

Leveraging Human Production-Interpretation Asymmetries to Test LLM Cognitive Plausibility

Suet-Ying Lam^{1*}, Qingcheng Zeng^{2*}, Jingyi Wu², Rob Voigt²

¹UMass Amherst, ²Northwestern University

Abstract

Whether large language models (LLMs) process language similarly to humans has been the subject of much theoretical and practical debate. We examine this question through the lens of the production-interpretation distinction found in human sentence processing and evaluate the extent to which instruction-tuned LLMs replicate this distinction. Using an empirically documented asymmetry between pronoun production and interpretation in humans for implicit causality verbs as a testbed, we find that some LLMs do quantitatively and qualitatively reflect human-like asymmetries between production and interpretation. We demonstrate that whether this behavior holds depends upon both model size—with larger models more likely to reflect human-like patterns and the choice of meta-linguistic prompts used to elicit the behavior. Our codes and results are available [here](#).

1 Introduction and Background

The extent to which large language models (LLMs) are “cognitively plausible,” that is, replicate human-like behaviors in language processing, has been the subject of ongoing debate (Dentella et al., 2023; Hu et al., 2024; Futrell and Mahowald, 2025; Kuribayashi et al., 2025). Existing research on the linguistic capabilities of LLMs has predominantly focused on their performance in language interpretation, e.g., pragmatic understanding (Hu et al., 2023), sentence acceptability judgment (Warstadt et al., 2020), garden-path effect (Futrell et al., 2019), reference resolution (Lam et al., 2023). In this study, we examine a previously unexplored dimension of cognitive plausibility: whether LLMs reflect human-like distinctions between **production** and **interpretation** in language processing.

Production and interpretation were traditionally treated as two independent processes in human language: for instance, in neurolinguistics the “classic” Lichtheim–Broca–Wernicke model assumes distinct anatomical pathways associated with production and interpretation (see Ben Shalom and Poeppel 2008). While this extreme dichotomy has been rejected recently (see Pickering and Garrod 2013), humans do exhibit different underlying biases in language processing between production and interpretation even in very similar tasks. Whether such distinctions carry over into LLMs is of particular interest when we consider that the fundamental unit of LLM computation is $P(\text{token}|\text{context})$, which is inherently ambiguous between production and interpretation and is practically applied towards both types of tasks.

The present study probes into this question using reference processing as a test case. Consider the following examples:

- (1) A production task: next-mentioned bias
 - a. John infuriated Bill. ... [IC1]
 - b. John praised Bill. ... [IC2]
- (2) A interpretation task: ambiguous pronoun resolution
 - a. John infuriated Bill. He ... [IC1]
 - b. John praised Bill. He ... [IC2]

When asked to continue the story in (1), speakers usually describe events that happened to one of the two mentioned characters. This is a production task, and psycholinguistic research has investigated how the preceding context affects the **next-mention bias** of the character, i.e., how likely a character will be referred to in the continued story $P(\text{referent}|\text{context})$. In (1-a), ‘John’ has a higher next-mention bias than ‘Bill’, because he is the implicit cause of the event. Verbs like ‘infuriate’ are therefore called the subject-biased im-

*Both authors contributed equally. Correspondence to qcz@u.northwestern.edu

PLICIT CAUSALITY (IC1) verbs. In contrast, ‘Bill’ has a higher next-mention bias than ‘John’ in (1-b), because ‘Bill’ implies an implicit cause for ‘John’ to praise him. Verbs like ‘praise’ are therefore called object-biased implicit causality (IC2) verbs (e.g., [Stevenson et al. 1994](#)).

A similar implicit causality bias can also be found in (2). Although participants are asked to perform the same sentence continuation task as in (1), note that an additional interpretation element is involved here: since an additional pronoun is provided after the context sentence, participants must first resolve the ambiguous pronoun “he” $P(\text{referent}|\text{pronoun})$ before providing a reasonable continuation.

Studies have shown that human participants were more likely to resolve the ambiguous pronoun to the subject ‘John’ than the object ‘Bill’ with an IC1 verb. That is, ‘John’ has a higher interpretation bias than ‘Bill’ under this condition. This interpretation bias flips when the verb is an IC2 verb: ‘Bill’ were more likely to be the antecedent of the ambiguous pronoun ‘he’ than ‘John’ in (2-b) (e.g., [Crinean and Garnham 2006](#)).

Given the additional interpretation element, henceforth we refer to the sentence continuation task with a given pronoun such as (2) as an **interpretation task** for measuring pronoun interpretation bias; and the one without a given pronoun such as (1) as the **production task** for measuring next-mentioned bias, despite the fact that the sentence continuation task is in principle a production task.

Crucially, psycholinguistic research has revealed an asymmetry between these two biases. In the interpretative case of pronoun interpretation bias, humans are robustly more likely to show a preference for the subject than in the production case of next-mention bias cross-linguistically (English: [Rohde and Kehler 2014](#); Mandarin: [Zhan et al. 2020](#); [Lam and Hwang 2024](#); German: [Patterson et al. 2022](#); Catalan: [Mayol 2018](#)). That is, participants were more likely to resolve an ambiguous pronoun towards the subject than choose the subject as the next referent, despite the same context.

For humans, this extra subject bias in interpretation comes from the bias of using pronouns for subject antecedent, i.e., $P(\text{he}|\text{subject})$. That is, when they see an ambiguous pronoun, they do not only consider the next-mention bias of which antecedent is more likely to be mentioned, but also why a pronoun is used. It is unknown whether and how LLMs can handle this difference, as they do

not generate such a conditional probability based on the choice of next referent instead of the context.

We therefore probe this dimension of LLM cognitive plausibility using this task, asking (1) whether the IC verb-type effect is reflected by LLMs in both production and interpretation; and (2) whether a human-like asymmetry between the two biases exists. This is our first set of questions: do LLMs show human-like interpretation and production biases, and if so under what conditions? Do human-like effects scale with parameter count?

Evaluating LLM in metalinguistic prompts

[Hu and Levy \(2023\)](#) demonstrated that direct probability-based measures in general outperformed meta-linguistic prompting in assessing plausibility and syntactic processing tasks. However, not all language processing tasks can be effectively quantified using probability-based measures, and for some tasks meta-linguistic prompts are the only possible method to measure processing. This is exactly our case: in the ambiguous pronoun resolution task, the bias towards the subject ‘John’ or the object ‘Bill’ is represented by the same term, i.e., $P(\text{he}|\text{context})$. Metalinguistic prompting is thus necessary to elicit meaningful results.

One might expect LLM performance to vary within different metalinguistic prompts, but it is unclear which type of metalinguistic prompts would perform better. Although [Hu and Levy \(2023\)](#) demonstrated that metalinguistic prompts that are more similar to the direct probability baseline perform better, we would like to explore whether probability measurements obtained via metalinguistic prompting also provide greater reliability. This constitutes the second aim of this paper: across different metalinguistic prompting strategies, which elicit more human-like language processing behaviors?

Our findings are summarized below: (1) in the most cases, LLMs cannot capture the difference between the production and the interpretation task in reference processing; if any, the asymmetry is limited to certain meta-linguistic prompts and rarely reach the magnitude found in human participants; and (2) the choice of meta-linguistic prompts matters in evaluating LLMs: most LLMs perform in a more human-like manner with one specific prompt that is unrelated to probability measure, namely the Yes/No prompt (details explained in Section 2.2).

2 Methodology

2.1 Stimuli

We constructed stimuli in the frame of “[Character A] IC-verb [Character B]” without pronouns for the next-mention production task (as in (1)), and with a pronoun with an ambiguous referent for the pronoun resolution task (as in (2)). We selected 137 IC1 verbs and 134 IC2 verbs from the original study that found the production-interpretation asymmetry (Rohde and Kehler, 2014), and the English IC verb corpus (Ferstl et al., 2011). For each verb, we heuristically created two items by assigning a pair of male names and a pair of female names randomly selected from 13 unambiguously male and 13 unambiguously female names from Rohde (2008). The congruence of the gender of the characters ensures the ambiguity of the pronoun. This results in 541 items in each task.

2.2 Models and Metalinguistic Prompting

Our experiments evaluate four representative LLMs from open-source LLMs ranging from 8B to 70B and one proprietary LLM: LLaMA3.1Instruct-8B (Grattafiori et al., 2024), QWen2.5Instruct-32B (Qwen et al., 2025), LLaMA3.3Instruct-70B (Grattafiori et al., 2024), and GPT-4o (OpenAI et al., 2024). We focus on instruction-tuned model¹ as they allow the effective use of metalinguistic prompting, and varying parameter counts also allow us to assess the impact of model scaling on human-like language processing behavior. Specifically, we constantly use greedy decoding in our generation. We employ four metalinguistic prompt strategies to assess LLM behavior:

- (i) Binary choice prompting: The model is prompted to select between subject and object.
- (ii) Continuation prompting: The model is instructed to extend the sentence by continuing with either the subject or the object.
- (iii) Yes/No prompting: The model is asked whether the following sentence (or the existing pronoun) will begin with the *subject*, requiring a binary response.
- (iv) Yes/No probability prompting: Similar to (iii), but instead of a categorical response, we ex-

¹We also experiment with base LLMs. However, even for continuation prompting, base LLMs could not follow instructions reliably.

tract the probability assigned to the *Yes* token as a quantitative measure.

Three authors of the paper manually verified all model outputs to confirm subject/object choice and exclude ambiguous, nonsensical, plural responses, and responses with repeated pronoun in the interpretation task. Table 2 reported the distribution of the excluded responses over the 1082 responses of each model in the two tasks.

Table 1: The distribution of excluded responses in continuation prompting

	Ambiguous	Non-sensical	Plural	Repeated pronoun
LLaMA-8B	54	21	17	0
LLaMA-70B	35	9	0	0
Qwen	17	6	0	0
GPT-4o	22	1	8	5

The specific prompts used in our experiments are provided in Appendix A, and the details of the exclusion criteria in Appendix B.1.

2.3 Evaluation

We evaluated model behavior from two perspectives. We first consider whether the IC verb type effect and the production-interpretation asymmetry exists in LLMs, which we both visualize to observe broad trends and verify with statistical tests like those run in human experiments. We then consider the magnitude of the effect found in LLMs, as psycholinguistic research has found that language models often fail to replicate the magnitude of effects found in human participants even when the directionality is similar. Below we only report effects that are verified by statistical evidence, the details of which be found in Appendix B.

3 Results

Result #1: Implicit causality biases are sometimes replicated by LLMs, depending on prompts and models. Figure 1 presents the task performance across different metalinguistic prompts for each model. The implicit causality (IC) verb effect — where subjects are chosen more frequently after IC1 verbs than IC2 verbs — was observed in at least one metalinguistic prompt for all models, though the strength and consistency of this effect varied by both model and prompt type.

For LLaMA models, the IC verb effect was present in both production and interpretation when using Yes/No and Yes/No probability prompts, indicating a broader sensitivity to IC biases across tasks. In contrast, the IC verb effect was limited

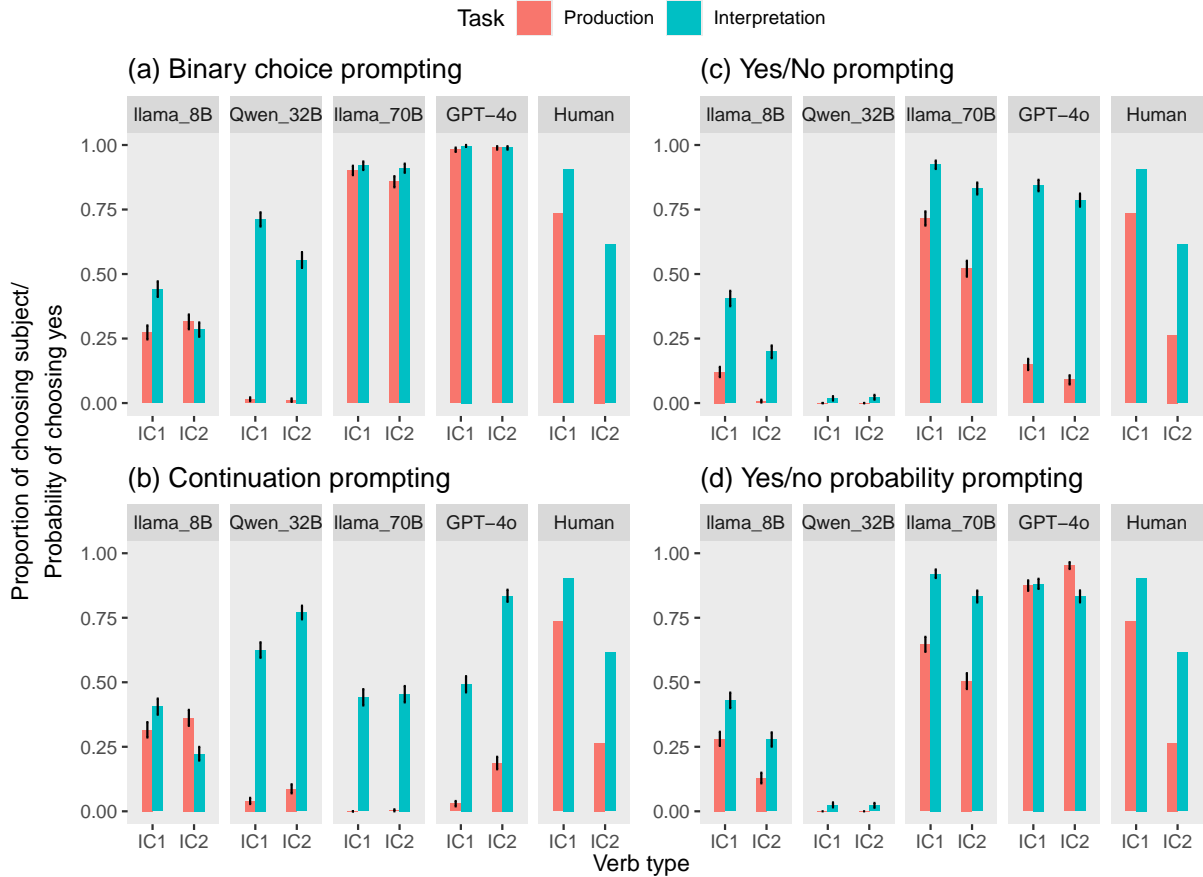


Figure 1: Model behavior as proportion of subject/yes choice as the antecedent by prompting strategy. Human behavior – rightmost facet in each subfigure for reference, from Rohde and Kehler (2014) – tends to reflect higher subject bias for IC1 over IC2 verbs, but with asymmetry between production and interpretation tasks.

to Yes/No prompting for GPT-4o and even more restricted for Qwen, in which the effect was only observed in interpretation and limited to binary choice prompt. This suggests that the model’s sensitivity to IC verbs may depend on task framing.

Result #2: The production-interpretation asymmetry is limitedly captured. Recall that human participants are more subject-biased in interpretation than in production. This difference is only elicited in LLaMA models using Yes/No and Yes/No probability prompts, and GPT-4o using Yes/No prompt and continuation prompt. Qwen was entirely unable to capture this asymmetry because it failed to predict the IC verb effect in production with all four metalinguistic prompts.

Another unexpected production-interpretation asymmetry is that LLMs generally are less likely to capture the IC verb bias in production than in interpretation. For instance, while the LLaMA-8B model was able to predict an IC verb effect in interpretation when using binary choice and continuation prompts, it predicted a reverse verb type effect in

production using these prompts. The same pattern can also be seen in GPT-4o using Yes/No probability prompts. This unexpected asymmetry is even clearer in the performance of the Qwen model, in which the IC verb effect is only found in interpretation. This suggests that LLMs are generally unable to recognize the difference between production and interpretation like humans.

Result #3: LLMs do not align with human behavior in effect magnitude. Although there was a verb type effect and a production/interpretation task effect on LLMs’ responses, the magnitude of these effects is different from humans even in the most similar case: with Yes/No prompting, LLaMA-70B predicted an IC verb effect difference of 19.4% in the production task and 9.2% in the interpretation task. Yet, these two differences were 47.2% and 28.8% respectively in human participants (Rohde and Kehler, 2014). In other words, even in the best scenario, LLMs underestimate the magnitude of the IC verb effect.

Surprisingly, the magnitude of the production-

interpretation asymmetry predicted by LLMs is well-aligned with humans in one case: in human participants, this asymmetry was 16.8% and 35.2% respectively for IC1 and IC2 verbs. With Yes/No prompting, LLaMA-70B predicted a 20.8% and 31% difference respectively for the two verb types. Yet, this seems to be an isolated case. Among the six cases (two LLaMA models in Yes/No and Yes/No probability promptings; GPT-4o in Yes/No and continuation promptings) that exhibit a trend of the production-interpretation asymmetry, this is the only scenario that the magnitude of the asymmetry is similar between LLM and humans. We therefore do not tend to generalize that LLM aligns with human behavior in terms of effect magnitude.

Result #4: Scaling matters. Overall, LLaMA-70B shows better performance than LLaMA-8B across all four prompts. With Binary choice prompts, LLaMA-8B predicted a reverse IC verb effect in the production task and a reverse production-interpretation asymmetry in IC2 verbs, while both effects were captured by LLaMA-70B. With Yes/No and Yes/No probability prompts, although both LLaMA-8B and LLaMA-70B reflected an IC verb effect and production-interpretation asymmetry, the magnitude in human participants was better approximated by LLaMA-70B.

Result #5: Overall, Yes/No prompts align the best with human performance. Among the four promptings used in our study, most models best align with human performance with the Yes/No prompting. This is the only prompting that LLaMA models and GPT-4o capture the IC verb effect and the production-interpretation asymmetry. However, for Qwen, the best performing prompt is the binary choice prompting. This shows that models still differ in their sensitivity to metalinguistic prompts.

4 Discussion

Our study examines the asymmetry between interpretation and interpretation in humans within the context of LLMs, showing that, under specific prompting strategies, certain LLMs can approximate human-like asymmetry. Among the four models we tested, while LLaMA models and GPT-4o performed the best with the Yes/No prompting, Qwen performed the best with the binary choice prompting.

Our findings partially align with [Hu and Levy \(2023\)](#), which suggests that the format of met-

alinguistic prompts influences the behavior of the model. In our study, the four distinct prompting strategies did lead to considerable variation in LLM performance, highlighting the sensitivity of model responses to prompt formulation. The different preferences for prompting strategies between model families also reinforce the need for extensive experimentation across multiple prompting approaches before evaluating a model using metalinguistic prompts.

Yet, different from their findings, we did not see prompting strategies that are more similar to direct probability measures better capture human preferences. Based on [Hu and Levy \(2023\)](#)'s observations, one might expect that metalinguistic prompts that incorporate explicit probability measures would be more reliably approximate human biases. However, our results show otherwise: although models differ in their best-performing prompts, the Yes/No prompting without probability performed best in three out of the four models in our setting. The probability-based format has never yielded superior alignment with human responses. This indicates that probability cues are not universally beneficial in prompting and that their effectiveness may depend on the specific task or model.

Another worth mentioning pattern in our results is that continuation prompting keeps performing the worst in capturing human-like performance, despite being the task originally performed by human participants. One possible explanation is that instruction-tuned LLMs, having been fine-tuned with instruction data or preference optimization objectives, may develop constrained response patterns ([Lin et al., 2023](#)), limiting their flexibility in generating diverse continuations. Additionally, such fine-tuning can reduce conceptual diversity ([Murthy et al., 2024](#)), making LLMs less sensitive to implicit biases in language processing. This suggests that continuation prompting may not be well-suited for probing human-like asymmetries in interpretation and production.

In conclusion, our study reveals that while some instruction-tuned LLMs can approximate human-like asymmetries between production and interpretation under specific conditions, their sensitivity and alignment to human behavior are highly dependent on model scale and the choice of metalinguistic prompting strategy. These findings highlight the nuanced and prompt-contingent nature of evaluating cognitive plausibility in LLMs.

Limitations

Our work has two primary limitations. First, our experiments are conducted on three English-centric LLMs (LLaMA and GPT-4o) and one multilingual-oriented LLM (Qwen). This selection may introduce biases into the models' performance, potentially limiting the generalizability of our findings across other LLMs. Second, our study focuses solely on the asymmetry between ambiguous pronoun resolution and production in English, without exploring cross-linguistic variations. Future research could address these limitations by incorporating a more diverse set of LLMs and broadening the scope of languages analyzed.

References

- Dale J Barr, Roger Levy, Christoph Scheepers, and Harry J Tily. 2013. Random effects structure for confirmatory hypothesis testing: Keep it maximal. *Journal of memory and language*, 68(3):255–278.
- Dorit Ben Shalom and David Poeppel. 2008. Functional anatomic models of language: assembling the pieces. *The Neuroscientist*, 14(1):119–127.
- Marcelle Crinean and Alan Garnham. 2006. Implicit causality, implicit consequentiality and semantic roles. *Language and Cognitive Processes*, 21(5):636–648.
- Vittoria Dentella, Fritz Günther, and Evelina Leivada. 2023. Systematic testing of three language models reveals low language accuracy, absence of response stability, and a yes-response bias. *Proceedings of the National Academy of Sciences*, 120(51):e2309583120.
- Evelyn C Ferstl, Alan Garnham, and Christina Manouildou. 2011. Implicit causality bias in english: A corpus of 300 verbs. *Behavior Research Methods*, 43:124–135.
- Richard Futrell and Kyle Mahowald. 2025. How linguistics learned to stop worrying and love the language models. *Preprint*, arXiv:2501.17047.
- Richard Futrell, Ethan Wilcox, Takashi Morita, Peng Qian, Miguel Ballesteros, and Roger Levy. 2019. Neural language models as psycholinguistic subjects: Representations of syntactic state. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 32–42, Minneapolis, Minnesota. Association for Computational Linguistics.
- Andrew Gelman, Aleks Jakulin, Maria Grazia Pittau, and Yu-Sung Su. 2008. A weakly informative default prior distribution for logistic and other regression models. Aaron Grattafiori, Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Alex Vaughan, Amy Yang, Angela Fan, Anirudh Goyal, Anthony Hartshorn, Aobo Yang, Archi Mitra, Archie Sravankumar, Artem Korenev, Arthur Hinsvark, Arun Rao, Aston Zhang, Aurelien Rodriguez, Austen Gregerson, Ava Spataru, Baptiste Roziere, Bethany Biron, Binh Tang, Bobbie Chern, Charlotte Caucheteux, Chaya Nayak, Chloe Bi, Chris Marra, Chris McConnell, Christian Keller, Christophe Touret, Chunyang Wu, Corinne Wong, Cristian Canton Ferrer, Cyrus Nikolaidis, Damien Al-lonsius, Daniel Song, Danielle Pintz, Danny Livshits, Danny Wyatt, David Esiobu, Dhruv Choudhary, Dhruv Mahajan, Diego Garcia-Olano, Diego Perino, Dieuwke Hupkes, Egor Lakomkin, Ehab AlBadawy, Elina Lobanova, Emily Dinan, Eric Michael Smith, Filip Radenovic, Francisco Guzmán, Frank Zhang, Gabriel Synnaeve, Gabrielle Lee, Georgia Lewis Anderson, Govind Thattai, Graeme Nail, Gregoire Mialon, Guan Pang, Guillem Cucurell, Hailey Nguyen, Hannah Korevaar, Hu Xu, Hugo Touvron, Iliyan Zarov, Imanol Arrieta Ibarra, Isabel Kloumann, Ishan Misra, Ivan Evtimov, Jack Zhang, Jade Copet, Jaewon Lee, Jan Geffert, Jana Vranes, Jason Park, Jay Mahadeokar, Jeet Shah, Jelmer van der Linde, Jennifer Billock, Jenny Hong, Jenya Lee, Jeremy Fu, Jianfeng Chi, Jianyu Huang, Jiawen Liu, Jie Wang, Jiecao Yu, Joanna Bitton, Joe Spisak, Jongsoo Park, Joseph Rocca, Joshua Johnstun, Joshua Saxe, Junteng Jia, Kalyan Vasuden Alwala, Karthik Prasad, Kartikeya Upasani, Kate Plawiak, Ke Li, Kenneth Heafield, Kevin Stone, Khalid El-Arini, Krithika Iyer, Kshitiz Malik, Kuenley Chiu, Kunal Bhalla, Kushal Lakhotia, Lauren Rantala-Yeary, Laurens van der Maaten, Lawrence Chen, Liang Tan, Liz Jenkins, Louis Martin, Lovish Madaan, Lubo Malo, Lukas Blecher, Lukas Landzaat, Luke de Oliveira, Madeline Muzzi, Mahesh Pasupuleti, Mannat Singh, Manohar Paluri, Marcin Kardas, Maria Tsimpoukelli, Mathew Oldham, Mathieu Rita, Maya Pavlova, Melanie Kam-badur, Mike Lewis, Min Si, Mitesh Kumar Singh, Mona Hassan, Naman Goyal, Narjes Torabi, Nikolay Bashlykov, Nikolay Bogoychev, Niladri Chatterji, Ning Zhang, Olivier Duchenne, Onur Çelebi, Patrick Alrassy, Pengchuan Zhang, Pengwei Li, Petar Vasic, Peter Weng, Prajjwal Bhargava, Pratik Dubal, Praveen Krishnan, Punit Singh Koura, Puxin Xu, Qing He, Qingxiao Dong, Ragavan Srinivasan, Raj Ganapathy, Ramon Calderer, Ricardo Silveira Cabral, Robert Stojnic, Roberta Raileanu, Rohan Maheswari, Rohit Girdhar, Rohit Patel, Romain Sauvestre, Ronnie Polidoro, Roshan Sumbaly, Ross Taylor, Ruan Silva, Rui Hou, Rui Wang, Saghar Hosseini, Sahana Chennabasappa, Sanjay Singh, Sean Bell, Seohyun Sonia Kim, Sergey Edunov, Shaoliang Nie, Sharan Narang, Sharath Rapparthi, Sheng Shen, Shengye Wan, Shruti Bhosale, Shun Zhang, Simon Vandenhende, Soumya Batra, Spencer Whitman, Sten Sootla, Stephane Collet, Suchin Gururangan, Sydney Borodinsky, Tamar Herman, Tara Fowler, Tarek Sheasha, Thomas Georgiou, Thomas Scialom, Tobias Speckbacher, Todor Mihaylov, Tong Xiao, Ujjwal

Karn, Vedanuj Goswami, Vibhor Gupta, Vignesh Ramanathan, Viktor Kerkez, Vincent Gonguet, Virginie Do, Vish Vogeti, Vitor Albiero, Vladan Petrovic, Weiwei Chu, Wenhan Xiong, Wenyin Fu, Whitney Meers, Xavier Martinet, Xiaodong Wang, Xiaofang Wang, Xiaoqing Ellen Tan, Xide Xia, Xinfeng Xie, Xuchao Jia, Xuwei Wang, Yaelle Goldschlag, Yashesh Gaur, Yasmine Babaei, Yi Wen, Yiwen Song, Yuchen Zhang, Yue Li, Yuning Mao, Zacharie Delpierre Coudert, Zheng Yan, Zhengxing Chen, Zoe Papakipos, Aaditya Singh, Aayushi Srivastava, Abha Jain, Adam Kelsey, Adam Shajnfeld, Adithya Gangidi, Adolfo Victoria, Ahuva Goldstand, Ajay Menon, Ajay Sharma, Alex Boesenberg, Alexei Baevski, Allie Feinstein, Amanda Kallet, Amit Sangani, Amos Teo, Anam Yunus, Andrei Lupu, Andres Alvarado, Andrew Caples, Andrew Gu, Andrew Ho, Andrew Poulton, Andrew Ryan, Ankit Ramchandani, Annie Dong, Annie Franco, Anuj Goyal, Aparajita Saraf, Arkabandhu Chowdhury, Ashley Gabriel, Ashwin Bharambe, Assaf Eisenman, Azadeh Yazdan, Beau James, Ben Maurer, Benjamin Leonhardi, Bernie Huang, Beth Loyd, Beto De Paola, Bhargavi Paranjape, Bing Liu, Bo Wu, Boyu Ni, Braden Hancock, Bram Wasti, Brandon Spence, Brani Stojkovic, Brian Gamido, Britt Montalvo, Carl Parker, Carly Burton, Catalina Mejia, Ce Liu, Changhan Wang, Changkyu Kim, Chao Zhou, Chester Hu, Ching-Hsiang Chu, Chris Cai, Chris Tindal, Christoph Feichtenhofer, Cynthia Gao, Damon Civin, Dana Beaty, Daniel Kreymer, Daniel Li, David Adkins, David Xu, Davide Testuggine, Delia David, Devi Parikh, Diana Liskovich, Didem Foss, Dingkan Wang, Duc Le, Dustin Holland, Edward Dowling, Eissa Jamil, Elaine Montgomery, Eleonora Presani, Emily Hahn, Emily Wood, Eric-Tuan Le, Erik Brinkman, Esteban Arcaute, Evan Dunbar, Evan Smothers, Fei Sun, Felix Kreuk, Feng Tian, Filippos Kokkinos, Firat Ozgenel, Francesco Caggioni, Frank Kanayet, Frank Seide, Gabriela Medina Florez, Gabriella Schwarz, Gada Badeer, Georgia Swae, Gil Halpern, Grant Herman, Grigory Sizov, Guangyi, Zhang, Guna Lakshminarayanan, Hakan Inan, Hamid Shojanazeri, Han Zou, Hannah Wang, Hanwen Zha, Haroun Habeeb, Harrison Rudolph, Helen Suk, Henry Aspegren, Hunter Goldman, Hongyuan Zhan, Ibrahim Damlaj, Igor Molybog, Igor Tufanov, Ilias Leontiadis, Irina-Elena Veliche, Itai Gat, Jake Weissman, James Geboski, James Kohli, Janice Lam, Japhet Asher, Jean-Baptiste Gaya, Jeff Marcus, Jeff Tang, Jennifer Chan, Jenny Zhen, Jeremy Reizenstein, Jeremy Teboul, Jessica Zhong, Jian Jin, Jingyi Yang, Joe Cummings, Jon Carvill, Jon Shepard, Jonathan McPhie, Jonathan Torres, Josh Ginsburg, Junjie Wang, Kai Wu, Kam Hou U, Karan Saxena, Kartikay Khanelwal, Katayoun Zand, Kathy Matosich, Kaushik Veeraraghavan, Kelly Michelena, Keqian Li, Kiran Jagadeesh, Kun Huang, Kunal Chawla, Kyle Huang, Lailin Chen, Lakshya Garg, Lavender A, Leandro Silva, Lee Bell, Lei Zhang, Liangpeng Guo, Licheng Yu, Liron Moshkovich, Luca Wehrstedt, Madian Khabsa, Manav Avalani, Manish Bhatt, Martynas Mankus, Matan Hasson, Matthew Lennie, Matthias Reso, Maxim Groshev, Maxim Naumov,

Maya Lathi, Meghan Keneally, Miao Liu, Michael L. Seltzer, Michal Valko, Michelle Restrepo, Mihir Patel, Mik Vyatskov, Mikayel Samvelyan, Mike Clark, Mike Macey, Mike Wang, Miquel Jubert Hermoso, Mo Metanat, Mohammad Rastegari, Munish Bansal, Nandhini Santhanam, Natascha Parks, Natasha White, Navyata Bawa, Nayan Singhal, Nick Egebo, Nicolas Usunier, Nikhil Mehta, Nikolay Pavlovich Laptev, Ning Dong, Norman Cheng, Oleg Chernoguz, Olivia Hart, Omkar Salpekar, Ozlem Kalinli, Parkin Kent, Parth Parekh, Paul Saab, Pavan Balaji, Pedro Rittner, Philip Bontrager, Pierre Roux, Piotr Dollar, Polina Zvyagina, Prashant Ratanchandani, Pritish Yuvraj, Qian Liang, Rachad Alao, Rachel Rodriguez, Rafi Ayub, Raghotham Murthy, Raghu Nayani, Rahul Mitra, Rangrabhu Parthasarathy, Raymond Li, Rebekkah Hogan, Robin Battey, Rocky Wang, Russ Howes, Ruty Rinott, Sachin Mehta, Sachin Siby, Sai Jayesh Bondu, Samyak Datta, Sara Chugh, Sara Hunt, Sargun Dhillon, Sasha Sidorov, Satadru Pan, Saurabh Mahajan, Saurabh Verma, Seiji Yamamoto, Sharadh Ramaswamy, Shaun Lindsay, Shaun Lindsay, Sheng Feng, Shenghao Lin, Shengxin Cindy Zha, Shishir Patil, Shiva Shankar, Shuqiang Zhang, Shuqiang Zhang, Sinong Wang, Sneha Agarwal, Soji Sajuyigbe, Soumith Chintala, Stephanie Max, Stephen Chen, Steve Kehoe, Steve Satterfield, Sudarshan Govindaprasad, Sumit Gupta, Summer Deng, Sungmin Cho, Sunny Virk, Suraj Subramanian, Sy Choudhury, Sydney Goldman, Tal Reamad, Tamar Glaser, Tamara Best, Thilo Koehler, Thomas Robinson, Tianhe Li, Tianjun Zhang, Tim Matthews, Timothy Chou, Tzook Shaked, Varun Vontimitta, Victoria Ajayi, Victoria Montanez, Vijai Mohan, Vinay Satish Kumar, Vishal Mangla, Vlad Ionescu, Vlad Poenaru, Vlad Tiberiu Mihalescu, Vladimir Ivanov, Wei Li, Wenchen Wang, Wenwen Jiang, Wes Bouaziz, Will Constable, Xiao Cheng Tang, Xiaojian Wu, Xiaolan Wang, Xilun Wu, Xinbo Gao, Yaniv Kleinman, Yanjun Chen, Ye Hu, Ye Jia, Ye Qi, Yenda Li, Yilin Zhang, Ying Zhang, Yossi Adi, Youngjin Nam, Yu, Wang, Yu Zhao, Yuchen Hao, Yundi Qian, Yunlu Li, Yuzi He, Zach Rait, Zachary DeVito, Zef Rosnbrick, Zhaoduo Wen, Zhenyu Yang, Zhiwei Zhao, and Zhiyu Ma. 2024. [The llama 3 herd of models](#). *Preprint*, arXiv:2407.21783.

Jennifer Hu, Sammy Floyd, Olessia Jouravlev, Evelina Fedorenko, and Edward Gibson. 2023. [A fine-grained comparison of pragmatic language understanding in humans and language models](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 4194–4213, Toronto, Canada. Association for Computational Linguistics.

Jennifer Hu and Roger Levy. 2023. [Prompting is not a substitute for probability measurements in large language models](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 5040–5060, Singapore. Association for Computational Linguistics.

Jennifer Hu, Kyle Mahowald, Gary Lupyman, Anna Ivanova, and Roger Levy. 2024. [Language mod-](#)

- els align with human judgments on key grammatical constructions. *Proceedings of the National Academy of Sciences*, 121(36):e2400917121.
- Tatsuki Kuribayashi, Yohei Oseki, Souhaib Ben Taieb, Kentaro Inui, and Timothy Baldwin. 2025. [Large language models are human-like internally](#). *Preprint*, arXiv:2502.01615.
- Suet-Ying Lam and Heeju Hwang. 2024. Pronoun interpretation is more subject-biased than expected by the bayesian model. *Language, Cognition and Neuroscience*, pages 1–20.
- Suet-Ying Lam, Qingcheng Zeng, Kexun Zhang, Chenyu You, and Rob Voigt. 2023. [Large language models are partially primed in pronoun interpretation](#). In *Findings of the Association for Computational Linguistics: ACL 2023*, pages 9493–9506, Toronto, Canada. Association for Computational Linguistics.
- Bill Yuchen Lin, Abhilasha Ravichander, Ximing Lu, Nouha Dziri, Melanie Sclar, Khyathi Chandu, Chandra Bhagavatula, and Yejin Choi. 2023. [The unlocking spell on base llms: Rethinking alignment via in-context learning](#). *Preprint*, arXiv:2312.01552.
- Laia Mayol. 2018. Asymmetries between interpretation and production in catalan pronouns. *Dialogue Discourse*. 2018 Dec; 9 (2): 1-34.
- Sonia K. Murthy, Tomer Ullman, and Jennifer Hu. 2024. [One fish, two fish, but not the whole sea: Alignment reduces language models’ conceptual diversity](#). *Preprint*, arXiv:2411.04427.
- OpenAI, :, Aaron Hurst, Adam Lerer, Adam P. Goucher, Adam Perelman, Aditya Ramesh, Aidan Clark, AJ Ostrow, Akila Welihinda, Alan Hayes, Alec Radford, Aleksander Mądry, Alex Baker-Whitcomb, Alex Beutel, Alex Borzunov, Alex Carney, Alex Chow, Alex Kirillov, Alex Nichol, Alex Paino, Alex Renzin, Alex Tachard Passos, Alexander Kirillov, Alexi Christakis, Alexis Conneau, Ali Kamali, Allan Jabri, Allison Moyer, Allison Tam, Amadou Crookes, Amin Tootoochian, Amin Tootoochian, Ananya Kumar, Andrea Vallone, Andrej Karpathy, Andrew Braunstein, Andrew Cann, Andrew Codispoti, Andrew Galu, Andrew Kondrich, Andrew Tulloch, Andrey Mishchenko, Angela Baek, Angela Jiang, Antoine Pelisse, Antonia Woodford, Anuj Gosalia, Arka Dhar, Ashley Pantuliano, Avi Nayak, Avital Oliver, Barret Zoph, Behrooz Ghorbani, Ben Leimberger, Ben Rossen, Ben Sokolowsky, Ben Wang, Benjamin Zweig, Beth Hoover, Blake Samic, Bob McGrew, Bobby Spero, Bogo Giertler, Bowen Cheng, Brad Lightcap, Brandon Walkin, Brendan Quinn, Brian Guarraci, Brian Hsu, Bright Kellogg, Brydon Eastman, Camillo Lugaresi, Carroll Wainwright, Cary Bassin, Cary Hudson, Casey Chu, Chad Nelson, Chak Li, Chan Jun Shern, Channing Conger, Charlotte Barette, Chelsea Voss, Chen Ding, Cheng Lu, Chong Zhang, Chris Beaumont, Chris Hallacy, Chris Koch, Christian Gibson, Christina Kim, Christine Choi, Christine McLeavey, Christopher Hesse, Claudia Fischer, Clemens Winter, Coley Czarnecki, Colin Jarvis, Colin Wei, Constantin Koumouzelis, Dane Sherburn, Daniel Kappler, Daniel Levin, Daniel Levy, David Carr, David Farhi, David Mely, David Robinson, David Sasaki, Denny Jin, Dev Valladares, Dimitris Tsipras, Doug Li, Duc Phong Nguyen, Duncan Findlay, Edede Oiwoh, Edmund Wong, Ehsan Asdar, Elizabeth Proehl, Elizabeth Yang, Eric Antonow, Eric Kramer, Eric Peterson, Eric Sigler, Eric Wallace, Eugene Brevdo, Evan Mays, Farzad Khorasani, Felipe Petroski Such, Filippo Raso, Francis Zhang, Fred von Lohmann, Freddie Sulit, Gabriel Goh, Gene Oden, Geoff Salmon, Giulio Starace, Greg Brockman, Hadi Salman, Haiming Bao, Haitang Hu, Hannah Wong, Haoyu Wang, Heather Schmidt, Heather Whitney, Heewoo Jun, Hendrik Kirchner, Henrique Ponde de Oliveira Pinto, Hongyu Ren, Huiwen Chang, Hyung Won Chung, Ian Kivlichan, Ian O’Connell, Ian O’Connell, Ian Osband, Ian Silber, Ian Sohl, Ibrahim Okuyucu, Ika Lan, Ilya Kostrikov, Ilya Sutskever, Ingmar Kanitscheider, Ishaan Gulrajani, Jacob Coxon, Jacob Menick, Jakub Pachocki, James Aung, James Betker, James Crooks, James Lennon, Jamie Kiros, Jan Leike, Jane Park, Jason Kwon, Jason Phang, Jason Teplitz, Jason Wei, Jason Wolfe, Jay Chen, Jeff Harris, Jenia Varavva, Jessica Gan Lee, Jessica Shieh, Ji Lin, Jiahui Yu, Jiayi Weng, Jie Tang, Jieqi Yu, Joanne Jang, Joaquin Quinonero Candela, Joe Beutler, Joe Landers, Joel Parish, Johannes Heidecke, John Schulman, Jonathan Lachman, Jonathan McKay, Jonathan Uesato, Jonathan Ward, Jong Wook Kim, Joost Huizinga, Jordan Sitkin, Jos Kraaijeveld, Josh Gross, Josh Kaplan, Josh Snyder, Joshua Achiam, Joy Jiao, Joyce Lee, Juntang Zhuang, Justyn Harriman, Kai Fricke, Kai Hayashi, Karan Singhal, Katy Shi, Kavin Karthik, Kayla Wood, Kendra Rimbach, Kenny Hsu, Kenny Nguyen, Keren Gu-Lemberg, Kevin Button, Kevin Liu, Kiel Howe, Krithika Muthukumar, Kyle Luther, Lama Ahmad, Larry Kai, Lauren Itow, Lauren Workman, Leher Pathak, Leo Chen, Li Jing, Lia Guy, Liam Fedus, Liang Zhou, Lien Mamitsuka, Lillian Weng, Lindsay McCallum, Lindsey Held, Long Ouyang, Louis Feувrier, Lu Zhang, Lukas Kondraciuk, Lukasz Kaiser, Luke Hewitt, Luke Metz, Lyric Doshi, Mada Aflak, Maddie Simens, Madelaine Boyd, Madeleine Thompson, Marat Dukhan, Mark Chen, Mark Gray, Mark Hudnall, Marvin Zhang, Marwan Aljube, Mateusz Litwin, Matthew Zeng, Max Johnson, Maya Shetty, Mayank Gupta, Meghan Shah, Mehmet Yatbaz, Meng Jia Yang, Mengchao Zhong, Mia Glaese, Mianna Chen, Michael Janner, Michael Lampe, Michael Petrov, Michael Wu, Michele Wang, Michelle Fradin, Michelle Pokrass, Miguel Castro, Miguel Oom Temudo de Castro, Mikhail Pavlov, Miles Brundage, Miles Wang, Minal Khan, Mira Murati, Mo Bavarian, Molly Lin, Murat Yesildal, Nacho Soto, Natalia Gimelshein, Natalie Cone, Natalie Staudacher, Natalie Summers, Natan LaFontaine, Neil Chowdhury, Nick Ryder, Nick Stathas, Nick Turley, Nik Tezak, Niko Felix, Nithanth Kudige, Nitish Keskar, Noah Deutsch, Noel Bundick, Nora Puckett, Ofir Nachum, Ola Okelola, Oleg Boiko, Oleg Murk, Oliver Jaffe, Olivia Watkins, Olivier Godement, Owen Campbell-Moore, Patrick

Chao, Paul McMillan, Pavel Belov, Peng Su, Peter Bak, Peter Bakkum, Peter Deng, Peter Dolan, Peter Hoeschele, Peter Welinder, Phil Tillet, Philip Pronin, Philippe Tillet, Prafulla Dhariwal, Qiming Yuan, Rachel Dias, Rachel Lim, Rahul Arora, Rajan Troll, Randall Lin, Rapha Gontijo Lopes, Raul Puri, Reah Miyara, Reimar Leike, Renaud Gaubert, Reza Zamani, Ricky Wang, Rob Donnelly, Rob Honsby, Rocky Smith, Rohan Sahai, Rohit Ramchandani, Romain Huet, Rory Carmichael, Rowan Zellers, Roy Chen, Ruby Chen, Ruslan Nigmatullin, Ryan Cheu, Saachi Jain, Sam Altman, Sam Schoenholz, Sam Toizer, Samuel Miserendino, Sandhini Agarwal, Sara Culver, Scott Ethersmith, Scott Gray, Sean Grove, Sean Metzger, Shamez Hermani, Shantanu Jain, Shengjia Zhao, Sherwin Wu, Shino Jomoto, Shiron Wu, Shuaiqi, Xia, Sonia Phene, Spencer Papay, Srinivas Narayanan, Steve Coffey, Steve Lee, Stewart Hall, Suchir Balaji, Tal Broda, Tal Stramer, Tao Xu, Tarun Gogineni, Taya Christianson, Ted Sanders, Tejal Patwardhan, Thomas Cunninghamman, Thomas Degry, Thomas Dimson, Thomas Raoux, Thomas Shadwell, Tianhao Zheng, Todd Underwood, Todor Markov, Toki Sherbakov, Tom Rubin, Tom Stasi, Tomer Kaftan, Tristan Heywood, Troy Peterson, Tyce Walters, Tyna Eloundou, Valerie Qi, Veit Moeller, Vinnie Monaco, Vishal Kuo, Vlad Fomenko, Wayne Chang, Weiyi Zheng, Wenda Zhou, Wesam Manassra, Will Sheu, Wojciech Zaremba, Yash Patil, Yilei Qian, Yongjik Kim, Youlong Cheng, Yu Zhang, Yuchen He, Yuchen Zhang, Yujia Jin, Yunxing Dai, and Yury Malkov. 2024. [Gpt-4o system card](#). *Preprint*, arXiv:2410.21276.

Clare Patterson, Petra B Schumacher, Bruno Nicenboim, Johannes Hagen, and Andrew Kehler. 2022. A bayesian approach to german personal and demonstrative pronouns. *Frontiers in psychology*, 12:672927.

Martin J Pickering and Simon Garrod. 2013. An integrated theory of language production and comprehension. *Behavioral and brain sciences*, 36(4):329–347.

Qwen, :, An Yang, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chengyuan Li, Dayiheng Liu, Fei Huang, Haoran Wei, Huan Lin, Jian Yang, Jianhong Tu, Jianwei Zhang, Jianxin Yang, Jiayi Yang, Jingren Zhou, Junyang Lin, Kai Dang, Keming Lu, Keqin Bao, Kexin Yang, Le Yu, Mei Li, Mingfeng Xue, Pei Zhang, Qin Zhu, Rui Men, Runji Lin, Tianhao Li, Tianyi Tang, Tingyu Xia, Xingzhang Ren, Xuancheng Ren, Yang Fan, Yang Su, Yichang Zhang, Yu Wan, Yuqiong Liu, Zeyu Cui, Zhenru Zhang, and Zihan Qiu. 2025. [Qwen2.5 technical report](#). *Preprint*, arXiv:2412.15115.

Hannah Rohde. 2008. *Coherence-driven effects in sentence and discourse processing*. University of California, San Diego.

Hannah Rohde and Andrew Kehler. 2014. Grammatical and information-structural influences on pronoun production. *Language, Cognition and Neuroscience*, 29(8):912–927.

Rosemary J Stevenson, Rosalind A Crawley, and David Kleinman. 1994. Thematic roles, focus and the representation of events. *Language and cognitive processes*, 9(4):519–548.

Alex Warstadt, Alicia Parrish, Haokun Liu, Anhad Mohananey, Wei Peng, Sheng-Fu Wang, and Samuel R Bowman. 2020. Blimp: The benchmark of linguistic minimal pairs for english. *Transactions of the Association for Computational Linguistics*, 8:377–392.

Meilin Zhan, Roger Levy, and Andrew Kehler. 2020. Pronoun interpretation in mandarin chinese follows principles of bayesian inference. *PloS one*, 15(8):e0237012.

A Prompt Template

The following shows the metalinguistic prompt design for LLMs.

System Prompt: You are a helpful assistant.

production-binary-choice-template:

In the following sentence, who is more likely to be the subject of the next sentence? {} or {}? Please ONLY return the name without any explanation or extra words.

Sentence: {}

Answer:

production-yes-no-template:

In the following sentence, judge whether the pronoun of the next sentence will refer to {}. Please ONLY answer with 'Yes' or 'No'.

Sentence: {}

production-continuation-template:

Please reasonably continue the sentence with one of the mentioned characters. You should start a new sentence rather than a clause. Please ONLY return the continuation.

Sentence: {}

interpretation-binary-choice-template:

In the following sentence, who is more likely to be the referent of the pronoun? {} or {}? Please ONLY return the name without any explanation or extra words.

Sentence: {}

Answer:

interpretation-yes-no-template:

In the following sentence, judge whether the pronoun refers to {}. Please ONLY answer with 'Yes' or 'No'.

Sentence: {}

interpretation-continuation-template:

Please reasonably continue the sentence following the pronoun. Please ONLY return the continuation.

Sentence: {}

B Statistical analyses and results

We focus on how the bias of IC verbs (IC1 vs. IC2) and the task (production vs. interpretation) affect the outcomes of LLMs in each type of metalinguistic prompts.

B.1 Data annotation

For binary prompts, we annotated based on the name answered by the LLM. For yes-no prompts, we annotated "yes" as "Subject" .

For continuation prompts, we manually annotated the choice of referent (Subject vs. Object) in the production task and the choice of antecedent (Subject vs. Object) in the ambiguous pronoun resolution task based on the meaning of the generated continuation. For instance, for a sentence like "Nick offended Steve. He decided to apologize and clear the air before things escalated further", we annotated the outcome as "Subject". Ambiguous references (e.g., "Zack divorced Paul. He later moved to a new city to start his life over."), non-sensical continuations (e.g., "Janet wanted Kate. She to join her at the party that night, but Kate had already made other plans."), continuations with a plural antecedent (e.g., "Claire played with Jane. They were building a sandcastle on the beach"), and continuations starting with a repeated pronoun in an interpretation task (e.g., "Claire played with Jane. She... she") were excluded. The distribution of the excluded item has been provided in Table 2 in the main text, repeated below:

Table 2: The distribution of excluded responses in continuation prompting

	Ambiguous	Non-sensical	Plural	Repeated pronoun
LLaMA-8B	54	21	17	0
LLaMA-70B	35	9	0	0
Qwen	17	6	0	0
GPT-4o	22	1	8	5

As can be seen, LLaMA-8B made more ambiguous and non-sensical continuations than all other models. Besides, it also provided continuations with "they", which violates the requirement posed in the prompts. This might suggest that scaling affects LLMs' ability to follow the instructions. More models should be evaluated to test this hypothesis.

B.2 Analyses

LLaMA models For the results of LLaMA models, we ran a mixed-effects Bayesian bernoulli regression model using the R package brms for the binary outcome resulted from continuation, binary, and yes-no prompts (Subject = 1; Object = 0); and a mixed-effects Bayesian linear regression model for the continuous probability outcome from yes-no probability prompt. Note that the dataset used in statistical model for continuation outcome is slightly different from that used in other models, as

some of the responses are excluded, as introduced in Appendix B.1.

Each model was fitted using 4 chains, each with 5000 iterations. The first 1000 were warm-up to calibrate the sampler. This results in 12000 posterior samples. They were all built with fixed predictors of IC verbs (sum-coded: IC1 = 0.5; IC2 = -0.5), task type (sum-coded: interpretation = 0.5; production = -0.5), and their interaction. A maximal random structure justified by design is implemented (Barr et al., 2013). For logistic regression models, we used weakly informative priors, i.e., a Cauchy distribution with a center of 0 and a scale of 2.5 for fixed effects following Gelman et al. (2008), and the default setting of the package for the other parameters. For linear regression models, we used a gaussian distribution with a mean of 0 and a standard deviation of 1 as the weakly informative prior for fixed predictors. When there is an interaction effect, we further ran nested models for pairwise comparison.

The Bayesian statistics framework does not use the p-value. We consider the 95% credible interval (CrI) as the evidence for an effect: if the 95% CrI does not include a zero, i.e., it is all positive or negative, we consider there is evidence for an effect. Below we report the estimate and the 95% CrI for each effect.

Qwen For Qwen, we only ran analysis for continuation prompts, and limited to the effect of the IC verb only. This is because responses of Qwen in the production task is so extreme that no statistical model can be successfully fitted. The settings of the mixed-effects Bayesian bernoulli regression model are the same as those used for LLaMA models.

GPT-4o Like LLaMA models, we ran mixed-effects Bayesian bernoulli regression model for binary outcomes from continuation and yes-no prompts. We did not run analysis for binary prompts because responses of GPT-4o were extremely subject-biased in all conditions. As a result, no statistical model can be fitted. We also ran linear regression model for probability outcomes from yes-no probability prompts. The settings of the statistical models are the same as in analyses for LLaMA models.

B.3 Results

We bold the predictor in which the effect is supported the statistical evidence, i.e., the 95% CrI does not contain a zero.

B.3.1 LLaMA-3.1-8B Model

Binary prompting Table 3 shows that there was an interaction effect between the IC verb type and the task type.

Formula: binary \sim verb * task + (1 + task itemID)			
	Estimate	Est. Error	95% CrI
Intercept	-1.80	0.74	[-3.78, -0.99]
verb	0.64	0.48	[-0.12, 1.80]
task	0.67	0.75	[-0.65, 2.47]
verb:task	2.06	0.99	[0.78, 4.61]

Table 3: Summary of the Bayesian logistic regression model of LLaMA-3.1-8B model, binary choice prompting.

Nested analyses further reveal that the production-interpretation asymmetry is only found in IC1 verbs (Table 4), and the verb type effect is only found in interpretation (Table 5).

Formula: binary \sim verb/task + (1 + task itemID)			
	Estimate	Est. Error	95% CrI
Intercept	-1.84	0.76	[-3.81, -0.99]
verb	0.63	0.47	[-0.14, 1.76]
verbIC1:task	1.76	0.96	[0.55, 4.24]
verbIC2:task	-0.40	0.82	[-2.30, 1.14]

Table 4: Pairwise comparison of the task type effect within IC1 and within IC2 conditions using binary choice prompting in LLaMA-3.1-8B model.

Formula: binary \sim task/verb + (1 + prompt itemID)			
	Estimate	Est. Error	95% CrI
Intercept	-1.80	0.79	[-3.85, -0.99]
task	0.65	0.73	[-0.68, 2.37]
taskProduction:verb	-0.43	0.50	[-1.57, 0.42]
taskInterpretation:verb	1.69	0.89	[0.61, 4.01]

Table 5: Pairwise comparison of the IC verb type effect within the production and within the interpretation task using binary choice prompting in LLaMA-3.1-8B model.

Yes-no prompting As shown in Table 6, there was an interaction effect between the verb type and task type for the responses of the model.

Formula: yes_no \sim verb * task + (1 itemID)			
	Estimate	Est. Error	95% CrI
Intercept	-4.06	0.61	[-5.44, -3.04]
verb	3.12	0.67	[1.97, 4.61]
task	4.31	0.70	[3.12, 5.89]
verb:task	-1.99	0.93	[-4.01, -0.39]

Table 6: Summary of the Bayesian logistic regression model of LLaMA-3.1-8B model, yes-no prompting.

Nested analysis in Table 7 further reveals that the model did give more "yes" (or referring to subject) with IC1 verbs than IC2 verbs (as indicated by the positive intercept of the verb predictor). Also, the production-interpretation asymmetry is found within both IC1 and IC2 verbs, such that the model chose more subjects in interpretation than production.

Formula: yes_no \sim verb/task + (1 itemID)			
	Estimate	Est. Error	95% CrI
Intercept	-4.06	0.62	[-5.50, -3.04]
verb	3.18	0.70	[2.00, 4.76]
verbIC1:task	3.22	0.55	[2.28, 4.43]
verbIC2:task	5.45	1.14	[3.58, 8.04]

Table 7: Pairwise comparison of the task effect within the IC1 and within the IC2 verbs using yes-no choice prompting in LLaMA-3.1-8B model.

Continuation prompting As shown in Table 8, there was an interaction effect between the verb type and the task type for the responses of the model.

Formula: cont \sim verb * task + (1 itemID)			
	Estimate	Est. Error	95% CrI
Intercept	-0.81	0.09	[-0.99, -0.65]
verb	0.36	0.16	[0.05, 0.67]
task	-0.15	0.14	[-0.44, 0.13]
verb:task	1.14	0.30	[0.57, 1.73]

Table 8: Summary of the Bayesian logistic regression model of LLaMA-3.1-8B model, continuation prompting.

Nested analyses in Table 9 show that the verb type effect is only found in interpretation.

Formula: cont ~ task/verb + (1 itemID)			
	Estimate	Est. Error	95% CrI
Intercept	-0.81	0.09	[-0.99, -0.64]
task	-0.15	0.14	[-0.44, 0.13]
taskProduction:verb	-0.22	0.21	[-0.64, 0.18]
taskInterpretation:verb	0.93	0.22	[0.51, 1.37]

Table 9: Pairwise comparison of the IC verb type effect within the production and within the interpretation task using continuation prompting in LLaMA-3.1-8B model.

The pairwise comparison in Table 10 shows that the production-interpretation asymmetry differs in direction between the two verb types: while interpretation is more subject-biased than production in IC1 verbs, production is more subject-biased than interpretation in IC2 verbs.

Formula: cont ~ verb/task + (1 itemID)			
	Estimate	Est. Error	95% CrI
Intercept	-0.81	0.09	[-0.99, -0.65]
verb	0.36	0.15	[0.06, 0.66]
verbIC1:task	0.42	0.20	[0.05, 0.81]
verbIC2:task	-0.73	0.21	[-1.14, -0.32]

Table 10: Pairwise comparison of the task type effect within IC1 and within IC2 verbs using continuation prompting in LLaMA-3.1-8B model.

Yes/No probability prompting As shown in Table 11, the model did generate a higher probability of ‘Yes’ (= subject) following IC1 verbs than IC2 verbs, and in interpretation task than in production task.

Formula: subject_yes_probability ~ verb * task + (1 itemID)			
	Estimate	Est. Error	95% CrI
Intercept	0.28	0.01	[0.27, 0.29]
verb	0.15	0.01	[0.12, 0.18]
task	0.15	0.01	[0.13, 0.17]
verb: task	-0.00	0.02	[-0.03, 0.03]

Table 11: Summary of the Bayesian linear regression model of LLaMA-3.1-8B model, Yes/No probability prompting.

B.3.2 LLaMA-3.3-70B

Binary prompting Table 12 clearly shows that the model only reveals the production-interpretation asymmetry, such that it chose more subject in interpretation than in production. There was no clear evidence for the IC verb type effect.

Formula: binary ~ verb * task + (1 itemID)			
	Estimate	Est. Error	95% CrI
Intercept	14.61	3.61	[9.24, 23.34]
verb	1.36	1.42	[-1.19, 4.45]
task	2.21	0.73	[0.98, 3.84]
verb:task	-1.24	1.07	[-3.44, 0.77]

Table 12: Summary of the Bayesian logistic regression model of LLaMA-3.3-70B model, binary prompting.

Yes-no prompting Table 13 shows that the model is able to capture the IC verb type effect and the production-interpretation asymmetry, such that it responded ‘Yes’ (=subject) more for IC1 verbs than IC2 verbs and the interpretation task than the production task.

Formula: yes_no ~ verb * task + (1 itemID)			
	Estimate	Est. Error	95% CrI
Intercept	2.24	0.25	[1.79, 2.78]
verb	1.46	0.34	[0.83, 2.15]
task	2.63	0.31	[2.05, 3.29]
verb:task	-0.10	0.44	[-0.95, 0.76]

Table 13: Summary of the Bayesian logistic regression model of LLaMA-3.3-70B model, Yes/No prompting.

Continuation prompting Like in binary choice prompting, Table 14 shows that the model only reveals the production-interpretation asymmetry, such that it chose more subject in interpretation than in production. There was no clear evidence for the IC verb type effect.

Formula: cont ~ verb * task + (1 itemID)			
	Estimate	Est. Error	95% CrI
Intercept	-3.66	0.66	[-5.23, -2.68]
verb	-0.61	0.88	[-2.67, 0.86]
task	6.83	1.27	[4.89, 9.82]
verb:task	1.10	1.74	[-1.78, 5.19]

Table 14: Summary of the Bayesian logistic regression model of LLaMA-3.3-70B model, continuation prompting.

Yes/No probability prompting Table 15 shows an interaction effect between verb and task type.

Formula: subject_yes_probability \sim verb * task + (1 itemID)			
	Estimate	Est. Error	95% CrI
Intercept	0.73	0.01	[0.70, 0.75]
verb	0.12	0.02	[0.07, 0.16]
task	0.30	0.01	[0.27, 0.33]
verb:task	-0.05	0.03	[-0.11, -0.00]

Table 15: Summary of the Bayesian linear regression model of LLaMA-3.3-70B model, Yes/No probability prompting.

Pairwise comparisons in Table 16 and 17 further show that the verb type effect can be found in both production and interpretation task, and the asymmetry can be found in both IC1 and IC2 conditions.

Formula: subject_yes_probability \sim verb/task + (1 itemID)			
	Estimate	Est. Error	95% CrI
Intercept	0.73	0.01	[0.71, 0.75]
verb	0.12	0.02	[0.07, 0.16]
verbIC1:task	0.27	0.02	[0.24, 0.31]
verbIC2:task	0.33	0.02	[0.29, 0.37]

Table 16: Pairwise comparison of the verb type effect within the production and the interpretation task using Yes/No probability prompting in LLaMA-3.1-70B model.

Formula: subject_yes_probability \sim task/verb + (1 itemID)			
	Estimate	Est. Error	95% CrI
Intercept	0.73	0.01	[0.70, 0.75]
task	0.30	0.01	[0.27, 0.33]
taskProduction:verb	0.14	0.03	[0.09, 0.19]
taskInterpretation:verb	0.09	0.03	[0.04, 0.14]

Table 17: Pairwise comparison of the task type effect within IC1 and within IC2 verbs using Yes/No probability prompting in LLaMA-3.1-70B model.

B.3.3 QWen

Note that we only analyzed the responses of the continuation prompting in the interpretation task for Qwen, because models cannot be converged in other case. As can be seen below in Table 18, there is an opposite IC verb type effect such that the model referred to more subjects following IC2 verbs than IC1 verbs.

Formula: cont \sim verb + (1 itemID)			
	Estimate	Est. Error	95% CrI
Intercept	-0.60	0.07	[-0.73, -0.47]
verb	-0.37	0.13	[-0.63, -0.11]

Table 18: Summary of the Bayesian logistic regression model of Qwen model, continuation prompting.

B.3.4 GPT-4o

Recall that due to the extreme pattern of GPT-4o with binary promptings, no statistical model could be fitted. Below we only report results of statistical analyses for outputs from Yes/No, Yes/No probability, and continuation promptings only.

Table 19 shows the results of outcomes from Yes/No prompting of GPT-4o. As can be seen, there was a main effect of both verb type and task type, such that IC1 verbs and interpretation task elicited more choices towards subject antecedent. This aligns with the human performance.

Formula: yes_no \sim verb * task + (1 itemID)			
	Estimate	Est. Error	95% CI
Intercept	-0.39	0.15	[-0.69, -0.12]
verb	0.68	0.29	[0.15, 1.28]
task	5.51	0.83	[4.17, 7.42]
verb:task	-0.23	0.46	[-1.15, 0.67]

Table 19: Summary of the Bayesian logistic regression model of GPT-4o, Yes/No prompting.

Table 20 shows the results of outcomes from Yes/No probability prompting of GPT-4o. Given statistical evidence for the interaction effect between verb and task type, we further ran a nested analysis for the effect of the task type within each level of verb, as shown in Table 11.

Formula: subject_yes_probability \sim verb * task + (1 itemID)			
	Estimate	Est. Error	95% CI
Intercept	0.89	0.01	[0.87, 0.90]
verb	-0.01	0.01	[-0.04, 0.01]
task	-0.06	0.01	[-0.08, -0.03]
verb:task	0.13	0.02	[0.08, 0.17]

Table 20: Summary of the Bayesian linear regression model of GPT-4o, Yes/No probability prompting.

As can be seen, there was only a task effect limited to IC2 verbs and in an opposite way to human performance: the interpretation task was even less subject-biased than the production task,

as shown in the negative sign of the estimate of verbIC2:task.

Formula: subject_yes_probability ~ verb/prompt + (1 itemID)			
	Estimate	Est. Error	95% CI
Intercept	0.89	0.01	[0.87, 0.90]
verb	-0.01	0.01	[-0.04, 0.01]
verbIC1:task	0.01	0.02	[-0.03, 0.04]
verbIC2:task	-0.12	0.02	[-0.15, -0.08]

Table 21: Pairwise comparison of the task type effect within IC1 and within IC2 verbs using of GPT-4o, Yes/No probability prompting.

Lastly, Table 22 shows statistical results from continuation promptings of GPT-4o. The main effect of the task type shows that the model captures the production-interpretation asymmetry that the interpretation task was more subject-biased than the production task. However, the negative sign of the verb type effect shows a reversed IC verb effect, such that IC2 verbs were more subject-biased than IC1 verbs.

Formula: cont ~ verb * task + (1 itemID)			
	Estimate	Est. Error	95% CI
Intercept	-1.43	0.30	[-2.12, -0.95]
verb	-3.11	0.63	[-4.57, -2.10]
task	5.85	1.04	[4.26, 8.28]
verb:task	0.10	0.60	[-1.05, 1.33]

Table 22: Summary of the Bayesian linear regression model of GPT-4o, continuation prompting.