

Evaluation of Deontic Conditional Reasoning in Large Language Models: The Case of Wason’s Selection Task

Hirohiko Abe¹, Kentaro Ozeki^{1,2}, Risako Ando¹, Takanobu Morishita¹,
Koji Mineshima¹, Mitsuhiro Okada¹

¹Keio University, Tokyo, Japan

²University of Tokyo, Tokyo, Japan

{hirohiko-abe, risakochaan, morishita}@keio.jp kentaro.ozeki@gmail.com

{minesima, okada}@abelard.flet.keio.ac.jp

Abstract

As large language models (LLMs) advance in linguistic competence, their reasoning abilities are gaining increasing attention. In humans, reasoning often performs well in domain specific settings, particularly in normative rather than purely formal contexts. Although prior studies have compared LLM and human reasoning, the domain specificity of LLM reasoning remains underexplored. In this study, we introduce a new Wason Selection Task dataset that explicitly encodes deontic modality to systematically distinguish deontic from descriptive conditionals, and use it to examine LLMs’ conditional reasoning under deontic rules. We further analyze whether observed error patterns are better explained by confirmation bias (a tendency to seek rule-supporting evidence) or by matching bias (a tendency to ignore negation and select items that lexically match elements of the rule). Results show that, like humans, LLMs reason better with deontic rules and display matching-bias-like errors. Together, these findings suggest that the performance of LLMs varies systematically across rule types and that their error patterns can parallel well-known human biases in this paradigm.

1 Introduction

As large language models (LLMs) advance in linguistic capabilities, their reasoning abilities are gaining increasing attention. Among various forms of reasoning, inference based on conditional rules of the form “if p then q ,” known as **conditional reasoning**, is especially important. Such reasoning underlies domains such as moral and legal judgment and has significant real-world implications.

In cognitive science, empirical studies have shown that humans tend to be more proficient in conditional reasoning within specific domains, especially those involving norms, compared to purely formal reasoning (Griggs and Cox, 1982; Cheng and Holyoak, 1985; Manktelow and Over, 1991).

This phenomenon, known as **domain specificity**, suggests that human reasoning relies on different cognitive mechanisms across domains (Davies et al., 1995; Fiddick, 2004). Such characteristics have been explained as outcomes of evolutionary adaptation (Cosmides, 1989; Cosmides and Tooby, 1992). Whether LLMs trained with large-scale neural networks exhibit domain specificity in their reasoning abilities remains an open question.

Research in cognitive science has also pointed out that human errors in conditional reasoning stem from *matching bias* (Evans and Lynch, 1973; Evans, 1998), a tendency to focus on the words appearing in the rule itself, rather than from confirmation bias (Wason, 1966), a tendency to *confirm* a given rule. Although prior research compared human and LLM reasoning based on insights from cognitive science (Seals and Shalin, 2024; Lampinen et al., 2024; Ozeki et al., 2024), it remains unclear which bias best accounts for LLM errors. To our knowledge, confirmation bias and matching bias have not been systematically compared within a unified experimental design.

In this paper, we evaluate LLMs’ conditional reasoning using the Wason Selection Task (Wason, 1966), a standard paradigm in cognitive psychology. We use a new dataset that explicitly encodes deontic modality, enabling systematic comparisons between deontic and descriptive rules.¹ LLMs achieve higher accuracy on deontic rules, paralleling patterns of human performance. Moreover, on deontic rules, their error patterns are more consistent with matching bias than with confirmation bias. These results suggest domain-specificity effects in LLM conditional reasoning and error patterns that align with matching bias.

Our contributions are summarized as follows:

1. We introduce a new Wason Selection Task

¹The dataset is publicly available at <https://github.com/kmineshima/NeuBAROCO>.

dataset with explicit deontic encoding.

2. We provide a comprehensive and up-to-date evaluation of LLMs using this dataset.
3. We clarify parallels in human and LLM performance across deontic versus descriptive rules.
4. We compare confirmation and matching bias as explanations of LLM errors, finding stronger support for matching bias.

2 Selection Task and Domain Specificity

The Wason Selection Task is a widely studied paradigm in cognitive science for probing the gap between formal logic and human intuitive reasoning about conditionals. It assesses whether a rule expressed in the form “if p then q ” holds true. Participants are presented with four cards that represent the affirmation or negation of the rule’s antecedent and consequent, and asked to select which cards must be turned over to verify the rule.

In the example below (Rule 1), the logically correct answer is to select the cards “7” and “d,” which have the potential to falsify the rule. However, experimental studies in cognitive psychology have widely observed that fewer than 10% of participants choose this correct combination (Evans et al., 1993, Chap. 4).

Rule 1 (Descriptive): If a card has an odd number on one side, then the opposite side has an uppercase letter.

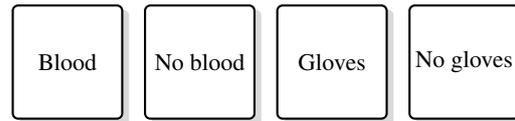


In our study, we refer to rules whose antecedent and consequent are arbitrarily paired and lack a meaningful commonsense connection (as in the original Wason task) as *descriptive*, in contrast to the *deontic* rules discussed below.

It is well established that participants perform better on tasks with concrete content, especially when the content involves normative concepts such as obligations or permissions, rather than abstract or general materials (Cheng and Holyoak, 1985; Manktelow and Over, 1991). For instance, in Rule 2, participants are more likely to select the correct cards (“blood” and “no gloves”) than in Rule 1.

Although the precise mechanisms behind this phenomenon remain debated, a leading hypothesis suggests that human reasoning operates through domain-specificity mechanisms that are specialized

Rule 2 (Deontic): If blood is spilled, the nurse must wear gloves.



according to the type of content involved (Cosmides, 1989; Cosmides and Tooby, 1992).

Recent NLP research has also focused on conditional reasoning using the Wason Selection Task (Lampinen et al., 2024; Seals and Shalin, 2024). While these studies investigate LLM performance on the Wason Selection Task, they do not explicitly frame their analyses in terms of domain specificity in the cognitive-science sense we adopt.

In Lampinen et al. (2024), the authors primarily investigate content effects (belief biases) in LLMs. They do not address the structural distinction between reasoning about deontic versus descriptive rules, which is central to our definition of domain specificity.

Seals and Shalin (2024) compare LLM performance on the Wason Selection Task with *social* and *non-social* rules. They report that the evaluated models tend to perform better when reasoning about social rules, though this advantage is smaller than would be expected in humans. However, they do not provide an explicit criterion for systematically distinguishing social from non-social categories.

We extend this line of research by focusing on the **modality** expressed in the rules. Specifically, to control the type of rule and enable a more systematic evaluation, we classify rules into two categories: *descriptive* and *deontic*.

Deontic rules express norms that regulate actions, such as obligations and prohibitions. For example, the rules containing obligation markers (e.g., “must” and “need to”) or prohibition markers (e.g., “must not” and “not allowed”) are classified as deontic. In contrast, *descriptive* rules (e.g., Rule 1) state generalizations without reference to actions or deontic modal expressions. This modality-based distinction helps define a clearer boundary between deontic and descriptive rules and supports more consistent classification.

3 Confirmation Bias and Matching Bias

A typical response to the task based on Rule 1 is to select the cards “7” and “D.” Schematically, for a

Rule Polarity	Pattern	Example Rule	TA	FA	TC	FC
Pos-Pos	$p \rightarrow q$	If the number is prime, then the letter is lowercase.	7	8	b	B
Pos-Neg	$p \rightarrow \neg q$	If the number is prime, then the letter is not lowercase.	7	8	B	b
Neg-Pos	$\neg p \rightarrow q$	If the number is not prime, then the letter is lowercase.	8	7	b	B
Neg-Neg	$\neg p \rightarrow \neg q$	If the number is not prime, then the letter is not lowercase.	8	7	B	b

Table 1: List of patterns involving negation. Pos: affirmative, Neg: negated. **TA**: case where the antecedent is true (True Antecedent), **FA**: case where the antecedent is false (False Antecedent), **TC**: case where the consequent is true (True Consequent), **FC**: case where the consequent is false (False Consequent). Cells highlighted in blue indicate logically correct selections. \square indicates options predicted by matching bias. Confirmation bias predicts selections of both **TA** and **TC**.

rule in the form “if p , then q ,” the logically correct answer is to choose cases where p is true and q is false; that is, cases that would *falsify* the rule. However, people commonly select cases where both p and q are true, i.e., cases that *confirm* the rule. Wason’s original hypothesis was that this tendency stems from a confirmation bias (Wason, 1966). On the other hand, later studies have suggested that the error pattern observed in human responses to the Wason Selection Task may be better explained by *matching bias*, rather than *confirmation bias* (Evans and Lynch, 1973; Evans, 1998). This can be tested by using the same card arrangement as in Rule 1, with the rule replaced by its *negated* form.

Rule 3 (Descriptive-negation): If a card has an odd number on one side, then the opposite side is **not** an uppercase letter.



In this case, the “D” card represents not a confirming instance, but a falsifying one. Nevertheless, participants have still been observed to select the “7” and “D” cards in such tasks (Evans, 1998). This suggests that people tend to select cards that match the elements appearing in the rule regardless of whether those elements are negated, rather than choosing based on the truth of the antecedent and consequent (as in confirmation bias.) This tendency indicates *matching bias*.

To investigate whether LLMs replicate confirmation bias or matching bias, we designed four types of *rule polarity*, based on whether the antecedent and/or consequent includes negation (see Table 1). If a model exhibits confirmation bias, it should consistently select the cards where the antecedent is true (**TA**) and the consequent is true (**TC**), regard-

less of whether negation appears in the rule. If a model exhibits matching bias, it should prefer options that directly match the symbols p and q in the rule, rather than those corresponding to $\neg p$ or $\neg q$.

Matching bias can be characterized as a tendency to overlook negation and select cards that lexically match the terms mentioned in the rule. Accordingly, testing for matching bias amounts to testing sensitivity to negation in conditional reasoning. Such analyses therefore probe whether LLMs can correctly handle negation, a long-standing challenge for Transformer-based models (Ettinger, 2020). Incorporating negation into the Wason task allows for a more precise identification of the sources of reasoning biases.

4 Experiments

4.1 Dataset

To evaluate LLM performance on the Wason Selection Task across different modalities and polarities, we constructed an evaluation dataset in which each problem consists of a rule, a set of four numbered cards corresponding to the rule, the numbers of the two correct choices, and a gold label. The dataset contains a total of 160 problems, including 80 problems with *deontic* rules and 80 problems with *descriptive* rules. Each of the four rule polarity patterns (Pos-Pos, Pos-Neg, Neg-Pos, Neg-Neg) consists of 20 problems. See Appendix A for details of the dataset construction.

4.2 Models and Settings

In the experiment, we evaluate five families of open-weight models, including both reasoning and non-reasoning models. For reasoning models, we evaluate gpt-oss 20B/120B (OpenAI, 2025) and Qwen 3 14B/32B (Yang et al., 2025). For non-reasoning models, we evaluate Gemma 3 4B/12B/27B (Gemma Team, 2025), Llama 3.3

70B (Grattafiori et al., 2024), and OLMo 2 32B (Team OLMo, 2024). For all non-reasoning models, instruction-tuned variants are used. Compared with Seals and Shalin (2024), which evaluated earlier, relatively small non-reasoning models and reported low performance compared to humans, we evaluate more recent models, including models explicitly trained or tuned for reasoning.

The basic (Zero-Shot) prompt used in the experiment is provided in Appendix B. For the Few-Shot prompt, we include one exemplar question and answer for a deontic rule and one for a descriptive rule. The answers in these exemplars are intentionally different from the correct answers to the corresponding experimental questions. The CoT prompt adopts the Zero-Shot Chain-of-Thought method (Kojima et al., 2022). For non-reasoning models, the initial maximum output token limit is set to 10 for the Zero- and Few-Shot prompts, and to 6,000 for the CoT prompt. For reasoning models, the initial maximum output token limit is set to 6,000. When this limit is insufficient, it is incrementally increased by 1,000 (up to a maximum of 11,000 in total) and the request is retried. The temperature is set to 0 (more deterministic outputs) for non-reasoning models and to the recommended values for reasoning models. All other hyperparameters are kept at their default values.

We use accuracy as the evaluation metric, computed using an exact-match criterion for this multi-select task. A response is counted as correct if and only if the model selects all and only the correct options.

4.3 Results and Discussion

This section presents results on domain specificity and reasoning bias. Examples of observed reasoning errors are provided in Appendix C.

4.3.1 Domain Specificity

Table 2 shows the differences in performance between the deontic and descriptive rule problems. Across all models and prompt types, accuracy is higher on deontic rule problems than on descriptive rule problems, with improvements ranging from 5.0 to 41.2%. The degree of the difference between the two types of rule tasks varied across models. Overall, these findings suggest that LLMs exhibit domain-specificity effects in reasoning with deontic rules, though the extent of this effect appears to depend on factors including model size.

Model	Zero	Few	CoT
Reasoning Models			
gpt-oss-20b	67.5 (+27.5)	91.2 (+41.2)	n/a
gpt-oss-120b	100.0 (+30)	98.8 (+30)	n/a
qwen3-14b	97.5 (+45)	98.8 (+37.5)	n/a
qwen3-32b	98.8 (+38.8)	100.0 (+25)	n/a
Non-Reasoning Models			
gemma3-4b	17.5 (+5)	21.2 (+12.5)	38.8 (+20)
gemma3-12b	58.8 (+30.1)	78.7 (+41.2)	43.8 (+33.8)
gemma3-27b	77.5 (+38.7)	76.2 (+28.7)	62.5 (+25)
llama-3.3-70b	62.5 (+10)	78.7 (+16.2)	83.8 (+38.8)
olmo-2-0325-32b	36.2 (+15)	38.8 (+16.3)	57.5 (+16.3)

Table 2: Accuracy (%) by prompt type for deontic rule problems. Values in parentheses show the difference in accuracy from the descriptive rule condition (positive if deontic > descriptive). **Zero** = Zero-Shot, **Few** = Few-Shot, **CoT** = Chain-of-Thought.

Polarity	Pattern	TA	FA	TC	FC
pos-pos	$p \rightarrow q$	100.0	0.0	45.0	55.0
pos-neg	$p \rightarrow \neg q$	100.0	0.0	5.0	95.0
neg-pos	$\neg p \rightarrow q$	80.0	20.0	60.0	40.0
neg-neg	$\neg p \rightarrow \neg q$	90.0	10.0	10.0	90.0

Table 3: Selection percentages (%) for each option in the deontic rule problems (gpt-oss-20b, Zero-Shot setting). As with Table 1, cells highlighted in blue indicate logically correct selections. \square indicates options predicted by matching bias. Confirmation bias predicts selections of both **TA** and **TC**.

4.3.2 Confirmation Bias versus Matching Bias

Table 3 shows a model’s selection percentages for each option under each rule polarity, following the format of Table 1. For example, when the rule has the polarity of Pos-Pos ($p \rightarrow q$), the model selects the **TA** (True Antecedent) option in 100.0% of cases (top-left cell of Table 3). Within the same row (i.e., under the same rule pattern), the selection percentages for **FA**, **TC**, and **FC** are 0.0%, 45.0%, and 55.0%, respectively. Note that each option is selected independently of the others. Corresponding tables for all models are provided in the Appendix D.

As stated in Section 3, if confirmation bias were to be replicated, we would expect a high selection rate for the options corresponding to **TA** and **TC**, regardless of whether they are in affirmative or negative form. However, as shown in Table 3, while **TA** exceeds **FA**, **TC** does not exceed **FC**. Therefore, these results do not provide evidence for confirmation bias.

Model	TA		FA		TC		FC	
	p	$\neg p$	p	$\neg p$	q	$\neg q$	q	$\neg q$
Reasoning Models								
gpt-oss-20b	100.0	85.0	15.0	0.0	52.5	7.5	92.5	47.5
gpt-oss-120b	100.0	100.0	0.0	0.0	0.0	0.0	100.0	100.0
qwen3-14b	100.0	95.0	5.0	0.0	0.0	0.0	100.0	100.0
qwen3-32b	100.0	97.5	0.0	0.0	0.0	0.0	100.0	100.0
Non-Reasoning Models								
gemma-3-4b	100.0	95.0	2.5	17.5	15.0	2.5	32.5	2.5
gemma-3-12b	77.5	95.0	10.0	55.0	15.0	2.5	85.0	57.5
gemma-3-27b	100.0	85.0	25.0	5.0	12.5	7.5	92.5	75.0
llama-3.3-70b	100.0	85.0	20.0	0.0	62.5	0.0	100.0	47.5
olmo-2-0325-32b	70.0	65.0	0.0	0.0	85.0	2.5	100.0	25.0

(a) Deontic Rule

Model	TA		FA		TC		FC	
	p	$\neg p$	p	$\neg p$	q	$\neg q$	q	$\neg q$
Reasoning Models								
gpt-oss-20b	100.0	92.5	5.0	15.0	25.0	30.0	70.0	75.0
gpt-oss-120b	100.0	100.0	2.5	20.0	5.0	25.0	95.0	95.0
qwen3-14b	100.0	97.5	0.0	7.5	27.5	30.0	87.5	77.5
qwen3-32b	100.0	100.0	0.0	12.5	17.5	32.5	90.0	90.0
Non-Reasoning Models								
gemma-3-4b	100.0	47.5	87.5	0.0	47.5	0.0	37.5	5.0
gemma-3-12b	97.5	82.5	37.5	32.5	30.0	20.0	55.0	20.0
gemma-3-27b	100.0	90.0	45.0	7.5	70.0	2.5	80.0	5.0
llama-3.3-70b	100.0	92.5	7.5	37.5	50.0	15.0	77.5	45.0
olmo-2-0325-32b	87.5	27.5	50.0	0.0	67.5	20.0	80.0	32.5

(b) Descriptive Rule

Table 4: Percentage (%) of types of options selected by the models, relative to rule polarity (Zero-Shot setting). For the meanings of **TA**, **FA**, **TC**, **FC**, refer to Section 3 and Table 1. **Boldface** indicates the higher value between p and $\neg p$, or q and $\neg q$, in each respective column (when applicable).

To assess whether models exhibit matching bias, Table 4 reorganizes the data following the analytic framework presented in Evans (1998) for human experiments. Here, the matching bias is evaluated by comparing **relative preferences** among p vs. $\neg p$ and q vs. $\neg q$ **within each option**. For example, in the **TC** column in Table 3, identifying whether the model prefers q (as reflected in the Pos-Pos and Neg-Pos rows) over $\neg q$ (the Pos-Neg and Neg-Neg rows) is crucial for diagnosing matching bias. Table 4 highlights these contrasts concisely; accordingly, the larger of the two percentages is boldfaced, visually marking the preferred choice within each pair. Note that the labels p , $\neg p$, q , and $\neg q$ in the column headers indicate the polarity of the **selected option type** (i.e., which card was chosen), rather than the antecedent and consequent forms of the rules (e.g., $p \rightarrow q$ and $\neg p \rightarrow \neg q$).

As discussed in Section 3, matching bias predicts a preference for p over $\neg p$ and q over $\neg q$, irrespective of whether the antecedent or consequent in the options is true or false. As shown in Table 4, our results are broadly consistent with this prediction.

One exception arises for reasoning models in the descriptive condition, where $\neg p$ is consistently preferred to p in **FA** and $\neg q$ to q in **TC**. Since only **TA** and **FC** are required to check for violations, selecting **FA** or **TC** yields an incorrect response under our exact-match scoring, and these options are chosen only infrequently compared with **TA** and **FC**.

Finally, note that matching bias and logical correctness are not mutually exclusive. In particular, selecting q in **FC**, as predicted by matching bias, is logically valid. Thus, a matching-bias-consistent preference in **FC** does not entail irrational reasoning but can even increase accuracy.

5 Conclusion

This study systematically evaluated LLMs’ conditional reasoning using the Wason Selection Task. As shown in Section 4.3, the results suggest that LLMs exhibit domain-specificity effects and that their error tendencies are more consistent with matching bias than confirmation bias.

From the perspective of explaining this behavior in LLMs, further investigation into the specific domains of specialization and the mechanisms behind them remains an open question. In addition, mechanistic analysis to determine the causal origins of those behaviors is left for future work. Although this study centered on the Wason Selection Task, future work should extend the analysis to other forms of conditional reasoning and distinct types of reasoning beyond conditional reasoning, within the modality-based framework proposed in this study. Additionally, a systematic evaluation of LLMs’ reasoning about permission and obligation constitutes an important avenue for further investigation.

Limitations

While our study aims to provide a comprehensive analysis of deontic conditional reasoning in LLMs, several limitations should be noted.

First, this study examined the Wason Selection Task as a case study of conditional reasoning. However, whether these findings extend to conditional reasoning more broadly remains an open question for future work.

In addition, while our study focused on deontic modality to examine domain specificity, further research could also consider other ways of classifying domains. Relatedly, a systematic evaluation of LLMs' reasoning about permission and obligation remains an important direction for further investigation.

More broadly, while our study focused on characterizing the observable reasoning behavior of LLMs in a controlled variant of the Wason Selection Task, mechanistic analyses that determine the causal origins of these behaviors, such as architectural inductive biases, properties of the training data, or specific exposure to Wason type problems, are valuable topics for future work.

Finally, because LLMs are continuously evolving, our findings may not generalize to future models with improved reasoning mechanisms. Model performance can change with updates to training data, architecture, and fine tuning strategies. Moreover, differences in fine tuning, prompt engineering, and instruction following abilities may contribute to performance variation, which makes it challenging to isolate the effects of model architecture alone.

Acknowledgment

We thank the anonymous reviewers for their insightful comments and suggestions, which have improved the paper. This work is partially supported by JST CREST Grant Number JPMJCR2114, Keio University Global Research Institute (KGRI) Challenge Grant, and JSPS Kakenhi Grant Numbers JP24K00004, JP21K00016, JP21H00467, JP23K20416, and JP21K18339.

References

- Patricia W Cheng and Keith J Holyoak. 1985. Pragmatic reasoning schemas. *Cognitive psychology*, 17(4):391–416.
- Leda Cosmides. 1989. The logic of social exchange: Has natural selection shaped how humans reason? studies with the wason selection task. *Cognition*, 31(3):187–276.
- Leda Cosmides and John Tooby. 1992. Cognitive adaptations for social exchange. In *The Adapted Mind: Evolutionary Psychology and the Generation of Culture*. Oxford University Press.
- Paul Sheldon Davies, James H Fetzer, and Thomas R Foster. 1995. Logical reasoning and domain specificity: A critique of the social exchange theory of reasoning. *Biology and Philosophy*, 10:1–37.
- Allyson Ettinger. 2020. [What BERT is not: Lessons from a new suite of psycholinguistic diagnostics for language models](#). *Transactions of the Association for Computational Linguistics*, 8:34–48.
- J St BT Evans and J St Lynch. 1973. Matching bias in the selection task. *British Journal of Psychology*, 64(3):391–397.
- Jonathan St BT Evans. 1998. Matching bias in conditional reasoning: Do we understand it after 25 years? *Thinking & Reasoning*, 4(1):45–110.
- Jonathan St.B. T. Evans, Stephen E. Newstead, and Ruth M. J. Byrne. 1993. *Human Reasoning: The Psychology of Deduction*. Psychology Press.
- Laurence Fiddick. 2004. Domains of deontic reasoning: Resolving the discrepancy between the cognitive and moral reasoning literatures. *The Quarterly Journal of Experimental Psychology Section A*, 57(3):447–474.
- Gemma Team. 2025. Gemma 3 technical report. *arXiv preprint arXiv:2503.19786*.
- Aaron Grattafiori, Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Alex Vaughan, et al. 2024. The llama 3 herd of models. *arXiv preprint arXiv:2407.21783*.
- Richard A Griggs and James R Cox. 1982. The elusive thematic-materials effect in Wason's selection task. *British Journal of Psychology*, 73(3):407–420.
- Takeshi Kojima, Shixiang Shane Gu, Machel Reid, Yutaka Matsuo, and Yusuke Iwasawa. 2022. Large language models are zero-shot reasoners. *Advances in neural information processing systems*, 35:22199–22213.
- Andrew K Lampinen, Ishita Dasgupta, Stephanie C Y Chan, Hannah R Sheahan, Antonia Creswell, Dharshan Kumaran, James L McClelland, and Felix Hill. 2024. [Language models, like humans, show content effects on reasoning tasks](#). *PNAS Nexus*, 3(7):page233.
- Ken I Manktelow and David E Over. 1991. Social roles and utilities in reasoning with deontic conditionals. *Cognition*, 39(2):85–105.
- OpenAI. 2025. gpt-oss-120b & gpt-oss-20b model card. *arXiv preprint arXiv:2508.10925*.
- Kentaro Ozeki, Risako Ando, Takanobu Morishita, Hirohiko Abe, Koji Mineshima, and Mitsuhiro Okada. 2024. Exploring reasoning biases in large language models through syllogism: Insights from the NeuBAROCO dataset. In *Findings of the Association for Computational Linguistics: ACL 2024*.
- Spencer M. Seals and Valerie L. Shalin. 2024. [Evaluating the deductive competence of large language models](#). 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 8614–8630, Mexico City, Mexico. Association for Computational Linguistics.

Team OLMo. 2024. 2 OLMo 2 furious. *arXiv preprint arXiv:2501.00656*.

Peter C. Wason. 1966. Reasoning. In Peter C. Wason, editor, *New Horizons in Psychology*, pages 135–151. Penguin Books.

An Yang, Anfeng Li, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chang Gao, Chengen Huang, Chenxu Lv, et al. 2025. Qwen3 technical report. *arXiv preprint arXiv:2505.09388*.

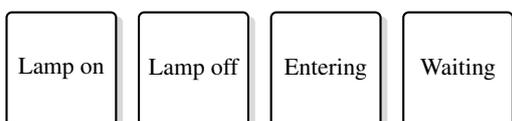
A Dataset Details

We created templates as described in Section 2. Then, we used Gemini 1.5 Pro to instantiate these templates. We manually checked and substantially adjusted the cases to ensure quality.

We created 20 *deontic* and 20 *descriptive* cases with a Pos-Pos polarity pattern. By manually adding negation to the antecedent and/or consequent, we generated the cases in other polarity patterns, namely Pos-Neg, Neg-Pos, Neg-Neg. Formally, the application of negation can be summarized as shown in Table 5.

Regarding the application of negation to the consequent in *deontic* cases, obligations were treated as affirmative whereas prohibitions (negations of permission) as negative. This is due to the fact that the rules expressing permission or unnecessary (negation of obligation) were excluded from the dataset because the cards to be selected differ in these modal types from obligation and prohibition (negation of permission). For example, when the rule is one of permission, such as Rule 4, there is no card that should be selected. This is because, even if the antecedent holds, there is no requirement to perform the action described in the consequent, and conversely, if the antecedent does not hold, it is not logically entailed that the action in the consequent must not be performed.

Rule 4 (Permission): If the lamp is on, then the visitor is allowed to enter the room.



The same applies to the unnecessary (negation of obligation), as in “If the lamp is on, then the visitor needs not enter the room.” In this case as well, there is no card that should be selected. In the cases of permission and unnecessary (negation of obligation), unlike in the cases of obligation and prohibition (negation of permission), the cards satisfying the antecedent and those negating the consequent are not the correct choices. To avoid inconsistency with respect to the correct answers, these types were excluded from the study.

B Prompts

An example of a basic (Zero-Shot) prompt is as follows:

Basic Prompt

To verify if the following rule holds true, which of the four cards below should you flip over? Please do not turn over more than required. Answer with zero or more of the numbers 1 to 4, separated by commas, and output nothing else.

Rule: If a symbol is a spiral, then the opposite side is a primary color.

Card 1: Spiral

Card 2: Triangle

Card 3: Red

Card 4: Green

Answer:1,4

C Error Examples

Examples of errors for the descriptive and deontic rule problems, including three polarity combinations (Pos-Pos, Neg-Pos, and Neg-Neg), are presented in Figures 1, 2, 3, and 4.

Figure 1 shows an example of a *deontic* rule with Pos-Pos polarity. Several models (e.g., gpt-oss-20b) incorrectly selected option 3, which corresponds to the consequent q , instead of option 4 (*Guard leaving*), which represents $\neg q$ and thus the actual violation when p holds. This error pattern is consistent with both confirmation bias and matching bias.

Figure 2 shows an example of a *deontic* rule with Neg-Pos polarity. Some models (e.g., gpt-oss-20b) again incorrectly selected option 3, which corresponds to the consequent q , instead of option 4

Rule Polarity	Pattern	Descriptive	Deontic
Pos-Pos	$p \rightarrow q$	If a number is prime, then the opposite side has a lowercase letter.	If the lamp is on, then the visitor must enter the room.
Pos-Neg	$p \rightarrow \neg q$	If a number is prime, then the opposite side does not have a lowercase letter.	If the lamp is on, then the visitor must not enter the room.
Neg-Pos	$\neg p \rightarrow q$	If a number is not prime, then the opposite side has a lowercase letter.	If the lamp is not on, then the visitor must enter the room.
Neg-Neg	$\neg p \rightarrow \neg q$	If a number is not prime, then the opposite side does not have a lowercase letter.	If the lamp is not on, then the visitor must not enter the room.

Table 5: Example rules of descriptive and deontic cases.

($\neg q$, *Document not filed*), which represents the actual violation when p holds. This error pattern is compatible with both confirmation bias and matching bias. On the other hand, gemma-3-27b and llama-3.3-70b incorrectly selected option 1 (*Car in front of the building*) instead of option 2 (*Car in the parking lot*). This error reflects a tendency to select options that lexically match the antecedent while ignoring negation, and is therefore more consistent with matching bias than with confirmation bias.

Figure 3 presents a *descriptive* rule with Pos-Pos polarity. Some models (e.g., gemma-3-27b) incorrectly picked option 3 (*Violin*), which matches the consequent q , even though the correct choice is to check option 4 (*Bicycle*, i.e., $\neg q$) to verify whether the rule is violated when the antecedent p holds. This error pattern, which also appears in Figure 1 for the *deontic* condition, is compatible with both confirmation bias and matching bias.

Figure 4 shows a *descriptive* rule with Neg-Neg polarity. The logically correct response is to select option 2 (%), which satisfies the antecedent, and option 3 (*Sad*), which negates the consequent. Some models, including gemma-3-4b and gemma-3-12b, selected options 1 (♣) and option 3 (*Sad*). Selecting option 1 instead of option 2 suggests that the model ignored negation in the antecedent and instead chose the lexically matching option. This behavior is therefore consistent with matching bias rather than confirmation bias. Note that selecting the negation of the consequent is required, so choosing option 3 ($\neg q$) is both logically correct and compatible with matching bias. Some models (gpt-oss-20b and gpt-oss-120b) additionally included option 4 (q) alongside option 3, yielding a mixed response pattern that is not straightforward to interpret as confirmation bias.

Type: Deontic Correct Answer: 1, 4	
Rule: If it is sunny, then the guard is obliged to unlock the door.	
<ol style="list-style-type: none"> 1. Sunny 2. Rainy 3. Guard unlocking the door 4. Guard leaving 	
Model	Prediction
gpt-oss-20b	✗ 1, 3
gpt-oss-120b	✓ 1, 4
qwen3-14b	✓ 1, 4
qwen3-32b	✓ 1, 4
gemma-3-4b	✗ 1
gemma-3-12b	✗ 1, 2
gemma-3-27b	✗ 1, 3
llama-3.3-70b	✗ 1, 3
olmo-2-0325-32b	✗ 3

Figure 1: An example of error for a deontic rule with Pos-Pos polarity.

Type: Deontic Correct Answer: 2, 4	
Rule: If the car is not in front of the building, then the lawyer needs to file the document with the court.	
<ol style="list-style-type: none"> 1. Car in front of the building 2. Car in the parking lot 3. Lawyer filing 4. Document not filed 	
Model	Prediction
gpt-oss-20b	✗ 2, 3
gpt-oss-120b	✓ 2, 4
qwen3-14b	✓ 2, 4
qwen3-32b	✓ 2, 4
gemma-3-4b	✓ 2, 4
gemma-3-12b	✓ 2, 4
gemma-3-27b	✗ 1, 4
llama-3.3-70b	✗ 1, 3, 4
olmo-2-0325-32b	✗ 3, 4

Figure 2: An example of error for a deontic rule with Neg-Pos polarity.

Type: Descriptive Correct Answer: 2, 3	
Rule: If a card does not show a musical note, then the opposite side does not contain a word describing a feeling.	
<ol style="list-style-type: none"> 1. ♪ 2. % 3. Sad 4. Fast 	
Model	Prediction
gpt-oss-20b	✗ 2, 3, 4
gpt-oss-120b	✗ 2, 3, 4
qwen3-14b	✓ 2, 3
qwen3-32b	✓ 2, 3
gemma-3-4b	✗ 1, 3
gemma-3-12b	✗ 1, 3
gemma-3-27b	✗ 1, 3
llama-3.3-70b	✓ 2, 3
olmo-2-0325-32b	✗ 3, 4

Figure 4: An example of error for a descriptive rule with Neg-Neg polarity.

Type: Descriptive Correct Answer: 1, 4	
Rule: If a number is a multiple of 3, then the opposite side shows a musical instrument.	
<ol style="list-style-type: none"> 1. 9 2. 11 3. Violin 4. Bicycle 	
Model	Prediction
gpt-oss-20b	✓ 1, 4
gpt-oss-120b	✓ 1, 4
qwen3-14b	✓ 1, 4
qwen3-32b	✓ 1, 4
gemma-3-4b	✗ 1
gemma-3-12b	✗ 2, 4
gemma-3-27b	✗ 1, 3
llama-3.3-70b	✗ 1, 2, 4
olmo-2-0325-32b	✗ 1, 3

Figure 3: An example of error for a descriptive rule with Pos-Pos polarity.

D Supplemental Results

Tables 6 to 28 present the selection percentages (%) by each model under different prompting conditions (Zero-Shot, Few-Shot, CoT), for all rule polarity patterns in both deontic and descriptive rule tasks. See Table 3 for the description of the table format.

Polarity	Pattern	TA	FA	TC	FC
pos-pos	$p \rightarrow q$	100.0	0.0	45.0	55.0
pos-neg	$p \rightarrow \neg q$	100.0	0.0	5.0	95.0
neg-pos	$\neg p \rightarrow q$	80.0	20.0	60.0	40.0
neg-neg	$\neg p \rightarrow \neg q$	90.0	10.0	10.0	90.0

(a) Deontic

Polarity	Pattern	TA	FA	TC	FC
pos-pos	$p \rightarrow q$	100.0	25.0	10.0	70.0
pos-neg	$p \rightarrow \neg q$	100.0	5.0	0.0	55.0
neg-pos	$\neg p \rightarrow q$	95.0	10.0	40.0	80.0
neg-neg	$\neg p \rightarrow \neg q$	90.0	0.0	60.0	85.0

(b) Descriptive

Table 6: gpt-oss-20b (Zero-Shot)

Polarity	Pattern	TA	FA	TC	FC
pos-pos	$p \rightarrow q$	100.0	0.0	0.0	100.0
pos-neg	$p \rightarrow \neg q$	100.0	0.0	15.0	100.0
neg-pos	$\neg p \rightarrow q$	100.0	0.0	5.0	100.0
neg-neg	$\neg p \rightarrow \neg q$	100.0	0.0	15.0	100.0

(a) Deontic

Polarity	Pattern	TA	FA	TC	FC
pos-pos	$p \rightarrow q$	100.0	35.0	5.0	85.0
pos-neg	$p \rightarrow \neg q$	100.0	5.0	10.0	85.0
neg-pos	$\neg p \rightarrow q$	90.0	15.0	35.0	95.0
neg-neg	$\neg p \rightarrow \neg q$	90.0	0.0	60.0	95.0

(b) Descriptive

Table 7: gpt-oss-20b (Few-Shot)

Polarity	Pattern	TA	FA	TC	FC
pos-pos	$p \rightarrow q$	100.0	0.0	0.0	100.0
pos-neg	$p \rightarrow \neg q$	100.0	0.0	0.0	100.0
neg-pos	$\neg p \rightarrow q$	100.0	0.0	0.0	100.0
neg-neg	$\neg p \rightarrow \neg q$	100.0	0.0	0.0	100.0

(a) Deontic

Polarity	Pattern	TA	FA	TC	FC
pos-pos	$p \rightarrow q$	100.0	40.0	0.0	95.0
pos-neg	$p \rightarrow \neg q$	100.0	0.0	5.0	100.0
neg-pos	$\neg p \rightarrow q$	100.0	5.0	10.0	95.0
neg-neg	$\neg p \rightarrow \neg q$	100.0	0.0	45.0	90.0

(b) Descriptive

Table 8: gpt-oss-120b (Zero-Shot)

Polarity	Pattern	TA	FA	TC	FC
pos-pos	$p \rightarrow q$	100.0	0.0	0.0	100.0
pos-neg	$p \rightarrow \neg q$	100.0	0.0	0.0	100.0
neg-pos	$\neg p \rightarrow q$	100.0	0.0	0.0	95.0
neg-neg	$\neg p \rightarrow \neg q$	100.0	0.0	0.0	100.0

(a) Deontic

Polarity	Pattern	TA	FA	TC	FC
pos-pos	$p \rightarrow q$	100.0	45.0	0.0	95.0
pos-neg	$p \rightarrow \neg q$	100.0	0.0	5.0	95.0
neg-pos	$\neg p \rightarrow q$	100.0	15.0	15.0	90.0
neg-neg	$\neg p \rightarrow \neg q$	100.0	0.0	35.0	100.0

(b) Descriptive

Table 9: gpt-oss-120b (Few-Shot)

Polarity	Pattern	TA	FA	TC	FC
pos-pos	$p \rightarrow q$	100.0	0.0	0.0	100.0
pos-neg	$p \rightarrow \neg q$	100.0	0.0	0.0	100.0
neg-pos	$\neg p \rightarrow q$	95.0	5.0	0.0	100.0
neg-neg	$\neg p \rightarrow \neg q$	95.0	5.0	0.0	100.0

(a) Deontic

Polarity	Pattern	TA	FA	TC	FC
pos-pos	$p \rightarrow q$	100.0	10.0	5.0	65.0
pos-neg	$p \rightarrow \neg q$	100.0	5.0	10.0	85.0
neg-pos	$\neg p \rightarrow q$	95.0	0.0	50.0	90.0
neg-neg	$\neg p \rightarrow \neg q$	100.0	0.0	50.0	90.0

(b) Descriptive

Table 10: qwen3-14b (Zero-Shot)

Polarity	Pattern	TA	FA	TC	FC
pos-pos	$p \rightarrow q$	100.0	0.0	0.0	100.0
pos-neg	$p \rightarrow \neg q$	100.0	0.0	0.0	100.0
neg-pos	$\neg p \rightarrow q$	95.0	0.0	0.0	100.0
neg-neg	$\neg p \rightarrow \neg q$	100.0	0.0	0.0	100.0

(a) Deontic

Polarity	Pattern	TA	FA	TC	FC
pos-pos	$p \rightarrow q$	100.0	10.0	0.0	90.0
pos-neg	$p \rightarrow \neg q$	100.0	5.0	10.0	85.0
neg-pos	$\neg p \rightarrow q$	90.0	0.0	55.0	100.0
neg-neg	$\neg p \rightarrow \neg q$	90.0	0.0	55.0	100.0

(b) Descriptive

Table 11: qwen3-14b (Few-Shot)

Polarity	Pattern	TA	FA	TC	FC
pos-pos	$p \rightarrow q$	100.0	0.0	0.0	100.0
pos-neg	$p \rightarrow \neg q$	100.0	0.0	0.0	100.0
neg-pos	$\neg p \rightarrow q$	100.0	0.0	0.0	100.0
neg-neg	$\neg p \rightarrow \neg q$	95.0	0.0	0.0	100.0

(a) Deontic

Polarity	Pattern	TA	FA	TC	FC
pos-pos	$p \rightarrow q$	100.0	25.0	5.0	85.0
pos-neg	$p \rightarrow \neg q$	100.0	0.0	10.0	85.0
neg-pos	$\neg p \rightarrow q$	100.0	0.0	30.0	95.0
neg-neg	$\neg p \rightarrow \neg q$	100.0	0.0	55.0	95.0

(b) Descriptive

Table 12: qwen3-32b (Zero-Shot)

Polarity	Pattern	TA	FA	TC	FC
pos-pos	$p \rightarrow q$	100.0	0.0	0.0	100.0
pos-neg	$p \rightarrow \neg q$	100.0	0.0	0.0	100.0
neg-pos	$\neg p \rightarrow q$	100.0	0.0	0.0	100.0
neg-neg	$\neg p \rightarrow \neg q$	100.0	0.0	0.0	100.0

(a) Deontic

Polarity	Pattern	TA	FA	TC	FC
pos-pos	$p \rightarrow q$	100.0	5.0	5.0	95.0
pos-neg	$p \rightarrow \neg q$	100.0	0.0	20.0	100.0
neg-pos	$\neg p \rightarrow q$	85.0	0.0	30.0	90.0
neg-neg	$\neg p \rightarrow \neg q$	100.0	0.0	50.0	100.0

(b) Descriptive

Table 13: qwen3-32b (Few-Shot)

Polarity	Pattern	TA	FA	TC	FC
pos-pos	$p \rightarrow q$	100.0	15.0	20.0	0.0
pos-neg	$p \rightarrow \neg q$	100.0	20.0	0.0	50.0
neg-pos	$\neg p \rightarrow q$	95.0	5.0	10.0	5.0
neg-neg	$\neg p \rightarrow \neg q$	95.0	0.0	5.0	15.0

(a) Deontic

Polarity	Pattern	TA	FA	TC	FC
pos-pos	$p \rightarrow q$	100.0	0.0	55.0	0.0
pos-neg	$p \rightarrow \neg q$	100.0	0.0	0.0	40.0
neg-pos	$\neg p \rightarrow q$	45.0	85.0	40.0	10.0
neg-neg	$\neg p \rightarrow \neg q$	50.0	90.0	0.0	35.0

(b) Descriptive

Table 14: gemma3-4b (Zero-Shot)

Polarity	Pattern	TA	FA	TC	FC
pos-pos	$p \rightarrow q$	100.0	40.0	0.0	5.0
pos-neg	$p \rightarrow \neg q$	100.0	55.0	5.0	5.0
neg-pos	$\neg p \rightarrow q$	90.0	35.0	5.0	65.0
neg-neg	$\neg p \rightarrow \neg q$	95.0	30.0	55.0	15.0

(a) Deontic

Polarity	Pattern	TA	FA	TC	FC
pos-pos	$p \rightarrow q$	95.0	10.0	60.0	10.0
pos-neg	$p \rightarrow \neg q$	90.0	30.0	35.0	20.0
neg-pos	$\neg p \rightarrow q$	85.0	70.0	15.0	30.0
neg-neg	$\neg p \rightarrow \neg q$	90.0	75.0	25.0	15.0

(b) Descriptive

Table 15: gemma3-4b (Few-Shot)

Polarity	Pattern	TA	FA	TC	FC
pos-pos	$p \rightarrow q$	100.0	30.0	65.0	0.0
pos-neg	$p \rightarrow \neg q$	100.0	10.0	0.0	85.0
neg-pos	$\neg p \rightarrow q$	90.0	15.0	25.0	20.0
neg-neg	$\neg p \rightarrow \neg q$	90.0	5.0	5.0	55.0

(a) Deontic

Polarity	Pattern	TA	FA	TC	FC
pos-pos	$p \rightarrow q$	95.0	15.0	80.0	0.0
pos-neg	$p \rightarrow \neg q$	90.0	30.0	0.0	70.0
neg-pos	$\neg p \rightarrow q$	75.0	65.0	35.0	10.0
neg-neg	$\neg p \rightarrow \neg q$	80.0	80.0	0.0	30.0

(b) Descriptive

Table 16: gemma3-4b (CoT)

Polarity	Pattern	TA	FA	TC	FC
pos-pos	$p \rightarrow q$	75.0	75.0	20.0	30.0
pos-neg	$p \rightarrow \neg q$	80.0	35.0	0.0	85.0
neg-pos	$\neg p \rightarrow q$	95.0	10.0	10.0	85.0
neg-neg	$\neg p \rightarrow \neg q$	95.0	10.0	5.0	85.0

(a) Deontic

Polarity	Pattern	TA	FA	TC	FC
pos-pos	$p \rightarrow q$	95.0	35.0	40.0	5.0
pos-neg	$p \rightarrow \neg q$	100.0	30.0	5.0	65.0
neg-pos	$\neg p \rightarrow q$	90.0	30.0	20.0	35.0
neg-neg	$\neg p \rightarrow \neg q$	75.0	45.0	35.0	45.0

(b) Descriptive

Table 17: gemma3-12b (Zero-Shot)

Polarity	Pattern	TA	FA	TC	FC
pos-pos	$p \rightarrow q$	95.0	20.0	0.0	75.0
pos-neg	$p \rightarrow \neg q$	75.0	0.0	20.0	95.0
neg-pos	$\neg p \rightarrow q$	90.0	15.0	5.0	90.0
neg-neg	$\neg p \rightarrow \neg q$	90.0	5.0	15.0	90.0

(a) Deontic

Polarity	Pattern	TA	FA	TC	FC
pos-pos	$p \rightarrow q$	75.0	30.0	20.0	40.0
pos-neg	$p \rightarrow \neg q$	90.0	20.0	25.0	75.0
neg-pos	$\neg p \rightarrow q$	75.0	25.0	10.0	90.0
neg-neg	$\neg p \rightarrow \neg q$	80.0	30.0	80.0	65.0

(b) Descriptive

Table 18: gemma3-12b (Few-Shot)

Polarity	Pattern	TA	FA	TC	FC
pos-pos	$p \rightarrow q$	90.0	85.0	5.0	10.0
pos-neg	$p \rightarrow \neg q$	90.0	45.0	5.0	55.0
neg-pos	$\neg p \rightarrow q$	100.0	0.0	5.0	55.0
neg-neg	$\neg p \rightarrow \neg q$	100.0	0.0	15.0	75.0

(a) Deontic

Polarity	Pattern	TA	FA	TC	FC
pos-pos	$p \rightarrow q$	100.0	40.0	0.0	5.0
pos-neg	$p \rightarrow \neg q$	100.0	55.0	10.0	15.0
neg-pos	$\neg p \rightarrow q$	95.0	30.0	10.0	30.0
neg-neg	$\neg p \rightarrow \neg q$	95.0	20.0	40.0	20.0

(b) Descriptive

Table 19: gemma3-12b (CoT)

Polarity	Pattern	TA	FA	TC	FC
pos-pos	$p \rightarrow q$	100.0	10.0	15.0	75.0
pos-neg	$p \rightarrow \neg q$	100.0	0.0	0.0	100.0
neg-pos	$\neg p \rightarrow q$	85.0	30.0	10.0	75.0
neg-neg	$\neg p \rightarrow \neg q$	85.0	20.0	15.0	85.0

(a) Deontic

Polarity	Pattern	TA	FA	TC	FC
pos-pos	$p \rightarrow q$	100.0	10.0	90.0	0.0
pos-neg	$p \rightarrow \neg q$	100.0	5.0	0.0	95.0
neg-pos	$\neg p \rightarrow q$	95.0	45.0	50.0	10.0
neg-neg	$\neg p \rightarrow \neg q$	85.0	45.0	5.0	65.0

(b) Descriptive

Table 20: gemma3-27b (Zero-Shot)

Polarity	Pattern	TA	FA	TC	FC
pos-pos	$p \rightarrow q$	100.0	10.0	0.0	90.0
pos-neg	$p \rightarrow \neg q$	100.0	0.0	0.0	100.0
neg-pos	$\neg p \rightarrow q$	70.0	35.0	30.0	65.0
neg-neg	$\neg p \rightarrow \neg q$	75.0	25.0	0.0	100.0

(a) Deontic

Polarity	Pattern	TA	FA	TC	FC
pos-pos	$p \rightarrow q$	100.0	85.0	5.0	10.0
pos-neg	$p \rightarrow \neg q$	100.0	10.0	0.0	90.0
neg-pos	$\neg p \rightarrow q$	100.0	35.0	25.0	40.0
neg-neg	$\neg p \rightarrow \neg q$	70.0	35.0	15.0	80.0

(b) Descriptive

Table 21: gemma3-27b (Few-Shot)

Polarity	Pattern	TA	FA	TC	FC
pos-pos	$p \rightarrow q$	100.0	5.0	65.0	30.0
pos-neg	$p \rightarrow \neg q$	100.0	0.0	0.0	100.0
neg-pos	$\neg p \rightarrow q$	85.0	30.0	30.0	55.0
neg-neg	$\neg p \rightarrow \neg q$	90.0	15.0	5.0	90.0

(a) Deontic

Polarity	Pattern	TA	FA	TC	FC
pos-pos	$p \rightarrow q$	100.0	10.0	90.0	0.0
pos-neg	$p \rightarrow \neg q$	100.0	5.0	0.0	95.0
neg-pos	$\neg p \rightarrow q$	90.0	55.0	55.0	0.0
neg-neg	$\neg p \rightarrow \neg q$	85.0	45.0	0.0	70.0

(b) Descriptive

Table 22: gemma3-27b (CoT)

Polarity	Pattern	TA	FA	TC	FC
pos-pos	$p \rightarrow q$	100.0	0.0	85.0	30.0
pos-neg	$p \rightarrow \neg q$	100.0	0.0	0.0	100.0
neg-pos	$\neg p \rightarrow q$	85.0	25.0	40.0	65.0
neg-neg	$\neg p \rightarrow \neg q$	85.0	15.0	0.0	100.0

(a) Deontic

Polarity	Pattern	TA	FA	TC	FC
pos-pos	$p \rightarrow q$	100.0	60.0	70.0	15.0
pos-neg	$p \rightarrow \neg q$	100.0	15.0	5.0	85.0
neg-pos	$\neg p \rightarrow q$	90.0	15.0	30.0	75.0
neg-neg	$\neg p \rightarrow \neg q$	95.0	0.0	25.0	70.0

(b) Descriptive

Table 23: llama-3.3-70b (Zero-Shot)

Polarity	Pattern	TA	FA	TC	FC
pos-pos	$p \rightarrow q$	100.0	0.0	25.0	80.0
pos-neg	$p \rightarrow \neg q$	100.0	0.0	0.0	100.0
neg-pos	$\neg p \rightarrow q$	100.0	5.0	55.0	40.0
neg-neg	$\neg p \rightarrow \neg q$	100.0	0.0	0.0	100.0

(a) Deontic

Polarity	Pattern	TA	FA	TC	FC
pos-pos	$p \rightarrow q$	100.0	40.0	40.0	25.0
pos-neg	$p \rightarrow \neg q$	95.0	10.0	0.0	95.0
neg-pos	$\neg p \rightarrow q$	100.0	5.0	5.0	90.0
neg-neg	$\neg p \rightarrow \neg q$	95.0	10.0	45.0	55.0

(b) Descriptive

Table 24: llama-3.3-70b (Few-Shot)

Polarity	Pattern	TA	FA	TC	FC
pos-pos	$p \rightarrow q$	100.0	0.0	20.0	90.0
pos-neg	$p \rightarrow \neg q$	100.0	0.0	0.0	100.0
neg-pos	$\neg p \rightarrow q$	95.0	15.0	20.0	95.0
neg-neg	$\neg p \rightarrow \neg q$	85.0	15.0	0.0	100.0

(a) Deontic

Polarity	Pattern	TA	FA	TC	FC
pos-pos	$p \rightarrow q$	100.0	10.0	45.0	15.0
pos-neg	$p \rightarrow \neg q$	100.0	5.0	10.0	75.0
neg-pos	$\neg p \rightarrow q$	80.0	15.0	0.0	80.0
neg-neg	$\neg p \rightarrow \neg q$	85.0	20.0	35.0	50.0

(b) Descriptive

Table 25: llama-3.3-70b (CoT)

Polarity	Pattern	TA	FA	TC	FC
pos-pos	$p \rightarrow q$	80.0	0.0	90.0	10.0
pos-neg	$p \rightarrow \neg q$	60.0	0.0	0.0	100.0
neg-pos	$\neg p \rightarrow q$	75.0	0.0	80.0	40.0
neg-neg	$\neg p \rightarrow \neg q$	55.0	0.0	5.0	100.0

(a) Deontic

Polarity	Pattern	TA	FA	TC	FC
pos-pos	$p \rightarrow q$	90.0	0.0	70.0	15.0
pos-neg	$p \rightarrow \neg q$	85.0	0.0	10.0	75.0
neg-pos	$\neg p \rightarrow q$	40.0	40.0	65.0	50.0
neg-neg	$\neg p \rightarrow \neg q$	15.0	60.0	30.0	85.0

(b) Descriptive

Table 26: olmo-2-0325-32b (Zero-Shot)

Polarity	Pattern	TA	FA	TC	FC
pos-pos	$p \rightarrow q$	0.0	0.0	100.0	95.0
pos-neg	$p \rightarrow \neg q$	0.0	0.0	60.0	100.0
neg-pos	$\neg p \rightarrow q$	90.0	0.0	15.0	95.0
neg-neg	$\neg p \rightarrow \neg q$	95.0	0.0	25.0	75.0

(a) Deontic

Polarity	Pattern	TA	FA	TC	FC
pos-pos	$p \rightarrow q$	10.0	5.0	85.0	90.0
pos-neg	$p \rightarrow \neg q$	10.0	5.0	80.0	85.0
neg-pos	$\neg p \rightarrow q$	75.0	0.0	25.0	100.0
neg-neg	$\neg p \rightarrow \neg q$	75.0	0.0	95.0	30.0

(b) Descriptive

Table 27: olmo-2-0325-32b (Few-Shot)

Polarity	Pattern	TA	FA	TC	FC
pos-pos	$p \rightarrow q$	100.0	0.0	80.0	20.0
pos-neg	$p \rightarrow \neg q$	90.0	0.0	5.0	100.0
neg-pos	$\neg p \rightarrow q$	90.0	5.0	55.0	50.0
neg-neg	$\neg p \rightarrow \neg q$	80.0	5.0	15.0	100.0

(a) Deontic

Polarity	Pattern	TA	FA	TC	FC
pos-pos	$p \rightarrow q$	95.0	5.0	85.0	15.0
pos-neg	$p \rightarrow \neg q$	95.0	0.0	10.0	95.0
neg-pos	$\neg p \rightarrow q$	80.0	30.0	40.0	50.0
neg-neg	$\neg p \rightarrow \neg q$	40.0	65.0	10.0	90.0

(b) Descriptive

Table 28: olmo-2-0325-32b (CoT)