Beyond Memorization: Distinguishing between Reductive and Epistemic Reasoning in LLMs using Classic Logic Puzzles

Anonymous ACL submission

Abstract

Epistemic reasoning requires agents to infer the state of the world from partial observations and information about other agents' knowledge. Prior work evaluating LLMs on canonical epistemic puzzles interpreted their behavior through a dichotomy between epistemic reasoning and brittle memorization. We argue that this framing is incomplete: in recent models, memorization is better understood as a special case of *reduction*, where a new instance is mapped onto a known problem. Instead, we introduce a *reduction ladder*, a sequence of modifications that progressively move instances away from a canonical epistemic puzzle, making reduction increasingly difficult while preserving the underlying logic. We find that while some large models succeed via reduction, other models fail early, and all models struggle once epistemic reasoning is required.

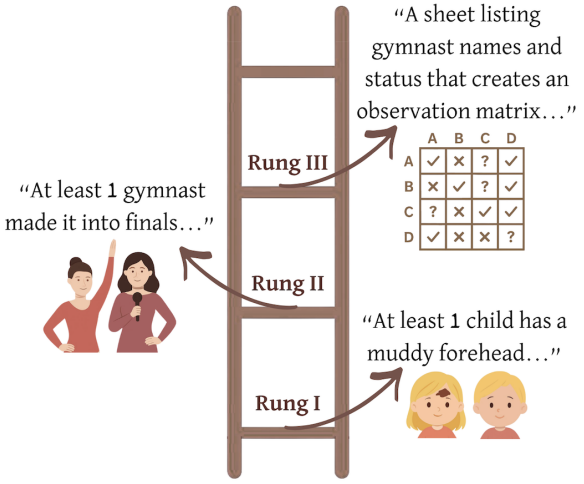

Figure 1: **The reduction ladder.** A progression of puzzle modifications designed to separate epistemic reasoning from reductive reasoning, by gradually obscuring the structure of the classic *Muddy Children* problem and making reduction-based solutions harder to apply.

1 Introduction

Problems requiring epistemic reasoning in multi-agent systems involve inferring the state of the world by integrating one's own observations with information about other agents' knowledge (Fagin et al., 1995). Such challenges arise naturally in collaborative settings where each agent has only a partial view of the world, such as legal proceedings, education, or medical applications (Rabinowitz et al., 2018; Li et al., 2023).

Previous work evaluating LLMs' epistemic reasoning has centered on canonical problems (e.g., the *Muddy Children* problem), with the goal of distinguishing between two possible solution strategies: *epistemic reasoning* and *memorization*. In the former, models reason over the specific parameters of the instance to derive a solution; in the latter, they match the task to a familiar pattern seen during pretraining and reproduce the corresponding answer. Prior findings suggest that earlier models often over-rely on memorization, which frequently breaks under superficial perturbations (Sileo and Lernould, 2023; Jiang et al., 2024).

In this work, we argue that recent advances in LLMs warrant a more nuanced treatment of memorization, viewing it as a special case of a broader and often effective strategy also employed by humans: *reduction* between problems (Gentner, 1983). Given an instance x of a problem A, a reduction is a function f that maps x to an instance $f(x)$ of another problem B. In contrast to epistemic reasoning, the central challenge in reduction is not reasoning over the specific parameters of the instance itself, but identifying the appropriate target problem B and the correspondence between the elements of the two problems. Under this view, memorization is the case in which $f(x) = x$.

We therefore ask: *Can reduction and epistemic reasoning be distinguished in LLMs, and can performance under each be quantified?* We show that canonical epistemic problems provide a natural framework for investigating this question. Instead

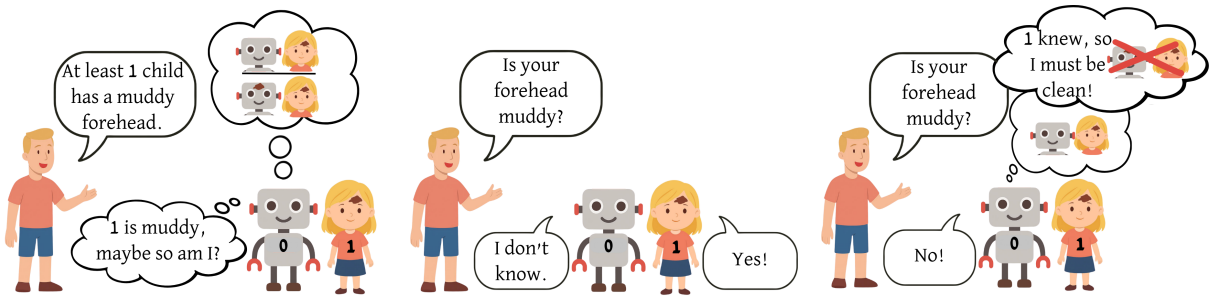

(A) Public Announcement: Following the announcement, child 1 eliminates the possibility that both children are clean.

(B) Round 1: Child 0 cannot rule out possibilities based on the public announcement, while child 1 knows she is muddy.

(C) Round 2: Child 0 knows he is clean, since if he were muddy, child 1 would have remained uncertain.

Figure 2: An example of the *Muddy Children* puzzle with two children, one of whom is muddy, where each child observes the other child but not their own forehead. Thought bubbles depict the possibilities the agent considers.

of seeking specific perturbations that "break" a model, we propose a *reduction ladder* (Figure 1) consisting of progressively larger changes that move instances farther from the original problem. For example, we change the classic *Muddy Children* problem to deal with Olympic gymnasts and introduce non-symmetric agent observations. Measuring performance along this ladder, together with trends in chain-of-thought reasoning, allows us to characterize how models behave as reduction becomes increasingly difficult to apply and success depends more heavily on epistemic reasoning.

We evaluate various LLMs and find that some large models can solve most of these problems by reduction, whereas other models fail at earlier steps. Once reduction ceases to be effective, all models struggle to solve problems through epistemic reasoning. Beyond the benchmark itself, this distinction provides an interpretable lens on model behavior, helping reveal not only how well models perform, but also *how* they succeed or fail.

Our contributions are threefold: (1) we reinterpret "memorization" as a form of *reduction*; (2) we introduce a *reduction ladder* to separate reduction from epistemic reasoning; and (3) we show that while some large LLMs solve many variants through reduction, all models struggle when epistemic reasoning is required.

2 Background: Dynamic Epistemic Logic and the *Muddy Children* Puzzle

Dynamic Epistemic Logic (DEL) is a formal framework for multi-agents systems, representing what they know, which possibilities they consider, and how this changes when new information is revealed. In this framework, *agents* are the entities whose knowledge is being modeled, and *public announcements* rule out possibilities that are inconsistent with newly available information, unfolding across multiple *rounds* of communication (Baltag et al., 1998; Van Ditmarsch et al., 2008).

We focus on the *Muddy Children* puzzle, a simple setting that captures many central ideas of DEL, illustrated for two children in Figure 2. Each child can see the other's forehead but not their own, so their uncertainty is represented by the possible worlds in the thought bubbles. In panel 2A, both remain uncertain about themselves; the father's announcement that at least one child is muddy rules out the all-clean world. This lets child 1 infer that she is muddy, and after she announces this in panel 2B, child 0 can further narrow the possibilities and reach his own conclusion in panel 2C.

3 Distinguishing Between Reductive and Epistemic Reasoning

We aim to distinguish between two ways in which LLMs may determine the correct response of an agent in DEL-style settings: *reductive reasoning* and *epistemic reasoning*. Under reductive reasoning, the model identifies the current epistemic state as a variant of a familiar problem and applies the corresponding response pattern; for instance, if it encountered the same puzzle during pretraining, it may retrieve the answer from parametric knowledge. Under epistemic reasoning, by contrast, the model derives the agent's response directly from the observations, public announcements, and prior answers in the interaction rounds.

To distinguish between reductive and epistemic reasoning in LLMs, we construct a "reduction ladder", illustrated in Figure 1, and elaborated below. The first rung in the ladder is seeded with a well-known epistemic problem, which LLMs have likely

encountered many times in their pretraining data. Each successive rung aims to move away from the previous one in terms of problem structure, while keeping the underlying logic. Therefore, if a model solves a problem via epistemic reasoning, its performance should not vary along the ladder. If, on the other hand, a model relies on reduction to a known problem, its performance may degrade the further we move from the base rung.

3.1 Task Definition

To evaluate LLMs' epistemic reasoning, we place the LLM in the role of an agent within a simulated world (e.g., a child in the muddy children puzzle). The model's goal is to infer its own status by reasoning over the information provided in its prompt.

Model Input. The model receives a prompt containing the public announcement, observations of the other agents, and the interaction history consisting of all answers provided by agents up to the particular round in question.

Model Output. The model is asked a question regarding its own status (e.g., "Is your forehead muddy?"). It must respond with exactly one of *"Yes"*, *"No"*, or *"I don't know"*. For any given round, exactly one response is correct based on the agent's current knowledge. To verify the model's response, we develop a symbolic solver grounded in Kripke possible worlds semantics (Kripke, 1959).

3.2 The Reduction Ladder

Our ladder consists of three rungs of problems, each moving away from the base problem.

Rung I: Seed problem. The first rung uses the classic *Muddy Children* problem. Our assumption is that LLMs have seen this puzzle many times during pretraining and hence will likely default to using a memorized solution, rather than solving it as a novel problem, via epistemic reasoning.

Rung II: Different settings which conflict with world knowledge. The goal of the second rung is to preserve the underlying epistemic reasoning, while moving away from the familiar setup of the *Muddy Children* puzzle. Specifically, We recast the puzzle in an *Olympic Games* setting: children are replaced by gymnasts, and muddy versus clean is replaced by qualified versus not qualified for the Olympic final. All other problem elements remain: gymnasts only know if others have qualified, and there is a public announcement declaring a lower bound for how many have qualified.

Inspired by (Jiang et al., 2024), we assign the queried agent a known name in a way that requires the model to go against its parametric knowledge. For example, in cases where the queried agent is not supposed to have qualified, we use the name famous gymnasts, so that solving the puzzle requires concluding that they did not qualify.

Rung III: Non-symmetric agents. The third rung reconsiders how agents observe one another and introduces a public announcement declaring an upper bound for how many have qualified. In the classic puzzle, each agent observes all others except themselves; here, we generalize this setup using an *observation matrix* $O \in \{0,1\}^{n \times n}$, inspired by Sileo and Lernould (2023), where $O_{i,j} = 1$ if and only if agent i observes the status of agent j. We sample each entry of O i.i.d. from $\{0,1\}$, while ensuring that the queried agent does not observe its own status, to prevent immediate inference.

Model	# Parameters
ChatGPT 5	?
ChatGPT 5-nano	?
Gemini 2.5-Pro	?
Gemini 2.5-Flash-Lite	?
Qwen 2.5-VL-Instruct	72B
Mistral Small-Instruct	24B
OLMo 3 Instruct	7B
OLMo 3.1 Think	32B

Table 1: Evaluated models. Parameter counts are listed only when publicly disclosed by the provider.

4 Experiments

In this section, we evaluate LLMs across the rungs of our ladder. In each rung, the model is asked to answer from the perspective of an agent in the problem, using the public announcement, the agents' observations, and the public history of the interaction up to a queried round.

4.1 Experimental setup

Models. We evaluate 8 models from diverse families, including ChatGPT (Ouyang et al., 2022), Gemini (Gemini Team, 2023), Mistral (Jiang et al., 2023), Qwen (Qwen Team, 2025), and OLMo (Groeneveld et al., 2024), as summarized in Table 1. We use the default temperature of 1 for the GPT-5 models and a temperature of 0 for all other models.

[Premise] There are n children… each child can clearly see every other child, but not themselves. The father then publicly announces that at least q children have muddy foreheads. Then, in repeated rounds, the father asks each child whether they know if their own forehead is muddy. You are one of the children…
[Example] For example, if there are 2 children…
[Hypothesis] Suppose there are 10 children… at least 2 children have muddy foreheads. In round 1, assuming that you see 2 children with muddy foreheads and the rest are clean; is your forehead muddy? Answer with "Yes", "No", or "I don't know".

Figure 3: **Reduced prompt example in the *Muddy Children* setup.** Each prompt contains a description of the protocol, a short example illustrating the required response format, and a test instance specifying the public interaction history and the target agent's observations.

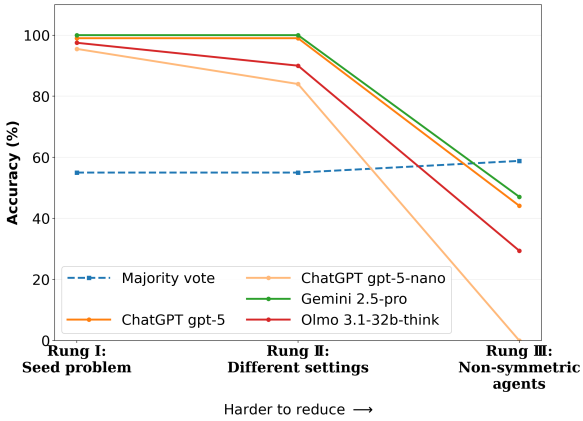

Figure 4: **Accuracy (%) across the epistemic complexity rungs.** The results illustrate an overall degradation in accuracy as the tasks move away from the *Muddy Children* problem and hence require harder reduction.

4.2 Results

We evaluate 1,320 puzzles for each of Rungs I and II, and 374 puzzles for Rung III (due to the technical complexity of generating solvable observation matrices), by uniformly sampling different puzzle parameters (number of agents, their status, the bound in the public announcement). Since some models fail to outperform the majority vote on Rung I and mostly reuse the same response (e.g., Olmo 3-7B-Instruct answers "No" in 96.7% of instances), we exclude them from the main analysis. As illustrated in Figure 3, We use a shared prompt template across all rungs.[1] Several observations can be drawn from our results.

Models seem to rely on reductive reasoning when solving epistemic problems. Figure 4 shows that, unlike earlier findings (Jiang et al., 2024), some LLMs can solve the problem under substantial surface changes in rung II. Still, accuracy significantly declines for all models as we move up the ladder. Since the rungs preserve the underlying epistemic logic while stripping away familiar cues, this suggests that early success depends partly on reductive rather than epistemic reasoning.

[1]See Appendix A for prompt examples and full results.

	Accuracy (%)	Reductive CoT(%)
Rung I	97.6%	62%
Rung II	96.7%	54%
Rung III	66.8%	15%

Table 2: **Manual analysis of Gemini 2.5 Pro across 100 CoT samples per rung**, showing correlation between use of reductive reasoning and accuracy.

Manual Analysis of CoT suggests reductive rather than epistemic reasoning. To better understand Gemini 2.5 Pro's reasoning,[2] we manually analyze 100 chain-of-thought (CoT) samples from each rung, and mark them as either reductive or epistemic. Table 2 shows strong performance in Rungs I and II is correlated with reductive reasoning: the model often explicitly identifies the task as a variant of a classic puzzle, even when its surface form differs substantially. In Rung III, however, reductive language largely disappears and is replaced by formal epistemic terminology, including references to Kripke models (Kripke, 1963) and the elimination of possible worlds, yet overall accuracy drops sharply. This suggests that while the model can solve DEL-style tasks by reduction, it struggles when success depends on epistemic reasoning.

Scale helps within model families but does not fully explain performance. Figure 4 shows a consistent within-family advantage for larger models (e.g., ChatGPT 5 is better than ChatGPT 5-nano). However, scale alone does not explain performance: despite being smaller, Olmo 3.1-32B-think outperforms Qwen 2.5-VL-72B-Instruct.

5 Conclusion

We revisited the evaluation of LLMs on epistemic puzzles, arguing that the contrast between epistemic reasoning and memorization is too coarse. We showed that strong performance seems to come from reducing to familiar problem structures rather than performing epistemic reasoning.

[2]We choose to focus on Gemini has it provides easy access to its chain-of-thought.

6 Limitations

A significant challenge lies in interpreting the internal mechanisms of LLMs. While we utilize Chain-of-Thought (CoT) to distinguish between reductive and epistemic reasoning, there may exist a disconnect between the model's verbalized reasoning and its actual computation; it is possible that LLMs perform reduction while mentioning terms that suggest epistemic reasoning in their CoT, and vice versa. This indication is in line with recent work that addresses the disconnect between CoT and the internal computation of the model (Levy et al., 2025b; Pfau et al., 2024; Pal et al., 2026; Levy et al., 2025a).

References

Alexandru Baltag, Lawrence S. Moss, and Slawomir Solecki. 1998. The logic of public announcements, common knowledge, and private suspicions. In *Proceedings of the 7th Conference on Theoretical Aspects of Rationality and Knowledge (TARK VII)*, pages 43–56.

Ronald Fagin, Joseph Y Halpern, Yoram Moses, and Moshe Y Vardi. 1995. *Reasoning About Knowledge*. MIT Press, Cambridge, Massachusetts.

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

Dedre Gentner. 1983. Structure-mapping: A theoretical framework for analogy. *Cogn. Sci.*, 7:155–170.

Dirk Groeneveld, Iz Beltagy, Pete Walsh, Akshita Bhagia, Rodney Kinney, Oyvind Tafjord, Ananya Harsh Jha, Hamish Ivison, Ian Magnusson, Yizhong Wang, and 1 others. 2024. Olmo: Accelerating the science of language models. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 15789–15809.

Albert Q. Jiang, Alexandre Sablayrolles, Arthur Mensch, Chris Bamford, Devendra Singh Chaplot, Diego de las Casas, Florian Bressand, Gianna Lengyel, Guillaume Lample, Lucile Saulnier, Lélio Renard Lavaud, Marie-Anne Lachaux, Pierre Stock, Teven Le Scao, Thibaut Lavril, Thomas Wang, Timothée Lacroix, and William El Sayed. 2023. Mistral 7b. *Preprint*, arXiv:2310.06825.

Bowen Jiang, Yangxinyu Xie, Zhuoqun Hao, Xiaomeng Wang, Tanwi Mallick, Weijie J. Su, Camillo J. Taylor, and Dan Roth. 2024. A peek into token bias: Large language models are not yet genuine reasoners. In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 4722–4756, Miami, Florida, USA. Association for Computational Linguistics.

Saul Kripke. 1963. Semantical considerations on modal logic. *Acta Philosophica Fennica*, 16:83–94.

Saul A Kripke. 1959. A completeness theorem in modal logic1. *The journal of symbolic logic*, 24(1):1–14.

Mosh Levy, Zohar Elyoseph, and Yoav Goldberg. 2025a. Humans perceive wrong narratives from ai reasoning texts. *arXiv preprint arXiv:2508.16599*.

Mosh Levy, Zohar Elyoseph, Shauli Ravfogel, and Yoav Goldberg. 2025b. State over tokens: Characterizing the role of reasoning tokens. *arXiv preprint arXiv:2512.12777*.

Huao Li, Yu Chong, Simon Stepputtis, Joseph Campbell, Dana Hughes, Charles Lewis, and Katia Sycara. 2023. Theory of mind for multi-agent collaboration via large language models. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 180–192, Singapore. Association for Computational Linguistics.

Long Ouyang, Jeff Wu, Xu Jiang, Diogo Almeida, Carroll L. Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, John Schulman, Jacob Hilton, Fraser Kelton, Luke Miller, Maddie Simens, Amanda Askell, Peter Welinder, Paul F. Christiano, Jan Leike, and Ryan Lowe. 2022. Training language models to follow instructions with human feedback. In *Advances in Neural Information Processing Systems*, volume 35, pages 27730–27744.

Koyena Pal, David Bau, and Chandan Singh. 2026. Do explanations generalize across large reasoning models? *arXiv preprint arXiv:2601.11517*. Under review.

Jacob Pfau, William Merrill, and Samuel R. Bowman. 2024. Let's think dot by dot: Hidden computation in transformer language models. *arXiv preprint arXiv:2404.15758*.

Qwen Team. 2025. Qwen3 technical report. *Preprint*, arXiv:2505.09388.

Neil Rabinowitz, Frank Perbet, Francis Song, Chiyuan Zhang, S. M. Ali Eslami, and Matthew Botvinick. 2018. Machine theory of mind. In *Proceedings of the 35th International Conference on Machine Learning*, volume 80 of *Proceedings of Machine Learning Research*, pages 4218–4227. PMLR.

Damien Sileo and Antoine Lernould. 2023. MindGames: Targeting theory of mind in large language models with dynamic epistemic modal logic. In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 4570–4577, Singapore. Association for Computational Linguistics.

Hans Van Ditmarsch, Wiebe van Der Hoek, and Barteld Kooi. 2008. *Dynamic epistemic logic*. Springer.

A Prompt Examples and Additional Results

A.1 Prompt examples

Figures 6-8 show example prompts from the three rungs of the reduction ladder.

A.2 Experimental Setup and Additional Results

Problem Instance structure. We define a problem instance using the tuple (n, k, t, q, j, O). In this formulation, n represents the total number of agents (e.g, the number of children in the *Muddy Children* puzzle), and k denotes the count of agents possessing the hidden status (e.g., being "muddy"). The parameter t represents the bound type (lower or upper) and q represents the bound value in the public announcement (e.g., "at least q are muddy"), while j specifies the target reasoning round at which the agent's knowledge is evaluated. Finally, $O \in \{0,1\}^{n \times n}$ is an observation matrix where the entry $O_{i,i'} = 1$ if agent i observes the status of agent i', and $O_{i,i'} = 0$ otherwise.

An example of this parameterization is illustrated in Figure 2. There, the number of agents is $n = 2$, the number of agents possessing the hidden status (muddy) is $k = 1$, the boundary is a lower bound of value $q = 1$, and the observations are as in the matrix $O = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}$, meaning each child observes the other child, but not themselves.

Our experiments use a fixed number of agents, $n = 10$, across all rungs. For each configuration (n, k, t, q, j, O), we evaluate agent 0 - the queried agent - under both possible hidden states (e.g., "muddy" or "clean"), provided that the assigned state is consistent with the specified number k of agents holding the hidden status. The remaining parameters are generated for each rung as follows:

Rung I: Seed problem and Rung II: Different settings. This configuration covers all valid combinations of marked agents $k \in [0, 10]$, bound value q (lower $q \in [0, k]$ or upper $q \in [k, 10]$), and reasoning rounds $j \in [1, 11]$. The diagonal of the observation matrices is 0 ($O_{i,i} = 0$) and the rest is filled with 1 ($O_{i,j} = 1, i \neq j$). This results in 1320 problem instances per rung.

Rung III: Non-symmetric agents. This dataset targets increased epistemic depth with 374 problem instances. It uses an upper bound $q = 2$, reasoning rounds $j \in [1, 11]$, marked agents $k \in [0, 2]$, and observation matrices that exclude self-observation ($O_{0,0} = 0$), sampled randomly from $\{0,1\}^{n \times n}$.

Post-process and verification. In post-processing, we evaluate only the first non-empty line of each response, since the prompt instructs the model to provide its answer on the first line and any explanation on a new line. We accept exactly one of "Yes", "No", and "I don't know", while allowing minor formatting variation such as surrounding whitespace, a final period, and apostrophe variants in "don't".

Results. Figure 5 reports the full results for all eight models and complements the main-text figure, which focuses on the four strongest models. The broader pattern remains the same: the accuracy declines as we move up the reduction ladder, with the sharpest drop on Rung III. The appendix figure also makes explicit why four models were omitted from the main analysis. Already on the seed problem, these models either fail to beat the majority vote of 41.7% or only marginally exceed it, while repeatedly falling back on the same response across many instances.

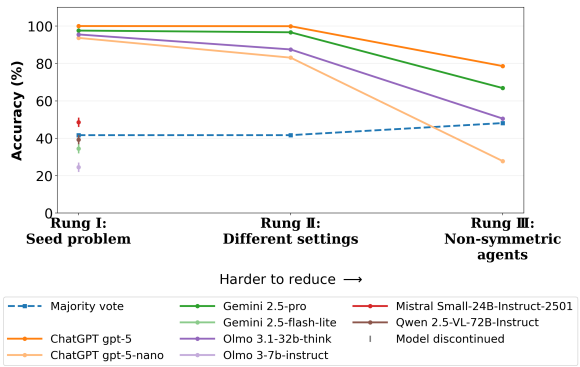

Figure 5: **Accuracy (%) across the epistemic complexity rungs.** The results illustrate an overall degradation in accuracy as the tasks move away from the *Muddy Children* problem and hence require harder reduction.

Rung I: Seed problem

[Premise] The following is a problem with parameters that may vary between instances:
There are n children playing in the mud. Their father calls them to the house and arranges them in a semicircle, so that each child can clearly see every other child, but not themselves. The father then publicly announces that at least q children have muddy foreheads. Then, in repeated rounds, the father asks each child whether they know if their own forehead is muddy. In each round, the children may rely only on the round number, the answers from the previous rounds, and the observed muddy/clean status of the other children.
You are one of the children, and in round j, you need to answer the question "Is your forehead muddy?" based on what you see, the round number, and the answers from the previous rounds. Answer with "Yes", "No", or "I don't know". Write your answer without additional text and, on a new line, explain your reasoning.
[Example] For example, if there are 2 children, and the father announces that at least 2 have muddy foreheads, then in round 1, you should respond "Yes" for both children.
Answer the following question:
[Hypothesis] Suppose there are 10 children and the father publicly announces that at least 2 children have muddy foreheads. In round 1 the muddy children answered that they knew if they had mud on their forehead, while the clean children answered that they did not know. In round 2, assuming that you see 1 child with a muddy forehead and the rest are clean; is your forehead muddy? Answer with "Yes", "No", or "I don't know". Write your answer without additional text and, on a new line, explain your reasoning.

Figure 6: **Prompt example for Rung I (Seed problem).** The canonical *Muddy Children* puzzle used as a seed problem in our reduction ladder.

Rung II: Different settings which conflict with world knowledge

[Premise] The following is a problem with parameters that may vary between instances:
There are n gymnasts competing in an Olympic qualification round. At the end of the round, the coach gives each gymnast a sheet listing every other gymnast's name and status - advanced to the final or not - excluding their own status. The gymnasts can't share the given information with others. The coach then publicly announces that at least q gymnasts made it into the final. Then, in repeated rounds, the coach asks each gymnast whether they know if they made it into the final. In each round, the gymnasts may rely only on the round number, the answers from the previous rounds, and the observed status of the other gymnasts.
You are one of the gymnasts, and in round j, you need to answer the question "Did you make it into the final?" based on what you see, the round number, and the answers from the previous rounds. Answer with "Yes", "No", or "I don't know". Write your answer without additional text and, on a new line, explain your reasoning.
[Example] For example, if there are 2 gymnasts, and the coach announces that at least 2 gymnasts made it into the final, then in round 1, you should respond "Yes" for both gymnasts.
Answer the following question:
[Hypothesis] Suppose there are 10 gymnasts and the coach publicly announces that at least 2 gymnasts made it into the final. In round 1 the gymnasts who made it into the final answered that they knew if they made it into the final, while the gymnasts who did not make it answered that they did not know. In round 2, assuming that you are the famous singer Lizzo, you see 1 gymnast that made it into the final and the rest didn't; did you make it into the final? Answer with "Yes", "No", or "I don't know". Write your answer without additional text and, on a new line, explain your reasoning.

Figure 7: **Prompt example for Rung II (Different setting conflicting with world knowledge).** A prompt that preserves the epistemic structure of the seed problem while changing its surface form. Children are replaced by Olympic gymnasts, and muddy versus clean is replaced by qualified versus not qualified for the final. The queried agent is assigned a familiar name, forcing the model to rely on the context rather than its parametric knowledge.

Rung III: Non-symmetric agents

[Premise] The following is a problem with parameters that may vary between instances:
There are n gymnasts competing in an Olympic qualification round. At the end of the round, the coach gives each gymnast a sheet listing some gymnast's name and status - advanced to the final or not - that creates an observation matrix. The gymnasts can't share the given information with others. Each entry (i,k) in the observation matrix indicates whether gymnast i knows the status of gymnast k. A value of 1 means gymnast i knows gymnast k's status, while a value of 0 means they do not know. The coach then publicly announces that at most q gymnasts made it into the final. Then, in repeated rounds, the coach asks each gymnast whether they know if they made it into the final. In each round, the gymnasts may rely only on the round number, the answers from the previous rounds, and the observed status of the other gymnasts.
You are one of the gymnasts, and in round j, you need to answer the question "Did you make it into the final?" based on what you see, the round number, and the answers from the previous rounds. Answer with "Yes", "No", or "I don't know". Write your answer without additional text and, on a new line, explain your reasoning.
[Example] For example, if there are 2 gymnasts, and the coach announces that at most 0 gymnasts made it into the final, then in round 1, given the following observation matrix:
[0,1]
[1,0]
you should respond "No" for both gymnasts.
Answer the following question:
[Hypothesis] Suppose there are 10 gymnasts indexed by 0-9, and the coach publicly announces that at most 2 gymnasts made it into the final. In round 1, gymnasts 1, 3, 8, and 9 answered that they knew if they made it into the final, while the rest answered that they did not know if they made it into the final. In round 2, assuming that you are gymnast 0, the famous Olympic gymnast Simone Biles, and observe the world as described in the following observation matrix:
[0, 0, 0, 1, 0, 0, 0, 0, 0, 1]
[1, 1, 0, 0, 0, 0, 1, 1, 0, 1]
[1, 1, 0, 0, 0, 1, 0, 1, 0, 0]
[1, 0, 1, 1, 0, 1, 1, 1, 0, 0]
[1, 0, 1, 0, 0, 1, 0, 1, 0, 1]
[1, 0, 0, 0, 1, 0, 1, 0, 1, 0]
[0, 1, 1, 1, 1, 1, 0, 1, 0, 0]
[1, 1, 0, 0, 1, 0, 1, 0, 0, 1]
[1, 1, 1, 0, 1, 1, 0, 1, 1, 0]
[0, 0, 0, 1, 0, 0, 0, 1, 1, 1]
Specifically, you see that gymnasts 3 and 9 did not make it into the final; did you make it into the final? Answer with "Yes", "No", or "I don't know". Write your answer without additional text and, on a new line, explain your reasoning.

Figure 8: **Prompt example for Rung III (Non-symmetric agents).** A prompt illustrating the most challenging rung of the reduction ladder. Instead of the standard symmetric visibility pattern, the prompt provides an observation matrix indicating which gymnast knows the status of which other gymnast, together with a public upper-bound announcement on how many gymnasts qualified. Solving the instance requires reasoning from asymmetric observations and the public history of earlier rounds, making straightforward reduction to the classic *Muddy Children* puzzle more difficult.