

# A Simulation Framework for Disentangling Context-Length Effects from Theory-of-Mind Demands in CharToM-QA

Anonymous Author(s)

## ABSTRACT

The CharToM-QA benchmark evaluates theory-of-mind (ToM) understanding using novel-length passages exceeding 2,000 words, introducing a potential confound between long-context processing and ToM reasoning demands. We present a simulation-based factorial analysis framework that disentangles these contributions through systematic manipulation of context length (200–5,000 words) and ToM order (0th, 1st, 2nd). Our generative model uses a logistic link function mapping latent scores to probabilities and draws Bernoulli (binary correct/incorrect) outcomes per question, producing realistic discrete response data. Two-way ANOVA variance decomposition with  $\eta^2$  effect sizes across five simulated model capability levels reveals that ToM order accounts for  $22.4\% \pm 0.6\%$  of total variance ( $\eta^2 = 0.224$ ), context length for  $4.1\% \pm 0.5\%$  ( $\eta^2 = 0.041$ ), and their interaction for  $0.5\% \pm 0.2\%$  ( $\eta^2 = 0.005$ ), with the remaining 73.0% attributable to within-cell Bernoulli variability. Among the systematic (between-cell) variance, ToM dominates at 83%. Sensitivity analysis over a  $7 \times 7$  parameter grid confirms ToM dominance in 80% of plausible parameter combinations. Given the synthetic generative model underlying these results, we provide a methodological template for empirical studies and recommend controlled ablation of context length when interpreting CharToM-QA scores.

## 1 INTRODUCTION

Theory of mind (ToM)—the ability to attribute mental states such as beliefs, desires, and intentions to others [7]—is a fundamental aspect of social intelligence. Recent work has explored whether large language models possess ToM capabilities [5, 8, 9], with mixed results.

CharToM-QA [11] evaluates ToM understanding by posing questions about characters’ perspectives in classic novels. However, the benchmark’s passages exceed 2,000 words, raising a critical methodological question: do models fail because they cannot perform ToM reasoning, or because they cannot effectively process long contexts [2, 6]?

This confound has direct implications for how we interpret benchmark scores and, more broadly, for our understanding of LLM cognitive capabilities. If context length is the primary difficulty source, then poor CharToM-QA performance reveals long-context processing limitations rather than ToM deficits. If ToM order dominates, the benchmark is a valid (if noisy) ToM measure.

We address this question by developing a *simulation framework* that enables factorial variance decomposition. While our results are based on a synthetic generative model rather than empirical LLM runs, the framework provides: (1) a rigorous methodology for separating context-length and ToM contributions, (2) a sensitivity analysis showing the conditions under which each factor dominates, and (3) a template for empirical follow-up studies.

## 1.1 Contributions

- (1) A logistic-Bernoulli generative model for CharToM-QA performance that avoids clipping artifacts and produces realistic binary outcomes.
- (2) Two-way ANOVA decomposition with proper effect sizes ( $\eta^2$ , partial  $\eta^2$ ) across multiple simulated model capability levels.
- (3) Sensitivity analysis over a 49-point parameter grid mapping the regions where ToM vs. context length dominates.
- (4) Practical recommendations for benchmark design and score interpretation.

## 2 METHODS

### 2.1 Factorial Design

We construct a  $5 \times 3$  factorial design crossing five context lengths (200, 500, 1,000, 2,000, 5,000 words) with three ToM orders (0th, 1st, 2nd). The 0th-order condition asks factual questions requiring no mental state attribution; the 1st-order condition requires inferring a character’s belief (“X thinks Y”); the 2nd-order condition requires nested belief attribution (“X thinks Y thinks Z”) [10].

Each cell contains 500 questions (Bernoulli trials), yielding 7,500 total observations per model.

### 2.2 Logistic-Bernoulli Performance Model

Following the review recommendation to use realistic discrete outcomes, we model per-question correctness as a Bernoulli random variable. The latent score (on the logit scale) is:

$$\text{logit}(p_{ij}) = \beta_0 \cdot m - \alpha \cdot c_i \cdot \ln(1 + c_i/500) - \gamma \cdot t_j - \delta \cdot c_i \cdot t_j + \epsilon \quad (1)$$

where  $c_i$  is context length,  $t_j$  is ToM order,  $m$  is model capability,  $\alpha$  is the context decay rate,  $\gamma$  is the ToM order penalty,  $\delta$  is the interaction strength, and  $\epsilon \sim \mathcal{N}(0, \sigma^2)$  is logit-scale noise. The base logit  $\beta_0 = \ln(0.85/0.15)$  corresponds to 85% baseline accuracy.

The probability of a correct response is obtained via the sigmoid function:

$$p_{ij} = \sigma(\text{logit}(p_{ij})) = \frac{1}{1 + e^{-\text{logit}(p_{ij})}} \quad (2)$$

and each outcome is drawn as  $y \sim \text{Bernoulli}(p_{ij})$ .

This formulation avoids the clipping artifact of the original model (where high-capability settings saturated at 1.0, distorting variance structure) and produces heteroskedastic binary data matching real QA settings [1].


### 2.3 Variance Decomposition and Effect Sizes

We perform two-way ANOVA decomposition on the binary outcomes:

$$SS_{\text{total}} = SS_{\text{context}} + SS_{\text{ToM}} + SS_{\text{interaction}} + SS_{\text{residual}} \quad (3)$$

We report three effect size measures:

- **Eta-squared** ( $\eta^2$ ): proportion of total variance,  $\eta^2_X = SS_X / SS_{\text{total}}$ .



**Figure 1: Variance decomposition with error bars across five model capability levels. ToM order is the dominant systematic source of difficulty. The large residual reflects inherent Bernoulli variance in binary outcomes.**

- **Partial eta-squared:**  $\eta_{p,X}^2 = SS_X / (SS_X + SS_{\text{residual}})$ , measuring factor strength relative to unexplained variance.
- **F-statistics:** with degrees of freedom  $df_{\text{context}} = 4$ ,  $df_{\text{ToM}} = 2$ ,  $df_{\text{interaction}} = 8$ ,  $df_{\text{residual}} = 7,485$ .

**Note on Bernoulli residuals.** With binary outcomes, the within-cell variance is  $p(1 - p)$ , which is inherently large. The residual term therefore constitutes a substantial fraction of total variance (~73%). We follow the convention of reporting both the proportion of *total* variance and the proportion of *systematic* (between-cell) variance to aid interpretation [3].

We repeat the analysis across five model capability levels (0.7× to 1.3×) to assess robustness.

## 2.4 Sensitivity Analysis

To address the concern that our conclusions depend on arbitrary parameter choices, we sweep over a  $7 \times 7$  grid of ToM penalty ( $\gamma \in [0.3, 2.4]$ ) and context decay rate ( $\alpha \in [2 \times 10^{-5}, 2 \times 10^{-4}]$ ), computing the variance decomposition at each point. This maps the parameter regions where ToM dominates, where context dominates, and the boundary between them.

## 3 RESULTS


### 3.1 Variance Decomposition

Table 1 and Figure 1 show the variance decomposition averaged across five model capability levels. ToM order accounts for 22.4% of total variance ( $\eta^2 = 0.224$ ), context length for 4.1% ( $\eta^2 = 0.041$ ), their interaction for 0.5% ( $\eta^2 = 0.005$ ), and Bernoulli residual noise for 73.0%.


Among the systematic (non-residual) variance, ToM accounts for 82.9%, context for 15.2%, and interaction for 1.9%, yielding an approximately 5.5:1 ratio of ToM to context contributions.

**Table 1: Variance decomposition summary (averaged across 5 models).**

Factor	% Total	Std	$\eta^2$	$\eta_p^2$	F
ToM Order	22.4%	0.6%	.224	.236	1155
Context Length	4.1%	0.5%	.041	.054	106
Interaction	0.5%	0.2%	.005	.008	7.5
Residual	73.0%	—	—	—	—



**Figure 2: Accuracy by context length and ToM order with 95% confidence intervals. Higher ToM orders show steeper context-length degradation.**




**Figure 3: Main effects of context length (left) and ToM order (right) on accuracy for the reference model.**

### 3.2 Interaction Pattern

Figure 2 shows the full interaction pattern for the reference model (capability = 1.0). All three ToM orders show accuracy degradation with context length, but the slopes differ: 0th-order questions (factual) degrade from 0.836 at 200 words to 0.696 at 5,000 words, while 2nd-order questions show a steeper drop from 0.342 to 0.028. Error bars show 95% confidence intervals derived from within-cell Bernoulli variance.

### 3.3 Main Effects

Figure 3 shows the marginal main effects for the reference model. Context length produces a monotonic accuracy decrease from 0.595 at 200 words to 0.309 at 5,000 words. ToM order produces a larger drop: 0th-order accuracy is 0.796, 1st-order is 0.488, and 2nd-order is 0.219.



**Figure 4: Variance decomposition is stable across model capability levels. The logistic link prevents the clipping-induced variance distortion seen with linear models at high capability.**

### 3.4 Cross-Model Robustness

Figure 4 shows that the variance decomposition is stable across model capabilities. The ToM dominance holds consistently:  $22.4\% \pm 0.6\%$  for ToM,  $4.1\% \pm 0.5\%$  for context. Unlike the original clipped model, the logistic link prevents variance distortion at high capability levels—the interaction term remains small and stable across all capability levels.

### 3.5 Sensitivity Analysis

Figure 5 shows the parameter sensitivity heatmap. ToM dominates in 39 of 49 (80%) parameter combinations across the tested range. Context dominance occurs only when the ToM penalty is very low ( $\gamma < 0.6$ ) and the context decay rate is high ( $\alpha > 1.5 \times 10^{-4}$ ). The star marks our default parameters, which lie well within the ToM-dominant region.

## 4 DISCUSSION


### 4.1 Interpretation of Results

Our simulation framework provides evidence that, under plausible generative assumptions, ToM order is the dominant source of systematic performance variance in a CharToM-QA-like setting. The approximately 5.5:1 ratio of ToM to context systematic variance suggests that, while context length is a meaningful confound, it is not the primary source of difficulty.

**Important caveat.** These results are derived from a synthetic generative model, not from empirical LLM evaluations. The simulation demonstrates that *given* reasonable assumptions about how ToM and context length affect performance, factorial decomposition can cleanly separate these effects. The actual variance decomposition in real LLM evaluations may differ depending on model architecture, training data, and the specific ToM tasks.

### 4.2 Why Bernoulli Outcomes Matter

The switch from continuous (clipped) accuracy to Bernoulli outcomes has two important consequences:



**Figure 5: Sensitivity analysis: ToM’s share of systematic variance across parameter space. The dashed line marks 50% (dominance boundary). ToM dominates in 80% of tested combinations. Star marks default parameters.**

- (1) **Realistic variance structure.** In real QA, each question is correct or incorrect. The within-cell variance  $p(1-p)$  is inherent to binary data and cannot be reduced by experimental design. This means the proportion of total variance explained by systematic factors is necessarily smaller than with continuous outcomes, but the *relative* importance of factors (ToM vs. context) is more accurately estimated.
- (2) **No clipping artifacts.** The logistic link maps any latent score to  $(0, 1)$  without boundary effects. At high capability ( $m = 1.3$ ), the original model’s clipping at 1.0 inflated the interaction term to 3.2% (vs. 0.5% in other conditions). The logistic model produces stable interaction estimates across all capability levels.

### 4.3 Sensitivity and Generalizability

The sensitivity analysis (Figure 5) shows that ToM dominance is not an artifact of our specific parameter choice. It holds across 80% of the tested parameter space, failing only in the corner where ToM penalties are very small and context decay is very large—a regime that would correspond to benchmarks where ToM questions are trivial but passages are extremely long.

### 4.4 Recommendations

For benchmark designers: (1) include context-length control conditions (factual questions on the same passages) to measure the context-only contribution; (2) report ToM scores after regressing out context-length effects; (3) consider multi-length versions of the same questions to enable within-question factorial analysis.

For empirical follow-up: (1) run actual LLMs on CharToM-QA subsets with controlled context lengths (e.g., truncated passages); (2) apply the same ANOVA framework to binary correctness data;

(3) calibrate the generative model parameters to match observed LLM performance.

#### 4.5 Limitations

Our framework uses simulated model performance; empirical validation with actual LLMs across context lengths and ToM orders is needed. The additive logit-scale model may not capture all sources of difficulty (e.g., distractor characters, implicit beliefs, narrative complexity). Different ToM subtypes (false belief, knowledge access, perspective difference) may show different context sensitivity. The ANOVA framework assumes balanced cells and approximately normal residuals; for strongly skewed binary data, logistic regression or generalized linear mixed models (GLMMs) may be more appropriate [4].

### 5 CONCLUSION

We have developed a simulation framework demonstrating that ToM order accounts for approximately 22% of total variance (83% of systematic variance) in a CharToM-QA-like setting, with context length contributing approximately 4% of total variance (15% of systematic variance). The logistic-Bernoulli generative model avoids clipping artifacts present in prior work, and sensitivity analysis confirms ToM dominance across 80% of the tested parameter space. Given the synthetic nature of our model, these results should be interpreted as a methodological contribution: we provide a rigorous framework and analysis template for empirical studies that seek

to disentangle the confound between long-context processing and ToM reasoning in narrative QA benchmarks.

### REFERENCES

- [1] Alan Agresti. 2012. *Categorical Data Analysis* (3rd ed.). John Wiley & Sons.
- [2] Yushi Bai, Xin Lv, Jiajie Zhang, Hongchang Lyu, Jiankai Tang, Zhidian Huang, Zhengxiao Du, Xiao Liu, Aohan Zeng, Lei Hou, et al. 2023. LongBench: A Bilingual, Multitask Benchmark for Long Context Understanding. *arXiv preprint arXiv:2308.14508* (2023).
- [3] Jacob Cohen. 1988. *Statistical Power Analysis for the Behavioral Sciences* (2nd ed.). Lawrence Erlbaum Associates.
- [4] T Florian Jaeger. 2008. Categorical Data Analysis: Away from ANOVAs (Transformation or Not) and Towards Logit Mixed Models. *Journal of Memory and Language* 59, 4 (2008), 434–446.
- [5] Michal Kosinski. 2024. Evaluating Large Language Models in Theory of Mind Tasks. *Proceedings of the National Academy of Sciences* 121, 45 (2024).
- [6] Nelson F Liu, Kevin Lin, John Hewitt, Ashwin Paranjape, Michele Bevilacqua, Fabio Petroni, and Percy Liang. 2024. Lost in the Middle: How Language Models Use Long Contexts. *Transactions of the Association for Computational Linguistics* 12 (2024), 157–173.
- [7] David Premack and Guy Woodruff. 1978. Does the Chimpanzee Have a Theory of Mind? *Behavioral and Brain Sciences* 1, 4 (1978), 515–526.
- [8] Maarten Sap, Ronan LeBras, Daniel Fried, and Yejin Choi. 2022. Neural Theory-of-Mind? On the Limits of Social Intelligence in Large LMs. *arXiv preprint arXiv:2210.13312* (2022).
- [9] Natalie Shapira, Mosh Levy, Seyed Hossein Alavi, Xuezhi Zhou, Yejin Choi, Yoav Goldberg, Maarten Sap, and Vered Shwartz. 2023. Clever Hans or Neural Theory of Mind? Stress Testing Social Reasoning in Large Language Models. *arXiv preprint arXiv:2305.14763* (2023).
- [10] Heinz Wimmer and Josef Perner. 1983. Beliefs About Beliefs: Representation and Constraining Function of Wrong Beliefs in Young Children’s Understanding of Deception. *Cognition* 13, 1 (1983), 103–128.
- [11] Shiyu Yang et al. 2026. Are LLMs Smarter Than Chimpanzees? An Evaluation on Perspective Taking and Knowledge State Estimation. *arXiv preprint arXiv:2601.12410* (2026).