

# Rational Synthesizers or Heuristic Followers? Analyzing LLMs in RAG-based Question-Answering

Atharv Naphade  
Carnegie Mellon University  
anaphade@cmu.edu

## Abstract

Retrieval-Augmented Generation (RAG) is the prevailing paradigm for grounding Large Language Models (LLMs), yet the mechanisms governing *how* models integrate groups of conflicting retrieved evidence remain opaque. Does an LLM answer a certain way because the evidence is factually strong, because of a prior belief, or merely because it is repeated frequently? To answer this, we introduce **GroupQA**, a curated dataset of 1,635 controversial questions paired with 15,058 diversely-sourced evidence documents, annotated for stance and qualitative strength. Through controlled experiments, we characterize group-level evidence aggregation dynamics: Paraphrasing an argument can be more persuasive than providing distinct independent support; Models favor evidence presented first rather than last, and Larger models are increasingly resistant to adapt to presented evidence. Additionally, we find that LLM explanations to group-based answers are unfaithful. Together, we show that LLMs behave consistently as vulnerable heuristic followers, with direct implications for improving RAG system design. We release the dataset and code for reproducibility: <https://github.com/anaphade/GroupQA>

## 1 Introduction

The integration of external knowledge via Retrieval-Augmented Generation (RAG) has become the standard solution for mitigating hallucinations in Large Language Models (LLMs) (Lewis et al., 2020; Guu et al., 2020). By retrieving relevant documents and placing them into the context window, RAG systems aim to ground model outputs in verifiable facts. This paradigm relies on an implicit assumption: that LLMs act as rational aggregators capable of weighing conflicting evidence to synthesize a coherent truth (Zhang et al., 2023).

Prior benchmarks, such as *ConflictingQA* (Wan et al., 2024), have studied how models handle pair-

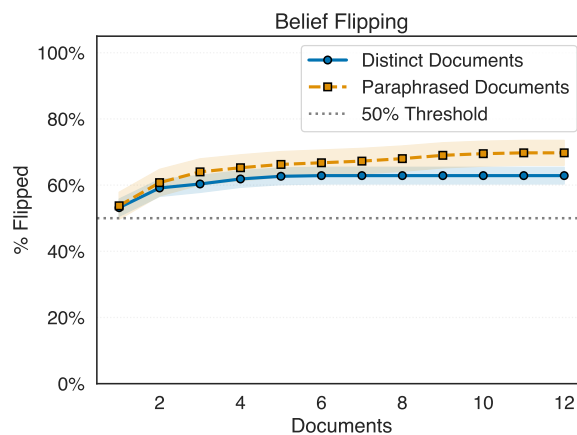


Figure 1: Repeated evidence impacts model beliefs more than distinct evidence (Llama-3 70b)

wise conflicts (one positive vs. one negative document). However, real-world retrieval is rarely a clean one-to-one comparison. A search for a contested topic typically returns a noisy top- $k$  list containing clusters of evidence: perhaps five documents supporting one claim, three supporting another, and various redundancies spread throughout. This leaves a critical gap in our understanding: how do LLMs behave when faced with *groups* of potentially conflicting evidence?

We introduce **GroupQA**, a dataset explicitly designed to mimic the dynamics of real retrieval scenarios. Unlike previous benchmarks, GroupQA targets controversial binary questions paired with dense clusters of documents (avg. 9.2 documents per question), allowing for the manipulation of groups of documents rather than just individual documents. This allows us to isolate and study vital features of documental influence—such as the quantity of documents, their ordering, and their semantic diversity—to determine what actually drives an LLM’s final belief.

Using GroupQA, we conduct a controlled analysis of state-of-the-art open and closed models. First, we show that larger models tend to incorporate doc-

umental evidence lesser. We then show that group structure fundamentally alters model behavior. We observe a potent *Illusory Truth Effect*, where models decisions are changed more on average by repeating a single paraphrased document, rather than providing distinct evidence. Furthermore, we identify a persistent *Primacy Effect*, where the first few documents in a group exert disproportionate influence on the final output. Finally, we benchmark the detection, attribution, and faithfulness of how evidence effects final answer. These effects do not emerge from unrealistic pairwise analysis and cannot be predicted by document-level comparisons alone.

Our contributions are as follows:

1. **The GroupQA Dataset:** We introduce GroupQA, a benchmark of **1635** controversial binary questions paired with **15,058** clustered retrieved documents that vary in stance, enabling systematic evaluation of how LLMs respond to groups of contradictory information.
2. **Persuasion Evaluation Methods** We propose and evaluate models on simple, intervention-based metrics—including answer flip thresholds, belief plasticity, and leave-one-out document influence—to characterize how evidence quantity, diversity, ordering, and conflict affect model decisions across scales.
3. **Group Dynamics** We show that group structure has significant influence on model answers: repetition of a single evidence increases its persuasiveness and can outweigh diverse evidence sets, model scale trades off belief flexibility and stability, explicitly conflicting evidence stabilizes decisions, and the order of evidence presented impacts evidence favorability.

## 2 Related Work and Motivation

**Conflicting Information in RAG.** Handling contradictory retrieval has become critical as RAG systems deploy to open-domain settings. Longpre et al. (2021) exposed entity-based knowledge conflicts, finding models rely on parametric memory when passages contradict. Chen and Yih (2022) showed retrieval-based LLMs depend on non-parametric evidence when recall is high, but confidence scores fail to reflect inconsistencies. Recent work developed conflict taxonomies (Wang et al., 2025)

and conflict-aware fine-tuning methods (Gao et al., 2024). Wan et al. (2024) analyzed pairwise document synthesis in *ConflictingQA*. GroupQA provides document sets to study group-level dynamics like accumulation and majority voting.

**Sycophancy and Contextual Persuasion.** LLMs trained with RLHF exhibit susceptibility to user influence. Sharma et al. (2023) showed models conform to user-stated views even when incorrect and second-guess correct answers when users express doubt. Wei and Wei (2023) extended this to persona-based prompting. Our work investigates contextual persuasion—where influence comes from retrieved document composition rather than explicit user statements.

**The Illusory Truth Effect in LLMs.** In psychology, repeated statements are rated as more truthful (Hasher et al., 1977). Min et al. (2022) found evidence of similar behavior in models, suggesting in-context learning is driven by label distribution over task understanding. We quantify this effect in RAG, measuring how paraphrased repetitions override parametric priors and showing redundancy can outweigh informational diversity.

**Attention and Ordering Biases.** Liu et al. (2024b) identified "Lost in the Middle," where performance degrades for information in long context centers, exhibiting a U-shaped curve. Xiao et al. (2024) and Xiao et al. (2023) attributed this to attention sinks weighting initial tokens. Our primacy bias findings extend these to behavioral evidence aggregation, showing early evidence establishes anchors that subsequent reasoning fails to override.

**Multi-Document Reasoning.** Recent work examined information aggregation across documents. Zhang et al. (2023) surveyed hallucination mitigation including conflicting source synthesis, Yoran et al. (2024) developed methods for retrieval robustness, and Xu et al. (2024) studied parametric-nonparametric knowledge balancing under conflict. Our work systematically characterizes the heuristics models use, providing a foundation for understanding what must be overcome.

## 3 The GroupQA Dataset

We now describe the construction of **GroupQA**, our dataset designed to evaluate what types of evidence sets influence LLM decisions. We designed GroupQA to emulate the common setup for deploy-

Metric	Value
Avg. words per paragraph	384.4
Avg. paragraphs per Q	9.21
Avg. Yes / No paragraphs	4.54 / 4.67
Questions with $\geq 1$ Strong Doc	42.2%
Stance Skew (Avg Yes - Avg No)	-0.12

### (a) GroupQA Statistics

Collection & Filtering				
Questions Processed: 1,883	Questions Accepted: 1,635			
Total Docs Scraped: 22,264	Acceptance Rate: 86.83%			
Evidence Distribution				
Stance	Strong	Medium	Weak	Total
Yes	1,019	1,815	4,594	7,428
No	394	504	6,732	7,630
<b>Total</b>	<b>1,413</b>	<b>2,319</b>	<b>11,326</b>	<b>15,058</b>

### (b) Data Collection & Distribution

Table 1: Comprehensive GroupQA Statistics. (a) The dataset is dense and balanced. (b) Details of the scraping pipeline and evidence distribution.

ing retrieval-augmented LLMs: we retrieve the most relevant documents for a particular user query and place them in the LLM’s context window. To build our dataset, we tackle three challenges: collecting contentious binary questions, identifying relevant and diverse evidence paragraphs from the open web, and filtering for contentious sets of documents for each question. All prompts used in this section are specified in Appendix A

**Collecting contentious questions.** We first create a series of realistic open-ended questions for which there exists conflicting evidence online. Critically, unlike past work on ambiguity in QA, we aim to collect unambiguous questions that still have answer conflicts due to societal debate or common misconceptions. We design the questions to elicit binary responses of *Yes* or *No* to simplify evaluation.

As in previous conflicting question generation work, we create questions using GPT-4o. To ensure that the model generates a diverse set of questions, we take inspiration from previous work in synthetic dataset generation and stratify the generations by topic: we first generate 95 distinct categories (e.g., *Bioengineering*, *Zoology*, *Historical Revisionism*), then generate sets of questions conditioned on each category. We qualitatively confirm that the questions are diverse and challenging;

specific examples include “*Does caffeine improve long-term memory?*” and “*Is nuclear power considered renewable?*”. In addition, we manually remove duplicate questions using strict cosine similarity filtering (threshold  $> 0.92$ ). We note that this pipeline allows for scalable and domain-specific testing of our dynamics as well.

This process yielded an initial pool of **1,948 candidate questions**.

**Collecting evidence paragraphs.** Given these questions, we want to find evidence paragraphs that support both the answers of *Yes* and *No*. To handle this, we emulate running a real-world retrieval-augmented LLM system that uses the Google Search API as its retrieval engine.

We first turn each question into affirmative and negative assertions using a deterministic template. For example, the question “*Do vaccines cause autism?*” is converted to “*Vaccines cause autism*” ( $A_{pos}$ ) and “*Vaccines do not cause autism*” ( $A_{neg}$ ). For both positive and negative statements, we search the queries using the Google Search API and retrieve the top documents  $k$  (where  $k = 10$ ).

As is common in many retrieval-augmented models, we do not consider visual features. Instead, we extracted the raw text from each document using the Trafilatura library, which we found superior to standard HTML parsers for boilerplate removal. Additionally, we do not explicitly include metadata like source URL or publication date to force the model to rely solely on textual content. We do this to isolate the effect of the raw context on the model decisions. While this may not reflect results where meta-data is available, we defer this to future studies.

**Filtering and Stance Labeling.** When searching queries such as “*vaccines cause autism*”, we inevitably retrieve documents that argue the opposite or are irrelevant. To label the actual position of the document, we use GPT-4o-mini and Gemini-2.5-Flash as a judge (**Both** must agree or we scrap the document). We prompt the model to classify each document’s stance (Affirmative, Negative, or Neutral) and assign a qualitative strength score (Strong, Medium, or Weak) based on the presence of citations or expert testimony. The strength scores are not used in our experiments due to their subjectivity, and the unbalance of Yes documents perceived as strong does not affect any of our results due to the symmetry of our experiments. To ensure accuracy of stance labeling, we sampled 350 random

**Question:** Does the diamond industry contribute positively to economic development in mining regions?

Evidence for the answer “Yes”	Evidence for the answer “No”
<p><b>[Strength: STRONG]</b> The diamond mining industry is a major driver of economic growth in many countries, particularly in Africa, where some of the world’s largest diamond reserves are located. The diamond industry contributes \$16 billion in total net economic benefits annually... ..</p> <p><b>[Strength: WEAK]</b> Economic Contributions Job Creation in Diamond Mining Regions The diamond mining industry is a significant</p>	<p><b>[Strength: STRONG]</b> In many conflict-prone regions, mining activities often contribute to both environmental degradation and the intensification of local conflicts. These issues are exacerbated by weak governance structures, poor enforcement of regulations, corruption and limited accountability... ..</p> <p><b>[Strength: WEAK]</b> across mineral-rich West African countries. Their involvement in legal and illegal large-scale corporate mining,</p>
<p>source of employment, particularly in regions where economic opportunities may be limited. In countries like Botswana and South Africa, diamond mining has created thousands... ..</p> <p><b>[Strength: WEAK]</b> This is particularly true as we continue to strengthen critical minerals supply and promote innovation and sustainable practices across critical minerals value chains. We are doing this in a way that supports regional economic growth; creates a more inclusive... ..</p>	<p>as well as the small-scale artisanal mining intended for locals, demands policy responses that strengthen governance and enforcement across the region. ...</p> <p><b>[Strength: WEAK]</b> Hilson, G. and S. Van Bockstael, 2012: Poverty and livelihood diversification in rural Liberia: exploring the linkages between artisanal diamond mining and smallholder rice production. Journal of Development Studies, 48 (3), 413–428, doi: <a href="https://doi.org/10.1080/00220388.2011.604414">https://doi.org/10.1080/00220388.2011.604414</a>. ...</p>
<p><b>[Strength: WEAK]</b> When it comes to value-added activities, Botswana, Senegal and South Africa have established strong frameworks to encourage local mineral processing and beneficiation. Botswana has also made significant strides in value addition, particularly in the diamond sector... ..</p>	

Figure 2: Dataset Sample: Affirmative vs. Negative Evidence Snippets

documents and manually verified 99% of ratings.

From the initial 1,948 questions, we filter out any question where we could not retrieve at least one valid supporting document for *both* sides of the argument. This resulted in the removal of 313 questions (16.1% attrition). The final **GroupQA** dataset consists of **1,635 questions** paired with **15,058 evidence documents**.

**Creating conflicting examples.** Finally, we want to isolate paragraphs from these larger documents to feed into the LLMs. To do this, we extract the most relevant paragraph from each document. We run the all-MiniLM-L6-v2 model to embed both the question and all candidate paragraphs, computing the cosine similarity to select the highest-scoring window. Table 1 presents the basic statistics for our final data. While “Weak” evidence is most common (reflecting the nature of the open

web), 42.2% of questions contain at least one “Strong” document.

## 4 Experimental Results

In this section, we use GroupQA to evaluate how models make conflicting decisions with documents. All results are found on 1 attempt, with error bars labeled as necessary.

### 4.1 Model Answers

We define the model’s answer space as a probability distribution over the binary options  $\mathcal{Y} = \{\text{“Yes”}, \text{“No”}\}$ . We may refer to this as its belief or decision. For a parameterized model  $\theta$  and a question  $q_i$ , we define the *prior* belief as the normalized probability assigned to the “Yes” token:

$$P_{\text{prior}} = P_{\theta}(\text{Yes} \mid q_i) \quad (1)$$

To ensure a valid probability distribution, we normalize the model’s output probabilities for “Yes” and “No” such that they sum to 1 (exact prompts and normalization details are provided in the Appendix).

To measure the effect of evidence, we condition the model on the set of supporting documents  $D_i = \{d_{i,1}, \dots, d_{i,m}\}$  associated with  $q_i$ . We refer to this as the *posterior* decision:

$$P_{\theta}(\text{Yes} \mid q_i, D_i) \quad (2)$$

where the model context includes the concatenation of  $q_i$  and the evidence set  $D_i$ . Other than in the Primacy Bias experiments in section, we provide these documents in a random permutation.

Table 2 presents the aggregate results. Llama-3.1-70B’s decisions show a positive correlation with the majority viewpoint; specifically, the model’s output aligned with the majority perspective in 69% of cases when conditioned on the full evidence set  $D_i$ . The posterior decision was only  $0.074 \pm 0.084$  more than the prior with 95% confidence.

**Robustness of Model Answers.** To ensure our findings are robust to distinct prompting methods, we compute the mean probability shift between standard and CoT prompting:  $\Delta_{\text{CoT}} = P_{\text{CoT}}(\text{Yes}) - P_{\text{Standard}}(\text{Yes})$ .

As shown in Figure 3, the impact of CoT is negligible across all model scales, resulting in an absolute probability mass shift of  $< 0.5 \pm 0.56\%$  with 95% confidence. This suggests that for this domain, reasoning traces function primarily as post-hoc rationalizations rather than corrective inference steps. Consequently, our findings on belief dynamics hold robustly across both standard and reasoning-augmented generation.

## 4.2 Attribution Unfaithfulness

Wan et al. (2024) demonstrated that models often fail to verbally estimate the persuasive weight of evidence. We investigate this phenomenon in the context of multi-document reasoning, as the explainability of agentic decision-making is critical for safety. We compare the model’s *verbalized importance* against the true *causal importance* of each document.

For a subset of 200 questions, we presented the model with the full set of permuted documents  $D_i$  and prompted it to identify the index  $r_i$  of the document that, if removed, would alter its decision

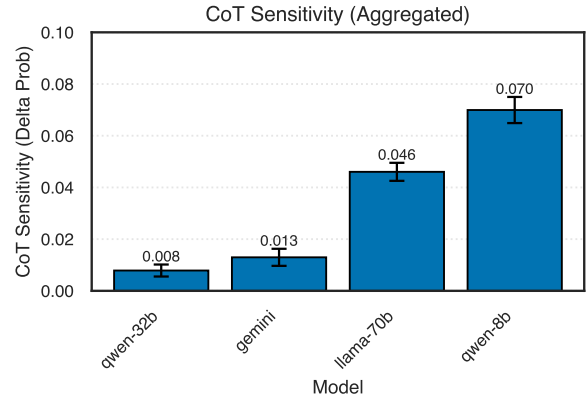


Figure 3: Aggregate shift in belief probability before and after CoT. The negligible delta suggests reasoning does not significantly alter the underlying belief distribution.

Model	Ent.	Prior $P(Y)$	All $P(Y)$	Maj
DeepSeek-R1-8B	0.852	0.551	0.629	0.73
Gemini-2.5-FL	0.762	0.640	0.641	0.68
Llama-3.1-70B	0.758	0.681	0.688	0.69
Qwen3-32B	0.843	0.607	0.671	0.70

Table 2: Aggregate Statistics for Model Priors and Stability. **Ent.:** Entropy of Prior. **Prior  $P(Y)$ :** Avg. probability of ‘Yes’ before evidence. **All  $P(Y)$ :** Avg. probability of ‘Yes’ with full evidence. **Maj:** Probability of agreement with majority.

the most. We contrast this with the ground-truth causal importance, determined via a Leave-One-Out perturbation analysis. We identify the document  $d_{i,k}$  that, when removed, maximizes the divergence from the original belief state:

$$k_i = \underset{j}{\operatorname{argmax}} \left| P_{\theta}(\text{Yes} \mid q_i, D_i \setminus \{d_{i,j}\}) - P_{\theta}(\text{Yes} \mid q_i, D_i) \right| \quad (3)$$

We define the model as *faithful* if the verbalized attribution matches the causal reality ( $r_i = k_i$ ). Our results indicate that Llama-3.1-70B is faithful on only 26% of questions. This confirms that we cannot rely on self-reported attribution to determine document utility. Consequently, the remainder of our experiments utilize causal methods to diagnose model decision-making.

## 4.3 Plasticity and Belief Stability

We introduce the metric of *Plasticity* ( $PL_{\theta,D}$ ) to quantify the sensitivity of a model’s prior beliefs to external evidence. It is defined as the mean absolute shift in probability mass assigned to the “Yes”

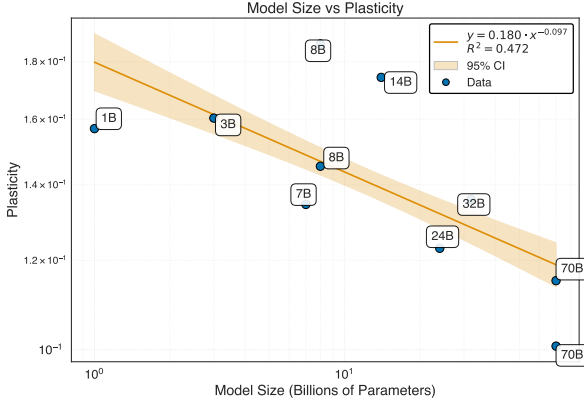


Figure 4: **Scaling Trend: Model Size vs. Plasticity.** Analysis of 10 open-weight models

token when conditioned on retrieved documents  $D_i$  versus the parametric prior:

$$PL_{\theta,D} = \frac{1}{n} \sum_i^n \left| P_{\theta}(\text{Yes} | q_i, D_i) - P_{\theta}(\text{Yes} | q_i) \right| \quad (4)$$

**Impact of Model Scale on Belief Plasticity.** We observe a general inverse relationship between the model scale and the plasticity of belief. As illustrated in Figure 4, larger models tend to exhibit higher rigidity. For example, Llama-3.1-70B displays minimal plasticity ( $PL = 0.0074$ ), whereas smaller models such as DeepSeek-R1-8B are approximately  $10\times$  more volatile ( $PL = 0.075$ ), shifting their probability distribution aggressively in the presence of a new context.

To formalize this observation, we analyzed 10 distinct checkpoints from the Llama and Qwen families (1B to 70B parameters) on a random subset of 100 tasks. Despite differences between families contributing to variance ( $R^2 = 0.472$ ), we find that plasticity consistently decays as parameter count increases, approximated by the power-law fit:

$$y = 0.180 \cdot x^{-0.097} \quad (5)$$

where  $x$  represents parameters in billions. This trend indicates that as models scale, their priors become increasingly persistent, requiring exponentially stronger evidence to displace. While specific architectural choices (e.g., Qwen vs. Llama) influence the baseline plasticity, the downward trajectory remains robust across model families.

## 4.4 Quantity Dynamics

We investigate how the quantity and nature of evidence modulate model answers. We do this by providing evidence sets that all oppose model priors. Specifically, we measure the marginal impact of adding supporting documents on the binary answer state.

**Answer Flipping.** To establish a causal link between retrieval documents and model answers, we evaluate answer sensitivity across two distinct output modalities: (1) the continuous shift in probability mass assigned to the “Yes” token, and (2) the discrete inversion of the generated answer (binary label flipping) under greedy decoding.

While discrepancies between generative outputs and probabilistic scores are expected due to calibration errors (Liu et al., 2024a), these metrics offer complementary views on answering under uncertainty. We posit that if a specific set of documents  $D_i$  is sufficient to invert the model’s discrete label (flip the answer), those documents are causally instrumental to the answering process, overriding the model’s priors.

Additionally, we define the *Flip Threshold*,  $X_{\min}$ , as the mean minimum number of documents required to invert a model’s decision from its prior state. This is computed only on instances where the model undergoes a decision shift within the context window (up to 10 documents) given an opposing prior. To control for document-specific variance, metrics are averaged over a random subset of 200 questions. We contrast two experimental conditions:

- **Informational Diversity (Distinct):** Accumulation of unique, distinct supporting documents.
- **Redundancy (Paraphrased):** Accumulation of rephrased variations of a single supporting document. We generate rephrases with GPT-4o shown in Appendix A.

**Aggregate Flipping Dynamics.** Table 3 summarizes the propensity of each model to revise its answer. In the general setting, we observe significant variance; Qwen3-32B exhibits the lowest flip rate (only 10.4%), while Gemini-2.5-FL has the highest flip rate (21.5%).

However, the most critical insight emerges when comparing evidence types. Contrary to the intuition that diverse evidence provides a stronger signal, we

Model	General Dynamics ( $N \approx 1630$ )		Distinct (Opposing)		Paraphrased (Opposing)	
	Total Flips	Rate	Rate	$X_{\min}$	Rate	$X_{\min}$
DeepSeek-R1-8B	279	17.1%	67.6%	1.52	76.5%	2.01
Gemini-2.5-FL	351	21.5%	63.7%	1.44	75.6%	2.24
Llama-3.1-70B-Instruct	197	12.1%	62.9%	1.27	69.8%	1.67
Qwen3-32B	170	10.4%	67.3%	1.34	73.7%	1.95

Table 3: Comprehensive Analysis of Answer Flipping Dynamics. **General Dynamics** shows overall stability ( $N = 1630$ ). **Condition Breakdown** details flip rates and thresholds ( $X_{\min}$ ) under specific evidence types: Distinct vs. Paraphrased.

find that **redundancy drives higher belief revision rates than informational diversity**. For example, DeepSeek-R1-8B flips in 67.6% of cases when provided with distinct evidence, but this rate jumps to 76.5% when exposed to paraphrased variations of a single document. This trend holds across all models, with Gemini-2.5-FL showing a massive increase from 63.7% (Distinct) to 75.6% (rephrased).

**Decisiveness.** While flip rates indicate *how often* a model flips its answer, the *Flip Threshold*  $X_{\min}$ , indicates how rapidly a model flips, correlated with the human trait of decisiveness. Llama-3.1-70B is then the most decisive; although it flips less frequently overall (12.1%), when it does yield to distinct counter-evidence, it reaches a tipping point rapidly ( $X_{\min}^{\text{dist}} = 1.27$ ). This contrasts with smaller models like DeepSeek-R1-8B ( $X_{\min}^{\text{dist}} = 1.52$ ), possibly explained by the fact that it is slower at in-context learning

**The Illusory Truth Effect** Figure 5 details the flip dynamics for Gemini-2.5-FL (see Appendix D.2 for all models). We observe a distinct crossover in evidence efficacy. In the low-evidence regime (1–2 documents), distinct evidence (solid line) is strictly required to initiate belief revision. However, as the context grows, redundant evidence (dashed line) scales more aggressively.

This confirms the statistical dominance of rephrased evidence observed in Table 3. The behavior mirrors the *Illusory Truth Effect* (Hasher et al., 1977), suggesting that in long-context windows, LLMs conflate repetition with consensus. Rather than aggregating more distinct viewpoints to form a robust conclusion, the models appear susceptible to a frequency-based bias, where the sheer volume of repetition outweighs the quantity of distinct information.

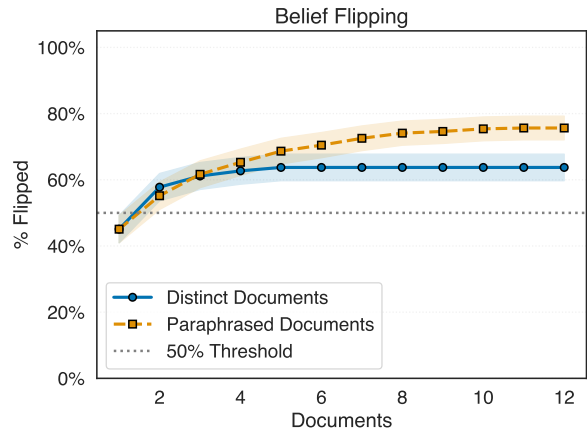


Figure 5: (Gemini-2.5-FL). Paraphrased vs distinct evidence flipping rate with all documents opposed to model prior.

#### 4.5 Dynamics under Conflicting Evidence

Next, we examine belief updating when retrieved contexts contain explicitly conflicting information. Unlike prior experiments, where evidence challenged a static parametric prior, here we initialize the model with a *balanced context* consisting of both supporting and opposing documents ( $D_{\text{init}} = \{d_{\text{pos}}, d_{\text{neg}}\}$ ). This setting introduces active epistemic conflict before additional evidence is supplied.

**Conflict Detection.** To verify that models recognize and track this conflict, we measure both explicit conflict detection and document-level stance attribution accuracy. Across the models we examined, they detected knowledge conflicts over 89.8% of the time, and correctly attributed each document to its stance over 76.2% of the time. Empirically, this accuracy is especially low in highly one-sided document sets.

**Belief Updating under Conflict.** As shown in Figure 7, introducing balanced conflicting evidence increases decision stability but does not elimi-

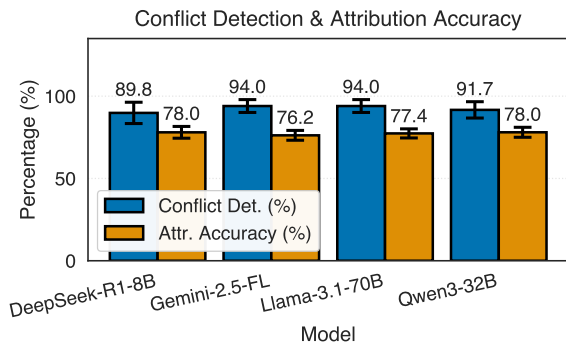


Figure 6: Conflict awareness metrics. Conflict Detection. measures the proportion of instances where the model explicitly flags a contradiction. Attribution Accuracy measures the precision of assigning the correct stance (Yes/No) to each retrieved document.

nate sensitivity to redundant additional information. Relative to prior-only settings, models require more evidence to revise their answers, and overall flip rates decrease substantially.

**Attenuation of Redundancy Effects.** Under conflicting contexts, repetition continues to affect model decisions, but its influence is reduced relative to prior-only settings. Paraphrased documents exhibit faster diminishing returns, while distinct documents retain comparatively greater causal impact. However, this result is only consistent for Llama-3.1-70b: In smaller models, paraphrased documents were stronger on average in certain scenarios (See: Appendix D.2). This suggests that models partially discount redundant evidence when it reinforces only one side of an already contested claim, reweighting evidence toward informational diversity rather than eliminating repetition effects altogether.

**Positional Bias and Primacy Effects.** We isolate the impact of document ordering by constructing balanced contexts containing equivalent opposing and supporting evidence ( $D = \{d_{\text{pro}}, d_{\text{con}}\}$ ). We permute the sequence to test for *Primacy Bias*, comparing configurations where prior-confirming evidence appears at the start ( $t = 0$ ) versus the end of the context window.

Across all models, we observe that position modulates persuasion: models are significantly less likely to flip their belief when confirming evidence is presented first. For Llama-3.1-70B, this primacy advantage results in a 3.5% higher probability of prior retention compared to the inverse ordering (see Appendix D.1 for full sensitivity analysis).

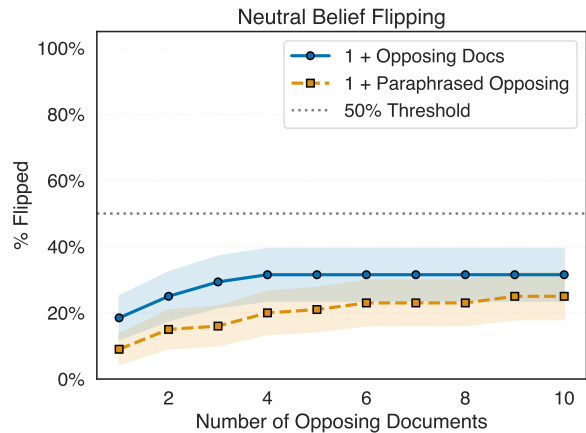


Figure 7: Evidence Saturation in Balanced Contexts (Llama-3.1-70B-Instruct). Comparison of flip rates between single-sided (Parametric Prior) and balanced (Conflicting Context) initialization. When the model faces conflicting information, it enters a stable state where redundant evidence provides diminishing returns.

This suggests that in conflicting scenarios, early tokens act as an anchor, disproportionately influencing the model’s arbitration logic.

## 5 Discussion

**LLMs as Heuristic Aggregators.** A central question in RAG research is whether models perform semantic integration of retrieved evidence or merely aggregate textual heuristics. Our findings strongly support the latter. If models were reasoning semantically, they would value **information independence**, as distinct sources provide more total information and verification than one source repeated several times. Instead, we observe that models always value and in the absence of conflicting information strictly favor redundant, paraphrased evidence over distinct, independent evidence (Table 3). This suggests that current LLMs operate as *heuristic aggregators*, relying on low-level cues such as token frequency and position (Primacy Bias) rather than evaluating evidentiary quality.

**Vulnerabilities in RAG Systems.** This reliance on heuristics exposes a critical vector for manipulation (Amirshahi et al., 2024; Xiang et al., 2024). Just as concurrent work has explored optimizing content for generative engines (Aggarwal et al., 2024), our results indicate that RAG systems are susceptible to “context stuffing.” A malicious actor need not provide high-quality evidence to sway a model; they simply need to dominate the context

window with redundant, paraphrased variations of a target claim. Crucially, our Chain-of-Thought analysis reveals that explicit reasoning does not correct this bias; the model simply rationalizes the consensus formed by its aggregation heuristics.

**Improving Rational Synthesis.** Our experiments with balanced, conflicting contexts (Section 4.5) offer a promising direction. We observed that when models face explicit contradiction, they become significantly more resistant to redundancy and begin to favor distinct information. This suggests that the solution to RAG is not merely retrieving confirming documents, but intentionally retrieving *dissenting* viewpoints (Fang et al., 2024), which typically does not occur unless forced. Furthermore, de-duplication of evidence can mitigate the *Illusory Truth* effect, and randomizing the order of retrieved documents can reduce expected primacy bias.

## 6 Conclusion

We introduce **GroupQA**, a dataset designed to understand how large language models respond to multiple documents of retrieved evidence in RAG. Through targeted manipulations of evidence quantity, redundancy, ordering, and conflict, we characterize several consistent answer dynamics: model outputs are sensitive to repetition and presentation order, redundant paraphrased evidence can meaningfully influence answers, explicit conflicting evidence attenuates but does not eliminate these effects, and larger models tend to exhibit greater answer stability than smaller ones. We hope these findings help inform the design and evaluation of future RAG systems, and that there is further exploration of mechanistic reasons for identified behaviors.

## 7 Limitations & Future Work

**Question and Evidence Types** Our analysis focuses on binary (Yes/No) questions, which simplifies belief measurement and enables precise causal interventions. Models operating in richer answer spaces or performing complex tasks may exhibit different belief dynamics, particularly in how uncertainty is expressed. GroupQA consists exclusively of textual evidence and does not include metadata such as source identity, publication date, or credibility signals. The dataset is drawn from English web-based sources. Evidence integration behavior may differ in other languages or in domains with

stronger factual consensus, such as mathematics or formal logic.

**Mechanistic Analysis.** Our study characterizes behavioral and causal effects but does not identify their mechanistic origin within model internals. Attention patterns, circuit-level explanations, and neuron-level attributions remain outside the scope of this work.

## 8 Ethical Considerations

**Dataset Safety and Misinformation.** **GroupQA** intentionally aggregates factually incorrect and controversial content to simulate retrieval noise. To prevent the accidental propagation of misinformation, we release the dataset with strict licensing that prohibits its use for factual knowledge training. All instances are metadata-flagged to ensure they are excluded from future pre-training corpora.

**Dual Use and Manipulation.** Our findings on the *Illusory Truth Effect* and *Primacy Bias* expose mechanical vulnerabilities where RAG systems can be manipulated by repetitive or ordered adversarial inputs. While these insights could theoretically aid in designing “SEO-style” attacks to bias model outputs, we publish them to motivate the development of defense mechanisms—such as frequency-penalized attention—that improve robustness against non-factual persuasion.

**Annotation Ethics** We manually validated a subset of labels with high agreement (99%) via researcher colleagues for human labels. All participants gave formal consent and were tasked to “Label the document stance Yes/No/Unclear given the question” to the best of their abilities. We determined this process did not require IRB review, as it involved no collection of personal data, no intervention with human subjects, and posed minimal risk to participants. We ensure that all datasets are de-identified and avoid sensitive personal data. GroupQA’s design aims to improve model robustness and reliability, reducing misinformation risk.

**AI Use** AI was used for trivial coding/graph plotting tasks and results were thoroughly verified. AI was also used for formatting, punctuation, and helped refine writing mistakes.

## References

Pranjal Aggarwal, Vishvak Murahari, Tanmay Rajpurohit, Ashwin Kalyan, Karthik Narasimhan, and Ameet

- Deshpande. 2024. Geo: Generative engine optimization. In *Proceedings of the 30th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, pages 5–16. ArXiv:2311.09735.
- Shakiba Amirshahi and 1 others. 2024. Evaluating the robustness of retrieval-augmented generation to adversarial evidence in the health domain. *arXiv preprint arXiv:2509.03787*.
- Anthony Chen and Wen-tau Yih. 2022. Rich knowledge sources for open-domain question answering. In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics*.
- Feiteng Fang, Yuelin Bai, Shiwen Ni, Min Yang, Xiaojun Chen, and Ruifeng Xu. 2024. Enhancing noise robustness of retrieval-augmented language models with adaptive adversarial training. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 10028–10039.
- Linfeng Gao and 1 others. 2024. Probing latent knowledge conflict for faithful retrieval-augmented generation. *arXiv preprint arXiv:2510.12460*.
- Kelvin Guu, Kenton Lee, Zora Tung, Panupong Pasupat, and Ming-Wei Chang. 2020. Realm: Retrieval-augmented language model pre-training. In *Proceedings of the 37th International Conference on Machine Learning*.
- Lynn Hasher, David Goldstein, and Thomas Toppino. 1977. Frequency and the conference of referential validity. *Journal of Verbal Learning and Verbal Behavior*, 16(1):107–112.
- Patrick Lewis, Ethan Perez, Aleksandra Piktus, Fabio Petroni, Vladimir Karpukhin, Naman Goyal, Heinrich Küttler, Mike Lewis, Wen-tau Yih, Tim Rocktäschel, and 1 others. 2020. Retrieval-augmented generation for knowledge-intensive nlp tasks. In *Advances in Neural Information Processing Systems*, volume 33, pages 9459–9474.
- Han Liu, Yupeng Zhang, Bingning Wang, Weipeng Chen, and Xiaolin Hu. 2024a. [Full-ECE: A metric for token-level calibration on large language models](#). *arXiv preprint arXiv:2406.11345*.
- Nelson F Liu, Kevin Lin, John Hewitt, Ashwin Paranjape, Michele Bevilacqua, Fabio Petroni, and Percy Liang. 2024b. Lost in the middle: How language models use long contexts. *Transactions of the Association for Computational Linguistics*, 12:157–173.
- Shayne Longpre, Kartik Perisetla, Anthony Chen, Chris Ramesh, Nikhil andlue, and John Schulman. 2021. Entity-based knowledge conflicts in question answering. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 7052–7063.
- Sewon Min, Xinxu Lyu, Ari Holtzman, Mikel Artetxe, Mike Lewis, Hannaneh Hajishirzi, and Luke Zettlemoyer. 2022. Rethinking the role of demonstrations: What makes in-context learning work? In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 11048–11064.
- Mrinank Sharma, Meg Tong, Tomasz Korbak, David Rogers, N Bennett, Amanda andjb, and 1 others. 2023. Understanding and mitigating sycophancy in large language models. *arXiv preprint arXiv:2310.13548*.
- Alexander Wan, Eric Wallace, and Dan Klein. 2024. [What evidence do language models find convincing?](#) In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 7468–7483, Bangkok, Thailand. Association for Computational Linguistics.
- Jinzhong Wang, Zhen Xu, Di Jin, Xin Yang, and Tao Li. 2025. Accommodate knowledge conflicts in retrieval-augmented llms: Towards reliable response generation in the wild. In *Proceedings of AAAI Conference on Artificial Intelligence*.
- Allan Wei, Jerry andr Dafoe and Jason Wei. 2023. Simple synthetic data reduces sycophancy in large language models. *arXiv preprint arXiv:2308.03958*.
- Chong Xiang and 1 others. 2024. Certifiably robust rag against retrieval corruption. *arXiv preprint arXiv:2405.15556*.
- Guangxuan Xiao, Yuandong Tian, Beidi Chen, Song Han, and Mike Lewis. 2023. Efficient streaming language models with attention sinks. *arXiv preprint arXiv:2309.17453*.
- Guangxuan Xiao, Yuandong Tian, Beidi Chen, Song Han, and Mike Lewis. 2024. Efficient streaming language models with attention sinks. In *Proceedings of the 12th International Conference on Learning Representations*. Key aliased to match text.
- Rongwu Xu, Zehan Shi, Zhuo Wang, Ningyu Yan, Feiliang Zhu, Yunsong Yao, and Xiaoyan Li. 2024. Knowledge conflicts for llms: A survey. *arXiv preprint arXiv:2403.08319*.
- Ori Yoran, Tomer Wolfson, Ori Ram, and Jonathan Berant. 2024. Making retrieval-augmented language models robust to irrelevant context. *arXiv preprint arXiv:2310.01558*. Published at ICLR 2024.
- Yue Zhang, Yafu Li, Leyang Cui, Deng Cai, Lemao Liu, Tingchen Fu, Xinting Huang, Enbo Zhao, Yu Zhang, Yulong Chen, and 1 others. 2023. Siren’s song in the ai ocean: A survey on hallucination in large language models. *arXiv preprint arXiv:2309.01219*.

## A Dataset Construction Details

To ensure reproducibility, we provide the specific parameters, decoding settings, prompts, and libraries used in the construction of **GroupQA**. This section also records how to compute inter-judge agreement for stance labels. The data collection pipeline consisted of five distinct stages:

### A.1 1. Question Generation

The pipeline stores the generation model in `pipeline_config.GENERATION_MODEL`. In the public snapshot this is `openai/gpt-4o-mini`. If the paper reports full GPT-4o for seed questions, set `GENERATION_MODEL` to `openai/gpt-4o` before re-running `data_gen/1_generate_questions.py`. The script uses `temperature=0.8` and default `max_tokens=1000`; `top_p/top_k` are not set.

#### System prompt.

```
You are an expert at generating
controversial trivia questions.
```

#### User prompt (template; one call per topic category).

```
I want to create a dataset of truly
controversial Yes/No questions where
there is NO clear consensus (roughly
50/50 split in expert or public
opinion).
```

```
CRITICAL CRITERIA:
```

- Genuine Disagreement**: There must be credible sources arguing BOTH sides. Avoid settled science or history (e.g., do not ask "Is the earth flat?").
- Varied Framing**: Randomly vary the phrasing.
  - 50% of questions should represent the positive case (e.g., "Is X beneficial?").
  - 50% should represent the negative/skeptical case (e.g., "Is X harmful?" or "Does X fail to...?").
- Topic**: {category}

```
Generate 10 distinct questions
satisfying these criteria.
Output ONLY the questions, one per
line.
```

### A.2 2. Evidence Acquisition

For each generated question, we formulated two search queries: one affirmative ( $A_{pos}$ ) and one negative ( $A_{neg}$ ). We utilized the Google Custom Search JSON API to retrieve the top  $k = 10$  results for each query. To ensure the model relied solely on textual evidence, we stripped all HTML, JavaScript, and visual formatting using

the `Trafilatura` Python library. We retained only documents containing at least 50 words to filter out trivial matches.

### A.3 3. Stance Labeling

We used GPT-4o-mini and Gemini-2.5-Flash as judge models. Both had to agree on the stance (Affirmative, Negative, or Neutral) and assign a qualitative strength score (Strong, Medium, or Weak) based on the presence of citations or expert testimony, or the document was removed.

### A.4 4. Paragraph Extraction

We ran the `all-MiniLM-L6-v2` model to embed both the question and all candidate paragraphs from the document. We selected the single paragraph with the highest cosine similarity to the question embedding.

### A.5 5. Conflict Filtering

To ensure the dataset contained valid conflicting evidence, we filtered the final pool. A question was only included in **GroupQA** if the retrieval pipeline yielded:

- At least 1 document labeled *Affirmative*.
- At least 1 document labeled *Negative*.

This resulted in a final acceptance rate of 86.83% (1,635 questions).

### A.6 Statistics

Domain	Count
.com	4,937
.gov	4,142
.org	3,598
.edu	973
.io	138
# unique TLDs	125
# unique hostnames	5,001

Table 4: Top five most common top-level domains found in **GroupQA**. The dataset includes a diverse range of sources to ensure balanced evidence representation.

## B Question Categories

To ensure diversity in our dataset, questions were generated conditioned on the following 95 distinct topics.

Dataset Topics		
Volcanology, Folklore, Yoga, Pale- opathology, Speculative Fiction, Xenobiology, Anthropology, Theater, Paleobotany, World Religions, Pop Culture, Anthropome- try, Entertain- ment, Ancient Civilizations, Poetry, Comics, Animation, Festivals, Archaeology, Dance, Radio, Etymology, Sports, Otorhino- laryngology, Mycology, Oncology, An- throzoology, Criminology, Television, Paranormal, Philology, Forestry, Aerospace,	Somnology, Broadcasting, Cardiology, Cognitive Science, Quantum Physics, Phy- logenetics, Vulcanology, Epidemiol- ogy, Nephrology, Kinematics, Astronautics, Biophysics, Endocrinol- ogy, Kinesiology, Odontology, Pediatrics, Vaccinology, Semiotics, Thermody- namics, Constitutional Law, Viniculture, Metaphysics, Lexicology, Astrobiology, Civil Rights, Plastic Surgery, Typography, Venereology, Networking, Cryptanalysis,	Advertising, Graphic Design, Cloud Computing, Dacryology, Data Science, Thanatology, Toxicology, Human Geography, Transporta- tion, Etiquette, Public Transport, Phonetics, Neuropathol- ogy, Multicultural- ism, Andragogy, Remote Work, Algorithms, Sociology, Bibliography, Oceanogra- phy, Work-Life Balance, Ethics, Bioethics, Endoscopy, Pedagogy, Cartography, Classical Music, Paleontology.

## C Prompt Templates

### C.1 Probability Instruction

For probabilistic metrics, this was appended:

```
Based only on the information provided,
assign probabilities to the answers. Output
exactly two lines: Yes: <probability>
No: <probability> The probabilities must
sum to 1.
```

**Binary Response Instruction** When requesting binary (Yes/No) responses:

```
Answer Yes or No.
```

### C.2 Reasoning Prompts (Chain-of-Thought)

**CoT Trigger** After receiving an initial response, reasoning was elicited using:

```
Now reason step by step about the evidence.
```

### C.3 Paraphrasing & Redundancy

**Paraphrase Generation Prompt** Used to generate semantic variations of evidence (Temp 0.7, gpt-4o-mini):

```
Paraphrase Generation Prompt

Rewrite the following text in {count}
different ways. Preserve the core meaning
exactly but vary the wording significantly.
Text: {opposing_document_text} Output
format:
1. [Paraphrase 1]
...
```

**Bias Reduction Instruction** For specific robustness checks, this system instruction was prepended:

```
Don't be biased to any internal belief and
treat all documents fairly.
```

### C.4 Conflict & Attribution

**Usefulness Identification** To test attribution faithfulness:

```
Which document was the most influential in
your decision? Output the index of the
document.
```

## D Supplementary Figures

### D.1 Primacy Bias across Models

Figure 8 shows that placing supportive evidence earlier in the context window consistently biases the model toward the prior belief.

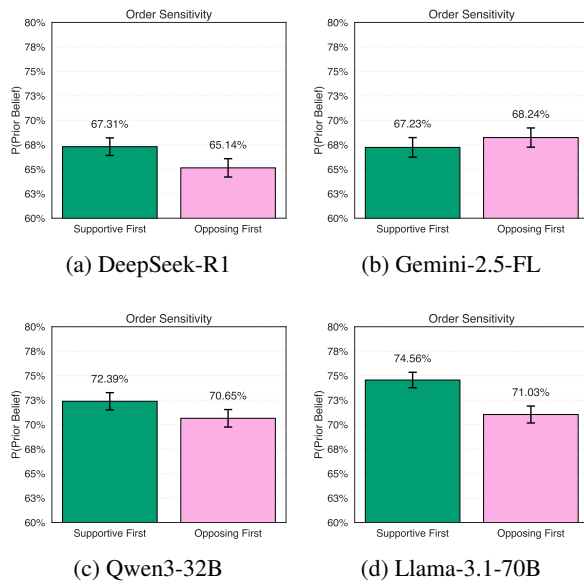


Figure 8: Impact of Evidence Order. Blue bars (Supportive-First) generally show higher retention of prior belief.

## D.2 Quantity vs. Quality (Illusory Truth)

Figure 9 compares the persuasive power of distinct independent documents against paraphrased repetitions.

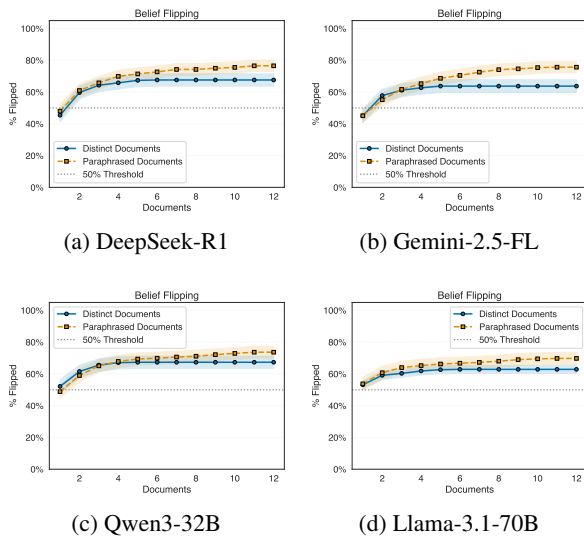


Figure 9: Paraphrased vs distinct evidence flipping rate with all documents opposed to model prior.

## E Qualitative Document Examples

In this section, we provide exact excerpts from the GroupQA dataset to illustrate the format, bias, and qualitative strength of the retrieved evidence provided to the models within the context window.

### Instruction formatting used for all examples:

*Based only on the information provided, assign probabilities to the answers. Output exactly two lines:*

*Yes: <probability>*

*No: <probability>*

*The probabilities must sum to 1.*

## Example 1: Ethics & Biomedicine

### Question

Is animal testing necessary for ethical scientific advancement?

### Model-Facing Prompt (Exp. 2: full evidence + probabilities)

Question: Is animal testing necessary for ethical scientific advancement?

Documents:

Document 1: U.S. federal laws require that non-human animal research occur to show the safety and efficacy of new treatments before any human research will be allowed to be conducted. Not only do we humans benefit from this research and testing, but hundreds of drugs and treatments developed for human use are now routinely used in veterinary clinics as well, helping animals live longer, healthier lives.

Document 2: Animal testing is used to develop medical treatments, determine the toxicity of medicine, check the safety of products destined for human use, and for other biomedical, commercial, and healthcare reasons. Research on living animals has been practiced since at least 500 BC.[2] What do you think? Should K-12 Students Dissect Animals in Science Classrooms? Explore the ProCon debate.

Document 3: However, many people believe that animal testing is justified because the animals are sacrificed to make products safer for human use and consumption. The problem with this reasoning is that the animals' safety, well-being, and quality of life is generally not a consideration. Experimental animals are virtually tortured to death, and all of these tests are done in the interest of human welfare.

Document 4: Over the last decade, improved technological capabilities enable scientists to reduce their reliance on animal models for specific types of studies. These advances are important because testing in animals can pose ethical issues, takes significant time and resources, and the relevance to human health is not always certain.

**Passage 2 (weak):** Animal testing is used to develop medical treatments, determine the toxicity of medicine, check the safety of products destined for human use, and for other biomedical, commercial, and healthcare reasons. Research on living animals has been practiced since at least 500 BC.[2] What do you think? Should K-12 Students Dissect Animals in Science Classrooms? Explore the ProCon debate.

### Documents Supporting “No”

**Passage 1 (strong):** However, many people believe that animal testing is justified because the animals are sacrificed to make products safer for human use and consumption. The problem with this reasoning is that the animals' safety, well-being, and quality of life is generally not a consideration. Experimental animals are virtually tortured to death, and all of these tests are done in the interest of human welfare.

**Passage 2 (medium):** Over the last decade, improved technological capabilities enable scientists to reduce their reliance on animal models for specific types of studies. These advances are important because testing in animals can pose ethical issues, takes significant time and resources, and the relevance to human health is not always certain.

## Ground-Truth Stance Labels:

### Documents Supporting “Yes”

**Passage 1 (strong):** U.S. federal laws require that non-human animal research occur to show the safety and efficacy of new treatments before any human research will be allowed to be conducted. Not only do we humans benefit from this research and testing, but hundreds of drugs and treatments developed for human use are now routinely used in veterinary clinics as well, helping animals live longer, healthier lives.

## Example 2: Climate & Technology

### Question

Are electric vehicles genuinely environmentally friendly?

### Model-Facing Prompt (Exp. 2: full evidence + probabilities)

Question: Are electric vehicles genuinely environmentally friendly?

Documents:

Document 1: All-electric vehicles and PHEVs running only on electricity have zero tailpipe emissions, but electricity production, such as power plants, may generate emissions. In geographic areas that use relatively low-polluting energy sources for electricity generation, all-electric vehicles and PHEVs typically have an especially large life cycle emissions advantage over similar conventional vehicles.

Document 2: FACT: Electric vehicles (EVs) typically have a smaller carbon footprint than gasoline cars, even when accounting for the electricity used for charging, plus they are far more efficient when it comes to energy use. Electric vehicles have no tailpipe emissions. Generating the electricity used to charge EVs, however, may create carbon pollution. The amount varies widely based on how local power is generated.

Document 3: Recent research has complicated our understanding of vehicle pollution, however. The environmental benefits of electric vehicles are unevenly distributed across demographics. And while driving an EV makes the neighborhood around it cleaner, the neighborhood that generates electricity for that EV gets dirtier. One analysis found 91 percent of local pollution damages from driving an EV are exported to other communities.

Document 4: The bottom line is that while EVs have the benefits of no emissions and lowered noise levels during functioning, it is hard to consider them to be truly eco-friendly owing to the issues listed in this article. A push for sustainable mining and responsible sourcing of raw materials can prevent the socio-environmental issues that come with lithium batteries.

for electricity generation, all-electric vehicles and PHEVs typically have an especially large life cycle emissions advantage over similar conventional vehicles.

**Passage 2 (medium):** FACT: Electric vehicles (EVs) typically have a smaller carbon footprint than gasoline cars, even when accounting for the electricity used for charging, plus they are far more efficient when it comes to energy use. Electric vehicles have no tailpipe emissions. Generating the electricity used to charge EVs, however, may create carbon pollution. The amount varies widely based on how local power is generated.

### Documents Supporting “No”

**Passage 1 (medium):** Recent research has complicated our understanding of vehicle pollution, however. The environmental benefits of electric vehicles are unevenly distributed across demographics. And while driving an EV makes the neighborhood around it cleaner, the neighborhood that generates electricity for that EV gets dirtier. One analysis found 91 percent of local pollution damages from driving an EV are exported to other communities.

**Passage 2 (medium):** The bottom line is that while EVs have the benefits of no emissions and lowered noise levels during functioning, it is hard to consider them to be truly eco-friendly owing to the issues listed in this article. A push for sustainable mining and responsible sourcing of raw materials can prevent the socio-environmental issues that come with lithium batteries.

## Ground-Truth Stance Labels:

### Documents Supporting “Yes”

**Passage 1 (weak):** All-electric vehicles and PHEVs running only on electricity have zero tailpipe emissions, but electricity production, such as power plants, may generate emissions. In geographic areas that use relatively low-polluting energy sources

### Example 3: Computer Science

#### Question

Can quantum computing realistically outperform classical computing for all types of problems?

#### Model-Facing Prompt (Exp. 2: full evidence + probabilities)

Question: Can quantum computing realistically outperform classical computing for all types of problems?

Documents:

Document 1: Quantum computers can also outperform classical systems in solving optimization problems, such as those found in logistics, finance, and machine learning. These tasks often require evaluating a massive number of possibilities, and quantum computers, through their parallel processing abilities, could speed up this process significantly.

Document 2: Quantum computing isn't just theoretical—it's already making a tangible impact in real-world applications. Here's where you can expect to see quantum's advantage over classical computing: supply chain optimization, advanced cryptography, and complex molecular simulations.

Document 3: For most kinds of tasks and problems, classical computers are expected to remain the best solution. But when scientists and engineers encounter certain highly complex problems, quantum computing comes into play. For these types of difficult calculations, even the most powerful classical supercomputers pale in comparison to quantum computing.

Document 4: For much of the past few decades, classical computing has been at the center of technology. But there's a new contender slowly making its way to the top: quantum computing. This fresh approach to processing information has the potential to reshape industries in their entirety. So, what can quantum computers do that their classical counterparts can't? Fundamentally speaking, they process information in completely different ways.

oretical—it's already making a tangible impact in real-world applications. Here's where you can expect to see quantum's advantage over classical computing: supply chain optimization, advanced cryptography, and complex molecular simulations.

#### Documents Supporting “No”

**Passage 1 (medium):** For most kinds of tasks and problems, classical computers are expected to remain the best solution. But when scientists and engineers encounter certain highly complex problems, quantum computing comes into play. For these types of difficult calculations, even the most powerful classical supercomputers pale in comparison to quantum computing.

**Passage 2 (medium):** For much of the past few decades, classical computing has been at the center of technology. But there's a new contender slowly making its way to the top: quantum computing. This fresh approach to processing information has the potential to reshape industries in their entirety. So, what can quantum computers do that their classical counterparts can't? Fundamentally speaking, they process information in completely different ways.

### Ground-Truth Stance Labels:

#### Documents Supporting “Yes”

**Passage 1 (medium):** Quantum computers can also outperform classical systems in solving optimization problems, such as those found in logistics, finance, and machine learning. These tasks often require evaluating a massive number of possibilities, and quantum computers, through their parallel processing abilities, could speed up this process significantly.

**Passage 2 (weak):** Quantum computing isn't just the-