning. Furthermore, our human annotators were recruited from a specific demographic (US-based, English-speaking users on Prolific). As common sense itself can be culturally situated, the impact of these rationales could vary across different cultural and linguistic groups. While we employed rigorous annotation protocols and made sure each instance was annotated by at least 5 annotators, there may still be instances where ambiguity in interpretation arises and hence, the impact a rationale has on the plausibility ratings would be different.

8 Acknowledgements

We would like to thank the anonymous reviewers for their valuable feedback on this paper. We would also like to thank Hal Daumé III, Marine Carpuat, Naomi Feldman, Srajit Sakhujia, Abhishashi Sancheti, Haozhe An, Neha Srikanth, Christabel Acquaye, Yu Hou, Nishant Balepur and other

members of the CLIP Lab for their helpful comments and suggestions. Rachel Rudinger and Shramay Palta are supported by NSF CAREER Award No. 2339746. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the National Science Foundation.

References

- Christabel Acquaye, Haozhe An, and Rachel Rudinger. 2024. [Susu box or piggy bank: Assessing cultural commonsense knowledge between Ghana and the US](#). In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 9483–9502, Miami, Florida, USA. Association for Computational Linguistics.
01. AI, :, Alex Young, Bei Chen, Chao Li, Chengen Huang, Ge Zhang, Guanwei Zhang, Guoyin Wang, Heng Li, Jiangcheng Zhu, Jianqun Chen, Jing Chang, Kaidong Yu, Peng Liu, Qiang Liu, Shawn Yue, Senbin Yang, Shiming Yang, and 14 others. 2025. [Yi: Open foundation models by 01.ai](#). *Preprint*, arXiv:2403.04652.
- Vicky Arnold, Philip A Collier, Stewart A Leech, and Steve G Sutton. 2000. The effect of experience and complexity on order and recency bias in decision making by professional accountants. *Accounting & Finance*, 40(2):109–134.
- Gagan Bansal, Tongshuang Wu, Joyce Zhou, Raymond Fok, Besmira Nushi, Ece Kamar, Marco Tulio Ribeiro, and Daniel Weld. 2021. [Does the whole exceed its parts? the effect of ai explanations on complementary team performance](#). In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*, CHI ’21, New York, NY, USA. Association for Computing Machinery.
- Rafael M. Batista and Thomas L. Griffiths. 2026. [A rational analysis of the effects of sycophantic ai](#). *Preprint*, arXiv:2602.14270.
- Chandra Bhagavatula, Ronan Le Bras, Chaitanya Malaviya, Keisuke Sakaguchi, Ari Holtzman, Hannah Rashkin, Doug Downey, Wen tau Yih, and Yejin Choi. 2020. [Abductive commonsense reasoning](#). In *International Conference on Learning Representations*.
- Mehar Bhatia and Vered Shwartz. 2023. [GD-COMET: A geo-diverse commonsense inference model](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 7993–8001, Singapore. Association for Computational Linguistics.
- Or Biran and Kathleen McKeown. 2017. Human-centric justification of machine learning predictions. In *Proceedings of the 26th International Joint Conference on Artificial Intelligence*, IJCAI’17, page 1461–1467. AAAI Press.
- Faeze Brahman, Vered Shwartz, Rachel Rudinger, and Yejin Choi. 2021. [Learning to rationalize for non-monotonic reasoning with distant supervision](#). *Proceedings of the AAAI Conference on Artificial Intelligence*, 35(14):12592–12601.
- Oana-Maria Camburu, Tim Rocktäschel, Thomas Lukasiewicz, and Phil Blunsom. 2018. [e-snli: Natural language inference with natural language explanations](#). In *Advances in Neural Information Processing Systems*, volume 31. Curran Associates, Inc.
- Chacha Chen, Shi Feng, Amit Sharma, and Chenhao Tan. 2023. [Machine explanations and human understanding](#). In *Proceedings of the 2023 ACM Conference on Fairness, Accountability, and Transparency*, FAccT ’23, page 1, New York, NY, USA. Association for Computing Machinery.
- Tongfei Chen, Zhengping Jiang, Adam Poliak, Keisuke Sakaguchi, and Benjamin Van Durme. 2020. [Uncertain natural language inference](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 8772–8779, Online. Association for Computational Linguistics.
- DeepSeek-AI. 2025. [Deepseek-r1: Incentivizing reasoning capability in llms via reinforcement learning](#). *Preprint*, arXiv:2501.12948.
- James N. Druckman. 2001. [Evaluating framing effects](#). *Journal of Economic Psychology*, 22(1):91–101.
- Laurette Dube-Rioux and J. Edward Russo. 1988. [An availability bias in professional judgment](#). *Journal of Behavioral Decision Making*, 1(4):223–237.
- Shi Feng and Jordan Boyd-Graber. 2019. [What can ai do for me? evaluating machine learning interpretations in cooperative play](#). In *Proceedings of the 24th International Conference on Intelligent User Interfaces*, IUI ’19, page 229–239, New York, NY, USA. Association for Computing Machinery.
- Raymond Fok and Daniel S Weld. 2024. [In search of verifiability:: Explanations rarely enable complementary performance in ai-advised decision making](#). *AI Magazine*, 45(3):317–332.
- Valerie S. Folkes. 1988. [The availability heuristic and perceived risk](#). *Journal of Consumer Research*, 15(1):13–23.
- Ana Valeria González, Gagan Bansal, Angela Fan, Yashar Mehdad, Robin Jia, and Srinivasan Iyer. 2021. [Do explanations help users detect errors in open-domain QA? an evaluation of spoken vs. visual explanations](#). In *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021*, pages 1103–1116, Online. Association for Computational Linguistics.

- Navita Goyal, Eleftheria Briakou, Amanda Liu, Connor Baumler, Claire Bonial, Jeffrey Micher, Clare Voss, Marine Carpuat, and Hal Daumé III. 2023. [What else do I need to know? the effect of background information on users' reliance on QA systems](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 3313–3330, Singapore. Association for Computational Linguistics.
- Aaron Grattafiori, Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Alex Vaughan, Amy Yang, Angela Fan, Anirudh Goyal, and 1 others. 2024. [The llama 3 herd of models](#). *Preprint*, arXiv:2407.21783.
- Herbert P Grice. 1975. Logic and conversation. In *Speech acts*, pages 41–58. Brill.
- Lisa Anne Hendricks, Zeynep Akata, Marcus Rohrbach, Jeff Donahue, Bernt Schiele, and Trevor Darrell. 2016. [Generating visual explanations](#). In *Proceedings of the 14th European Conference on Computer Vision (ECCV)*, Amsterdam, Netherlands.
- Daniel Hershcovich, Stella Frank, Heather Lent, Miryam de Lhoneux, Mostafa Abdou, Stephanie Brandl, Emanuele Bugliarello, Laura Cabello Piqueras, Ilias Chalkidis, Ruixiang Cui, Constanza Fierro, Katerina Margatina, Phillip Rust, and Anders Søgaard. 2022. [Challenges and strategies in cross-cultural NLP](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 6997–7013, Dublin, Ireland. Association for Computational Linguistics.
- Alon Jacovi, Ana Marasović, Tim Miller, and Yoav Goldberg. 2021. [Formalizing trust in artificial intelligence: Prerequisites, causes and goals of human trust in ai](#). In *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency, FAccT '21*, page 624–635, New York, NY, USA. Association for Computing Machinery.
- Albert Q. Jiang, Alexandre Sablayrolles, Antoine Roux, Arthur Mensch, Blanche Savary, Chris Bamford, Devendra Singh Chaplot, Diego de las Casas, Emma Bou Hanna, Florian Bressand, Gianna Lengyel, Guillaume Bour, and 1 others. 2024. [Mixtral of experts](#). *Preprint*, arXiv:2401.04088.
- Edward E Jones and Steven Bergias. 1976. A recency effect in attitude attribution. *Journal of Personality*, 44(3).
- Brihi Joshi, Ziyi Liu, Sahana Ramnath, Aaron Chan, Zhewei Tong, Shaoliang Nie, Qifan Wang, Yejin Choi, and Xiang Ren. 2023. [Are machine rationales \(not\) useful to humans? measuring and improving human utility of free-text rationales](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 7103–7128, Toronto, Canada. Association for Computational Linguistics.
- Arie W Kruglanski and Icek Ajzen. 1983. Bias and error in human judgment. *European Journal of Social Psychology*, 13(1):1–44.
- Mary Kynn. 2007. [The 'heuristics and biases' bias in expert elicitation](#). *Journal of the Royal Statistical Society Series A: Statistics in Society*, 171(1):239–264.
- Andrew Lampinen, Ishita Dasgupta, Stephanie Chan, Kory Mathewson, Mh Tessler, Antonia Creswell, James McClelland, Jane Wang, and Felix Hill. 2022. [Can language models learn from explanations in context?](#) In *Findings of the Association for Computational Linguistics: EMNLP 2022*, pages 537–563, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.
- Veronica Latcinnik and Jonathan Berant. 2020. Explaining question answering models through text generation. *arXiv preprint arXiv:2004.05569*.
- Mina Lee, Megha Srivastava, Amelia Hardy, John Thickstun, Esin Durmus, Ashwin Paranjape, Ines Gerard-Ursin, Xiang Lisa Li, Faisal Ladhak, Frieda Rong, Rose E Wang, Minae Kwon, Joon Sung Park, Hancheng Cao, Tony Lee, Rishi Bommasani, Michael S. Bernstein, and Percy Liang. 2023. [Evaluating human-language model interaction](#). *Transactions on Machine Learning Research*.
- Hector J. Levesque, Ernest Davis, and L. Morgenstern. 2011. [The winograd schema challenge](#). In *AAAI Spring Symposium: Logical Formalizations of Commonsense Reasoning*.
- Jiacheng Liu, Wenya Wang, Dianzhuo Wang, Noah Smith, Yejin Choi, and Hannaneh Hajishirzi. 2023. [Vera: A general-purpose plausibility estimation model for commonsense statements](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 1264–1287, Singapore. Association for Computational Linguistics.
- David R Mandel. 2014. Do framing effects reveal irrational choice? *Journal of Experimental Psychology: General*, 143(3):1185.
- Ana Marasovic, Iz Beltagy, Doug Downey, and Matthew Peters. 2022. [Few-shot self-rationalization with natural language prompts](#). In *Findings of the Association for Computational Linguistics: NAACL 2022*, pages 410–424, Seattle, United States. Association for Computational Linguistics.
- André C. R. Martins. 2006. [Probability biases as bayesian inference](#). *Judgment and Decision Making*, 1(2):108–117.
- Nasrin Mostafazadeh, Nathanael Chambers, Xiaodong He, Devi Parikh, Dhruv Batra, Lucy Vanderwende, Pushmeet Kohli, and James Allen. 2016. [A corpus and cloze evaluation for deeper understanding of commonsense stories](#). In *Proceedings of the 2016 Conference of the North American Chapter of the*

- Association for Computational Linguistics: Human Language Technologies*, pages 839–849, San Diego, California. Association for Computational Linguistics.
- Nadia Hanin Nazlan, Sarah Tanford, and Rhonda Montgomery. 2018. [The effect of availability heuristics in online consumer reviews](#). *Journal of Consumer Behaviour*, 17(5):449–460.
- OpenAI, :, Aaron Hurst, Adam Lerer, Adam P. Goucher, Adam Perelman, Aditya Ramesh, Aidan Clark, AJ Ostrow, Akila Welihinda, Alan Hayes, Alec Radford, Aleksander Mądry, Alex Baker-Whitcomb, Alex Beutel, and 1 others. 2024a. [Gpt-4o system card](#). *Preprint*, arXiv:2410.21276.
- OpenAI, :, Aaron Jaech, Adam Kalai, Adam Lerer, Adam Richardson, Ahmed El-Kishky, Aiden Low, Alec Helyar, Aleksander Madry, Alex Beutel, Alex Carney, Alex Iftimie, Alex Karpenko, Alex Tachard Passos, and 1 others. 2024b. [Openai o1 system card](#). *Preprint*, arXiv:2412.16720.
- OpenAI, Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altenschmidt, Sam Altman, Shyamal Anadkat, and 1 others. 2024c. [Gpt-4 technical report](#). *Preprint*, arXiv:2303.08774.
- Thorsten Pachur, Ralph Hertwig, and Florian Steinmann. 2012. How do people judge risks: availability heuristic, affect heuristic, or both? *Journal of Experimental Psychology: Applied*, 18(3):314.
- Shramay Palta, Nishant Balepur, Peter Rankel, Sarah Wiegrefe, Marine Carpuat, and Rachel Rudinger. 2024. [Plausibly problematic questions in multiple-choice benchmarks for commonsense reasoning](#). In *Findings of the Association for Computational Linguistics: EMNLP 2024*, pages 3451–3473, Miami, Florida, USA. Association for Computational Linguistics.
- Shramay Palta and Rachel Rudinger. 2023. [FORK: A bite-sized test set for probing culinary cultural biases in commonsense reasoning models](#). In *Findings of the Association for Computational Linguistics: ACL 2023*, pages 9952–9962, Toronto, Canada. Association for Computational Linguistics.
- Arjun Panickssery, Samuel R. Bowman, and Shi Feng. 2024. [LLM evaluators recognize and favor their own generations](#). In *The Thirty-eighth Annual Conference on Neural Information Processing Systems*.
- Bhargavi Paranjape, Julian Michael, Marjan Ghazvininejad, Hannaneh Hajishirzi, and Luke Zettlemoyer. 2021. [Prompting contrastive explanations for commonsense reasoning tasks](#). In *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021*, pages 4179–4192, Online. Association for Computational Linguistics.
- Alicia Parrish, Harsh Trivedi, Ethan Perez, Angelica Chen, Nikita Nangia, Jason Phang, and Samuel Bowman. 2022. [Single-turn debate does not help humans answer hard reading-comprehension questions](#). In *Proceedings of the First Workshop on Learning with Natural Language Supervision*, pages 17–28, Dublin, Ireland. Association for Computational Linguistics.
- Lianhui Qin, Vered Shwartz, Peter West, Chandra Bhagavatula, Jena D. Hwang, Ronan Le Bras, Antoine Bosselut, and Yejin Choi. 2020. [Back to the future: Unsupervised backprop-based decoding for counterfactual and abductive commonsense reasoning](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 794–805, Online. Association for Computational Linguistics.
- Nazneen Fatema Rajani, Bryan McCann, Caiming Xiong, and Richard Socher. 2019. [Explain yourself! leveraging language models for commonsense reasoning](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 4932–4942, Florence, Italy. Association for Computational Linguistics.
- Melissa Roemmele, Cosmin Adrian Bejan, and Andrew S Gordon. 2011. Choice of plausible alternatives: An evaluation of commonsense causal reasoning. In *2011 AAAI Spring Symposium Series*.
- Rachel Rudinger, Vered Shwartz, Jena D. Hwang, Chandra Bhagavatula, Maxwell Forbes, Ronan Le Bras, Noah A. Smith, and Yejin Choi. 2020. [Thinking like a skeptic: Defeasible inference in natural language](#). In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 4661–4675, Online. Association for Computational Linguistics.
- Maarten Sap, Hannah Rashkin, Derek Chen, Ronan Le Bras, and Yejin Choi. 2019. [Social IQa: Commonsense reasoning about social interactions](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 4463–4473, Hong Kong, China. Association for Computational Linguistics.
- Shlomi Sher and Craig RM McKenzie. 2008. Framing effects and rationality. *The probabilistic mind: Prospects for Bayesian cognitive science*, pages 79–96.
- Vered Shwartz, Peter West, Ronan Le Bras, Chandra Bhagavatula, and Yejin Choi. 2020. [Unsupervised commonsense question answering with self-talk](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 4615–4629, Online. Association for Computational Linguistics.
- Chenglei Si, Navita Goyal, Tongshuang Wu, Chen Zhao, Shi Feng, Hal Daumé Iii, and Jordan Boyd-Graber. 2024. [Large language models help humans verify](#)

- truthfulness – except when they are convincingly wrong. In *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, pages 1459–1474, Mexico City, Mexico. Association for Computational Linguistics.
- Dirk D Steiner and Jeffrey S Rain. 1989. Immediate and delayed primacy and recency effects in performance evaluation. *Journal of Applied Psychology*, 74(1):136.
- Alon Talmor, Jonathan Herzig, Nicholas Lourie, and Jonathan Berant. 2019. **CommonsenseQA: A question answering challenge targeting commonsense knowledge**. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4149–4158, Minneapolis, Minnesota. Association for Computational Linguistics.
- Amos Tversky and Daniel Kahneman. 1973. **Availability: A heuristic for judging frequency and probability**. *Cognitive Psychology*, 5(2):207–232.
- Su Wang, Greg Durrett, and Katrin Erk. 2018. **Modeling semantic plausibility by injecting world knowledge**. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 2 (Short Papers)*, pages 303–308, New Orleans, Louisiana. Association for Computational Linguistics.
- Mark E. Whiting and Duncan J. Watts. 2024. **A framework for quantifying individual and collective common sense**. *Proceedings of the National Academy of Sciences*, 121(4):e2309535121.
- Sarah Wiegrefe, Jack Hessel, Swabha Swayamdipta, Mark Riedl, and Yejin Choi. 2022. **Reframing human-AI collaboration for generating free-text explanations**. In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 632–658, Seattle, United States. Association for Computational Linguistics.
- Sarah Wiegrefe and Ana Marasovic. 2021. **Teach me to explain: A review of datasets for explainable natural language processing**. In *Thirty-fifth Conference on Neural Information Processing Systems Datasets and Benchmarks Track (Round 1)*.
- Sarah Wiegrefe, Ana Marasović, and Noah A. Smith. 2021. **Measuring association between labels and free-text rationales**. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 10266–10284, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Sheng Zhang, Rachel Rudinger, Kevin Duh, and Benjamin Van Durme. 2017. **Ordinal common-sense inference**. *Transactions of the Association for Computational Linguistics*, 5:379–395.

A Appendix

A.1 License for Artifacts

All datasets used in this work are publicly available and free to use on HuggingFace. All LLMs are used in accordance with the licenses released by their respective developing organizations.

A.2 Usage of AI Assistants

We only use AI Assistants to assist our writing to identify grammar errors, typos and rephrase terms for readability.

A.3 Computational Details

LLaMA-3.1-Instruct (8B, 70B), LLaMA-3.2-Instruct (3B), LLaMA-3.3-Instruct (70B), Mistral Instruct (8x22B), Yi-1.5-Chat (6B, 9B, 34B) and DeepSeek-R1 (Qwen 32B) were all run on eight NVIDIA:RTXA5000 GPUs and were allocated a total of 24 GPU hours to run all experiments. All OpenAI models were run on CPU and were allocated a total of 8 hours to run all experiments. Each LLM decodes with a maximum token generation length of 50, to ensure no extra verbiage is generated. We did not perform a hyperparameter search. All results are obtained from a single run.

For generating the rationales as mentioned in § 2, we set temperature to be 1 for GPT-4o. For generating plausibility ratings as discussed in § 4, we set temperature to be 0 and 0.1 for OpenAI models and Non-OpenAI models respectively to ensure greedy decoding.

A.4 Annotation Process Details

We used Prolific to collect the human annotations for all three rationale conditions. The annotators for our task were selected on the basis of the following criteria:

1. Must be located in the United States.
2. Primary language must be English.
3. Must not have any literacy difficulties.
4. Must have attained a minimum of an undergraduate level degree.
5. Must have an approval rate between 99 – 100% on Prolific.
6. We use a 50 – 50 split of male and female⁴ annotators to minimize the risk of any gender-specific biases creeping in.

⁴Gender as indicated on Prolific.

Dataset	Humans		OpenAI Models		Non-OpenAI Models	
	Gold Label	Distractor	Gold Label	Distractor	Gold Label	Distractor
SIQA	-0.0025	-0.0021	-0.0068	-0.0058	-0.0074	-0.0023
CQA	-0.0042	-0.0037	-0.0113	-0.0046	-0.0081	-0.0037

Table 5: OLS regression coefficients for the Rationale Length as a feature for ratings given by Humans, OpenAI and Non-OpenAI models. The dependent variable here is the change (Δ) in the plausibility rating after a rationale was presented. Numbers highlighted in **bold** represent coefficients that are statistically significant ($p < 0.05$) according to the t-test.

The total cost for our entire human annotations for all three rationales came out to be \$1000. We also received an exempt status from the IRB at our institution for this research. We do not collect any personally identifiable information from the annotators.

A.5 Prompts used for Rationale Generation and Plausibility Ratings

We release the prompts used in our work for generating the rationales and prompting the models for plausibility ratings below:

Prompt A.1: PRO Rationale Generation Prompt

Give your best rationale as to why the given choice is a plausible answer for the question. The rationale should not be more than 2 sentences long and should be concise. Your rationale has to be in favour of the plausibility of the choice, no matter what. Some of these situations may require creativity or thinking outside to box to come up with a rationale that best explains the plausibility of the answer choice. Feel free to be as creative as required.

Question: q

Choice: a

Generated Rationale: **PRO Rationale**

Prompt A.2: CON Rationale Generation Prompt

Give your best rationale as to why the given choice is an implausible answer for the question. The rationale should not be more than 2 sentences long and should be concise. Your rationale has to be in favour of the implausibility of the choice, no matter what. Some of these situations may require creativity or thinking outside to box to come up with a rationale that best explains the implausibility of the answer choice. Feel free to be as creative as required.

Question: q

Choice: a

Generated Rationale: **CON Rationale**

Prompt A.3: Plausibility Rating Without Rationales

I will give you a question and an answer choice. Your task is to rate the plausibility of the answer choice on the following Likert scale:

Likert Scale:

1 - Impossible

2 - Technically Possible

3 - Plausible

4 - Likely

5 - Very Likely

Your answer should just be a complete option from the Likert scale. Nothing else. No extra words or explanations.

Question: q

Choice: a

Generated Rating: r

Prompt A.4: Plausibility Rating With Pro Rationale

I will give you a question, an answer choice, and a rationale that best explains the plausibility of the answer choice. Your task is to rate the plausibility of the answer choice on the following Likert scale taking into consideration the rationale given.

Likert Scale:

1 - Impossible

2 - Technically Possible

3 - Plausible

4 - Likely

5 - Very Likely

Your answer should just be a complete option from the Likert scale. Nothing else. No extra words or explanations.

Question: q

Choice: a

Explanation: **PRO Rationale**

Generated Rating: r

Prompt A.5: Plausibility Rating With Con Rationale

I will give you a question, an answer choice, and a rationale that best explains the implausibility of the answer choice. Your task is to rate the plausibility of the answer choice on the following Likert scale taking into consideration the rationale given.

Likert Scale:

- 1 - Impossible
- 2 - Technically Possible
- 3 - Plausible
- 4 - Likely
- 5 - Very Likely

Your answer should just be a complete option from the Likert scale. Nothing else. No extra words or explanations.

Question: *q*

Choice: *a*

Explanation: Con Rationale

Generated Rating: r

Prompt A.6: Plausibility Rating With Pro+Con Rationales

I will give you a question, an answer choice, a rationale that best explains the plausibility of the answer choice, and a rationale that best explains the implausibility of the answer choice. Your task is to rate the plausibility of the answer choice on the following Likert scale taking into consideration the rationales given.

Likert Scale:

- 1 - Impossible
- 2 - Technically Possible
- 3 - Plausible
- 4 - Likely
- 5 - Very Likely

Your answer should just be a complete option from the Likert scale. Nothing else. No extra words or explanations.

Question: *q*

Choice: *a*

Explanation 1: Pro Rationale

Explanation 2: Con Rationale

Generated Rating: r

Agent Rating Change	Example
<p>Human Rating \uparrow</p> <p>LLM Rating \uparrow</p>	<p><i>Context:</i> Jenny wanted to find out what happened to Quinn’s leg and read Quinn’s book about his travels where he discussed his accident later in the first chapter.</p> <p><i>Question:</i> How would Jenny feel as a result?</p> <p><i>Choice:</i> satisfied (<i>gold label</i>)</p> <p><i>Pro Rationale:</i> Jenny would feel satisfied because she successfully found the information she was seeking about Quinn’s leg accident in his book, fulfilling her curiosity and providing her with the answers she needed.</p> <p>Human Rating Δ: +1.2</p> <p>LLM Rating Δ (OpenAI): +0.88</p> <p>LLM Rating Δ (Non-OpenAI): +0.67</p>
<p>Human Rating \uparrow</p> <p>LLM Rating \downarrow</p>	<p><i>Question:</i> I took my seat, the curtains drew back and I enjoyed the what?</p> <p><i>Choice:</i> movie (<i>distractor</i>)</p> <p><i>Pro Rationale:</i> The mention of "curtains drew back" is commonly associated with the start of a film in a theater setting, making "movie" a plausible choice as to what was enjoyed. Additionally, sitting down right before an experience occurs supports the idea of a movie theater viewing.</p> <p>Human Rating Δ: +0.2</p> <p>LLM Rating Δ (OpenAI): -0.12</p> <p>LLM Rating Δ (Non-OpenAI): -0.22</p>
<p>Human Rating \downarrow</p> <p>LLM Rating \uparrow</p>	<p><i>Context:</i> Quinn got a new job in a different town so they moved to a new house.</p> <p><i>Question:</i> What will Quinn want to do next?</p> <p><i>Choice:</i> reload his things (<i>distractor</i>)</p> <p><i>Pro Rationale:</i> Quinn may want to reload his things as part of the process of unpacking and arranging his belongings in his new home, ensuring everything is set up correctly and efficiently after the move. Additionally, reloading could refer to reorganizing or reconfiguring items like furniture and electronics to better suit the new living space.</p> <p>Human Rating Δ: -1.2</p> <p>LLM Rating Δ (OpenAI): +1</p> <p>LLM Rating Δ (Non-OpenAI): +1</p>
<p>Human Rating \downarrow</p> <p>LLM Rating \downarrow</p>	<p><i>Question:</i> The teacher played on the upright piano, she was explaining the song to all the students in the what?</p> <p><i>Choice:</i> music room (<i>gold label</i>)</p> <p><i>Pro Rationale:</i> The choice "music room" is plausible because it’s a common setting within a school where pianos are typically located and where a teacher would explain music to students, making it a logical location for the scenario.</p> <p>Human Rating Δ: -1</p> <p>LLM Rating Δ (OpenAI): -0.62</p> <p>LLM Rating Δ (Non-OpenAI): -0.33</p>

Table 6: Examples from SIQA and CQA showing mean plausibility rating changes for different agents for PRO Rationales.

Agent Rating Change	Example
<p>Human Rating ↑</p> <p>LLM Rating ↑</p>	<p><i>Question:</i> Where can many stores with clothing be found?</p> <p><i>Choice:</i> library (<i>distractor</i>)</p> <p><i>Con Rationale:</i> While a mall is a common place for multiple clothing stores, it is implausible to consider it the specific answer because a fashion district or online marketplaces might have far more clothing stores concentrated in one place.</p> <p>Human Rating Δ: +0.6</p> <p>LLM Rating Δ (OpenAI): 0</p> <p>LLM Rating Δ (Non-OpenAI): +0.33</p>
<p>Human Rating ↑</p> <p>LLM Rating ↓</p>	<p><i>Context:</i> Cameron wasted money on an expensive handbag.</p> <p><i>Question:</i> Why did Cameron do this?</p> <p><i>Choice:</i> resent her waste (<i>distractor</i>)</p> <p><i>Con Rationale:</i> It is implausible that Cameron would waste money on an expensive handbag with the intention of resenting her waste because typically, people do not make financial decisions to purposefully feel negative emotions like resentment. Such a motive is counterproductive and irrational, as purchases are usually made for positive reasons such as necessity, desire, or status, not to induce regret.</p> <p>Human Rating Δ: +1.6</p> <p>LLM Rating Δ (OpenAI): -0.62</p> <p>LLM Rating Δ (Non-OpenAI): -0.11</p>
<p>Human Rating ↓</p> <p>LLM Rating ↓</p>	<p><i>Context:</i> Riley layered down their arms with a blanket to keep warm.</p> <p><i>Question:</i> What does Riley need to do before this?</p> <p><i>Choice:</i> get a blanket from the closet (<i>gold label</i>)</p> <p><i>Con Rationale:</i> It is implausible that Riley needs to get a blanket from the closet because the sentence implies that the blanket is already in their possession, suggesting that any necessary retrieval must have already occurred. Additionally, there is no indication that Riley is not near a blanket or lacks access to one already placed conveniently nearby.</p> <p>Human Rating Δ: -3.2</p> <p>LLM Rating Δ (OpenAI): -1.38</p> <p>LLM Rating Δ (Non-OpenAI): -0.78</p>

Table 7: Examples from SIQA and CQA showing mean plausibility rating changes for different agents for CON Rationales. We observed no cases where there was a decrease in Human ratings, but decrease in LLM ratings.

Agent Rating Change	Example
<p>Human Rating ↑</p> <p>LLM Rating ↑</p>	<p><i>Context:</i> Aubrey was a police officer that pulled a driver over for speeding down the highway one day.</p> <p><i>Question:</i> What will Aubrey want to do next?</p> <p><i>Choice:</i> run away from the driver (<i>distractor</i>)</p> <p><i>Pro Rationale:</i> Aubrey might want to run away from the driver if the situation escalates into a dangerous confrontation or if they perceive an immediate threat to their safety, prompting them to seek cover or help. As a police officer, ensuring personal safety can sometimes take precedence in volatile situations.</p> <p><i>Con Rationale:</i> Running away from the driver is implausible because as a police officer, Aubrey’s duty is to address the traffic violation and ensure public safety, not to avoid the situation. Such an action would be unprofessional and contrary to law enforcement responsibilities.</p> <p>Human Rating Δ: +0.8</p> <p>LLM Rating Δ (OpenAI): +0.88</p> <p>LLM Rating Δ (Non-OpenAI): +1.11</p>
<p>Human Rating ↑</p> <p>LLM Rating ↓</p>	<p><i>Context:</i> Kendall took Skylar into account when making plans for the day.</p> <p><i>Question:</i> How would Skylar feel as a result?</p> <p><i>Choice:</i> Glad they were thought about (<i>gold label</i>)</p> <p><i>Pro Rationale:</i> Skylar would likely feel glad they were thought about because being considered in someone’s plans can lead to feelings of being valued and appreciated, reinforcing a sense of importance in their relationship with Kendall. This consideration shows that Kendall cares about Skylar’s preferences and needs, which can foster a positive emotional response.</p> <p><i>Con Rationale:</i> The choice is implausible because Skylar might have preferred to be surprised rather than having their preferences considered, as they enjoy spontaneity and unpredictability. Additionally, Skylar may have wanted a day without any plans, making the consideration unnecessary and possibly unwanted.</p> <p>Human Rating Δ: +0.2</p> <p>LLM Rating Δ (OpenAI): -1.12</p> <p>LLM Rating Δ (Non-OpenAI): -0.33</p>

Table 8: Examples from SIQA and CQA showing mean plausibility rating changes for different agents when PRO Rationales and CON Rationales are shown together.

Agent Rating Change	Example
<p>Human Rating ↓</p> <p>LLM Rating ↑</p>	<p><i>Question:</i> James loved to surf but he wasn't good at it. He would always do what?</p> <p><i>Choice:</i> enjoy yourself (<i>distractor</i>)</p> <p><i>Pro Rationale:</i> Even though James wasn't skilled at surfing, his passion and love for it ensured that he focused more on the experience and thrill of being in the water, allowing him to enjoy himself regardless of his performance. This positive attitude highlights that enjoyment does not solely depend on skill level, but on the enthusiasm and engagement in the activity.</p> <p><i>Con Rationale:</i> The phrase "enjoy yourself" implies a level of satisfaction and contentment, which contradicts the notion of James not being good at surfing, as his lack of skill is likely to cause frustration rather than enjoyment. Therefore, it is implausible that James would consistently enjoy himself given his admitted lack of proficiency.</p> <p>Human Rating Δ: -0.6</p> <p>LLM Rating Δ (OpenAI): $+0.75$</p> <p>LLM Rating Δ (Non-OpenAI): $+0.11$</p>
<p>Human Rating ↓</p> <p>LLM Rating ↓</p>	<p><i>Question:</i> Aside from water and nourishment what does your dog need?</p> <p><i>Choice:</i> lots of attention (<i>gold label</i>)</p> <p><i>Pro Rationale:</i> Dogs are social animals that thrive on interaction and companionship, making lots of attention essential for their emotional well-being and development. By providing attention, owners can strengthen their bond with their dog, enhance socialization skills, and reduce the risk of behavioral issues stemming from loneliness or neglect.</p> <p><i>Con Rationale:</i> While dogs may enjoy attention, it is not a fundamental need comparable to shelter or regular exercise, making "lots of attention" an implausible equivalent to basic survival necessities.</p> <p>Human Rating Δ: -0.4</p> <p>LLM Rating Δ (OpenAI): -1</p> <p>LLM Rating Δ (Non-OpenAI): -0.33</p>

Table 9: (Continued from Table 8) Examples from SIQA and CQA showing mean plausibility rating changes for different agents when PRO Rationales and CON Rationales are shown together.

For the given context, question and answer choice, choose the displayed rationale which best explains the plausibility of the answer choice, and is also the most believable rationale.
Context: Jan had a long intensive labor but gave birth to a daughter that day.

Question: How would you describe Jan?

Answer Choice: disappointed in the result

Rationale 1: Jan might be "disappointed in the result" if she was expecting a boy and had mentally prepared for that outcome, or if she had specific expectations about the birth that were not met. Additionally, the stress and exhaustion from a long labor might temporarily overshadow her joy, leading to feelings of disappointment.

Rationale 2: It's plausible that Jan might feel disappointed in the result if she had hoped for a different outcome, such as a son, especially after enduring a long and intensive labor. Additionally, the physical and emotional toll of a difficult delivery could amplify any feelings of letdown, regardless of a healthy birth.

Rationale 3: Jan might have been expecting a son based on various family traditions or personal preferences, leading her to feel initially disappointed by the outcome of having a daughter despite the joy of motherhood. This emotional reaction could be a temporary adjustment period, where she navigates her expectations versus reality.

Rationale 4: Jan could be described as disappointed in the result if she was strongly hoping for a son, possibly due to cultural or family expectations, and despite the successful delivery, her initial reaction is one of letdown. This emotional response might also stem from pre-birth preparations tailored specifically for a boy, making the unexpected outcome jarring for her.

Please leave any feedback about the above survey item (if you have any) below:

Figure 4: An example of the interface that annotators used while choosing the best rationale for an answer choice as described in § 2.1. Example taken from SIQA.

Kendall took Skylar into account when making plans for the day.

How would Skylar feel as a result?

Answer: Glad they were thought about

Please rate the plausibility of the **highlighted answer** on the following 5-point scale, taking into consideration the underlined explanation below.

1 - Impossible 2 - Technically Possible 3 - Plausible 4 - Likely 5 - Very Likely

Answer: Glad they were thought about

Explanation: Skylar would likely feel glad they were thought about because being considered in someone's plans can lead to feelings of being valued and appreciated, reinforcing a sense of importance in their relationship with Kendall. This consideration shows that Kendall cares about Skylar's preferences and needs, which can foster a positive emotional response.

Please leave any feedback about the above survey item (if you have any) below:

Figure 5: An example of the interface that annotators used while giving plausibility ratings to answer choices with a PRO rationale as described in § 3. Example taken from SIQA.

When you fail to finish something, you failed at doing what to it

Answer: completing

Please rate the plausibility of the **highlighted answer** on the following 5-point scale, taking into consideration the underlined explanation below.

1 - Impossible 2 - Technically Possible 3 - Plausible 4 - Likely 5 - Very Likely

Answer: completing

Explanation: The choice "completing" is implausible because failing to finish something inherently means that you did not complete it; thus, you cannot fail at completing something if completing is the very action you did not achieve.

Please leave any feedback about the above survey item (if you have any) below:

Figure 6: An example of the interface that annotators used while giving plausibility ratings to answer choices with a CON rationale as described in § 3. Example taken from CQA.

Aside from water and nourishment what does your dog need?

Answer: charm

Please rate the plausibility of the **highlighted answer** on the following 5-point scale, taking into consideration the underlined explanations below.

Answer: charm 1 - Impossible 2 - Technically Possible 3 - Plausible 4 - Likely 5 - Very Likely

○ ○ ○ ○ ○

Explanation 1: Charm can enhance a dog's interactions with humans and other pets, making it more likely to receive affection, playtime, and social engagement, all of which contribute to its emotional well-being. Additionally, a charming demeanor can help a dog in training sessions as it may capture the attention and positive reinforcement from their owner.

Explanation 2: It is implausible that a dog needs charm because charm is a subjective human quality rather than a physical or emotional necessity for a dog's well-being. Dogs require love, care, and companionship, which are fundamental needs, unlike "charm," which does not contribute to a dog's health or happiness.

Please leave any feedback about the above survey item (if you have any) below:

Figure 7: An example of the interface that annotators used while giving plausibility ratings to answer choices with both PRO and CON rationales as described in § 3. Example taken from CQA.