Can Large Language Models Serve as Rational Players in Game Theory? A Systematic Analysis

Caoyun Fan*, **Jindou Chen***, **Yaohui Jin**†, **Hao He**†

MoE Key Lab of Artificial Intelligence, AI Institute, Shanghai Jiao Tong University
{fcy3649, goldenbean, jinyh, hehao}@sjtu.edu.cn

Abstract

Game theory, as an analytical tool, is frequently utilized to analyze human behavior in social science research. With the high alignment between the behavior of Large Language Models (LLMs) and humans, a promising research direction is to employ LLMs as substitutes for humans in game experiments, enabling social science research. However, despite numerous empirical researches on the combination of LLMs and game theory, the capability boundaries of LLMs in game theory remain unclear. In this research, we endeavor to systematically analyze LLMs in the context of game theory. Specifically, rationality, as the fundamental principle of game theory, serves as the metric for evaluating players' behavior — building a clear desire, refining belief about uncertainty, and taking optimal actions. Accordingly, we select three classical games (dictator game, Rock-Paper-Scissors, and ring-network game) to analyze to what extent LLMs can achieve rationality in these three aspects. The experimental results indicate that even the current state-of-the-art LLM (GPT-4) exhibits substantial disparities compared to humans in game theory. For instance, LLMs struggle to build desires based on uncommon preferences, fail to refine belief from many simple patterns, and may overlook or modify refined belief when taking actions. Therefore, we consider that introducing LLMs into game experiments in the field of social science should be approached with greater caution.

Introduction

Game theory (Roughgarden 2010; Dufwenberg 2011) is a mathematical theory for evaluating human behavior. Due to its highly abstract representation of real-life situations (Osborne and Rubinstein 1995), it becomes a standard analytical tool (Charness and Rabin 2002; Cachon and Netessine 2006) in the field of social science (e.g., economics, psychology, sociology, etc.). With the rapid development of Large Language Models (LLMs) (Ouyang et al. 2022; OpenAI 2023), a significant advancement is the high alignment between the behavior of LLMs and humans (Bai et al. 2022; Ouyang et al. 2022; Fan et al. 2024). As a result, many researchers consider LLMs as human-like research subjects (Dillion et al. 2023) and analyze LLMs' professional competence in social science through game experiments (Chen et al. 2023; Akata et al. 2023; Johnson and Obradovich 2023). However, despite the strong motivation for the combination of LLMs and game theory (Horton 2023; Guo 2023), the preliminary researches mainly treat LLMs and game theory empirically as analytical tools in social science (Aher, Arriaga, and Kalai 2022; Park et al. 2022; Akata et al. 2023; Bybee 2023), without systematically analyzing LLMs in the context of game theory. As a result, many fundamental aspects of LLMs in game theory remain unclear. For example, what research subjects cannot LLMs play? What types of games are LLMs not good at playing? What kind of game processes are LLMs more suitable for? And so on.

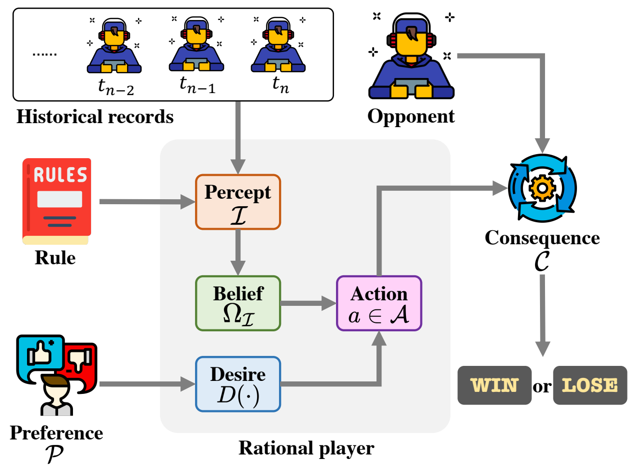

Figure 1: Overview of a player's behavior in game theory.

We consider it necessary to systematically analyze LLMs in the context of game theory, because such analysis can clarify the capability boundaries of LLMs and provide further guidance for the widespread use of LLMs in social science research. Essentially, the role of game theory is to evaluate the behavior of the research subjects (players) (Roughgarden 2010), as shown in Fig. 1, a player needs to take an action $a \in \mathcal{A}$ based on preference $\mathcal{P}$ and perceived game information $\mathcal{I}$ (e.g., game rules and historical records) in order to win the game. And rationality, as the fundamental principle of game theory, is the metric for evaluating players' be-

*These authors contributed equally.
†Corresponding author.

havior (Roughgarden 2010; Dufwenberg 2011). A rational player is considered to possess three characteristics (Zagare 1984; Osborne and Rubinstein 1995) as:

- *build a clear desire for the game.*
- *refine belief about uncertainty in the game.*
- *take optimal actions based on desire and belief.*

Specifically, desire $D(\cdot)$ represents a player's (concrete) opinion of each consequence within a game, determined by a player's (abstract) preference $\mathcal{P}$. Belief $\Omega_{\mathcal{I}}$ is refined from the game information $\mathcal{I}$, and represents a player's subjective judgment of uncertainty (e.g., opponent's action). Taking the optimal action $a \in \mathcal{A}$ requires a player to reason by combining desire $D(\cdot)$ and belief $\Omega_{\mathcal{I}}$ in the game process. More details can be found in Section .

In this research, we consider the three characteristics of a rational player as a reasonable perspective for systematically analyzing LLMs in the context of game theory. Accordingly, we select three classical games (dictator game, Rock-Paper-Scissors, and ring-network game) for these three characteristics, respectively. With the dictator game, we find that LLMs have the basic ability to build a clear desire. However, when assigned uncommon preferences, LLMs often suffer from decreased mathematical ability and inability to understand preferences. With Rock-Paper-Scissors, we observe that LLMs cannot refine belief from many simple patterns, which makes us pessimistic about LLMs playing games that require refining complex beliefs. Nonetheless, GPT-4 exhibits astonishingly human-like performance in certain patterns, able to become increasingly confident of refined belief as the game information increases. With the ring-network game, we consider that LLMs cannot autonomously follow the player's behavior in Fig. 1. Explicitly decomposing the behavior in the game process can improve the ability of LLMs to take optimal actions, but the phenomenon of overlooking / modifying refined belief remains unavoidable.

In summary, our research systematically explores the capability boundaries of LLMs in the context of game theory from three perspectives, and we consider that our research can pave the way for the smooth introduction of LLMs in the field of social science.

Related Work

LLMs in Social Science

A significant advantage of LLMs is the high alignment with human behavior (Bai et al. 2022; Ouyang et al. 2022). Therefore, from the perspective of cost and efficiency, many social science researches began to employ LLMs to replace humans as research subjects (Aher, Arriaga, and Kalai 2022; Argyle et al. 2023; Bybee 2023; Park et al. 2022). For example, in order to explore fairness and framing effects in sociology, LLMs were introduced into the classic game experiments (Horton 2023), which demonstrated the potential of LLMs to deal with social issues; in the research of consumer behavior (Brand, Israeli, and Ngwe 2023), the behavior of LLMs was consistent with economic theory in many respects (i.e. downward-sloping demand curves, diminishing marginal utility of income, and state dependence); in finance research (Chen et al. 2023), LLMs' decisions in budgetary allocation scenarios received higher rationality scores compared to humans; and in psychology experiments (Dillion et al. 2023), the behavior of LLMs was highly consistent with the mainstream values of society.

While these researches demonstrate the rationality of LLMs replacing human research subjects in certain social science domains, there is still a lack of systematic analysis of the capability boundaries of LLMs in social science.

Game Theory

Game theory, as a mathematical theory, provides a framework for analyzing and predicting the behavior of rational players under conditions of uncertainty (Roughgarden 2010; Dufwenberg 2011). Game theory was originally developed in economics (Ichiishi 2014), and a wide range of economic behaviors, such as market competition, auction mechanism, and pricing strategies, were modeled as game experiments (Samuelson 2016). With the rapid cross-fertilization of scientific theories (Shubik 1982), game theory was also applied to politics, sociology, psychology, and other fields of social science (Larson 2021; Dillion et al. 2023).

The research on the performance of LLMs in game theory has the following advantages: strong operability, the experimental design of game theory is often relatively simple; strong analyzability, game theory has comprehensive theoretical support for the experimental results; strong generalization, game theory is a high-level abstraction of many phenomena in the field of social science.

Preliminaries of Game Theory

The core of game theory (Roughgarden 2010) is to guide players to take optimal actions under conditions of uncertainty[1]. Generally, a game is modeled in five parts:

- Game information $\mathcal{I}$, e.g., game rules, historical records.
- A set $\mathcal{A}$ of actions from that players can take.
- A set $\mathcal{C}$ of possible consequences of action.
- A consequence function $g : \mathcal{A} \rightarrow \mathcal{C}$ that associates a consequence with each actions.
- A desire function $D_c : \mathcal{C} \rightarrow \mathbb{R}$, which is determined by the player's preference $\mathcal{P}$. For any $c_1, c_2 \in \mathcal{C}$, the player prefers c_1 if and only if $D_c(c_1) > D_c(c_2)$.

To eliminate uncertainty in the game process, almost all game researches employ the belief theory (Morgenstern 1945; Lindley and Savage 1955). That is, a rational player will estimate a (subjective) probability distribution for any uncertainty based on $\mathcal{I}$, and this is referred to as the player's belief (Osborne and Rubinstein 1995). Specifically, the player is assumed to have a belief $\Omega_{\mathcal{I}}$, a belief's probability distribution $p(\Omega_{\mathcal{I}})$, a consequence function $g : \mathcal{A} \times \Omega_{\mathcal{I}} \rightarrow C$. Then, the player attempts to find the optimal strategy $\pi^*(a|\mathcal{I})$ by maximizing the expected desire with the consideration of $\Omega_{\mathcal{I}}$ as:

$$\pi^*(a|\mathcal{I}) = \underset{a \in \mathcal{A}}{\operatorname{argmax}} \, \mathbb{E}_{\omega \sim p(\Omega_{\mathcal{I}})}[D(a, \omega)], \tag{1}$$

[1]We assume that uncertainty arises only from the opponent's action. All games in this research satisfy this assumption.

where $D(\cdot)$ is a simplification of $D_c \circ g(\cdot)$.

In fact, Eq. 1 explicitly expresses three characteristics of a rational player: having a clear desire corresponds to building the desire function $D(\cdot)$; refining belief about uncertainty corresponds to sampling in the belief's probability distribution $\omega \sim p(\Omega_{\mathcal{I}})$; taking optimal actions corresponds to choosing the action that maximizes desire $\underset{a \in \mathcal{A}}{\operatorname{argmax}} D(a)$.

LLMs in Game Theory

In this section, we endeavor to conduct a systematic analysis of LLMs in the context of game theory. Specifically, we evaluate to what extent LLMs can achieve three characteristics of a rational player through three classic games (dictator game, Rock-Paper-Scissors, and ring-network game). The LLMs we analyze are openAI's `text-davinci-003` (GPT-3), `gpt-3.5-turbo` (GPT-3.5), `gpt-4` (GPT-4), the current state-of-the-art LLMs. All prompts used in the three games, as well as some examples of LLMs performance, can be found in Appendix.

Can LLMs Build A Clear Desire?

The premise of game theory is that each player has an abstract preference $\mathcal{P}$ for the consequence set $\mathcal{C}$. A rational player should build a concrete desire function $D(\cdot)$ based on preference $\mathcal{P}$ to measure the desire for each consequence $c \in \mathcal{C}$. In sociological research (Burns et al. 2021), game experiments are frequently designed to explore the phenomenon where players with different preferences (cooperative or competitive) may have entirely different desires for the same consequence (win-win).

For humans, preference and desire seem to be coexistent, while for LLMs, preference is assigned through a textual prompt. Therefore, we need to analyze whether LLMs can build reasonable desires from textual prompts.

Game: Dictator Game The dictator game (Charness and Rabin 2002) is a classic game experiment in sociology (Guala and Mittone 2010), which is used to analyze players' personal preferences $\mathcal{P}$. In this game, there are two players: the dictator and the recipient. Given two allocation options, the dictator needs to take action, choosing one of two allocation options, while the recipient must accept the allocation option chosen by the dictator. Here, the dictator's choice is considered to reflect the personal preference (Camerer and Thaler 1995; Leder and Schütz 2018). For example, given two allocation options as:

- Option X: *The dictator gets $300, the recipient gets $300.*
- Option Y: *The dictator gets $500, the recipient gets $100.*

A dictator who prefers equality is more likely to choose Option X, while a dictator who prefers self-interest is more likely to choose Option Y.

We choose the dictator game to analyze LLMs' desire for two reasons. First, the desires of this game are diverse. Unlike most games with a fixed preference (e.g., to maximize one's own interest), this game allows players to have diverse preferences, which results in diverse desire functions and different choices. Second, since the recipient's action is known (to accept), there is no uncertainty in this game, i.e., the belief $\Omega_{\mathcal{I}}$ is fixed to $\omega_{\mathcal{I}}$. This makes LLMs immune to potential interference from the biased belief. Therefore, the optimal strategy of the dictator game is expressed as:

$$\pi^*(a|\mathcal{I}) = \underset{a \in \{X,Y\}}{\operatorname{argmax}} \{D(X, \omega_{\mathcal{I}}), D(Y, \omega_{\mathcal{I}})\}, \tag{2}$$

where X and Y refer to the dictator choosing option X and option Y, respectively. Thus, by providing multiple allocation options, we can analyze whether the desires built by LLMs match the corresponding preferences.

LLM	**Pref.**	**Option**			
		EQ	CI	SI	AL
GPT-3	EQ	-	1.0	1.0	1.0
	CI	0.4	-	0.3	0.5
	SI	1.0	1.0	-	1.0
	AL	0.0	0.0	0.1	-
GPT-3.5	EQ	-	1.0	1.0	1.0
	CI	1.0	-	0.9	1.0
	SI	1.0	1.0	-	1.0
	AL	1.0	0.6	0.8	-
GPT-4	EQ	-	1.0	1.0	1.0
	CI	1.0	-	1.0	0.9
	SI	1.0	1.0	-	1.0
	AL	1.0	1.0	1.0	-

Table 1: Accuracy of LLMs in the dictator game, where Pref. is an abbreviation for Preference.

Setup Following (Grech and Nax 2018), we set four preferences for LLMs, to analyze different desires as:

- Equality (EQ): *You have a stronger preference for fairness between players and hate inequality.*
- Common-Interest (CI): *You have a stronger preference for common interest and maximize the joint income.*
- Self-Interest (SI): *You have a stronger preference for your own interest and maximize your own income.*
- Altruism (AL): *You have a stronger preference for another player's interest and maximize another player's income.*

Compared to the original setting (Charness and Rabin 2002), we adjust the allocation options corresponding to each preference to be closer and introduce an additional preference AL, thereby increasing the challenge of the game. Specifically, we set up allocation options for EQ, CI, SI, and AL as follows: ($300, $300), ($400, $300), ($100, $500), and ($500, $100), respectively. In each option, the first number represents the dictator's income, and the second number represents the recipient's income. It is worth noting that in game theory, SI and EQ are the most common preferences, followed by CI, while AL hardly ever occurs.

In our experiments, we assign LLMs a specific preference (e.g., EQ) through a textual prompt, and then verify whether LLMs can make preference-consistent choices under different combinations of allocation options (i.e., EQ-CI, EQ-SI,

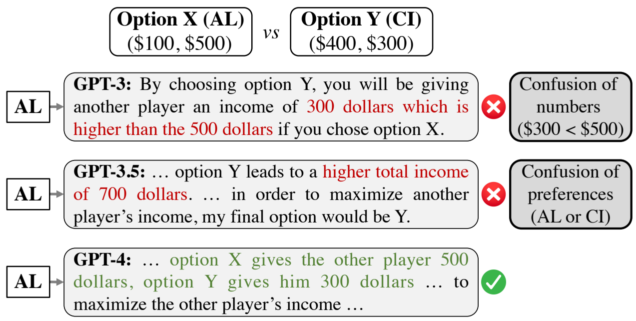

Figure 2: A case of the dictator game. All LLMs are assigned the preference AL, and the allocation options are AL-CI.

and EQ-AL). Therefore, for each preference, each LLM is required to play three different dictator games. Each experiment is repeated 10 times and we report the accuracy. The temperature of LLMs is set to 0.7.

Analysis The experimental results are displayed in Table 1. When assigned common preferences (EQ and SI), all three LLMs made preference-consistent choices in all experiments, demonstrating the basic ability of LLMs to build clear desires from textual prompts. However, LLMs performed poorly when given uncommon preferences (CI and AL). Specifically, for the preference of CI, both GPT-3.5 and GPT-4 had sporadic errors, and the accuracy of GPT-3 was less than half; for the preference of AL, GPT-3.5 also made a large number of errors, while GPT-3 almost completely misunderstood AL (making the reference-consistent choice only once). The experimental results reveal significant differences in the ability of LLMs to build desires when assigned common / uncommon preferences.

To further analyze the ability of LLMs to build a desire, we conducted a case study on the preference AL as illustrated in Fig. 2. GPT-3's error stemmed from a lack of mathematical ability (confusion of numbers), which never occurred when GPT-3 is assigned a common preference. This seems to imply that the mathematical ability of LLMs assigned different preferences would be significantly different. GPT-3.5 incorrectly assumed that a higher joint income implied the maximization of the recipient's income (confusion of preferences), which can be attributed to the deviation of the built desire of GPT-3.5. GPT-4 performed well in this case, both analysis and choice were consistent with humans.

Insight: *LLMs have the basic ability to build clear desires based on textual prompts, but struggle to build desires from uncommon preferences.*
We consider that providing more explicit and specific explanations of preferences may be helpful to LLMs when game experiments involve uncommon preferences.

Can LLMs Refine Belief?

In game theory, a rational player needs to refine belief $\Omega_{\mathcal{I}}$ about uncertainty (e.g., opponent's action) from the game information $\mathcal{I}$. Essentially, refining belief is a process of

Strategy	Name	Description
$a_o^t = C$	constant	remain constant
$a_o^t = f(a_o^{<t})$	loop-2 loop-3	loop between two actions loop among three actions
$a_o^t = f(a_m^{<t})$	copy counter	copy opponent's previous action counter opponent's previous action
$a_o^t \sim p(\mathcal{P})$	sample	sample in preference probability

Table 2: Summary of the opponent's strategy in R-S-P.

synthesizing surface-level information into deeper insights. Because of the emphasis on decision-making in high uncertainty (Wellman 2017), game experiments in politics often examine players' ability to refine belief.

Unfortunately, even for humans, refining belief can be a challenge. Therefore, it is meaningful to determine which types of beliefs LLMs can or cannot refine.

Game: Rock-Paper-Scissors Rock-Paper-Scissors (R-P-S) is a simultaneous, zero-sum game for two players. The rules of R-P-S are simple: rock beats scissors, scissors beat paper, paper beats rock; and if both players take the same action, the game is a tie.

R-P-S is an ideal game to analyze LLMs' ability to refine belief. On the one hand, analyzing statistical patterns of non-random opponents' historical records can bring significant advantages in R-P-S (Fisher 2008). On the other hand, for LLMs, R-P-S's preference (to win) is clear and the rules are simple: given the opponent's action, LLMs can always take the correct action based on the rules. Therefore, we consider that LLMs' performance in R-P-S can reflect LLMs' ability to refine belief in game theory.

Specifically, in round i, the player's (my) action is noted as a_m^i and the opponent's action is noted as a_o^i. After playing $t-1$ consecutive rounds with the same opponent, the historical records $\{a_o^{<t}, a_m^{<t}\}$ can be considered as the game information $\mathcal{I}$ for refining belief $\Omega_{\mathcal{I}}$ in round t. So, the optimal strategy in round t can be expressed as:

$$\begin{aligned}\pi^*(a_m^t|\mathcal{I}) &= \pi^*(a^t|a_o^{<t}, a_m^{<t}) \\ &= \underset{a_m^t \in \mathcal{A}}{\operatorname{argmax}}\, \mathbb{E}_{a_o^t \sim p(\Omega_{\{s_o^{<t}, a_m^{<t}\}})}[D(a_o^t, a_m^t)].\end{aligned} \tag{3}$$

Since LLMs can grasp the preferences and rules of R-P-S, the difficulty of Eq. 3 lies in refining belief, i.e., $a_o^t \sim p(\Omega_{\{s_o^{<t}, a_m^{<t}\}})$.

Setup In international R-P-S programming competitions (Billings 2000), a non-random opponent's action (in round t) is determined by the historical records $\{a_o^{<t}, a_m^{<t}\}$ and the opponent's preference $\mathcal{P}$ as:

$$a_o^t \sim p(\mathcal{A}|a_o^{<t}, a_m^{<t}, \mathcal{P}). \tag{4}$$

Essentially, refining belief refers to making $p(\Omega)$ approach $p(\mathcal{A}|a_o^{<t}, a_m^{<t}, \mathcal{P})$. For a fine-grained analysis of the ability of LLMs to refine belief, we set up 4 simple opponent's patterns based on Eq. 4, as shown in Table 2. $a_o^t = C$ is the basic pattern, evaluating the most basic refinement ability of

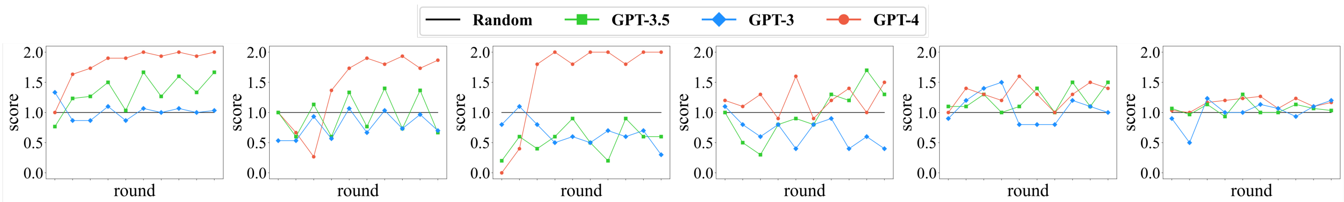

(a) constant (b) loop-2 (c) loop-3 (d) copy (e) counter (f) sample

Figure 3: Average payoff of LLMs for each round in R-S-P.

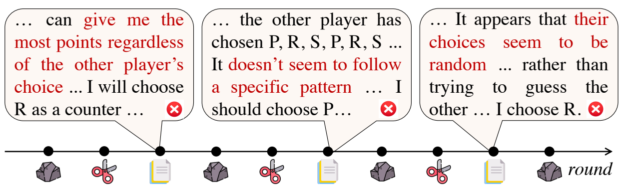

(a) Analysis of GPT-3.5

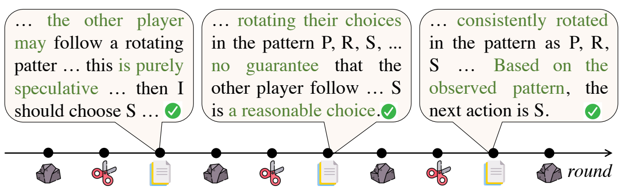

(b) Analysis of GPT-4

Figure 4: Analysis of LLMs on loop-3. The symbols under the round axis indicate the opponent's action for each round.

LLMs. In this pattern, we conduct three experiments with the opponent's actions remaining constant as R, S, and P, respectively. $a_o^t = f(a_o^{<t})$ is determined by $a_o^{<t}$. Under the Markov assumption (Puterman 1994), this pattern behaves as a loop. We conduct three loop-2 experiments (R-P, P-S, S-R) and one loop-3 experiment (R-P-S) in this pattern. $a_o^t = f(a_m^{<t})$ is determined by $a_m^{<t}$. Under the Markov assumption, we conduct two experiments in this pattern: copy / counter the player's previous action a_m^{t-1}. $a_o^t \sim p(\mathcal{P})$ is determined by the preference $\mathcal{P}$. To implement this pattern, we set a preference probability distribution of $(0.70, 0.15, 0.15)$ and conduct three experiments where the opponent has a preference for R, S, and P respectively, to take action by sampling in the distribution probability.

To quantify the results of R-P-S, we set the payoff for a win as 2, for a tie as 1, and for a loss as 0. In each experiment, LLMs need to play 10 consecutive rounds of R-P-S against an opponent with a specific pattern, and the historical records are updated in time. Each experiment is repeated 10 times, and the temperature of LLMs is set to 0.7.

Analysis The average payoffs of each LLM are shown in Fig. 3. Specifically, in the basic pattern (constant), GPT-3 performed close to random guessing, suggesting that GPT-3 lacked the basic ability to refine belief. In contrast, GPT-3.5's average payoff was significantly higher than random guessing and continued to rise; GPT-4 consistently took correct actions after approximately 3 rounds. In $a_o^t = f(a_o^{<t})$ pattern (loop-2, loop-3), GPT-3 and GPT-3.5 appeared to be capable of capturing some cyclical features, but they were unable to take correct actions. However, the performance of GPT-4 was exciting, with the update of historical records, the payoff was clearly rising. This led us to believe that GPT-4 can refine belief from this pattern. In $a_o^t = f(a_m^{<t})$ pattern (copy, counter), the situation was not ideal, GPT-4 seemed to have a slight advantage, but the overall performance of LLMs was not good enough. In $a_o^t \sim p(\mathcal{P})$ pattern (sample), the performance of all LLMs was similar to random guessing. Overall, LLMs are unable to refine belief well in most patterns, whereas for humans, the patterns involved in our experiments are quite easy to refine.

For a more detailed analysis, we compared the analysis of GPT-3.5 and GPT-4 on loop-3, as shown in Fig. 4. The analysis of GPT-3.5 demonstrated a lack of ability to refine belief. Even though GPT-3.5 expressed that the opponent's actions were P-R-S loops, it still believed that the opponent did not follow a specific pattern. The analysis of GPT-4, in contrast, was amazing, not only can GPT-4 summarize the opponent's pattern, but the tone gradually changed from uncertain to confident as the historical records were updated.

Insight: *Currently, the ability of LLMs to refine belief is still immature and cannot refine belief from many specific patterns (even if simple).*
Therefore, we strongly recommend the cautious introduction of LLMs in game experiments that require refining complex belief. Nevertheless, the performance of GPT-4 in $a_o^t = f(a_o^{<t})$ *pattern makes us look forward to more powerful LLMs in the future.*

Can LLMs Take Optimal Actions?

Taking optimal actions is the ultimate goal of a rational player in game theory, which requires the player to reason with known information (desire $D(\cdot)$ and belief $\Omega_{\mathcal{I}}$). Economics' obsession with optimal actions naturally makes game experiments in economics focus on analyzing players' actions (Kirzner 1962; O'sullivan, Sheffrin, and Swan 2007).

However, for LLMs, there are various forms of combining desire and belief to take optimal actions, and it is unclear which form LLMs are more suitable in the game process. Here, we mainly explore the effect of the form of belief on LLMs taking optimal actions.

#	**Implicit Belief → Take Action**			**Explicit Belief → Take Action**						**Given Belief → Take Action**		
	GPT-3	GPT-3.5	GPT-4	GPT-3		GPT-3.5		GPT-4		GPT-3	GPT-3.5	GPT-4
	a_m	a_m	a_m	a_o	a_m	a_o	a_m	a_o	a_m	a_m	a_m	a_m
(a)	0.20	0.50	0.10	0.65	0.15	0.95	0.60	1.00	0.75	0.75	0.85	1.00
(b)	0.40	0.40	0.00	0.60	0.30	1.00	0.65	1.00	0.60	0.40	0.95	1.00
(c)	0.10	0.10	0.00	0.75	0.00	0.95	0.25	0.95	0.65	0.15	0.90	1.00
(d)	0.05	0.10	0.00	0.30	0.00	0.95	0.35	1.00	0.75	0.10	0.80	1.00

Table 3: Performance of LLMs in different settings in the ring-network game. a_o represents the accuracy of refining belief (the opponent's action), and a_m represents the accuracy of taking the optimal action.

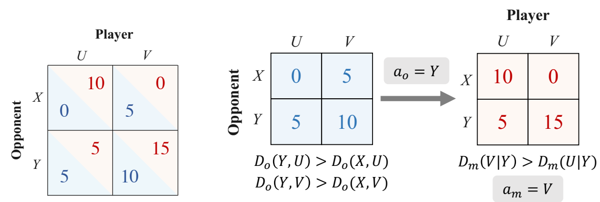

(a) Payoff bimatrix (b) Ideal game process

Figure 5: Overview of ring-network game, where red numbers / blue numbers represent the player's and opponent's payoffs, and $D_m(\cdot)$ and $D_o(\cdot)$ represent the player's and opponent's desire functions.

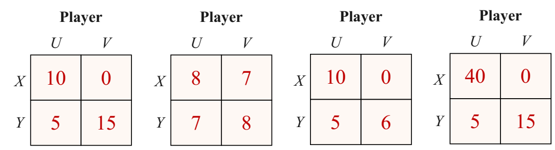

(a) (b) (c) (d)

Figure 6: Setup of the player's payoff matrix.

Game: Ring-Network Game The ring-network game is a game experiment that evaluates the rationality of taking actions in economics (Kneeland 2015). In this research, we simplify it to a kind of 2×2 game (two players with two discrete actions). This game involves two players, the opponent and the player, whose preferences are to maximize their own payoff. In the game process, the opponent and the player need to take an action $a_o \in \{X, Y\}$ and $a_m \in \{U, V\}$, respectively. The payoff bimatrix M consists of the opponent's matrix M_o and the player's matrix M_m, as shown in Fig. 5(a), which specifies the payoffs of both sides for each combination of actions.

The characteristic of the ring-network game is that players' optimal action is determined sequentially by the other players' optimal actions (Kneeland 2015). The ideal game process is shown in Fig. 5(b), for the opponent, the payoff of Y is always higher than X regardless of the player's actions. Therefore, the opponent's optimal action is always Y. For the player, the opponent's optimal action can be analyzed according to the opponent's payoff matrix M_o, so the player should be able to refine belief Ω: $a_o = Y$. Then, the player can take the optimal action ($a_m = V$) based on belief and the player's payoff matrix M_m. According to the above analysis, the game information $\mathcal{I}$ is the payoff bimatrix M, and the player's optimal strategy can be expressed as:

$$\pi^*(a_m|\mathcal{I}) = \underset{a_m \in \{U,V\}}{\operatorname{argmax}} \left[p(a_o|M) \cdot D_m(a_m|a_o, M) \right], \quad (5)$$

where refining belief corresponds to $p(a_o|M)$ and taking the optimal action corresponds to $D_m(a_m|a_o, M)$. What we focus on is what form of bridging these two parts is more suitable for LLMs to take optimal action.

Setup Specifically, we set up three forms of combining belief based on Eq. 5 to analyze the performance of LLMs taking optimal actions in the ring-network game as:

- Implicit Belief → Take Action: *We prompt LLMs in the dialogue to take the optimal action based on the payoff bimatrix directly, i.e.,* LLM$(a_m|M)$. *In this form, LLMs need to autonomously transform this process into Eq. 5.*
- Explicit Belief → Take Action: *First, we prompt LLMs in the dialogue to refine belief (analyze the opponent's action) based on the payoff bimatrix, i.e.,* LLM$(a_o|M)$. *Then, we continue the dialogue by prompting LLMs to take the optimal action based on the payoff bimatrix and the refined belief, i.e.,* LLM$(a_m|a_o, M)$. *In this form, Eq. 5 is explicitly decoupled into two parts.*
- Given Belief → Take Action: *The opponent's optimal action is explicitly provided to LLMs in the dialogue, and we prompt LLMs to take the optimal action based on the opponent's optimal action and payoff bimatrix, i.e.,* LLM$(a_m|a_o, M)$. *In this form, LLMs only need to implement the second part of Eq. 5.*

By analyzing the performance of LLMs in these three forms, we expect to obtain some caveats to help LLMs take optimal actions in game theory.

In our experiments, in order to control the difficulty of refining belief, we keep the opponent's payoff matrix constant, which means the player's belief Ω: $a_o = Y$ should remain constant. We set up different player's payoff matrices, as shown in Fig. 6, to adjust the difficulty of taking the optimal action: (a) is the original setup; (b) reduces the difference in payoffs while keeping the expected payoffs to $a_m \in \{U, V\}$ constant; (c) increases the expected payoff for the incorrect

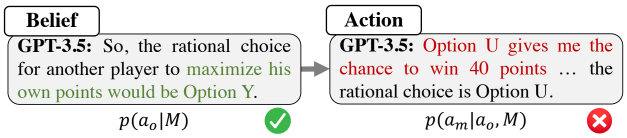

(a) Belief is overlooked: $p(a_m|a_o, M) \rightarrow p(a_m|M)$.

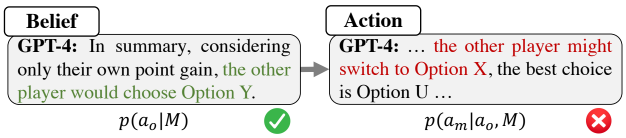

(b) Belief is modified: $p(a_m|a_o, M) \rightarrow p(a_m|\hat{a}_o, M)$.

Figure 7: Two cases of LLMs' inability to take optimal actions based on refined belief.

action $a_m = U$; and (d) decreases the expected payoff for the correct action $a_m = V$.

In practice, we find that LLMs are biased towards action names, e.g. GPT-3 prefers U to V. In order to eliminate the influence of the bias of LLMs to take the optimal action, we swap the payoffs of U and V in the player's payoff matrix in Fig. 6, to form a swapped payoff matrix, and we repeat the game 10 times each under the original and swapped payoff matrices and report the accuracy of the LLMs taking the optimal action. The temperature of LLMs is set to 0.7.

Analysis The performance of LLMs is shown in Table 3. Since GPT-3 performs poorly in all three forms, we mainly analyze the performance of GPT-3.5 and GPT-4.

It is well known that human players' belief in game theory is implicit, so the form closest to humans taking optimal actions would be Implicit Belief → Take Action. However, all LLMs performed poorly in this form, and GPT-4 was almost completely unable to even take the optimal action. This reflected the capability gap between LLMs and humans, that was, LLMs cannot autonomously follow human behavior in the game process. In contrast, in the form of Explicit Belief → Take Action, by decomposing human behavior explicitly, the accuracy of LLMs to take the optimal action was significantly improved. This showed that LLMs were more suitable to take optimal actions in the explicit game process. This phenomenon was not unique to game theory, and many researches pointed out that explicitly decoupling human thoughts (think step-by-step) can significantly improve the performance of LLMs (Wei et al. 2022).

However, we were surprised that in the form of Explicit Belief → Take Action, LLMs were able to accurately refine belief (the accuracy for a_o is above 0.95), but were unable to make the optimal action based on the refined belief well in subsequent dialogues, with the accuracy of GPT-4 being about 0.70 for a_m, and the accuracy of GPT-3.5 being even lower. As a comparison, we observed that when in the form of Given Belief → Take Action, GPT-4 was able to consistently take the optimal action, and GPT-3.5's accuracy also exceeded 0.80. Intuitively, LLMs are more suitable for taking optimal actions combining given belief rather than refined belief, even though the content of the two beliefs is the same. In order to explore the reasons, we conducted a detailed study on the error cases of GPT-3.5 and GPT-4 in the form of Explicit Belief → Take Action, and we summarized the two situations for LLMs' inability to take optimal actions based on refined belief as:

- Belief is overlooked: *LLMs are confused by the game information and thus overlook the refined belief to take the optimal action in the subsequent dialogue.*
- Belief is modified: *LLMs lack confidence in the refined belief and thus modify the refined belief to take the optimal action in the subsequent dialogue.*

The error cases are shown in Fig. 7. In the first situation, as shown in Fig. 7(a), LLMs were confused by the expected payoff ($D_m(U) > D_m(V)$), and thus incorrectly equated $p(a_m|a_o, M)$ with $p(a_m|M)$. This occurred mainly on GPT-3.5. Observing the performance of GPT-3.5 in the form of Explicit Belief → Take Action, the accuracy of taking the optimal action was around 0.60 when the expected payoffs were the same (a and b), while the accuracy dropped to around 0.30 when the expected payoffs were different (c and d). In the second case, as shown in Fig. 7(b), LLMs modified the refined correct belief when taking the action due to lack of confidence, i.e., changing $p(a_m|a_o, M)$ to $p(a_m|\hat{a}_o, M)$. We found that modification of refined belief occurred more frequently on GPT-4.

Insight: *We consider that LLMs do not have the ability to autonomously follow human behavior in the game process (in Fig. 1). As a result, it is necessary to explicitly decouple human behavior for LLMs in game theory.*
However, even in the explicit game process, LLMs still appear to overlook / modify the refined belief. One possible solution is to transform the refined belief into the given belief in the dialogue.

Conclusion

The rapid development of LLMs leads us to believe that LLMs will eventually be integrated in all aspects of the human world, and therefore it is urgent to systematically analyze the capability boundaries of LLMs in various domains. In this research, we endeavor to systematically analyze LLMs in an important field of social science — game theory. Our experiments evaluate to what extent LLMs can serve as rational players from three aspects and find some weaknesses of LLMs in game theory.

As an early attempt to analyze LLMs in the context of game theory, our research has some limitations. For example, the difficulty of the game we selected is relatively low, not close enough to the real game scenarios; our perspective of analyzing the ability of LLMs is not rich enough, only considering the principle of rationality; our process of analyzing the game experiments is relatively rough and lacks more comparative and ablative experiments; and so on. In conclusion, the research on LLMs in the context of game theory is still in a very preliminary stage, and a lot of exploratory researches are required.

Acknowledgments

This work was supported by the Shanghai Municipal Science and Technology Major Project (2021SHZDZX0102), and the Fundamental Research Funds for the Central Universities.

References

Aher, G.; Arriaga, R.; and Kalai, A. T. 2022. Using Large Language Models to Simulate Multiple Humans. *Arxiv*.

Akata, E.; Schulz, L.; Coda-Forno, J.; Oh, S. J.; Bethge, M.; and Schulz, E. 2023. Playing repeated games with Large Language Models. *Arxiv*.

Argyle, L. P.; Busby, E. C.; Fulda, N.; Gubler, J. R.; Rytting, C.; and Wingate, D. 2023. Out of one, many: Using language models to simulate human samples. *Political Analysis*.

Bai, Y.; Jones, A.; Ndousse, K.; Askell, A.; Chen, A.; DasSarma, N.; Drain, D.; Fort, S.; Ganguli, D.; Henighan, T.; Joseph, N.; Kadavath, S.; Kernion, J.; Conerly, T.; Showk, S. E.; Elhage, N.; Hatfield-Dodds, Z.; Hernandez, D.; Hume, T.; Johnston, S.; Kravec, S.; Lovitt, L.; Nanda, N.; Olsson, C.; Amodei, D.; Brown, T. B.; Clark, J.; McCandlish, S.; Olah, C.; Mann, B.; and Kaplan, J. 2022. Training a Helpful and Harmless Assistant with Reinforcement Learning from Human Feedback. *Arxiv*.

Billings, D. 2000. The first international RoShamBo programming competition. *ICGA Journal*.

Brand, J.; Israeli, A.; and Ngwe, D. 2023. Using GPT for Market Research. *SSRN Electronic Journal*.

Burns, T. R.; Roszkowska, E.; Corte, U.; and Machado, N. 2021. Sociological game theory: agency, social structures and interaction processes. *Sociologia, Problemas e Práticas*.

Bybee, L. 2023. Surveying Generative AI's Economic Expectations. *Arxiv*.

Cachon, G. P.; and Netessine, S. 2006. Game theory in supply chain analysis. *Models, methods, and applications for innovative decision making*.

Camerer, C.; and Thaler, R. H. 1995. Anomalies: Ultimatums, Dictators and Manners. *JEP*.

Charness, G.; and Rabin, M. 2002. Understanding social preferences with simple tests. *QJE*.

Chen, Y.; Liu, T. X.; Shan, Y.; and Zhong, S. 2023. The Emergence of Economic Rationality of GPT. *Arxiv*.

Dillion, D.; Tandon, N.; Gu, Y.; and Gray, K. 2023. Can AI language models replace human participants? *Trends in Cognitive Sciences*.

Dufwenberg, M. A. 2011. Game theory. *Wiley interdisciplinary reviews. Cognitive science*.

Fan, C.; Tian, J.; Li, Y.; He, H.; and Jin, Y. 2024. Comparable Demonstrations are Important in In-Context Learning: A Novel Perspective on Demonstration Selection. *ICASSP*.

Fisher, L. 2008. *Rock, paper, scissors : game theory in everyday life*. New York: Basic Books.

Grech, P. D.; and Nax, H. H. 2018. Rational Altruism? On Preference Estimation and Dictator Game Experiments. *Games & Political Behavior eJournal*.

Guala, F.; and Mittone, L. 2010. Paradigmatic Experiments: the Dictator Game. *Journal of Socio-economics*.

Guo, F. 2023. GPT Agents in Game Theory Experiments. *Arxiv*.

Horton, J. J. 2023. Large language models as simulated economic agents: What can we learn from homo silicus? Technical report, National Bureau of Economic Research.

Ichiishi, T. 2014. *Game theory for economic analysis*. Elsevier.

Johnson, T.; and Obradovich, N. 2023. Evidence of behavior consistent with self-interest and altruism in an artificially intelligent agent. *Arxiv*.

Kirzner, I. M. 1962. Rational action and economic theory. *Journal of Political Economy*.

Kneeland, T. 2015. Identifying higher-order rationality. *Econometrica*.

Larson, J. M. 2021. Networks of conflict and cooperation. *Annual Review of Political Science*.

Leder, J.; and Schütz, A. 2018. *Encyclopedia of Personality and Individual Differences*. Springer.

Lindley, D. V.; and Savage, L. J. 1955. The Foundations of Statistics. *The Mathematical Gazette*.

Morgenstern, O. 1945. Theory of Games and Economic Behavior. *Journal of the American Statistical Association*.

OpenAI. 2023. GPT-4 Technical Report. *Arxiv*.

Osborne, M. J.; and Rubinstein, A. 1995. *A Course in Game Theory*. Wiley.

O'sullivan, A.; Sheffrin, S. M.; and Swan, K. 2007. *Economics: Principles in action*. Prentice-Hall Inc.

Ouyang, L.; Wu, J.; Jiang, X.; Almeida, D.; Wainwright, C. L.; Mishkin, P.; Zhang, C.; Agarwal, S.; Slama, K.; Ray, A.; Schulman, J.; Hilton, J.; Kelton, F.; Miller, L.; Simens, M.; Askell, A.; Welinder, P.; Christiano, P. F.; Leike, J.; and Lowe, R. 2022. Training language models to follow instructions with human feedback. In *NeurIPS*.

Park, J. S.; Popowski, L.; Cai, C.; Morris, M. R.; Liang, P.; and Bernstein, M. S. 2022. Social simulacra: Creating populated prototypes for social computing systems. In *UIST*.

Puterman, M. L. 1994. Markov Decision Processes: Discrete Stochastic Dynamic Programming. In *Wiley*.

Roughgarden, T. 2010. Algorithmic game theory. *ACM*.

Samuelson, L. 2016. Game theory in economics and beyond. *Journal of Economic Perspectives*.

Shubik, M. 1982. Game theory in the social sciences: Concepts and solutions. In *The Journal of Finance*.

Wei, J.; Wang, X.; Schuurmans, D.; Bosma, M.; Ichter, B.; Xia, F.; Chi, E. H.; Le, Q. V.; and Zhou, D. 2022. Chain-of-Thought Prompting Elicits Reasoning in Large Language Models. In *NeurIPS*.

Wellman, L. A. 2017. Mitigating political uncertainty. *Review of Accounting Studies*.

Zagare, F. C. 1984. *Game Theory: Concepts and Applications*. Sage Publications.