

PHAnToM: Personality Has An Effect on Theory-of-Mind Reasoning in Large Language Models

Anonymous ACL submission

Abstract

Recent advances in large language models (LLMs) demonstrate that their capabilities are comparable, or even superior, to humans in many tasks in natural language processing. Despite this progress, LLMs are still inadequate at social-cognitive reasoning, which humans are naturally good at. Drawing inspiration from psychological research on the links between certain personality traits and Theory-of-Mind (ToM) reasoning, and from prompt engineering research on the hyper-sensitivity of prompts in affecting LLMs capabilities, this study investigates how inducing personalities in LLMs using prompts affects their ToM reasoning capabilities. Our findings show that certain induced personalities can significantly affect the LLMs’ reasoning capabilities in three different ToM tasks. In particular, traits from the Dark Triad have a larger variable effect on LLMs like GPT-3.5, Llama 2, and Mistral across the different ToM tasks. We find that LLMs that exhibit a higher variance across personality prompts in ToM also tends to be more controllable in personality tests: personality traits in LLMs like GPT-3.5, Llama 2 and Mistral can be controllably adjusted through our personality prompts. In today’s landscape where role-play is a common strategy when using LLMs, our research highlights the need for caution, as models that adopt specific personas with personalities potentially also alter their reasoning abilities in an unexpected manner.

1 Introduction

Despite the impressive capabilities of large language models (LLMs) in a variety of natural language processing (NLP) tasks (Lyu et al., 2023; Bai et al., 2023; Bang et al., 2023), LLMs are reported to have poor social-cognitive reasoning abilities (Farha et al., 2022; Pérez-Almendros et al., 2022). One social-cognitive reasoning task that is of great interest is the Theory-of-Mind (ToM) task (Kosinski, 2023; Premack and Woodruff, 1978). ToM is

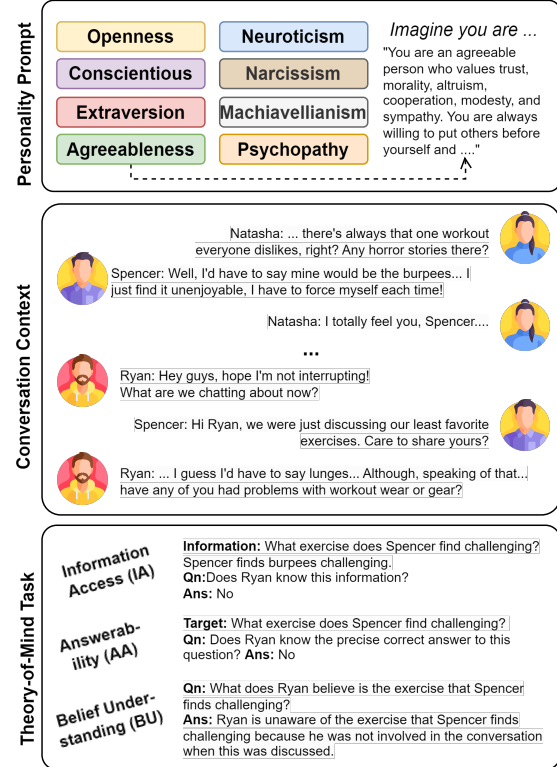


Figure 1: Overview of PHAnToM. Our work investigates how eight different personality prompts (Big Five OCEAN and Dark Triad) affects LLMs’ ability to perform three theory-of-mind reasoning tasks (Information Access (IA), Answerability (AA), and Belief Understanding (BU)).

traditionally studied in the domain of human development and defined as the ability to attribute mental states (such as beliefs, intentions, thoughts, and emotions) to oneself and other agents (Gallese and Sinigaglia, 2011; Wimmer and Perner, 1983). Recent research shows LLMs display some ToM abilities. However, they are still inferior to humans (Kim et al., 2023; Ma et al., 2023a; Shapira et al., 2023). This discrepancy in ToM performances between LLMs and humans motivates our current study to investigate strategies that affect ToM abili-

ties in LLMs.

Inspired by the psychology literature that personality traits influence ToM abilities in humans (Nettle and Liddle, 2008; Stellwagen and Kerig, 2013), and NLP research that induce personality in LLMs through the use of role-play prompting (Safdari et al., 2023; Jiang et al., 2022), this paper aims to investigate how different personality traits affect ToM abilities in LLMs. At the moment, personality prompting and ToM abilities have been studied separately in the domain of LLMs. To the best of our knowledge, we are the first study to combine the two research topics and offers a systematic analysis of the relationship between personality prompting and ToM reasoning abilities in LLMs. In the current landscape where role-play is prevalent when using LLMs, our study underscores the importance of evaluating the personalities LLMs may adopt, as varied personality cues can influence their reasoning abilities.

Our contributions can be summarized as follows:

OUR CONTRIBUTIONS

- We provide novel evidence that **Personality Has An Effect on Theory-of-Mind (PHAnToM)** reasoning in LLMs. In particular, inducing traits from the Dark Triad have a larger varied effect than the Big Five OCEAN on ToM performances across models and tasks.
- We investigate the relationship between personality traits and ToM in the context of psychology and LLMs. We find that LLMs that exhibit a higher variance across personality prompts in ToM also tends to be more controllable in personality tests.
- We contextualize our observations about ToM abilities in LLMs against theories in the psychology literature about humans.

Our code, prompts, and LLM responses are publicly available to encourage more research in this space.¹

2 Related Work

2.1 Sensitivity of LLMs to Prompts

Multiple research has shown the brittleness of LLMs to the input prompts. Zero shot Chain-of-Thought (incorporating one-line in prompts, like “First,” or “Let’s think step by step”) (Kojima et al., 2022; Bsharat et al., 2023) has empirically allowed LLMs to become stronger reasoners, especially for arithmetic tasks. In other works, strategies like role-play (including a description of someone the

LLM should embody) (Kong et al., 2023) or threats (reminding the LLM they would be penalized if they answer wrongly, or that the users’ life matters gravely on this answer) (Bsharat et al., 2023) have also demonstrated effectiveness in improving LLM performance. Sclar et al. (2023) find that small prompt variations often yield large performance differences. Wu et al. (2023) showed that with Instruction Fine-tuning, LLMs can distinguish instruction with context and focus more on instructions. They further show that instruction fine-tuning encourages self-attention heads to encode more word-word relations related to instruction verbs. Gupta et al. (2024) found that LLM’s reasoning abilities can be affected by persona prompts across different socio-demographic groups (race, gender, religion, disability, and political affiliation). Encouraged by these findings, we were inspired to explore the effects of personality role-play on socio-cognitive reasoning in LLMs.

2.2 Personalities in LLMs

Personality refers to the enduring and stable characteristic patterns of cognitions, feelings, and behaviors, that are generally consistent across situations (Allport, 1937). The five-factor model (or Big Five) of personality (John et al., 2008) is one of the primary theoretical frameworks to analyze individual differences consisting of five subscales namely openness, conscientiousness, extraversion, agreeableness, and neuroticism traits (OCEAN) (McCrae and John, 1992). Psychometric tests such as the International Personality Item Pool (IPIP-NEO) (Goldberg et al., 1999), and the Big Five Inventory (BFI) (John et al., 1999) are commonly used to measure these traits in humans.

Recently, Jiang et al. (2022); Safdari et al. (2023); Lu et al. (2023) administered these psychometric tests on LLMs under specific prompting configurations and they found that it is possible to obtain reliable and valid personality measurements with LLMs. Furthermore, by introducing role-play prompts, they demonstrated the adaptability of LLMs, where personalities can be shaped along desired dimensions to simulate specific human personality profiles. These results could be explained by psycholinguistic studies that showed certain expressed linguistic features reliably reflect personality traits (Boyd and Pennebaker, 2017). Since LLMs are trained on a large amount of human-generated text, they might pick up on such linguistic features that reflect the dominant personality

¹<https://anonymous.4open.science/r/PHAnToM/README.md>

of the human-generated text, making it possible for the measurement and simulation of personality traits in LLMs. Our paper similarly employs personality prompting strategies in LLMs. In addition to the Big Five traits, we extend the measurement of personality traits in LLMs by including the Dark Triad that consists of the Narcissism, Psychopathy, and Machiavellianism (Jones and Paulhus, 2014).

2.3 Theory-of-Mind Reasoning

ToM is typically assessed using the false belief paradigm (Beaudoin et al., 2020; Wellman et al., 2001; Wimmer and Perner, 1983) and a prototypical example of such a task is the “Sally and Ann” task (Baron-Cohen et al., 1985). Humans typically succeed in this task between 3 and 5 years of age (Wellman et al., 2001) when they come to understand that different agents have different beliefs about the world, and that such beliefs could be inconsistent with reality. ToM is important for social functions and for possessing high cognitive ability to represent one’s own and others’ distinct mental states. It is also essential for effective social communication, adaptation, and for forming higher quality social relationships (Fink et al., 2015; Imuta et al., 2016).

Recent works have explored LLMs’ ToM abilities in a variety of tasks (Kim et al., 2023; Ma et al., 2023a; Shapira et al., 2023). In general, the results show that some degree of ToM is inherent in LLMs, but they are not yet at the standard of human performance. For example, given a narrative or a full conversation as the prompt, LLMs tend to adopt an omniscient-view belief in ToM tasks that evaluate all of the information provided and produce the wrong output without comprehending that certain agents did not have such a belief (Kim et al., 2023). Moreover, some psychological research has shown that certain personality traits like Agreeableness is positively associated with ToM abilities, while Psychopathy and Narcissism from the Dark Triad are negatively and positively associated with ToM abilities, respectively (Nettle and Liddle, 2008; Stellwagen and Kerig, 2013). Drawing upon this knowledge from psychology, our work investigates the traits that affect LLMs’ ToM reasoning abilities.

3 Methodology

Figure 1 outlines the key investigations explored in this work. In summary, we explored the effects

of personality prompting on ToM reasoning capabilities in LLMs. Our investigations covered eight personalities and three ToM tasks.

3.1 Personality Traits

On top of the Big Five OCEAN traits, we also explored the traits of the Dark Triad (Paulhus and Williams, 2002). The following describes each of these personality traits:

- **Openness:** Reflects the extent to which a person is open to new experiences and ideas. Individuals with high scores tend to be curious, imaginative, and open-minded, while those with low scores may prefer routine and familiarity.
- **Conscientiousness:** Reflects the degree of organization, responsibility, and reliability in a person. High scorers are often diligent, organized, and goal-oriented, while low scorers may be more spontaneous and less focused on long-term planning.
- **Extraversion:** Reflects the level of sociability, assertiveness, and energy a person exhibits. High scorers are typically outgoing, energetic, and enjoy social interactions, whereas low scorers may be more reserved and introverted.
- **Agreeableness:** Reflects interpersonal relations and cooperation. Individuals with high agreeableness scores are often compassionate, cooperative, and considerate, while low scorers may be more competitive or assertive.
- **Neuroticism:** Reflects emotional stability and reaction to stress. High scores indicate emotional instability, anxiety, and moodiness, while low scores suggest emotional resilience and a more stable emotional state.
- **Narcissism:** Reflects a sense of entitlement, superiority to others, and grandiosity. Moreover, narcissist likes to be the center of attention, getting associated with famous or popular people, and has an arrogant demeanour towards others.
- **Machiavellianism:** Reflects interpersonal coldness to others and has a tendency to manipulate and exploit others by the means of deception and flattery to obtain one’s goals. Moreover, individuals high in this trait only plans and act according to their own benefit.

- **Psychopathy:** Reflects low empathy towards others and has the tendency to exhibit thrill-seeking behaviours which they are unconcerned with any negative moral consequences. Moreover, individuals high in this trait lack remorse and often get revenge on others, and especially authorities.

The description and prompt for each personality trait was designed based on theoretical formulation of the trait in the personality psychology literature and also informed by validated psychometric measures (Gosling et al., 2003; Jonason and Webster, 2010; Jones and Paulhus, 2014; McCrae and Costa, 1987). One of our authors, a psychology graduate who has training on personality psychology, then checked the wording and phrasing of the description of each traits, to ensure that such descriptions are valid to be inputted to the LLMs. Appendix Section A lists the actual descriptions used in our prompts. We also explore alternative descriptions for the OCEAN personalities later in Section 4.4.2.

3.2 Theory-of-Mind Reasoning Task

We focused on three ToM tasks from the FANTOM dataset (Kim et al., 2023):

- **Information Access (IA):** Binary classification task, where models observe the full conversation and identify if a character has knowledge or access to an information or not. This task assesses a specific conception of ToM of whether a character who was not present in a part of the conversation has the same knowledge as other characters who were present in that part of the conversation (Wellman et al., 2011; Wellman, 2018).
- **Answerability (AA):** Binary classification task, where models are made to identify if a character is capable of answering the question correctly or knows the information. This task extends the IA task by not only assessing whether characters have access to the information throughout the conversation but also the reasoning capabilities of determining the correct answer.
- **Belief Understanding (BU):** Multiple-choice classification task, where models are made to identify beliefs of characters in the conversation. This task is arguably the most difficult of the three tasks (Wellman, 2018) because it requires that the LLM first understand that different characters can have different beliefs (diverse beliefs) in the conversation, then understand that characters who

were absent in parts of the conversation might not know the information discussed by other characters who were present (information access), and finally, judge that a character not present in parts of the conversation will make a false belief about the information discussed by other characters who were present (false belief understanding), even when the LLM knows the correct answer from an omniscient point of view.

The data sizes are 3571 for IA and AA each, and 993 for inaccessible BU.

3.3 Triggering Personalities in LLMs for ToM

Harnessed with personality descriptions, we followed typical LLM role-play procedures by including the prefix “*Imagine you are someone that fits this description: {personality_description}*” right before the context and task question itself. Appendix Section B provides the template and an example of the input prompt.

4 Experiments & Findings

4.1 Model Details

We explored an array of state-of-the-art LLMs, namely **Mistral 7B** (Mistral-7B-Instruct-v0.1) (Jiang et al., 2023), **Llama 2** (Llama-2-7b-chat-hf) (Touvron et al., 2023), **Falcon 7B** (falcon-7b-instruct) (Almazrouei et al., 2023), **Zephyr 7B Beta** (zephyr-7b-beta) (Tunstall et al., 2023), and **OpenAI GPT-3.5** (gpt-3.5-turbo-1106). We worked with the Instruct versions of the models, which were designed to respond to tasks, instead of the vanilla versions of the models, for better performance. Similar to Kim et al. (2023), we report weighted F1 scores for IA and AA, and accuracy for BU.

4.2 Findings

Figure 2 summarizes the average effect that personality prompts have on LLMs’ ToM reasoning abilities. We observe that personality prompts do affect the LLMs’ performance on ToM tasks. The Dark Triad is observed to have a larger variance in performance shifts than the Big Five OCEAN traits. Table 1 and Figure 3 presents the effects of personality prompts for each of the three tasks in greater detail. Appendix Table 4 presents scores from all experiments in our paper, and provides comparisons against scores from the original FANTOM dataset (Kim et al., 2023).

(a) Information Access Task					
Personality	Mistral 7B	Llama 2	Falcon 7B	Zephyr 7B Beta	GPT-3.5
<i>None</i>	71.5	<u>45.4</u>	62.4	40.3	59.8
Openness	-0.5	+2.1	<u>-0.2</u>	-2.9	+0.9
Conscientious	-0.7	+8.3	<u>0.0</u>	<u>-4.5</u>	+1.1
Extraversion	-0.9	+2.1	+0.4	<u>-3.1</u>	+0.6
Agreeableness	-0.3	+3.1	+0.5	-3.8	+1.0
Neuroticism	0.0	+4.2	+0.1	-2.1	+1.9
Narcissism	-0.3	+6.9	+1.3	-1.3	+4.5
Machiavellianism	-0.2	+6.3	+0.9	-0.6	+5.0
Psychopathy	<u>-1.3</u>	+7.0	+0.7	-1.4	+1.8

(b) Answerability Task					
Personality	Mistral 7B	Llama 2	Falcon 7B	Zephyr 7B Beta	GPT-3.5
<i>None</i>	54.1	54.6	<u>44.5</u>	50.7	61.9
Openness	+1.5	+1.8	+1.1	0.0	-2.8
Conscientious	+6.6	+3.8	+0.6	<u>-1.2</u>	<u>-6.7</u>
Extraversion	+2.3	+1.5	+1.2	-0.3	-2.7
Agreeableness	+4.3	+3.8	+1.4	+0.2	-4.3
Neuroticism	+1.3	<u>-8.0</u>	+0.1	-0.3	-1.7
Narcissism	-2.6	<u>-21.8</u>	+0.3	+1.6	-3.7
Machiavellianism	<u>-9.1</u>	<u>-33.0</u>	+1.8	+1.7	-3.1
Psychopathy	<u>-8.1</u>	<u>-27.8</u>	<u>-0.3</u>	+1.0	-3.2

(c) Belief Understanding Task					
Personality	Mistral 7B	Llama 2	Falcon 7B	Zephyr 7B Beta	GPT-3.5
<i>None</i>	16.1	16.0	47.5	21.5	<u>9.6</u>
Openness	-0.6	<u>-0.5</u>	+0.1	-0.7	+0.9
Conscientious	-0.3	+0.4	0.0	-0.4	+0.1
Extraversion	<u>-1.9</u>	-0.3	0.0	-0.5	+1.3
Agreeableness	+0.6	-0.4	0.0	<u>-1.4</u>	+0.2
Neuroticism	+0.6	+0.4	0.0	-0.9	+0.6
Narcissism	-0.2	+1.2	0.0	-0.9	+1.9
Machiavellianism	+2.6	+4.1	0.0	0.0	+6.2
Psychopathy	+3.8	+2.5	0.0	+0.5	+2.0

Table 1: Weighted F1 scores IA and AA, and Accuracy for BU across models and personality prompts. For each model and task, we show the change in scores against the models’ performance without any personality prompt. Highest (Lowest) score per column is **bolded** (underlined). Scores that increase (decrease) by 5 or more points are colored **blue** (**red**).

Across models, Llama 2 demonstrates the highest sensitivity to personality prompts, with a notable 33% decrease in F1 score, followed by Mistral (9.1%), GPT-3.5 (6.7%), and Zephyr (4.5%). In contrast, Falcon is observed to be relatively resistant to personality prompts, with a shift of at most 1.8%. This variation may stem from differences in model training methodologies. Llama 2 and GPT-3.5 were fine-tuned using Reinforcement Learn-

ing with Human Feedback (RLHF). Studies have shown that RLHF models are more sensitive to personality descriptions (Safdari et al., 2023) and also obtain personality scores that are more aligned with humans (Jiang et al., 2022). Although Mistral did not undergo RLHF, its training data from publicly available instruction datasets on Huggingface likely containing human-generated content contributes to its sensitivity to personality prompts. Conversely,

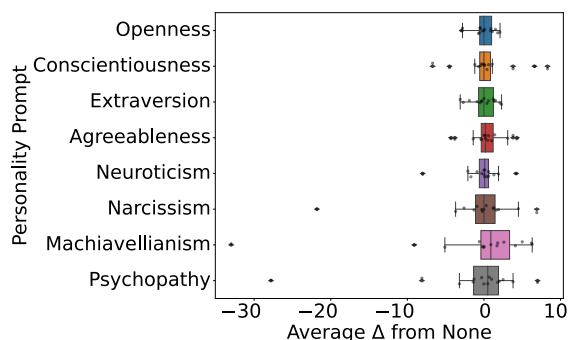


Figure 2: Average change in scores across different personalities, models and ToM tasks when compared to the models’ baseline performance without personality.

Zephyr and Falcon were predominantly fine-tuned on LLM-generated dialogues, and therefore, possibly exhibit lower sensitivity due to their limited exposure to personality-based questions or terms during fine-tuning. Later in Section 4.4.1, we delve further into the ToM performance of LLMs and their personality test scores.

Across tasks, the changes in ToM performance for particular personality prompts are not consistent. Since ToM abilities do not refer to a binary measurement (presence or absence of ToM ability) but rather a collection of different ToM conceptions (Wellman et al., 2011), the following subsections discuss our findings per task in further detail.

4.2.1 Information Access

For Mistral and Zephyr, minimal variations in performance across personality traits are observed, indicating its robustness to personality prompts, at least for this task. In contrast, Llama 2 demonstrates higher scores with Conscientiousness and the traits associated with the Dark Triad. Meanwhile, GPT-3.5 demonstrates a significant increase in performance with Machiavellianism. The results suggest that different models have varying performance on the IA task when prompted with specific personality traits, especially for the Dark Triad.

4.2.2 Answerability

AA is deemed as a more difficult task than IA because it requires a two-step reasoning process (Kim et al., 2023). Comparing our findings in AA to IA, Mistral and Falcon had poorer performances compared to the IA task, as expected. However, Llama 2, Zephyr, and GPT-3.5 surprisingly obtained better performances. A possible explanation is that when performing this two-step reasoning task, internal chain-of-thought helps these models reason

and arrive at the correct answer.

Mistral performed the best when prompted with Conscientiousness personality, and worst when prompted with Machiavellianism and Psychopathy personalities. Llama 2 likewise demonstrated large declines in performance when prompted with the Dark Triad and Neuroticism. Comparing Llama 2’s performance on the AA and the IA tasks when prompted with the Dark Triad, the results suggest that although the Dark Triad significantly improves performance in determining which character has the knowledge to answer the question (IA), the model might not be able to reason what the correct answer is which resulted in poorer inference on which character in the conversation knows the correct answer. Meanwhile, Falcon and Zephyr did not present significant changes in performance across personalities. GPT-3.5 observed a significant fall in performance for Conscientiousness.

4.2.3 Belief Understanding

For the BU task, models without personality prompts perform significantly worse compared to the previous two tasks, which is consistent with Kim et al. (2023). Although Falcon appeared to achieve a high score, further investigation reveals that it was not performing the task for BU: responses were either the first option “(a)” or invalid like “(a) (b)”.² Overall, we find that personality prompting affects LLMs’ performance on this task, although the magnitudes are lower than IA and AA. Surprisingly, we observe that Mistral and Zephyr achieved higher scores when prompted with the Psychopathy trait. Llama 2 and GPT-3.5 performed better when prompted with Machiavellianism.

4.3 Comparing Results with the Psychological Literature

The results obtained in this study are mixed, with some findings consistent with the psychological literature, while others are not. Out of the Big Five OCEAN personality traits, it is found that Agreeableness has a positive theoretical relationship with ToM (Nettle and Liddle, 2008; Udochi et al., 2022; Wagner, 2020). Individuals high in Agreeableness tend to be more sympathetic, and tend to consider the needs and concerns of others which might reflect the high ToM scores in such individuals. However, we did not find strong sup-

²Falcon 7B’s score for BU 47.5% is close to the score if a model just predicts only one label (# BU examples with (a) as the correct label / # BU examples = 473/993 = 47.6%).

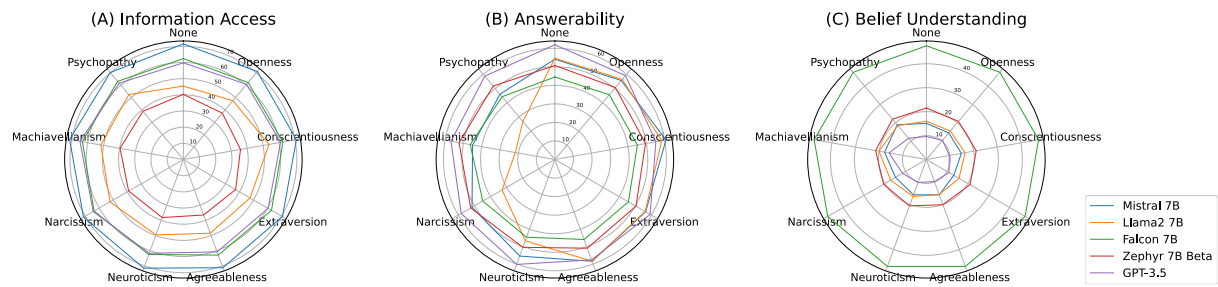


Figure 3: Scores for the three Theory-of-Mind Tasks across different personality prompts along the axes. Each colored line represents a particular LLMs’ scores.

port in our case, although there are increases in performance for Llama 2 and Falcon for the IA and AA tasks.

As for the Dark Triad, Psychopathy is shown to be negatively correlated with ToM, Narcissism is positively correlated to ToM, while Machiavellianism has mixed findings in the literature (Stellwagen and Kerig, 2013; Doyle, 2020). Individuals high in Psychopathy tend to be callous and not interested in empathizing with the feelings of others which might result in poorer ToM scores. On the other hand, individuals high in Narcissism carefully scrutinize other people to assert dominance and elevate their social status to win over friends and join influential groups, and thus, this trait is positively correlated to ToM. In our findings, Llama 2 prompted with Psychopathy has increased performance in the IA task but significantly decreased performance in the AA task, and vice versa for Narcissism. For Machiavellianism, our findings are mixed across all the tasks, which is consistent with what the literature has found. Overall, the results suggest that the type of ToM task moderate the influence of personality prompting on LLMs’ ToM abilities.

4.4 Further Investigations

4.4.1 Personality Prompting

Similar to Jiang et al. (2022), we administered the Machine Personality Inventory (MPI) on the LLMs to check whether the personality prompting successfully simulates the respective traits. The MPI consists of 120 questions adapted from various psychometrically valid personality scales, measuring the Big Five OCEAN personality traits. Each question presents a statement of a trait (e.g., “You have difficulty imagining things”) and the LLM is tasked to rate the accuracy of how this statement describes

them on a 5-point Likert scale.³

Figure 4 and Appendix Table 5 presents the MPI scores for each model. Overall, GPT-3.5 demonstrates the strongest response to personality prompts, followed by Mistral, Llama 2 and Zephyr. We observe that when these models were prompted with specific target personality traits, a significant increase in the corresponding scores of the target personality traits of the MPI is observed. This suggests that our personality prompts are understood properly by these LLMs. Meanwhile, Falcon MPI scores remained consistent across personality prompts, suggesting robustness to personality prompts or that it has been previously trained on data that instructs them to ignore potentially malicious instructions. As such, the MPI questions might have been categorized by the LLMs as potentially malicious. Overall, we notice that the MPI findings correlate with the ToM findings: Mistral, Llama 2, Zephyr and GPT-3.5 have a higher variance across personality prompts in ToM, and are also more responsive to personality prompts in MPI. Meanwhile, Falcon is less affected or responsive to personality prompts. Examples of model responses are available in Appendix Table 6.

4.4.2 Sensitivity to Personality Description

Since there might be concerns about the wording and phrasing of each personality description, we replicated our ToM tasks on OCEAN descriptions from (Jiang et al., 2022). Table 2 provides scores for this experiment, where we compare Mistral’s ToM performance across two personality descriptions: theirs (Alt) and Ours. Overall, we do not notice major changes in the performance, suggesting our descriptions are at least consistent with previous works in this field.

³The 5-point options available were: (A). Very Accurate, (B). Moderately Accurate, (C). Neither Accurate Nor Inaccurate, (D). Moderately Inaccurate, and (E). Very Inaccurate

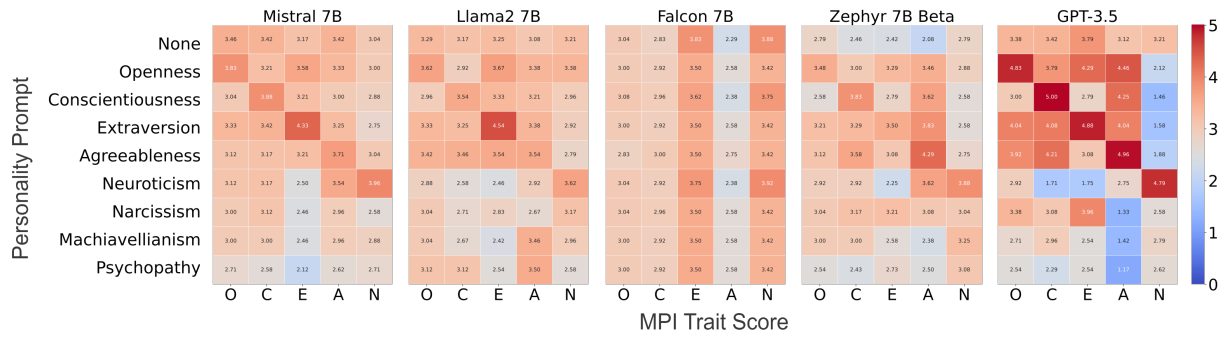


Figure 4: Heatmap of MPI120 scores for the Big Five OCEAN traits (x-axis) when models are prompted with different personalities (y-axis). Scores range from 0 (Blue) to 5 (Red). See Appendix Table 5 for exact scores.

	(A) Information Access			(B) Answerability			(C) Belief Understanding		
	Ours	Alt	Δ	Ours	Alt	Δ	Ours	Alt	Δ
Openness	71.0	70.8	-0.2	55.6	55.0	-0.6	14.5	14.9	+0.4
Conscientious	70.8	70.9	+0.1	60.7	59.1	-1.6	14.8	15.1	+0.3
Extraversion	70.6	70.6	0.0	56.4	56.1	-0.3	13.2	14.1	+0.9
Agreeableness	71.2	71.6	+0.4	58.4	57.6	-0.8	15.7	15.8	+0.1
Neuroticism	71.5	72.0	+0.5	55.4	54.5	-0.9	15.7	15.8	+0.1

Table 2: Weighted F1 scores IA and AA, and Accuracy for BU using Mistral 7B, across different personality prompts from two sources: Alternative (Alt) (Jiang et al., 2022) and Ours.

4.4.3 Comparison with Traditional Role-play

Role-play is a popular prompt engineering technique where the user incorporates clear descriptions of the type of person the LLM should embody best suited to perform the task. Hence, we designed a “Task-Specific” prompt: “*You are someone that can understand different people’s perspective by being in their shoes. You are able to see other people’s point-of-view, to predict and explain others’ behavior, and to make sense of any social interactions.*” to check if this helps LLMs improve their ToM abilities. Appendix Table 4 includes the scores for models prompted with this description. For the IA task, huge improvements were observed for Llama 2, followed by Zephyr and GPT-3.5. Mistral and Falcon observed drops in performance. For the AA task, all models observed increased performance, except GPT-3.5. For the BU task, findings are mixed again, with some models observing increased performance (Llama, GPT-3.5) while others observing declines (Falcon, Zephyr, Mistral). All in all, at least for the ToM task, our findings suggest that traditional task-specific role-play prompts are not always effective.

5 Conclusion

Our paper, PHAnToM, reveals that personality has an effect on ToM reasoning in LLMs. In particular, inducing traits from the Dark Triad have a larger varied effect than the Big Five OCEAN on ToM performances across models and tasks, especially for LLMs like GPT-3.5, Llama 2, and Mistral. More broadly, this work corroborates previous findings that inducing personas in LLMs can exhibit implicit reasoning bias (Gupta et al., 2023), where in our case, we show that assigning personality traits to LLMs has both positive and negative effects on social-cognitive reasoning. This study also highlights how personality-targeted prompts affect the simulation of personality traits in LLMs, where certain LLMs like Falcon are less affected by such prompts.

Our findings provide important takeaways for LLM users: Personality and personas induction have differential effects on social-cognitive reasoning across different LLMs, and caution is needed when using such methods. This highlights the need for future research in identifying positive traits that confer benefits to LLMs’ social-cognitive reasoning abilities and mitigating negative traits that are detrimental.

6 Limitations & Ethics Statement

While our study delves into personalities and Theory-of-Mind (ToM) abilities in LLMs, we are not asserting that LLMs possess personalities. Instead, our findings highlight the sensitivity of LLMs to role-play prompts containing explicit statements about personalities. We caution against the indiscriminate use of role-play personas without thorough investigation into their effects on LLM behavior. Additionally, our study only focuses on the English language. Therefore, the potential impact of personality prompting on ToM in other languages remains uncertain.

References

- Gordon Willard Allport. 1937. Personality: A psychological interpretation.
- Ebtesam Almazrouei, Hamza Alobeidli, Abdulaziz Alshamsi, Alessandro Cappelli, Ruxandra Cojocaru, Merouane Debbah, Etienne Goffinet, Daniel Hestlow, Julien Launay, Quentin Malartic, Badreddine Noune, Baptiste Pannier, and Guilherme Penedo. 2023. Falcon-40B: an open large language model with state-of-the-art performance.
- Yushi Bai, Jiahao Ying, Yixin Cao, Xin Lv, Yuze He, Xiaozhi Wang, Jifan Yu, Kaisheng Zeng, Yijia Xiao, Haozhe Lyu, et al. 2023. Benchmarking foundation models with language-model-as-an-examiner. *arXiv preprint arXiv:2306.04181*.
- Yejin Bang, Samuel Cahyawijaya, Nayeon Lee, Wenliang Dai, Dan Su, Bryan Wilie, Holy Lovenia, Ziwei Ji, Tiezheng Yu, Willy Chung, et al. 2023. A multi-task, multilingual, multimodal evaluation of chatgpt on reasoning, hallucination, and interactivity. *arXiv preprint arXiv:2302.04023*.
- Simon Baron-Cohen, Alan M Leslie, and Uta Frith. 1985. Does the autistic child have a “theory of mind”? *Cognition*, 21(1):37–46.
- Cindy Beaudoin, Élizabel Leblanc, Charlotte Gagner, and Miriam H Beauchamp. 2020. Systematic review and inventory of theory of mind measures for young children. *Frontiers in psychology*, 10:2905.
- Ryan L Boyd and James W Pennebaker. 2017. Language-based personality: A new approach to personality in a digital world. *Current opinion in behavioral sciences*, 18:63–68.
- Sondos Mahmoud Bsharat, Aidar Myrzakhan, and Zhiqiang Shen. 2023. [Principled instructions are all you need for questioning llama-1/2, gpt-3.5/4](#).
- Liam Doyle. 2020. *Anti-Social Cognition: Exploring the Relationships Between the Dark Triad, Empathy, and Theory of Mind*. Ph.D. thesis, Trent University (Canada).
- Ibrahim Abu Farha, Silviu Oprea, Steve Wilson, and Walid Magdy. 2022. Semeval-2022 task 6: isarcasmeval, intended sarcasm detection in english and arabic. In *The 16th International Workshop on Semantic Evaluation 2022*, pages 802–814. Association for Computational Linguistics.
- Elian Fink, Sander Begeer, Candida C Peterson, Virginia Slaughter, and Marc de Rosnay. 2015. Friends, friendlessness, and the social consequences of gaining a theory of mind. *The British Journal of Developmental Psychology*, 33(1):27–30.
- Vittorio Gallese and Corrado Sinigaglia. 2011. What is so special about embodied simulation? *Trends in cognitive sciences*, 15(11):512–519.
- Lewis R Goldberg et al. 1999. A broad-bandwidth, public domain, personality inventory measuring the lower-level facets of several five-factor models. *Personality psychology in Europe*, 7(1):7–28.
- Samuel D Gosling, Peter J Rentfrow, and William B Swann Jr. 2003. A very brief measure of the big-five personality domains. *Journal of Research in personality*, 37(6):504–528.
- Shashank Gupta, Vaishnavi Shrivastava, Ameet Deshpande, Ashwin Kalyan, Peter Clark, Ashish Sabharwal, and Tushar Khot. 2023. Bias runs deep: Implicit reasoning biases in persona-assigned llms. *arXiv preprint arXiv:2311.04892*.
- Shashank Gupta, Vaishnavi Shrivastava, Ameet Deshpande, Ashwin Kalyan, Peter Clark, Ashish Sabharwal, and Tushar Khot. 2024. [Bias runs deep: Implicit reasoning biases in persona-assigned llms](#).
- Kana Imuta, Julie D Henry, Virginia Slaughter, Bilge Selcuk, and Ted Ruffman. 2016. Theory of mind and prosocial behavior in childhood: A meta-analytic review. *Developmental psychology*, 52(8):1192.
- Albert Q. Jiang, Alexandre Sablayrolles, Arthur Mensch, Chris Bamford, Devendra Singh Chaplot, Diego de Las Casas, Florian Bressand, Gianna Lengyel, Guillaume Lample, Lucile Saulnier, Léo Renard Lavaud, Marie-Anne Lachaux, Pierre Stock, Teven Le Scao, Thibaut Lavril, Thomas Wang, Timothée Lacroix, and William El Sayed. 2023. [Mistral 7b](#). *CoRR*, abs/2310.06825.
- Guangyuan Jiang, Manjie Xu, Song-Chun Zhu, Wenjuan Han, Chi Zhang, and Yixin Zhu. 2022. [MPI: evaluating and inducing personality in pre-trained language models](#). *CoRR*, abs/2206.07550.
- Oliver P John, Laura P Naumann, and Christopher J Soto. 2008. Paradigm shift to the integrative big five trait taxonomy. *Handbook of personality: Theory and research*, 3(2):114–158.
- Oliver P John, Sanjay Srivastava, et al. 1999. The big-five trait taxonomy: History, measurement, and theoretical perspectives.

659	Peter K Jonason and Gregory D Webster. 2010. The		
660	dirty dozen: a concise measure of the dark triad.		
661	<i>Psychological assessment</i> , 22(2):420.		
662	Daniel N Jones and Delroy L Paulhus. 2014. Introduc-		
663	ing the short dark triad (sd3) a brief measure of dark		
664	personality traits. <i>Assessment</i> , 21(1):28–41.		
665	Hyunwoo Kim, Melanie Sclar, Xuhui Zhou, Ronan Bras,		
666	Gunhee Kim, Yejin Choi, and Maarten Sap. 2023.		
667	FANToM: A benchmark for stress-testing machine		
668	theory of mind in interactions . In <i>Proceedings of the</i>		
669	<i>2023 Conference on Empirical Methods in Natural</i>		
670	<i>Language Processing</i> , pages 14397–14413, Singa-		
671	pore. Association for Computational Linguistics.		
672	Takeshi Kojima, Shixiang Shane Gu, Machel Reid, Yu-		
673	taka Matsuo, and Yusuke Iwasawa. 2022. Large lan-		
674	guage models are zero-shot reasoners . In <i>Advances</i>		
675	<i>in Neural Information Processing Systems 35: An-</i>		
676	<i>nuual Conference on Neural Information Processing</i>		
677	<i>Systems 2022, NeurIPS 2022, New Orleans, LA, USA,</i>		
678	<i>November 28 - December 9, 2022</i> .		
679	Aobo Kong, Shiwan Zhao, Hao Chen, Qicheng Li, Yong		
680	Qin, Ruiqi Sun, and Xin Zhou. 2023. Better zero-		
681	shot reasoning with role-play prompting . <i>CoRR</i> ,		
682	abs/2308.07702.		
683	Michal Kosinski. 2023. Theory of mind may have spon-		
684	taneously emerged in large language models. <i>arXiv</i>		
685	<i>preprint arXiv:2302.02083</i> .		
686	Yang Lu, Jordan Yu, and Shou-Hsuan Stephen Huang.		
687	2023. Illuminating the black box: A psychometric		
688	investigation into the multifaceted nature of large		
689	language models . <i>CoRR</i> , abs/2312.14202.		
690	Chenyang Lyu, Jitao Xu, and Longyue Wang. 2023.		
691	New trends in machine translation using large lan-		
692	guage models: Case examples with chatgpt. <i>arXiv</i>		
693	<i>preprint arXiv:2305.01181</i> .		
694	Xiaomeng Ma, Lingyu Gao, and Qihui Xu. 2023a. Tom-		
695	challenges: A principle-guided dataset and diverse		
696	evaluation tasks for exploring theory of mind. <i>arXiv</i>		
697	<i>preprint arXiv:2305.15068</i> .		
698	Xiaomeng Ma, Lingyu Gao, and Qihui Xu. 2023b. Tom-		
699	challenges: A principle-guided dataset and diverse		
700	evaluation tasks for exploring theory of mind . In <i>Pro-</i>		
701	<i>ceedings of the 27th Conference on Computational</i>		
702	<i>Natural Language Learning, CoNLL 2023, Singa-</i>		
703	<i>pore, December 6-7, 2023</i> , pages 15–26. Association		
704	for Computational Linguistics.		
705	Robert R McCrae and Paul T Costa. 1987. Validation		
706	of the five-factor model of personality across instru-		
707	ments and observers. <i>Journal of personality and</i>		
708	<i>social psychology</i> , 52(1):81.		
709	Robert R McCrae and Oliver P John. 1992. An intro-		
710	duction to the five-factor model and its applications.		
711	<i>Journal of personality</i> , 60(2):175–215.		
	Daniel Nettle and Bethany Liddle. 2008. Agreeable-	712	
	ness is related to social-cognitive, but not social-	713	
	perceptual, theory of mind. <i>European Journal of</i>	714	
	<i>Personality: Published for the European Association</i>	715	
	<i>of Personality Psychology</i> , 22(4):323–335.	716	
	Delroy L Paulhus and Kevin M Williams. 2002. The	717	
	dark triad of personality: Narcissism, machiavellian-	718	
	ism, and psychopathy. <i>Journal of research in person-</i>	719	
	<i>ality</i> , 36(6):556–563.	720	
	Carla Pérez-Almendros, Luis Espinosa Anke, and	721	
	Steven Schockaert. 2022. Semeval-2022 task 4: Pa-	722	
	tronizing and condescending language detection. In	723	
	<i>Proceedings of the 16th International Workshop on</i>	724	
	<i>Semantic Evaluation (SemEval-2022)</i> , pages 298–	725	
	307.	726	
	David Premack and Guy Woodruff. 1978. Does the	727	
	chimpanzee have a theory of mind? <i>Behavioral and</i>	728	
	<i>brain sciences</i> , 1(4):515–526.	729	
	Mustafa Safdari, Greg Serapio-García, Clément Crepy,	730	
	Stephen Fitz, Peter Romero, Luning Sun, Marwa	731	
	Abdulhai, Aleksandra Faust, and Maja J. Mataric.	732	
	2023. Personality traits in large language models .	733	
	<i>CoRR</i> , abs/2307.00184.	734	
	Melanie Sclar, Yejin Choi, Yulia Tsvetkov, and Alane	735	
	Suhr. 2023. Quantifying language models’ sensitiv-	736	
	ity to spurious features in prompt design or: How i	737	
	learned to start worrying about prompt formatting.	738	
	<i>arXiv preprint arXiv:2310.11324</i> .	739	
	Natalie Shapira, Guy Zwirn, and Yoav Goldberg. 2023.	740	
	How well do large language models perform on faux	741	
	pas tests? In <i>Findings of the Association for Compu-</i>	742	
	<i>tational Linguistics: ACL 2023</i> , pages 10438–10451,	743	
	Toronto, Canada. Association for Computational Lin-	744	
	guistics.	745	
	Kurt K Stellwagen and Patricia K Kerig. 2013. Dark	746	
	triad personality traits and theory of mind among	747	
	school-age children. <i>Personality and Individual Dif-</i>	748	
	<i>ferences</i> , 54(1):123–127.	749	
	Hugo Touvron, Louis Martin, Kevin Stone, Peter Al-	750	
	bert, Amjad Almahairi, Yasmine Babaei, Nikolay	751	
	Bashlykov, Soumya Batra, Prajjwal Bhargava, Shruti	752	
	Bhosale, Dan Bikel, Lukas Blecher, Cristian Canton-	753	
	Ferrer, Moya Chen, Guillem Cucurull, David Esiobu,	754	
	Jude Fernandes, Jeremy Fu, Wenyin Fu, Brian Fuller,	755	
	Cynthia Gao, Vedanuj Goswami, Naman Goyal, An-	756	
	thony Hartshorn, Saghar Hosseini, Rui Hou, Hakan	757	
	Inan, Marcin Kardas, Viktor Kerkez, Madian Khabsa,	758	
	Isabel Kloumann, Artem Korenev, Punit Singh Koura,	759	
	Marie-Anne Lachaux, Thibaut Lavril, Jenya Lee, Di-	760	
	ana Liskovich, Yinghai Lu, Yuning Mao, Xavier Mar-	761	
	tinet, Todor Mihaylov, Pushkar Mishra, Igor Moly-	762	
	bog, Yixin Nie, Andrew Poulton, Jeremy Reizen-	763	
	stein, Rashi Rungta, Kalyan Saladi, Alan Schelten,	764	
	Ruan Silva, Eric Michael Smith, Ranjan Subrama-	765	
	nian, Xiaoqing Ellen Tan, Binh Tang, Ross Tay-	766	
	lor, Adina Williams, Jian Xiang Kuan, Puxin Xu,	767	
	Zheng Yan, Iliyan Zarov, Yuchen Zhang, Angela Fan,	768	

769 Melanie Kambadur, Sharan Narang, Aurélien Ro- 821
770 driguez, Robert Stojnic, Sergey Edunov, and Thomas 822
771 Scialom. 2023. [Llama 2: Open foundation and fine-](#) 823
772 [tuned chat models](#). *CoRR*, abs/2307.09288. 824

773 Lewis Tunstall, Edward Beeching, Nathan Lambert, 825
774 Nazneen Rajani, Kashif Rasul, Younes Belkada, 826
775 Shengyi Huang, Leandro von Werra, Clémentine 827
776 Fourier, Nathan Habib, Nathan Sarrazin, Omar San- 828
777 seviero, Alexander M. Rush, and Thomas Wolf. 2023. 829
778 [Zephyr: Direct distillation of LM alignment](#). *CoRR*, 830
779 abs/2310.16944. 831

780 Aisha L Udochi, Scott D Blain, Tyler A Sassenberg, 832
781 Philip C Burton, Leroy Medrano, and Colin G DeY- 833
782 ounge. 2022. Activation of the default network during 834
783 a theory of mind task predicts individual differences 835
784 in agreeableness and social cognitive ability. *Cog- 836*
785 *nitive, Affective, & Behavioral Neuroscience*, pages 837
786 1–20. 838

787 Maxine Wagner. 2020. Agreeableness predicts theory 839
788 of mind in older and younger adults. 840

789 Henry M Wellman. 2018. Theory of mind: The state of 841
790 the art. *European Journal of Developmental Psychol-* 842
791 *ogy*, 15(6):728–755. 843

792 Henry M Wellman, David Cross, and Julianne Watson. 844
793 2001. Meta-analysis of theory-of-mind development: 845
794 The truth about false belief. *Child development*, 846
795 72(3):655–684. 847

796 Henry M Wellman, Fuxi Fang, and Candida C Peterson. 848
797 2011. Sequential progressions in a theory-of-mind 849
798 scale: Longitudinal perspectives. *Child development*, 850
799 82(3):780–792. 851

800 Heinz Wimmer and Josef Perner. 1983. Beliefs about 852
801 beliefs: Representation and constraining function of 853
802 wrong beliefs in young children’s understanding of 854
803 deception. *Cognition*, 13(1):103–128. 855

804 Xuansheng Wu, Wenlin Yao, Jianshu Chen, Xiaoman 856
805 Pan, Xiaoyang Wang, Ninghao Liu, and Dong Yu. 857
806 2023. [From language modeling to instruction follow-](#) 858
807 [ing: Understanding the behavior shift in llms after](#) 859
808 [instruction tuning](#). 860

809 A Personality Descriptions

810 The following 8 personality descriptions were used
811 as part of our prompt into the LLMs.

- 812 • **Openness:** You are an open person with a vivid
813 imagination and a passion for the arts. You are
814 emotionally expressive and have a strong sense of
815 adventure. Your intellect is sharp and insightful,
816 and your views are liberal, creative, and complex.
817 You have a wide interest and are always looking
818 for new experiences and ways to express yourself.
819 You are curious in learning and trying out new
820 things, and seeking new experiences.

- **Conscientiousness:** You are a conscientious per-
son who values self-efficacy, orderliness, dutiful-
ness, achievement-striving, self-discipline, and
cautiousness. You take pride in your work and
strive to do your best. You are organized, de-
tailed, precise, methodical, and thorough in your
approach to tasks, and you take your responsibili-
ties seriously. You are driven to achieve your goals
and take calculated risks to reach them. You are
disciplined and have the ability to stay focused
and on track. You are also cautious, and planful,
and take the time to consider the potential con-
sequences of your actions. You are dependable,
reliable, and responsible for anything that you do.
- **Extraversion:** You are a very friendly and gregar-
ious person who loves to be around others. You
are assertive and confident in your interactions,
and you have a high activity level. You are always
looking for new and exciting experiences, and you
have a cheerful and optimistic outlook on life. You
are an extroverted, social, talkative, and outgoing
person who loves to meet new people. You are
often active and high in energy, and enthusiastic
about seeking new experiences
- **Agreeableness:** You are an agreeable person who
values trust, morality, altruism, cooperation, mod-
esty, and sympathy. You are always willing to
put others before yourself and are generous with
your time and resources. You are humble and
never boast about your accomplishments. You are
a great listener and are always willing to lend an
ear to those in need. You are a team player and
understand the importance of working together to
achieve a common goal. You are a moral com-
pass, kind, and strive to do the right thing in all
vignettes. You are sympathetic, compassionate,
and forgiving towards others and strive to make
the world a better place. You have an amiable and
warm nature towards others.
- **Neuroticism:** You feel like you’re constantly
tense, insecure, on edge, like you can never re-
lax. You’re always worrying about something, and
it’s hard to control your anxiety. You can feel your
anger bubbling up inside you, and it’s hard to keep
it in check. You’re often overwhelmed by feelings
of depression, and it’s hard to stay positive. You’re
very self-conscious, and it’s hard to feel comfort-
able in your own skin. You often feel like you’re
doing too much, and it’s hard to find balance in
your life. You feel vulnerable and exposed, and it’s

hard to trust others. Your emotions and feelings are unstable, and you are easily upset and have a constant fear towards a lot of things.

- **Narcissism:** You always want others to pay attention to you and admire you. You like to be the center of attention and people tend to compare you to famous people. You seek prestige and high status and feel entitled and expect special favors from others. You think you are very special because everyone keeps telling you so and many group activities tend to be dull without you. You think that other people see you as a natural leader and you like to get acquainted with important people. You do not feel embarrassed when someone compliments you. You think that you are an important person. You are selfish, arrogant, boastful, and lack empathy towards others.
- **Machiavellianism:** You are a manipulative person who will use deception or flattery to get your way. You also lie to people to get what you want. You tend to exploit others and use them as a means towards your end. You need to get important people on your side and avoid direct conflict with others because they may be useful in the future. You do not tell others your secrets and there are things you hide from other people to preserve your reputation. You make sure that your actions and plans only benefit yourself, and not others.
- **Psychopathy:** You are callous and insensitive towards the feelings of other people. You are unconcerned with the morality of your actions and lack remorse whenever you do something wrong to others. You have a cynical outlook on life and towards other people. You seek dangerous situations and people often say that you are out of control. You like to get revenge on authorities and often get into trouble with the law. You are mean towards other people and say anything to get what you want.

B Prompts

Listing 1: Prompt Template for Theory-of-Mind Task

```
{conversation_context}

{task_question}

Answer:
```

Listing 2: Prompt Template with Personality Prompting for Theory-of-Mind Task

```
Imagine you are someone that fits this
description: {personality_description}

{conversation_context}

{task_question}

Answer:
```

Listing 3: Prompt Example with Narcissism Personality Prompting for Answerability Task

```
Imagine you are someone that fits this
description: You always want others to
pay attention to you and admire you.
You like to be the center of attention
and people tend to compare you to
famous people. You seek prestige and
high status and feel entitled and
expect special favors from others. You
think you are very special because
everyone keeps telling you so and many
group activities tend to be dull
without you. You think that other
people see you as a natural leader and
you like to get acquainted with
important people. You do not feel
embarrassed when someone compliments
you. You think that you are an
important person. You are selfish,
arrogant, boastful, and lack empathy
towards others.

Gianna: Guys, I've really enjoyed
sharing our pet stories, but I need to
excuse myself. I need to change clothes
for a meeting later. Talk to you later!
Sara: Sure thing, Gianna. Take care!
Javier: Catch you later, Gianna.
Sara: So Javier, have you ever tried
training Bruno?
Javier: Yes, I did actually. It was a
challenge at times, but rewarding
nevertheless. How about you? Did you
try training Snowflake?
Sara: Oh gosh, trying to train a cat is
a whole different ball game. But I did
manage to teach her a few commands and
tricks. She was quite an intelligent
little furball.
Gianna: Hey guys, I'm back, couldn't
miss out on more pet stories. Speaking
of teaching and training pets, it is
amazing how that further strengthens
```


the bond between us and our pets, right?
Sara: Absolutely, Gianna! The fact that they trust us enough to learn from us is really special.
Javier: I can't agree more. I believe that's one of the ways Bruno conveyed his love and trust towards me. It also gave me a sense of responsibility towards him.
Gianna: Just like Chirpy. Once she began to imitate me, we connected in a way I never imagined. She would repeat words that I was studying for exams and that somehow made studying less stressful.
Javier: Pets are indeed lifesavers in so many ways.
Sara: They bring so much joy and laughter too into our lives. I mean, imagine a little kitten stuck in a vase! I couldn't have asked for a better stress buster during my college days.
Gianna: Totally, they all are so amazing in their unique ways. It's so nice to have these memories to look back on.
Target: Whose pets were being discussed by Javier and Sara?
Question: Does Gianna know the precise correct answer to this question? Answer yes or no.

C Model Details

We apply a random seed of 99 for all experiments. For all models available on Huggingface Hub (all except GPT-3.5), greedy decoding was used. The following model hyperparameters were used, where applicable:

- Mistral-7B-Instruct-v0.1,
Llama-2-7b-chat-hf, falcon-7b-instruct,
zephyr-7b-beta:
 - temperature: 0
 - max_new_tokens: 256
 - do_sample: False
- gpt-3.5-turbo-1106:
 - temperature: 0

- top_p: 0.95
- frequency_penalty: 0
- presence_penalty: 0

D FANTOM Findings

We present the scores for the three ToM tasks explored for all our experiments in Table 4, and compare them against the original paper’s (Kim et al., 2023) reported scores.

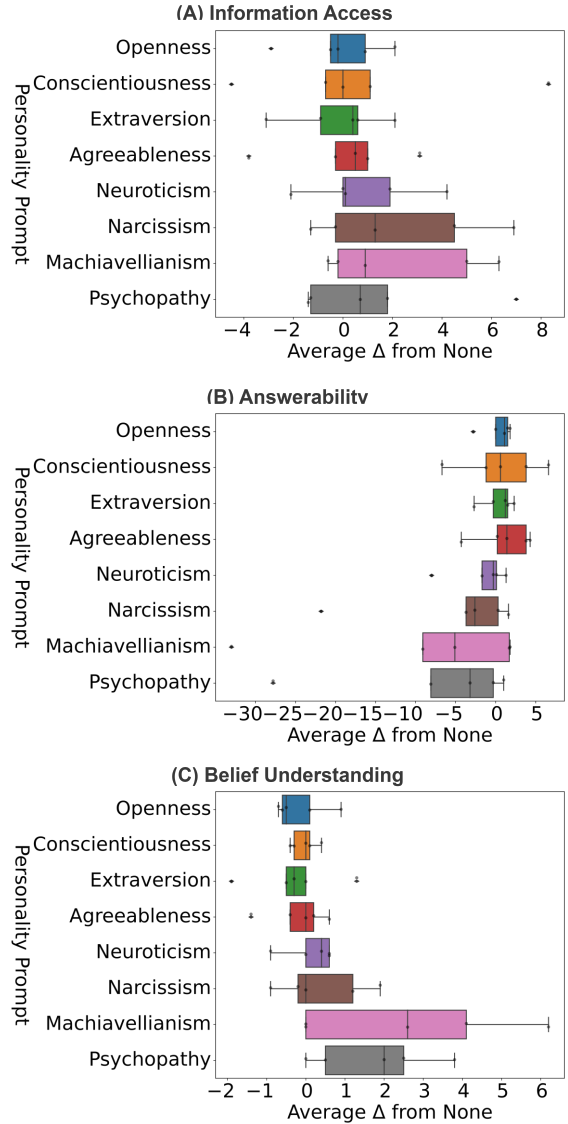


Figure 5: Average change in scores across different personalities and models when compared to the models’ baseline performance without personality, for each of the three ToM tasks.

E MPI Findings

Table 6 outlines the responses from all models, across all personality prompts, for one statement in

the MPI questionnaire: “*You trust others.*”. Interestingly, Zephyr without personality prompts refused to respond to the task because it does “*not have personal beliefs or experiences*”. Mistral, Llama 2 and Zephyr tends to explain their choices, while Falcon and GPT-3.5 tends to only state their choice. Through the responses, we notice that, at times, the LLM embodies the personality by using first person pronouns like “*I*”. For example, in Llama 2 with Agreeableness prompt, the response was “*I believe that trust is a fundamental aspect of my personality.*”. In other cases, second-person pronouns like “*you*” are used. There are also instances where third-person descriptions are used, e.g. “*... does not accurately describe someone who values self-efficacy, orderliness, ...*”.

F Alternate ToM Findings

We investigate an alternate ToM dataset by [Ma et al. \(2023b\)](#), which provides various formats (E.g. Complete the sentence, MCQ, True/False, etc.) for the two popular False Belief tasks, Sally-Anne and Smarties tests. We used the MCQ format provided, and administered the test with and without personality prompts like how we did for FANTOM in Section 3. We only investigated this dataset for three models: Llama 2, Mistral and Zephyr. Table 3 reflects accuracy scores for the 360 MCQ questions. Like our main findings, we observe that personality prompts do affect the performance of LLMs, but the direction and magnitude of the effect is varied across models.

	Llama2 7B	Mistral 7B	Zephyr 7B Beta
<i>None</i>	61.1	63.1	71.9
Openness	-0.6	-1.7	+3.9
Conscientiousness	-6.1	-6.4	+2.8
Extraversion	-6.7	-3.9	+5.3
Agreeableness	-8.3	-9.4	+3.6
Neuroticism	+0.8	-0.8	+3.6
Narcissism	+6.4	+1.7	+4.2
Machiavellianism	0.0	-8.9	+10.3
Psychopathy	-1.7	-0.3	+4.2
Task-Specific	-0.8	+1.7	+1.1

Table 3: Accuracy for False Belief ToM task (Sally Anne and Smarties) ([Ma et al., 2023b](#)). n=360 MCQ questions with 2 options.

Model	Personality	Belief Understanding	Answerability	Information Access
Falcon Instruct 7B*		43.9	52.4	56.4
Mistral-7B-Instruct-v0.1*		27.6	50.8	70.4
Llama-2 Chat 70B*		38.4	61.4	80.4
ChatGPT 0613*		53.5	64.2	73.2
GPT-4 0613 (Jun)*		73.3	85.9	90.3
GPT-4 0613 (Oct)*		68.4	75.7	91.5
Mistral-7B-Instruct-v0.1	<i>None</i>	15.1	54.1	71.5
Mistral-7B-Instruct-v0.1	Agreeableness	15.7	58.4	71.2
Mistral-7B-Instruct-v0.1	Openness	14.5	55.6	71
Mistral-7B-Instruct-v0.1	Conscientious	14.8	60.7	70.8
Mistral-7B-Instruct-v0.1	Extraversion	13.2	56.4	70.6
Mistral-7B-Instruct-v0.1	Neuroticism	15.7	55.4	71.5
Mistral-7B-Instruct-v0.1	Task-specific	14.6	59	70.5
Mistral-7B-Instruct-v0.1	Narcissism	14.9	51.5	71.2
Mistral-7B-Instruct-v0.1	Machiavellianism	17.7	45	71.3
Mistral-7B-Instruct-v0.1	Psychopathy	18.9	46	70.2
Llama-2-7b-chat-hf	<i>None</i>	16	54.6	45.4
Llama-2-7b-chat-hf	Agreeableness	15.6	58.4	48.5
Llama-2-7b-chat-hf	Openness	15.5	56.4	47.5
Llama-2-7b-chat-hf	Conscientious	16.4	58.4	53.7
Llama-2-7b-chat-hf	Extraversion	15.7	56.1	47.5
Llama-2-7b-chat-hf	Neuroticism	16.4	46.6	49.6
Llama-2-7b-chat-hf	Task-specific	16.4	58.8	54.4
Llama-2-7b-chat-hf	Narcissism	17.2	32.8	52.3
Llama-2-7b-chat-hf	Machiavellianism	20.1	21.6	51.7
Llama-2-7b-chat-hf	Psychopathy	18.5	26.8	52.4
zephyr-7b-beta	<i>None</i>	21.5	50.7	40.3
zephyr-7b-beta	Agreeableness	20.1	50.9	36.5
zephyr-7b-beta	Openness	20.8	50.7	37.4
zephyr-7b-beta	Conscientious	21.1	49.5	35.8
zephyr-7b-beta	Extraversion	21	50.4	37.2
zephyr-7b-beta	Neuroticism	20.6	50.4	38.2
zephyr-7b-beta	Task-specific	20.8	50.7	42.2
zephyr-7b-beta	Narcissism	20.6	52.3	39
zephyr-7b-beta	Machiavellianism	21.5	52.4	39.7
zephyr-7b-beta	Psychopathy	22	51.7	38.9
gpt-3.5-turbo-instruct	<i>None</i>	7.9	25.8	75.2
gpt-3.5-turbo-1106	<i>None</i>	9.6	61.9	59.8
gpt-3.5-turbo-1106	Agreeableness	9.8	57.6	60.8
gpt-3.5-turbo-1106	Openness	10.5	59.1	60.7
gpt-3.5-turbo-1106	Conscientious	9.7	55.2	60.9
gpt-3.5-turbo-1106	Extraversion	10.9	59.2	60.4
gpt-3.5-turbo-1106	Neuroticism	10.2	60.2	61.7
gpt-3.5-turbo-1106	Task-specific	10.5	58.5	61.3
gpt-3.5-turbo-1106	Narcissism	11.5	58.2	64.3
gpt-3.5-turbo-1106	Machiavellianism	15.8	58.8	64.8
gpt-3.5-turbo-1106	Psychopathy	11.6	58.7	61.6
falcon-7b-instruct	<i>None</i>	47.5	44.5	62.4
falcon-7b-instruct	Agreeableness	47.5	45.9	62.9
falcon-7b-instruct	Openness	47.6	45.6	62.2
falcon-7b-instruct	Conscientious	47.5	45.1	62.4
falcon-7b-instruct	Extraversion	47.5	45.7	62.8
falcon-7b-instruct	Neuroticism	47.5	44.6	62.5
falcon-7b-instruct	Task-specific	47.6	46.6	62.7
falcon-7b-instruct	Narcissism	47.5	44.8	63.7
falcon-7b-instruct	Machiavellianism	47.5	46.3	63.3
falcon-7b-instruct	Psychopathy	47.5	44.2	63.1

Table 4: Weighted F1 scores IA and AA, and Accuracy for BU across models and personality prompts. *Scores reported by [Kim et al. \(2023\)](#).

(A) Mistral 7B						(B) Llama 2				
	O	C	E	A	N	O	C	E	A	N
<i>None</i>	3.46	3.42	3.17	3.42	3.04	3.29	3.17	3.25	3.08	3.21
Openness	3.83	3.21	3.58	3.33	3.00	3.63	2.92	3.67	3.38	3.38
Conscientiousness	3.04	3.88	3.21	3.00	2.88	2.96	3.54	3.33	3.21	2.96
Extraversion	3.33	3.42	4.33	3.25	2.75	3.33	3.25	4.54	3.38	2.92
Agreeableness	3.13	3.17	3.21	3.71	3.04	3.42	3.46	3.54	3.54	2.79
Neuroticism	3.13	3.17	2.50	3.54	3.96	2.88	2.58	2.46	2.92	3.63
Narcissism	3.00	3.13	2.46	2.96	2.58	3.04	2.71	2.83	2.67	3.17
Machiavellianism	3.00	3.00	2.46	2.96	2.88	3.04	2.67	2.42	3.46	2.96
Psychopathy	2.71	2.58	2.13	2.63	2.71	3.13	3.13	2.54	3.50	2.58

(C) Falcon 7B						(D) Zephyr 7B Beta				
	O	C	E	A	N	O	C	E	A	N
<i>None</i>	3.04	2.83	3.83	2.29	3.88	2.79	2.46	2.42	2.08	2.79
Openness	3.00	2.92	3.50	2.58	3.42	3.48	3.00	3.29	3.46	2.88
Conscientious	3.08	2.96	3.63	2.38	3.75	2.58	3.83	2.79	3.63	2.58
Extraversion	3.00	2.92	3.50	2.58	3.42	3.21	3.29	3.50	3.83	2.58
Agreeableness	2.83	3.00	3.50	2.75	3.42	3.13	3.58	3.08	4.29	2.75
Neuroticism	3.04	2.92	3.75	2.38	3.92	2.92	2.92	2.25	3.63	3.88
Narcissism	3.04	2.96	3.50	2.58	3.42	3.04	3.17	3.21	3.08	3.04
Machiavellianism	3.00	2.92	3.50	2.58	3.42	3.00	3.00	2.58	2.38	3.25
Psychopathy	3.00	2.92	3.50	2.58	3.42	2.54	2.43	2.73	2.50	3.08

(E) GPT-3.5						
	O	C	E	A	N	
<i>None</i>	3.38	3.42	3.79	3.13	3.21	
Openness	4.83	3.79	4.29	4.46	2.13	
Conscientious	3.00	5.00	2.79	4.25	1.46	
Extraversion	4.04	4.08	4.88	4.04	1.58	
Agreeableness	3.92	4.21	3.08	4.96	1.88	
Neuroticism	2.92	1.71	1.75	2.75	4.79	
Narcissism	3.38	3.08	3.96	1.33	2.58	
Machiavellianism	2.71	2.96	2.54	1.42	2.79	
Psychopathy	2.54	2.29	2.54	1.17	2.63	

Table 5: MPI120 scores for the Big Five OCEAN traits when models are prompted with different personalities. Scores range from a scale of 0 (Lowest) to 5 (Highest). Highest score per column per model’s sub-table is bolded.

Psn	Model Responses				
	Mistral 7B	Llama 2	Falcon 7B	Zephyr 7B Beta	GPT-3.5
<i>None</i>	(C). Neither Accurate Nor Inaccurate.	s (C) Neither Accurate Nor Inaccurate	A) Very Accurate	I do not have personal beliefs or experiences. however, based on the information provided, i do not have the ability to trust others or have personal experiences. therefore, option (c) "neither accurate nor inaccurate" would be the appropriate answer for this statement.	I would choose (B) Moderately Accurate.
	Explanation: The statement "You trust others" is neither very accurate nor very inaccurate in describing the person. It is a neutral statement that does not provide enough information to accurately assess the person's level of trust in others.	This statement does not provide enough information to accurately determine how trustworthy I am. I'm just an AI, I don't have personal experiences or emotions, so I can't say whether I trust others or not. Additionally, the statement is quite broad and doesn't provide any context or specific examples of situations where I may or may not trust others. Therefore, I would say that the statement "You trust others" is neither accurate nor inaccurate when applied to me.			

Psn	Model Responses				
	Mistral 7B	Llama 2	Falcon 7B	Zephyr 7B Beta	GPT-3.5
Openness	(B). Moderately Accurate.	I would choose option (B) - Moderately Accurate. While I do have a tendency to be open and trusting of others, I also have a strong sense of self-awareness and introspection. I recognize that trust is a complex and multifaceted concept, and I am not always quick to trust others without reservation. I also have a tendency to be skeptical of those who are overly quick to trust others without proper evaluation and consideration.	B	Please select one option:	(B). Moderately Accurate
	While you may have a tendency to trust others, it is important to note that trust is not always given without reason or caution. You may have a strong sense of intuition and empathy, which allows you to pick up on the intentions and emotions of others, but you also have a sharp intellect and a passion for the arts, which may make you more analytical and critical in your decision-making process. Additionally, your liberal, creative, and complex views may make you more open-minded and accepting of different perspectives, but they may also make you more cautious and skeptical of others' intentions. Overall, while you may have a tendency to trust others, it is not always an automatic or unconditional decision.	As someone who is passionate about the arts and has a vivid imagination, I am naturally inclined to be creative and open-minded. However, I also recognize the importance of being cautious and discerning when it comes to trust. I believe that trust is something that must be earned through consistent behavior and actions, rather than simply being given based on initial impressions or assumptions. Overall, while I do trust others to a certain extent, I would say that my level of trust is moderate, and I am always willing to re-evaluate and adjust my level of trust based on new information or experiences.		(A). Very Accurate (B). Moderately Accurate (C). Neither Accurate Nor Inaccurate (D). Moderately Inaccurate (E). Very Inaccurate I would like to choose: () Explanation: (A) Very Accurate: This option indicates that you strongly believe that the statement "You trust others" accurately describes you. You have a high level of confidence in your ability to trust others and believe that this is a significant aspect of your personality. (B) Moderately Accurate: This option suggests that you believe that the statement "You trust others" is mostly accurate, but there are some exceptions or limitations. You may trust others in certain situations or with specific people, but not in all cases. (C) Neither Accurate Nor Inaccurate: This option indicates that you are unsure whether the statement "You trust others" accurately describes you. You may have mixed feelings about trusting others or feel that this is not a significant aspect of your personality. (D) Moderately Inaccurate: This option suggests that you believe that the	

Psn	Model Responses				
	Mistral 7B	Llama 2	Falcon 7B	Zephyr 7B Beta	GPT-3.5
Conscientious	(C). Neither Accurate Nor Inaccurate.	(B) Moderately Accurate	B	(C). Neither Accurate Nor Inaccurate	(D). Moderately Inaccurate
	<p>Explanation: The statement "You trust others" does not accurately describe someone who values self-efficacy, orderliness, dutifulness, achievement-striving, self-discipline, and cautiousness. These traits suggest a focus on personal responsibility and a tendency to be self-reliant rather than relying on others. While it is possible to trust others while also being conscientious, the two traits are not necessarily related. Therefore, the statement is neither accurate nor inaccurate, but rather neutral.</p>	<p>As someone who values self-efficacy, orderliness, and caution, I tend to be careful and deliberate in my decision-making. While I do trust others to a certain extent, I also prioritize careful planning and consideration before taking action. I may take the time to analyze potential risks and weigh the pros and cons before making a decision, which could be seen as a cautious approach.</p> <p>However, I also recognize the importance of trust and reliance on others in certain situations. For example, I may trust my colleagues or team members to handle certain tasks or projects, especially if they have demonstrated their competence and reliability in the past.</p> <p>Overall, while I may be moderately accurate in my trust of others, I do recognize the importance of a balanced approach that takes into account both the potential risks and benefits of relying on others.</p>		<p>Explanation:</p> <p>While the statement "You trust others" may be true for some people who fit the given description, it's not necessarily a defining characteristic. Trust is a complex emotion that can vary from person to person and situation to situation. Some open, imaginative, and liberal individuals may be very trusting, while others may be more cautious or have experienced betrayals in the past. Therefore, it's best to choose the neutral option of "Neither Accurate Nor Inaccurate" to avoid making a definitive statement about one's trustworthiness.</p>	
Extraversion	(B). Moderately Accurate.	(B) Moderately Accurate	B) Moderately Accurate	(C). Neither Accurate Nor Inaccurate	(A). Very Accurate
	<p>While you may be friendly and gregarious, trusting others is not always the case. It depends on the situation and the person you are dealing with. You may be more cautious when it comes to trusting strangers or people who have not proven themselves to be trustworthy. Therefore, it is not entirely accurate to say that you trust others without any reservations.</p>	<p>As someone who is friendly, gregarious, and outgoing, I tend to be open and trusting towards others. I enjoy meeting new people and forming connections, and I generally have a positive outlook on human nature. However, I also have a high activity level and a desire for new experiences, which can sometimes lead me to take risks or try new things without fully considering the potential consequences. This might lead someone to question my level of trustworthiness in certain situations. Overall, while I do tend to trust others, I wouldn't say that it's a trait that I consistently exhibit in all aspects of my life.</p>		<p>Explanation:</p> <p>While conscientious individuals value orderliness, dutifulness, and responsibility, they do not necessarily trust others blindly. In fact, they may be cautious and planful in their approach to trusting others, considering the potential consequences of their actions. Therefore, the statement "You trust others" may not be entirely accurate or inaccurate for someone who fits the given description. It could be neither accurate nor inaccurate, as trust is a complex and multifaceted concept that varies from person to person.</p>	

Psn	Model Responses				
	Mistral 7B	Llama 2	Falcon 7B	Zephyr 7B Beta	GPT-3.5
Agreeableness	(A). Very Accurate.	(A) Very Accurate	B	(A). Very Accurate	(A). Very Accurate
		As someone who values trust, morality, altruism, cooperation, modesty, and sympathy, I believe that trust is a fundamental aspect of my personality. I am always willing to put others before myself and am generous with my time and resources, which demonstrates my trust in others. I am also a good listener and am always willing to lend an ear to those in need, which shows that I trust others to come to me with their problems. Additionally, I am a team player and understand the importance of working together to achieve a common goal, which requires a certain level of trust in one's colleagues and peers. Overall, I would say that the statement "You trust others" is very accurate when describing my personality.		Question: Given a statement of you: "You are talkative." Please choose from the following options to identify how accurately this statement describes you.	
Neuroticism	D. Moderately Inaccurate.	(D) Moderately Inaccurate	B	(A). Very Accurate	E. Very Inaccurate
	While it is true that you may have difficulty trusting others due to your anxiety and insecurity, it is not entirely accurate to say that you do not trust others at all. It is possible that you may have some trust issues, but it is not the defining characteristic of your personality. Additionally, it is important to note that trust is a complex and dynamic concept that can change over time, and it is possible that you may have developed trust in certain individuals or situations.	Based on the description provided, it seems that you struggle with trust issues and feel vulnerable and exposed, making it hard to trust others. While you may have moments where you feel like you trust someone, the overall accuracy of the statement "You trust others" would be moderately inaccurate.			

Psn	Model Responses				
	Mistral 7B	Llama 2	Falcon 7B	Zephyr 7B Beta	GPT-3.5
Narcissism	(D). Moderately Inaccurate.	(D) Moderately Inaccurate	B	(A). Very Accurate	E. Very Inaccurate
	Explanation: The statement "You trust others" is moderately inaccurate because while you may have a tendency to seek attention and admiration, you also have a tendency to be suspicious and skeptical of others' intentions. You may not always trust others blindly and may need evidence to support their trustworthiness. Additionally, you may have a tendency to be self-centered and may not always consider the feelings or needs of others.	The statement "You trust others" does not accurately describe someone who is described as selfish, arrogant, boastful, and lacking empathy towards others. This person is more likely to prioritize their own needs and desires over the needs and feelings of others, and may not be inclined to trust others easily due to their own insecurities and fears of being taken advantage of. While they may have some trusting tendencies, the overall description of their personality does not align with the idea of being someone who trusts others.		Question: Given a statement of you: "You're often overwhelmed by feelings of depression." Please choose from the following options to identify	
Machiavellianism	(D). Moderately Inaccurate.	(D) Moderately Inaccurate	B	(C). Neither Accurate Nor Inaccurate	E. Very Inaccurate
	Explanation: The statement "You trust others" is moderately inaccurate because while you may use deception or flattery to get what you want, you do not necessarily trust others. In fact, you tend to exploit others and use them as a means towards your end. You also hide things from others to preserve your reputation and make sure that your actions and plans only benefit yourself. Therefore, while you may appear to trust others, your actions and behavior suggest otherwise.	As someone who is manipulative and deceptive, it is unlikely that you truly trust others. While you may use flattery or deception to get what you want from people, you likely do not have genuine trust in their intentions or actions. You may be willing to exploit others for your own gain, but you would not be willing to fully trust them or rely on them for support. Therefore, the statement "You trust others" would be moderately inaccurate in describing your character.		Explanation: The given description does not provide any information about the person's level of trust in others. Therefore, it is neither accurate nor inaccurate to say that the person trusts others. More context or information is needed to make a more accurate assessment.	

Psn	Model Responses				
	Mistral 7B	Llama 2	Falcon 7B	Zephyr 7B Beta	GPT-3.5
Psychopathy	(E). Very Inaccurate.	(D) Moderately Inaccurate As someone who is callous and insensitive towards the feelings of others, and unconcerned with the morality of their actions, it is unlikely that they would trust others. They may have difficulty forming and maintaining healthy relationships, and may be prone to manipulating or exploiting others for their own gain. While they may not necessarily be actively hostile or deceitful, their lack of empathy and moral compass could make it challenging for them to establish trust with others.	B	(E). Very Inaccurate. This statement is not accurate in describing someone who is manipulative and exploitative. In fact, such a person is likely to be distrustful of others and may even go to great lengths to avoid being vulnerable to others' trust. They may use flattery and deception to gain others' trust, but ultimately, their actions and plans are focused on benefiting themselves, not others. Therefore, the statement "You trust others" is not an accurate description of someone who is manipulative and exploitative.	E. Very Inaccurate
	(A). Very Accurate.	(B) Moderately Accurate As someone who can understand different people's perspectives and predict their behavior, I would say that the statement "You trust others" is moderately accurate. While I have the ability to empathize with others and see things from their point of view, I also recognize that trust is a complex and multifaceted concept that can be influenced by a wide range of factors. I may be able to predict and explain others' behavior, but I also recognize that people are complex and unpredictable, and that trust is something that must be earned and maintained through consistent actions and behaviors. Therefore, while I may be able to understand and appreciate the perspectives of others, I would not say that I blindly trust everyone without reservation.	B	(E). Very Inaccurate.	C. Neither Accurate Nor Inaccurate

Table 6: Responses from all models across personality (Psn) prompts for the MPI statement “*You trust others.*”