

Does ChatGPT have Theory of Mind?

Bart Holterman
Utrecht University

Kees van Deemter
Utrecht University

Abstract

Theory of Mind (ToM) is the ability to understand human thinking and decision-making, an ability that plays a crucial role in social interaction between people, including linguistic communication. This paper investigates to what extent recent Large Language Models in the ChatGPT tradition possess ToM. We posed six well-known problems that address biases in human reasoning and decision making to two versions of ChatGPT and we compared the results under a range of prompting strategies. While the results concerning ChatGPT-3 were somewhat inconclusive, ChatGPT-4 was shown to arrive at the correct answers more often than would be expected based on chance, although correct answers were often arrived at on the basis of false assumptions or invalid reasoning.

1 Introduction

ChatGPT is a large language model (LLM) chatbot based on the OpenAI ChatGPT-3 language model (Aydm and Karaarslan, 2022). A wealth of research articles are starting to chart ChatGPT’s strengths and weaknesses (Aljanabi et al., 2023).

We add to these insights by studying ChatGPT’s ability to comprehend human Theory of Mind (ToM). We chart ChatGPT 3 and 4’s ability to answer a range of well-known ToM problems, focussing on problems that address known human biases (henceforth, “fallacies”). Since ChatGPT is sensitive to subtle differences in the prompt (e.g., Moghaddam and Honey (2023)), we chart how ChatGPT responds to prompts with different levels of detail.

Since we wanted to cover a broad range of ToM abilities, we took a well-known overview of biases in human reasoning and decision making, namely Kahneman (2011), and selected a representative problem from each of its main chapters. Every prompt we gave to ChatGPT was based on a problem described in the book; ChatGPT’s answers

were compared with the correct answers. For example, suppose we ask, about a question Q , “Which answer to Q is most likely to be given by an arbitrary subject: answer A or answer B ?”; then if earlier experiments have shown overwhelmingly that more people chose A than B , then the correct way for ChatGPT to answer our question is to say that A is the most likely answer to Q . This means that, when subjects overwhelmingly tend to give answers to Q that are wrong then the ChatGPT answer we count as *correct* must mark that false answer as the most likely one.

2 Related Work

Theory of mind (ToM) is the ability to reason about mental states, such as beliefs, desires, and intentions, and to understand how mental states feature in everyday explanations and predictions of people’s behavior (Apperly, 2012). A variety of studies examine the presence and extent of ToM in large language models (LLM’s) (Kosinski, 2023; Borji, 2023; Dou, 2023; Brunet-Gouet et al., 2023; Moghaddam and Honey, 2023). These studies tend to employ false-belief tasks such as the “Sally-Anne Test”, a widely used test originally designed for assessing ToM capabilities in infants. This paper summarizes conducting an evaluation of ChatGPT’s understanding of human “fallacies”, looking at a range of mental states such as desires and intentions. Fallacies are a key aspect of ToM, because they show where humans behave in ways that deviate from what might be expected on purely “rational” grounds.

Our work bears some similarities to Binz and Schulz (2023)’s study of ChatGPT’s reasoning and decision making. These authors, however, asked questions of the form “what would you (ChatGPT) choose?”. By contrast, our paper poses questions of the form, “What would a third person be likely to choose”, thus focussing on ChatGPT’s understanding of human tendencies.

3 Experiments

Table 1 provides an overview of our approach to prompt engineering. To account for ChatGPT’s stochastic nature, each prompt was presented nine times ($n=9$), in such a way that it can answer ‘A’ one time and ‘B’ another time the same prompt is posed.

Detailedness	Id	Type of prompt that was used
low	p1	A basic version of the experiment, with minimal details.
medium	p2	Prompts are made more concrete, e.g., by giving people names.
	p3	More detail is added, making the scenario less abstract.
high	p4	Further extraneous details are added.
	p5	Like p4, but with a different choice of details.

Table 1: Our approach to prompt engineering

3.1 Experiments

We briefly describe each of the six problem areas that featured in our experiments. The concrete versions of these problem areas that were used can be found in Appendix A.

3.1.1 The Mental Shotgun

“When we intend to complete one task, we involuntarily compute others at the same time, slowing us down” (Kahneman, 2011). Kahneman and colleagues investigated this hypothesis by having participants listen to spoken word pairs; participants were asked to press a button if and only if the pair rhymes. It turned out that, for word pairs that rhyme, morphological similarity speeds up the time lag, so “vote-note” is recognized as rhyming more quickly than “vote-goat” (Seidenberg and Kahneman, 1979). Our study (see Appendix A.1) presented this experiment to ChatGPT, asking the system for which word pair a human subject would be likely to press the button faster. In prompts p1,...,p4, the same word pairs are presented as the ones used by Kahneman i.e. “vote-note” and “vote-goat”. In prompt p5, different word pairs are presented i.e. “boat-coat” and “boat-moat”.

3.1.2 The Anchoring Effect

“The anchoring effect occurs when people consider a particular value for an unknown quantity before estimating that quantity. What happens is one of the most reliable and robust results of experimental psychology: the estimates stay close to the number

that people considered” (Kahneman, 2011). Our experiment (Appendix A.2) is derived from the original one (Jacowitz and Kahneman, 1995). In prompts p1-4 we used the same experiment as Kahneman, but we used different names and ages. In p5, an entirely different instance of the problem was created.

3.1.3 The Linda Problem

The Linda problem addresses the human tendency to give undue weight to information that is in line with our pre-existing beliefs or expectations (Tversky and Kahneman, 1983). In earlier experiments, participants were presented with a scenario about a person called Linda, about whom they are told, “Linda is very active in equal rights movements”. The subject is asked to choose the most likely option from a set of possibilities. One of the options involves a conjunction of two traits, such as “Linda is a bank teller and a feminist”, while another option only mentions just one of these traits, e.g., “Linda is a bank teller”. Perhaps surprisingly, participants tend to choose the option involving two traits, thus violating some elementary laws of probability. In our experiment (see Appendix A.3), we modified the scenario so it does not involve the name Linda, nor the terms bank teller and feminist.

3.1.4 The Planning Fallacy

Humans frequently “make decisions based on delusional optimism rather than on a rational weight of gains, losses, and probabilities,” (Kahneman, 2011); see also Kahneman and Lovallo (2003). In our experiment (Appendix A.4), we used a kitchen remodeling scenario (p1-4) and a scenario of a project for a new subway line in Amsterdam (p5) as the prompts to pose to ChatGPT.

3.1.5 Relative Wealth

“Magically making a poor person’s portfolio worth a million dollars would be fabulous! Magically making a billionaire’s portfolio worth a million dollars would be agony!” (Kahneman, 2011). In other words, the subjective value, to a given individual, of that individual’s wealth is relative to this individual’s wealth in the recent past (see (Coombs et al., 1970) for extensive discussion). In versions p1-4 of our experiment (see Appendix A.5), the question posed to ChatGPT is whether a third person would think one of the two people was happier with an amount of money than the other, or if they are both equally happy. In version p5, financial wealth is

replaced by physical fitness. In p1-4 we presented the same experiment as described in by Kahneman. In p5, an entirely different instance of the problem was created.

3.1.6 Loss aversion

Human loss aversion is extremely well supported (e.g., Novemsky and Kahneman (2005)). The term describes how, “when directly weighted against each other, losses loom larger than gains” (Kahneman, 2011). This leads to a tendency to make choices that would be irrational when viewed in terms of expected value. In our experiment (see Appendix A.6), a scenario was created with a person presented with a bet with a positive expected value. According to Kahneman’s theory of loss aversion, an individual would decline this bet. We used the same scenario as Kahneman describes, however we changed the amount of money.

4 Results

GPT-3 outputs were obtained using ChatGPT version 3 with its last update on 24 March 2023. Outputs for GPT-4 were obtained using ChatGPT version 4 with its last update on 24 March 2023. We opted to use ChatGPT’s user interface which is also used by the majority of the public. Therefore, the parameter values used in the study are the default values provided in the OpenAI sample code (Motoki et al., 2023). An important parameter is the temperature, with a default value of 0.7 (Motoki et al., 2023). A total of 45 outputs were obtained for each individual problem and each model (n=45). In order to examine the impact of varying prompts, five prompts were used per ToM principle. Both ChatGPT-3 and ChatGPT-4 were presented with each prompt. For each ChatGPT version, every prompt was posed nine times (n=9). The results of each experiment will be presented separately in this section. In total, each model was tested 270 times on its ability to comprehend ToM, using varying experiments and prompts. Out of the 270 total questions, ChatGPT-3 answered 147 correctly, a binomial test returns a p-value of $p > 0.05$ which is in-significant. ChatGPT-4 answered 224 out of the 270 total questions correctly, a binomial test gives a p-value of ($p < 0.0001$), which is significant. To test the effect of the level of detail of the input prompts, five different input prompts were created per experiment. Figure 1 shows the average number of correct answers over all experiments, by model and prompt. ChatGPT-4 answered more

questions correctly than ChatGPT-3, as revealed by a two-proportion Z-Test ($p < 0.0001$). Results per individual experiment are shown in Appendix B.

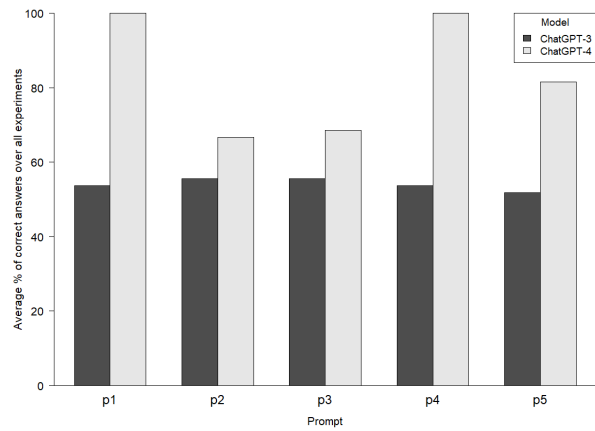


Figure 1: A bar plot showing the average number of correct answers over all experiments, by model and prompt. ChatGPT-3 has a similar correct answer average over all prompts, while ChatGPT-4 shows differences in the averages of correct answers over each prompt.

5 Conclusion

This paper reports on an investigation into ChatGPT’s knowledge and understanding of an important but complicated area of knowledge, namely the nature of human bias in reasoning and decision making. Summing up, ChatGPT-3 did not answer significantly more than half the total number of questions correctly ($p > 0.05$), although it performed somewhat better on the Linda problem, the planning fallacy and loss aversion. However, ChatGPT-4 performed much better, answering around 80% of the total number of questions correctly. The system performs particularly well on the mental shotgun, the Linda problem, the Anchoring effect, and the Planning Fallacy. For unknown reasons that may be explored in future research, ChatGPT-4 performed better with the low and high-detailed prompts compared to the mid-detailed prompts.

Thus, our results could be taken to suggest that ChatGPT does not, but ChatGPT-4 does, have a certain amount of comprehension of human biases in reasoning and decision making, a key facet of ToM. Broadly speaking, these results are in line with earlier findings (Borji, 2023) (Brunet-Gouet et al., 2023), which ascribed a certain amount of ToM to ChatGPT, particularly regarding false-belief tasks (as discussed in section 2). But although these re-

sults hold promise for a future in which Artificial Intelligence might reach a genuine understanding of people, we believe that our results, and earlier research on ChatGPT's ToM as well, (Kosinski, 2023)(Moghaddam and Honey, 2023), might flatter ChatGPT's actual abilities.

One reason is that systems can be "right for the wrong reasons" (McCoy et al., 2019). So although this was not the original focus of our work, we decided to examine the *explanations* (Appendix C) ChatGPT gave for its decisions, in order to test the quality of the systems' reasoning. Because the reasoning is not always clear (and occasionally completely *unclear*, as in the case of the Anchoring Effect, Appendix C2.1 and C2.2), such an analysis is not easy to perform, but in some cases, we detected flaws in the argumentation even in cases where the correct answer was chosen.

For example, in the answers to the Mental Shotgun, ChatGPT3 repeatedly explains its answer offering an incorrect assessment of whether two words rhyme. (Appendix C1.1 and C1.2). An interesting situation was encountered in the analysis of the Loss Aversion problem. In its first response (C6.1), the system provides the correct answer. However, the model incorrectly claims that the expected value of the bet is negative. In the second instance, the model provides an answer that does not take into account the human tendency to be loss averse and is therefore recorded as incorrect. Nonetheless, in contrast to the first instance, the system correctly calculates the expected value.

These observations lead on to a worry, with implications for all experimental evaluations of ChatGPT, about what may be called "parrotting". Despite our best efforts, it is conceivable that, when faced with a ToM problem, ChatGPT bases its answers on discussions of that same problem, or a closely related problem, in its training material. For instance, if experiments demonstrated that ChatGPT can prove Pythagoras's theorem, we wouldn't jump to the conclusion that ChatGPT can invent mathematical proofs in geometry; the system's proof may simply have been found in the texts on which ChatGPT was trained.

The study reported in this paper mitigated this problem by changing the original problems in all the prompts, and additionally varying the prompts very considerably, looking at 5 different versions of each problem. Except perhaps in the case of the Mental Shotgun (where the nature of the problem

make it challenging to vary the topic), especially version p5 of each problem was always completely different from versions of the problem that we were able to find in existing experiments. Yet Fig. 1 shows that performance on p5 was generally in line with performance on p1-4, suggesting that similarity to training material did not influence performance unduly. In future work, these issues should be investigated further, for example by perturbing the information in the prompts more systematically (e.g., by judiciously adding negations, cf. Niven and Kao (2019)) and probing how these perturbations affect ChatGPT's responses.

Far from considering these matters settled, we believe that deep questions are lurking in the background: How "novel" does a problem, in any area, have to be so we can count it as a fair test of ChatGPT's understanding of that area? Suppose, for example, a famous theme like Loss Aversion was discussed in the literature so extensively, in such detail and with so many examples, that a good proportion of ToM problem could be solved by learning from these discussions (e.g., using a state-of-the-art Natural Language Inference system (e.g., Schick and Schütze (2020))). Under these assumptions, the system would tend to perform above chance level on any reasonable test battery for Loss Aversion. If this happened, would this mean that ChatGPT "understands" Loss Aversion? Returning to the title of the paper, what does it mean to "have" Theory of Mind? These questions, which call to mind Searle's Chinese Room experiment (Searle, 1980), are all the more pertinent given that psychologists have learned to be cautious to distinguish between having a capability and actually using that capability. These considerations are all the more relevant because some of the most cogent experimental results underpinning the distinction between having and using a capability come precisely from the study of human Theory of Mind, where researchers have found strong limitations of the circumstances under which people use ToM (see e.g. (Keysar et al., 2003) for elaboration and experiments).

References

- Mohammad Aljanabi, Mohanad Ghazi, Ahmed Hussein Ali, Saad Abas Abed, et al. 2023. Chatgpt: Open possibilities. *Iraqi Journal For Computer Science and Mathematics*, 4(1):62–64.
- Ian A Apperly. 2012. What is “theory of mind”? concepts, cognitive processes and individual differences. *Quarterly Journal of Experimental Psychology*, 65(5):825–839.
- Ömer Aydın and Enis Karaarslan. 2022. Openai chatgpt generated literature review: Digital twin in healthcare. Available at SSRN 4308687.
- Marcel Binz and Eric Schulz. 2023. Using cognitive psychology to understand gpt-3. *Proceedings of the National Academy of Sciences*, 120(6):e2218523120.
- Ali Borji. 2023. A categorical archive of chatgpt failures. *arXiv preprint arXiv:2302.03494*.
- Eric Brunet-Gouet, Nathan Vidal, and Paul Roux. 2023. Do conversational agents have a theory of mind? a single case study of chatgpt with the hinting, false beliefs and false photographs, and strange stories paradigms.
- C.H. Coombs, R.M. Dawes, and A. Tversky. 1970. *Mathematical Psychology: An Elementary Introduction*. Prentice-Hall, New York.
- Zenan Dou. 2023. Exploring gpt-3 model’s capability in passing the sally-anne test a preliminary study in two languages.
- K.E. Jacowitz and D. Kahneman. 1995. Measures of anchoring in estimation tasks. *Personality and Social Psychology Bulletin*, 21:1161–66.
- D. Kahneman and D. Lovallo. 2003. Delusions of success: how optimism undermines executives’ decisions. *Harvard Business Review*, 81:56–63.
- Daniel Kahneman. 2011. *Thinking, fast and slow*. Farrar, Straus and Giroux, New York.
- Boaz Keysar, Shuhong Lin, and Dale J Barr. 2003. Limits on theory of mind use in adults. *Cognition*, 89(1):25–41.
- Michal Kosinski. 2023. Theory of mind may have spontaneously emerged in large language models. *arXiv preprint arXiv:2302.02083*.
- R Thomas McCoy, Ellie Pavlick, and Tal Linzen. 2019. Right for the wrong reasons: Diagnosing syntactic heuristics in natural language inference. *arXiv preprint arXiv:1902.01007*.
- Shima Rahimi Moghaddam and Christopher J Honey. 2023. Boosting theory-of-mind performance in large language models via prompting. *arXiv preprint arXiv:2304.11490*.
- Fabio Motoki, Valdemar Pinho Neto, and Victor Rodrigues. 2023. More human than human: Measuring chatgpt political bias. Available at SSRN 4372349.
- Timothy Niven and Hung-Yu Kao. 2019. Probing neural network comprehension of natural language arguments. *arXiv preprint arXiv:1907.07355*.
- N. Novemsky and D. Kahneman. 2005. The boundaries of loss aversion. *Journal of Marketing Research*, 42:119–28.
- Timo Schick and Hinrich Schütze. 2020. Exploiting cloze questions for few shot text classification and natural language inference. *arXiv preprint arXiv:2001.07676*.
- J. Searle. 1980. Minds, brains and programs. *Behavioral and Brain Sciences*, 3:417–57.
- M.S. Seidenberg and D. Kahneman. 1979. Orthographic effects on rhyme monitoring. *Journal of Experimental Psychology: Human Learning and Memory*, 5:546–54.
- A. Tversky and D. Kahneman. 1983. Extensional versus intuitive reasoning: the conjunction fallacy in probability judgment. *Psychological Review*, 90:293–315.

A Appendix: Prompts per ToM problem

A.1 Prompts for the Mental Shotgun problem

Detail	Id	Prompt
low	p1	What will a person recognize as being a rhyme faster? A: A word that ends with “oat” and a word that ends with “ote” or B: two words that end in “ote”.
medium	p2	A person is being asked to listen to two pairs of words. Option A: a word that ends with “oat” and a word that ends with “ote” or option B: two words that end with “ote”. Which pair will the person recognize faster as rhyming?
	p3	Jan is asked to listen to two pairs of words. While listening to the words Jan has to press a button when he thinks the pair of words rhyme. Pair A: consists of one word that ends with “ote” and one word that ends with “oat”. Pair B: consists of two words that both end in “ote”. For which pair will Jan press the button faster?
high	p4	Jan is asked to listen to pairs of words. While listening to the words Jan must press a button when he thinks the pair of words rhyme. Pair A: consists of the words “note” and “goat”. Pair B: consists of the words “note” and “vote”. For which pair will Jan press the button faster?
	p5	Jan is asked to listen to pairs of words. While listening to the words Jan must press a button when he thinks the pair of words rhyme. Pair A: consists of the words, “boat” and “mote”. Pair B: consists of the words, “boat” and “coat”. For which pair will Jan press the button faster?

A.2 Prompts for the Anchoring effect problem

Detail	Id	Prompt
low	p1	One person is asked if van Gogh died before or after he was age 105. Another person is asked if van Gogh died before or after he was age 55. They are both asked to guess van Gogh’s age. Which option is most likely? Option A: The first person guesses a higher number than the second person Option B: The second person guesses a higher number than the first person.
medium	p2	Two persons are participating in a questionnaire, the first person is asked if Vincent van Gogh died before or after he was age 105. The second person is asked if Vincent van Gogh died before or after he was age 55. The next question in the questionnaire is to guess the age at which Vincent van Gogh died. Which option is most likely? Option A: The first person guesses a higher number than the second person Option B: The second person guesses a higher number than the first person.
	p3	Two persons are participating in a questionnaire, the first person, called John, is asked if Vincent van Gogh died before or after he was age 105. The second person, called Rick, is asked if Vincent van Gogh died before or after he was age 55. The next question in the questionnaire is to guess the age at which Vincent van Gogh died. Which option is most likely? Option A: John guesses a higher number than Rick Option B: Rick guesses a higher number than John.
high	p4	Two persons are participating in a questionnaire, the first person’s name is John, John is a 27-year-old Graphic designer from Pittsburgh, he is asked if Vincent van Gogh died before or after he was age 105. The second person’s name is Rick, Rick is a 45-year-old lawyer from Seattle, he is asked if Vincent van Gogh died before or after he was age 55. The next question in the questionnaire is to guess the age at which Vincent van Gogh died. Which option is most likely? Option A: John guesses a higher number than Rick Option B: Rick guesses a higher number than John.
	p5	There are two criminals, their names are Paul and John. They have each individually committed tax fraud with which they have evaded 3 million dollars. For a crime like this, a normal sentence would be 5 years in prison. Paul and John have two different prosecutors. Paul’s prosecutor demands 7 years prison. John’s prosecutor demands 12 years in prison. Note, they have done the same crime. They both have the same judge. Which option has the highest probability? Option A: One of them will get a higher sentence. Option B: Paul and John will get the same sentence.

A.3 Prompts for the Linda Problem

Detail	Id	Prompt
low	p1	Robin is a person, who loves kids. A third person does not know Robin's employment or gender. Option A: Robin is a teacher and a woman or Option B: Robin is a teacher. Which one of these options is more likely to be chosen by the third person?
medium	p2	Robin is a person who loves kids, likes to play board games, and is very active in equal rights movements. Furthermore, a third person has no knowledge of Robin's employment or gender. Option A: Robin is a teacher and a woman or Option B: Robin is a teacher. Which one of these options is more likely to be chosen by the third person?
	p3	Robin is a person from the United Kingdom who loves kids, Robin is smart and likes sharing knowledge with others. Robin likes to play board games every Tuesday at a board game club. Additionally, Robin is very active in equal rights movements going to almost every protest in the nation (United Kingdom). Furthermore, Rick is a third person who has no knowledge of Robin's employment or gender. Option A: Robin is a teacher and a woman or Option B: Robin is a teacher. Which one of these options is more likely to be chosen by Rick?
high	p4	Robin is a person from the United Kingdom who has a deep love for children and enjoys spending time with them. In their free time, they volunteer at a local children's charity, where they organize and run educational activities for underprivileged children. Robin is smart and likes sharing knowledge with others. Robin likes to play board games every Tuesday at a board game club. Additionally, Robin is very active in equal rights movements going to almost every protest in the nation (United Kingdom). Furthermore, Rick is a third person who has no knowledge of Robin's employment or gender. Option A: Robin is a teacher and a woman or Option B: Robin is a teacher. Which one of these options is more likely to be chosen by Rick?
	p5	Jan is a man that has been all over the world, he has lived in 4 different countries, and visited 50 more countries. All this traveling is needed for his work. Furthermore, Jan has a wife and kids and loves to do crossword puzzles. Because Jan has lived in so many countries, he can speak 5 languages fluently. Robert doesn't know Jan's employment. Option A: Jan is a pilot and likes learning new languages. Option B: Jan is a pilot. Which one of these options is more likely to be chosen by Robert?

A.4 Prompts for the Planning Fallacy

Detail	Id	Prompt
low	p1	Someone wants to remodel their kitchen, he estimates it will be around a certain amount of money. Option A: the actual price will be around that amount of money. Option B: The actual amount of money will be higher. Which option is more likely?
medium	p2	Rick is an average American, he wants to remodel his kitchen, Rick estimates the price he thinks the remodeling will cost. Option A: the actual price will be around that price. Option B: The actual price will be higher. Which option is more likely?
	p3	Rick is an average American, he wants to remodel his kitchen by putting in a new counter and refrigerator, Rick estimates the price he thinks the remodeling will cost. Option A: the actual price will be around that price. Option B: The actual price will be higher. Which option is more likely?
high	p4	Rick, an average American homeowner, has plans to renovate his kitchen by installing a new countertop and refrigerator. After conducting some research, he comes up with an estimated budget for the remodeling project. Option A: the actual price will be around that price. Option B: The actual price will be higher. Which option is more likely?
	p5	The city of Amsterdam wants to build a new metro line that goes from east to west Amsterdam it is estimated to cost 20 million euros and will be done in 2025. Option A: the actual price will be around that price. Option B: The actual price will be higher. Which option is more likely?

A.5 Prompts for the Relative Wealth problem

Detail	Id	Prompt
low	p1	Two men have the same amount of money. The first man lost money while the second man gained money, but now they have the same amount. Option A: A third person will likely think one of the two men is happier with his amount of money. Option B: A third person will likely think both men are equally happy with their amount of money
medium	p2	Yesterday Rick gained 4 million dollars and Emma lost 4 million dollars, today they both have 5 million. A third person is asked who is happier with the amount of money they have, what will he likely say? Option A: one of the two is happier with the amount of money they have. Option B: They are both equally happy with the amount of money they have.
	p3	Rick as of yesterday had 1 million dollars, Emma on the other hand had 9 million dollars as of yesterday, due to volatility in the market they both have the same amount of money, namely 5 million dollars. Tim hears about this and is asked who is happier with the amount of money they have. What will Tim likely say? A: One of the two is happier with the amount of money they have. Option B: They are both equally happy with the amount of money they have
high	p4	Rick as of yesterday had 1 million dollars, Emma on the other hand had 9 million dollars as of yesterday. Emma lost 4 million dollars in a day while Rick gained 4 million dollars due to volatility in the market. This means they both have the same amount of money, namely 5 million dollars. Tim hears about this and is asked who is happier with the amount of money they have. Which option will Tim likely choose? A: one of the two is happier with the amount of money they have. Option B: They are both equally happy with the amount of money they have.
	p5	Devon and Marc regularly go to the gym to get stronger. They both bench press 80kg as a one rep max as of today. One month ago, Devon bench pressed 70kg while Marc bench pressed 90kg a month ago. Tim is a friend of theirs and is asked who is likely to be happier with their one rep max of 80kg today. Option A: One of them is happier with their one rep max of 80kg. Option B: They are both equally happy with their one rep max of 80kg.

A.6 Prompts for the Loss Aversion problem

Details	Id	Prompt
low	p1	A person is presented with a choice 50% chance to lose \$100 and a 50% chance to win \$125. What is more likely? Option A: They take the bet. Option B: They do not take the bet.
medium	p2	An average American is presented with a bet. The bet is a 50% chance to lose \$100 and a 50% chance to win \$125. Option A: They take the bet. Option B: They do not take the bet. Which option is most likely to be chosen by the average person?
	p3	An average American is presented with a bet, the bet involves flipping a coin, if it comes up heads lose \$100, but if it comes up tails win \$125. Option A: The average American takes the bet. Option B: The average American does not take the bet. Which option is most likely to be chosen by the average American?
high	p4	Julia is a 42-year-old accountant from New York City. One day, her friend offers her a bet. The bet involves flipping a coin, where if it comes up heads, she will lose \$100, but if it comes up tails, she will win \$125. She is given the option to take the bet or not. Option A: Julia takes the bet. Option B: Julia does not take the bet. Which option is most likely to be chosen by Julia?
	p5	Tom is a 28-year-old college student from Los Angeles. One day, his roommate offers him a bet. The bet involves flipping a coin, where if it comes up heads, he will lose \$200, but if it comes up tails, he will win \$250. He is given the option to take the bet or not. Option A: Tom takes the bet. Option B: Tom does not take the bet. Which option is most likely to be chosen by Tom?

B Appendix: Results per prompt, per ToM problem

B.1 Results for the Mental Shotgun problem

	Prompt	Correct (%)	p -value	Result
ChatGPT-3	p1	56	1.0000	Inconclusive
	p2	0	0.0039	Wrong
	p3	0	0.0039	Wrong
	p4	0	0.0039	Wrong
	p5	11	0.0391	Wrong
ChatGPT-4	p1	100	0.0039	Right
	p2	100	0.0039	Right
	p3	100	0.0039	Right
	p4	100	0.0039	Right
	p5	100	0.0039	Right

Table 2: Shows, for each model, for each prompt, the percentage of correct answers, the p -value from a binomial test, and the result that can be concluded given the p -value.

B.2 Results for the Anchoring effect problem

	Prompt	Correct (%)	p -value	Result
ChatGPT-3	p1	56	1.0000	Inconclusive
	p2	22	0.1797	Inconclusive
	p3	44	1.0000	Inconclusive
	p4	33	0.5078	Inconclusive
	p5	0	0.0039	Wrong
ChatGPT-4	p1	100	0.0039	Right
	p2	100	0.0039	Right
	p3	100	0.0039	Right
	p4	100	0.0039	Right
	p5	100	0.0039	Right

Table 3: Shows, for each model, for each prompt, the percentage of correct answers, the p -value from a binomial test, and the result that can be concluded given the p -value.

B.3 Results for the Linda Problem

	Prompt	Correct (%)	p -value	Result
ChatGPT-3	p1	11	0.0391	Wrong
	p2	100	0.0039	Right
	p3	89	0.0391	Right
	p4	100	0.0039	Right
	p5	33	0.5078	Inconclusive
ChatGPT-4	p1	100	0.0039	Right
	p2	100	0.0039	Right
	p3	67	0.5078	Inconclusive
	p4	100	0.0039	Right
	p5	0	0.0039	Wrong

Table 4: Shows, for each model, for each prompt, the percentage of correct answers, the p -value from a binomial test, and the result that can be concluded given the p -value.

B.4 Results for the Planning Fallacy

	Prompt	Correct (%)	p -value	Result
ChatGPT-3	p1	100	0.0039	Right
	p2	56	1.0000	Inconclusive
	p3	100	0.0039	Right
	p4	67	0.5078	Inconclusive
	p5	78	0.1797	Inconclusive
ChatGPT-4	p1	100	0.0039	Right
	p2	100	0.0039	Right
	p3	100	0.0039	Right
	p4	100	0.0039	Right
	p5	100	0.0039	Right

Table 5: Shows, for each model, for each prompt, the percentage of correct answers, the p -value from a binomial test, and the result that can be concluded given the p -value.

B.5 Results for the Relative Wealth problem

	Prompt	Correct (%)	<i>p</i> -value	Result
ChatGPT-3	p1	0	0.0039	Wrong
	p2	67	0.5078	Inconclusive
	p3	0	0.0039	Wrong
	p4	22	0.1797	Inconclusive
	p5	89	0.0391	Right
ChatGPT-4	p1	100	0.0039	Right
	p2	0	0.0039	Wrong
	p3	22	0.1797	Inconclusive
	p4	100	0.0039	Right
	p5	100	0.0039	Right

Table 6: Shows, for each model, for each prompt, the percentage of correct answers, the *p*-value from a binomial test, and the result that can be concluded given the *p*-value.

B.6 Results for the Loss Aversion problem

	Prompt	Correct (%)	<i>p</i> -value	Result
ChatGPT-3	p1	100	0.0039	Right
	p2	100	0.0039	Right
	p3	100	0.0039	Right
	p4	100	0.0039	Right
	p5	100	0.0039	Right
ChatGPT-4	p1	100	0.0039	Right
	p2	0	0.0039	Wrong
	p3	22	0.1797	Inconclusive
	p4	100	0.0039	Right
	p5	89	0.0391	Right


Table 7: Shows, for each model, for each prompt, the percentage of correct answers, the *p*-value from a binomial test, and the result that can be concluded given the *p*-value.

C Appendix: Examples of responses of ChatGPT

C.1 Responses for the Mental Shotgun problem

C.1.1

BA Jan is asked to listen to pairs of words. While listening to the words Jan must press a button when he thinks the pair of words rhyme. Pair A: consists of the words "boat" and "mote". Pair B: consists of the words, "boat" and "coat". For which pair will Jan press the button faster?


 Jan is likely to press the button faster for Pair B, which consists of the words "boat" and "coat".

This is because "boat" and "coat" have a higher degree of phonetic similarity and share a common vowel sound (/ou/), which makes them more likely to be perceived as rhyming compared to "boat" and "mote" which have different vowel sounds. The greater the phonetic similarity between the words, the easier it is for the brain to identify a rhyme and respond quickly.

Therefore, it is expected that Jan will press the button faster for Pair B ("boat" and "coat") compared to Pair A ("boat" and "mote").

C.1.2


BA Jan is asked to listen to pairs of words. While listening to the words Jan must press a button when he thinks the pair of words rhyme. Pair A: consists of the words "note" and "goat". Pair B: consists of the words "note" and "vote". For which pair will Jan press the button faster?

 Jan is likely to press the button faster for pair A, which consists of the words "note" and "goat", because these two words rhyme. On the other hand, pair B, which consists of the words "note" and "vote", do not rhyme as they have different vowel sounds in their endings. Rhyming words have the same ending sound, and since "note" and "goat" share the "oat" sound, Jan is likely to recognize the rhyme in pair A more quickly than in pair B. Therefore, he is expected to press the button faster for pair A.

C.2 Responses for the Anchoring Effect problem


C.2.1

BA Two persons are participating in a questionnaire, the first person, called John, is asked if Vincent van Gogh died before or after he was age 105. The second person, called Rick, is asked if Vincent van Gogh died before or after he was age 55. The next question in the questionnaire is to guess the age at which Vincent van Gogh died. Which option is most likely? Option A: John guesses a higher number than Rick Option B: Rick guesses a higher number than John.

 Option A is more likely. If John thinks Vincent van Gogh died after age 105 and Rick thinks Vincent van Gogh died after age 55, then John's guess may be much higher than Rick's. For example, Vincent van Gogh died at age 37, so John might guess 110 while Rick guesses 60. This would result in John's guess being higher than Rick's. However, it's also possible that both John and Rick guess too high or too low, so the correct answer is not certain based on the information given.

C.2.2


BA Two persons are participating in a questionnaire, the first person, called John, is asked if Vincent van Gogh died before or after he was age 105. The second person, called Rick, is asked if Vincent van Gogh died before or after he was age 55. The next question in the questionnaire is to guess the age at which Vincent van Gogh died. Which option is most likely? Option A: John guesses a higher number than Rick Option B: Rick guesses a higher number than John.

 Option B is most likely. Since Rick was asked if Vincent van Gogh died before or after he was 55 years old, it is possible that Rick may have some knowledge or interest in Vincent van Gogh's life and death. On the other hand, John was asked if Vincent van Gogh died before or after he was 105 years old, which seems like an unusual age to ask about and may indicate that John is not as familiar with Vincent van Gogh's life and death. Therefore, Rick is more likely to guess a higher age for Vincent van Gogh's death than John.

C.3 Responses for the Linda problem

C.3.1


BA Robin is a person, who loves kids. A third person does not know Robins employment or gender. Option A: Robin is a teacher and a woman or Option B: Robin is a teacher. Which one of these options is more likely to be chosen by the third person?

 Based solely on the information provided, it is more likely that the third person would choose Option A: "Robin is a teacher and a woman".

This is because the statement "Robin loves kids" is often associated with female gender stereotypes, particularly in the context of teaching or childcare professions. Additionally, the inclusion of the information that Robin is a teacher reinforces this stereotype, as teaching is a profession that is historically dominated by women.

C.3.2


BA Robin is a person, who loves kids. A third person does not know Robins employment or gender. Option A: Robin is a teacher and a woman or Option B: Robin is a teacher. Which one of these options is more likely to be chosen by the third person?

 Based solely on the information provided, it is equally likely that the third person will choose either Option A: "Robin is a teacher and a woman" or Option B: "Robin is a teacher." Without any additional context or clues about Robin's gender or profession, the third person's choice is purely a matter of personal assumption or preference.

C.4 Responses for the Planning Fallacy


C.4.1

BA Rick is an average American, he wants to remodel his kitchen, Rick makes an estimation of the price he thinks the remodelling will cost. Option A: the actual price will be around that price. Option B: The actual price will be higher. Which option is more likely?

 Option B is more likely. According to studies, people tend to underestimate the cost of home improvement projects. This is because they often overlook additional expenses such as labor costs, unforeseen issues, and materials. Therefore, it is more likely that the actual cost of remodeling Rick's kitchen will be higher than his initial estimation.

C.4.2


BA Rick, an average American homeowner, has plans to renovate his kitchen by installing a new countertop and refrigerator. After conducting some research, he comes up with an estimated budget for the remodeling project. Option A: the actual price will be around that price. Option B: The actual price will be higher. Which option is more likely?

 Option A, the actual price being around the estimated budget, is more likely if Rick's research and budgeting efforts accurately considered all the necessary expenses for the renovation. However, unexpected expenses or issues can arise during the renovation process, which can increase the cost beyond what was initially budgeted.

C.5 Responses for the Relative Wealth problem

C.5.1


BA Rick as of yesterday had 1 million dollars, Emma on the other hand had 9 million dollars as of yesterday, due to volatility in the market they both have the same amount of money, namely 5 million dollars. Tim hears about this and is asked who is happier with the amount of money they have. What will Tim likely say? A: One of the two is happier with the amount of money they have. Option B: They are both equally happy with the amount of money they have. You must choose one option

 B: They are both equally happy with the amount of money they have.

It is difficult to predict individual happiness based on financial gains or losses, as happiness is subjective and influenced by various factors. However, since both Rick and Emma now have the same amount of money, it can be assumed that they are equally happy with their current financial situation, at least in terms of the money they have.

C.5.2


BA Rick as of yesterday had 1 million dollars, Emma on the other hand had 9 million dollars as of yesterday. Emma lost 4 million dollars in a day while Rick gained 4 million dollars due to volatility in the market. This means they both have the same amount of money, namely 5 million dollars. Tim hears about this and is asked who is happier with the amount of money they have. Which option will Tim likely choose? A: one of the two is happier with the amount of money they have. Option B: They are both equally happy with the amount of money they have. You must choose one option

 Tim will likely choose Option A: one of the two is happier with the amount of money they have. This is because Rick gained 4 million dollars, while Emma lost 4 million dollars. Even though they both have the same amount of money now, their recent experiences with the market may have a significant impact on their happiness. Rick might be happier due to his recent gain, while Emma might be less happy due to her recent loss.

C.6 Responses for the Loss Aversion problem


C.6.1

BA Tom is a 28-year-old college student from Los Angeles. One day, his roommate offers him a bet. The bet involves flipping a coin, where if it comes up heads, he will lose \$200, but if it comes up tails, he will win \$250. He is given the option to take the bet or not. Option A: Tom takes the bet. Option B: Tom does not take the bet. Which option is most likely to be chosen by Tom?

 Tom is most likely to not take the bet (Option B). The potential loss of \$200 is greater than the potential gain of \$250, so the expected value of the bet is negative. In other words, on average, Tom would lose money by taking the bet. Therefore, it is rational for him to decline the bet.

C.6.2

BA Julia is a 42-year-old accountant from New York City. One day, her friend offers her a bet. The bet involves flipping a coin, where if it comes up heads, she will lose \$100, but if it comes up tails, she will win \$125. She is given the option to take the bet or not. Option A: Julia takes the bet. Option B: Julia does not take the bet. Which option is most likely to be chosen by Julia?

 From a financial perspective, the expected value of taking the bet is positive. The expected value is calculated by multiplying the probability of each outcome by its respective payoff and then summing the results. In this case, the expected value of taking the bet is:

$$\text{Expected value} = (0.5 * -\$100) + (0.5 * \$125) = \$12.50$$

This means that, on average, Julia stands to gain \$12.50 if she takes the bet. Even though there is a possibility of losing \$100, the potential gain of \$125 outweighs that risk, making it a financially rational decision for Julia to take the bet. Therefore, Julia is most likely to take the bet.