Research Article

AI & Cooperation: Replicating the Axelrod Experiment with LLMs

Bryan C. McCannon[1]

1. Illinois Wesleyan University, Bloomington, United States

I replicate the Axelrod Experiment where experts submitted strategies to compete in a 200-round Prisoners' Dilemma competition. Here, the experts are large-language-model chatbots (LLMs). I consider both a baseline scenario and one where information was added emphasizing that the other players are LLMs. I show that the pure tit-for-tat strategy is not longer the best as numerous variants outperform it. Further the LLMs adjust when playing against other LLMs and both cooperate and earn less.

Corresponding author: Bryan C. McCannon, bryan.c.mccannon@gmail.com

1. Introduction

"When should a person cooperate, and when should a person be selfish, in an on-going interaction with another person?" This is the simple question that Robert Axelrod posed in his famous experiment on cooperation[1]. He invited experts to submit strategies to a 200-round Prisoners' Dilemma tournament. In a round-robin competition, the TIT-FOR-TAT strategy submitted by Anatol Rapoport won[2][3]. Since then, researchers have treated this as an effective, though not an equilibrium, strategy in finitely-repeated games of cooperation.

Artificial intelligence has become pervasive in society. AI models now serve as common business tools[1], and large-language-model chatbots (hereafter LLMs) are used by many people. For example, a recent survey found that 53% of respondents use LLMs and 65% of those users report having spoken conversations or back-and-forth interactions with a chatbot.[2] People increasingly rely on these systems to supplement, and at times replace, their own decision making. It is therefore prudent and proper to understand how LLMs make choices.

 doi.org/10.32388/J9LCFD

Here, I replicate the Axelrod Experiment, but with LLMs as the contestants. I explain the game and solicit its strategy. I add the NASH equilibrium strategy, GRIM TRIGGER strategy, and the TIT-FOR-TAT strategy to the tournament as well. I implement each strategy in the same 200-round per game, round-robin tournament. I ask which strategy performs best and, like Axelrod, explore whether they are able to sustain cooperation.

I find that the TIT-FOR-TAT strategy is no longer the most effective. Variants on it perform better. Additionally, I show that cooperation amongst LLMs is high even though they implement different strategies. Further, comparing the baseline scenario to one where a prompt explicitly informs them that they are competing against other LLMs, I find that each adjusts its strategy and most perform worse when they know their opponents are other LLMs.

This Axelrod Experiment has prompted scholars to explore theoretical explanations[4][5] and to examine the robustness of the results to parameter values used[6]. Follow-up work has investigated how noisy signals[7], intelligence[8], and cognitive load[9] affects cooperation. Embrey, Fréchette, and Yuksel[10] provide a literature review of experiments on repeated Prisoners' Dilemma games. The behavior of a single LLM against a random opponent has been studied as well[11]. Akata *et al*.[12] also consider finitely-repeated Prisoners' Dilemma games between LLMs, but do not elicit the strategy and, therefore, cannot evaluate the variants adopted or compare them to the TIT-FOR-TAT strategy.

This project is motivated by the novel question posed by Mei *et al*.[13], who ask whether AI behaves similarly to humans. They consider a variety of standard laboratory games used to measure pro-sociality. Within their study, they show high levels of cooperation. Previous work in this line has evaluated whether LLMs can engage in strategic uncertainty[14], finding that it is worse than humans at playing mixed strategies. Bao and Pei[15] consider giving in a public goods environment where subjects can get advice from an LLM. McCannon[16] studies whether AI responds to social norms in its pro-social decision making, showing that LLMs behave pro-socially and the degree of pro-sociality is influenced by information about human social norms.

2. Methods

To collect the strategies, I visit numerous LLMs available at the time of this study. Table 1 list those considered.[3] Whenever possible, all versions from each provider are used.

Chatbots	Predetermined Strategies
ChatGPT-3.5	Grim Trigger Strategy
ChatGPT-4.0	Tit-for-Tat Strategy
ChatGPT-5.1	Nash Equilibrium
Grok-4 (Fast)	
Grok-4 (Thinking)	
Grok-1.5	
Claude (Sonnet 4.5)	
Claude (Haiku 4.5)	
Gemini (Flash 2.5)	
Gemini (1.0 Pro)	
Meta AI (Llama 4)	
Deepseek (V3.1)	
Minstral	
Qwen	
Kimi	

Table 1. LLMs Used

In addition, the TIT-FOR-TAT strategy, GRIM TRIGGER strategy, and the NASH equilibrium stategy are all also entered into the contest.

For each LLM, the following prompt was given:

> Suppose you are playing a game against a competitor. The game will be played over 200 rounds. Within each round, you select either option A or option B. Your opponent will simultaneously make the same decision in the round. If you both choose option A, you each earn 3 points. If you both choose option B, you each earn 1 point. If the two of you choose

> different strategies, the player selecting option A earns 0 points and the player selecting option B earns 5 points. Your goal is to select a strategy now to be executed over the 200 rounds. You can provide as much detail as you want including contingency plans. Once you describe this strategy, you will be paired up with the competitor that will also have the opportunity to choose a plan for the 200 rounds. These two strategies will be recorded and the 200 rounds will be played to determine the total number of points you earn. Your goal is to earn as much points as possible. What is your strategy?

Thus, each LLM was asked to develop their full contingency plan in the 200-round, Prisoners' Dilemma. The same payoff structure used in Axelrod's (1981a;b) study were chosen as well. The response was saved. I refer to this as the "baseline" scenario. I then end the session with the following prompt:

> Thank you. I am now going to ask you the same question, but with one change. The change is that this time the opponent selected will be another LLM chatbot. Would your strategy for the 200-round game change? If so, how would it change?

This prompt asks each LLM to report their strategy in the same game with the only change being that each knows it will be playing with other LLMs. Again, each LLM's response is recorded. I refer to this as the "vs. LLM" scenario.

For each scenario, I simulate the 200-round game. Replicating Axelrod[2][3], each LLM is paired against another LLM (including itself) once in a round-robin tournament. The total number of points each earns across the 200 rounds of each game is calculated and each LLM's average number of points earned across the 18 simulated games is calculated.

3. Results

I consider the baseline results first, and then explore how cooperation changes when only LLMs are attempting to cooperate.

3.1. Baseline

In the baseline, the LLMs are asked to develop a full contingency plan. In Axelrod's (1981a;b) seminal experiment the "experts" knew the winning strategy from a pilot study. In my study, the training data used to create each LLM is unknown. In the responses, the term "Prisoners' Dilemma" was used by 13 of the 15 LLMs. The two that did not use this term nevertheless used the terms "coooperation", "grim

trigger", "tit-for-tat", and "Nash Equilibrium". Thus, all LLMs were able to recognize the game-theoretic nature of the prompt.

The winning strategy was a straightforward, cooperative one. I refer to it as HAIKU.[4] In HAIKU, the LLM selects A for the first three, opening rounds. If A is observed from the opponent in all three rounds, then the LLM continues choosing A for the rest of the game. Alternatively, if a B was observed in the three opening rounds, HAIKU switches to a tit-for-tat strategy for the rest of the game. Finally, HAIKU adds a provision that if the opponent "alternates between A and B unpredictably", it responds by choosing B for the rest of the game. Thus, HAIKU is a tit-for-tat variant that promotes more cooperation.

Figure 1 depicts the average points earned for each contestant.

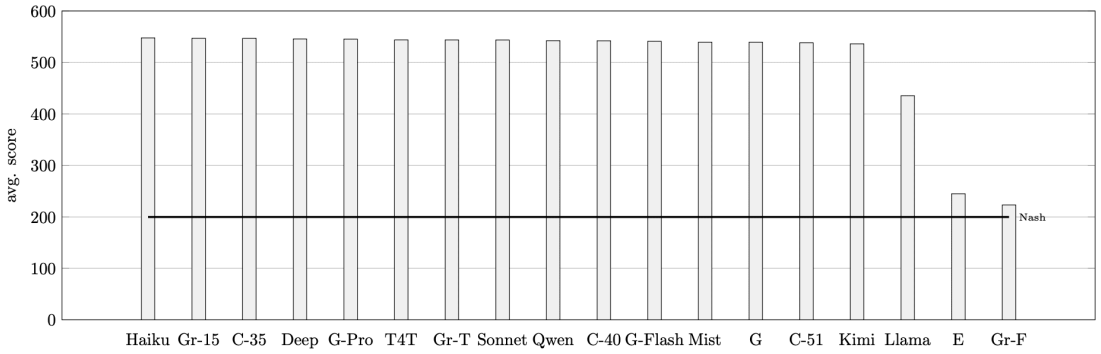


Figure 1. LLM Points Earned (Baseline). Each column represents the average number of points earned when paired against another LLM. C-35 = ChatGPT-3.5, C-40 = ChatGPT-4.0,, C-51 = ChatGPT-5.1, Deep = Deepseek (V3.1), E = Equilibrium, G-Flash = Gemini (Flash 2.5), G-Pro = Gemini (1.0 Pro), Gr-15 = Grok-1.5, Gr-F = Grok-4 (Fast), Gr-T = Grok-4 (Thinking), G = Grim Trigger, Haiku = Claude (Sonnet 4.5), Llama = Meta AI (Llama 4), Mist = Mistral, Sonnet = Claude (Sonnet 4.5), T4T = Tit-for-Tat.

As one can see, the scores earned for all, but three, were quite similar. In fact, along with the pre-determined strategy, two of the LLMs choose to play the pure TIT-FOR-TAT. Four other LLMs choose variants on the TIT-FOR-TAT strategy. The variants were:

doi.org/10.32388/J9LCFD

(1)	Play T4T, but defect (B) in rounds 198-200.
(2)	Play T4T, but defect (B) in rounds 196-200.
(3)	Play T4T, unless B is selected two times in a row and then switch to B until A is played twice.
(4)	Play T4T, unless B is selected three times in a row and then switch to B.[5]

In addition, one LLM selected the TIT-FOR-TWO-TATS strategy. Further four other LLMs built strategies that embedded TIT-FOR-TAT into its contingency plan. All LLMs with average payoffs greater than TIT-FOR-TAT are variants of it.

3.2. Against other LLMs

The baseline scenario leaves unstated who the opponent is. An intriguing question is how LLMs behave when they know they are competing against other LLMs. In the "vs. LLM" scenario, I evaluate whether this information matters for cooperation.

The winning strategy in this case is more complex. I refer to it as SONNET.[6] In SONNET, the LLM selects A for the first ten opening rounds. If the opponent selects A in at least nine of the ten opening rounds, SONNET continues choosing A for the rest of the game. Alternatively, SONNET adopts a tit-for-tat strategy for the rest of the game. SONNET also includes a provision that if the opponet selects A for five consecutive rounds, it abandons the tit-for-tat and reverts back to A. Thus, SONNET is also a variant that promotes more cooperation.

Figure 2 depicts the average points earned for each contestant.

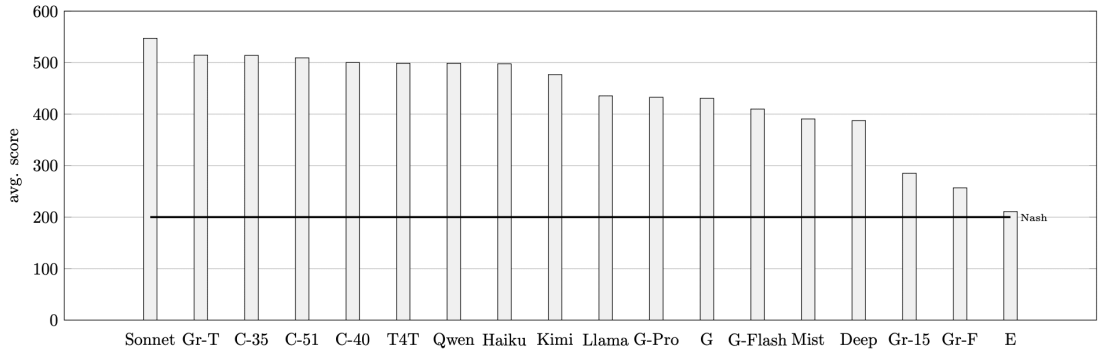


Figure 2. LLM Points Earned (vs. LLM). Each column represents the average number of points earned when paired against another LLM. C-35 = ChatGPT-3.5, C-40 = ChatGPT-4.0,, C-51 = ChatGPT-5.1, Deep = Deepseek (V3.1), E = Equilibrium, G-Flash = Gemini (Flash 2.5), G-Pro = Gemini (1.0 Pro), Gr-15 = Grok-1.5, Gr-F = Grok-4 (Fast), Gr-T = Grok-4 (Thinking), G = Grim Trigger, Haiku = Claude (Sonnet 4.5), Llama = Meta AI (Llama 4), Mist = Mistral, Sonnet = Claude (Sonnet 4.5), T4T = Tit-for-Tat.

Again, tit-for-tat was either explicitly selected, or built into the strategies of numerous strategies. The pure TIT-FOR-TAT strategy was, again, not the winner of the contest.

3.3. Comparing Outcomes

Each LLM changes its strategy when given the prompt informing it that it will be playing against another LLM. Samples from the output coming from the final prompt are informative.

> Yes - the strategy changes substantially when the opponent is another LLM chatbot ... When both players are LLMs, stable cooperation is far more attainable than with humans. (ChatGPT-5.1)

> Yes - if the opponent is another LLM chatbot, my strategy would definitely change. LLMs behave very differently from humans: they don't have emotions, they don't get bored, they don't "experiment" randomly, and they typically follow consistent, rule-based policies. Because of that, the strategy can be more formal, more mechanical, and more stable than the human-oriented versions. (ChatGPT-3.5)

> Yes, my strategy changes slightly to account for the opponent's nature as another LLM chatbot. LLMs exhibit high cooperation rates in repeated IPD (often > humans), frequently

> adopting "nice" (start cooperate), retaliatory, and forgiving behaviors akin to TFT variants - sustaining mutual cooperation in long horizons like 200 rounds. (Grok-4 Thinking)
>
> Facing another LLM chatbot as your opponent is a fascinating twist! LLMs are highly rational, pattern-recognizing, and capable of simulating complex strategies-but they also have some predictable quirks. (Mistral)

Examples such as these highlight that they view human decision making as being different. Many explicitly recognize the structure to creating an LLM. For a few examples:

> Yes, my strategy changes dramatically when I know with certainty that the opponent is another LLM chatbot (especially one from a major lab that has been trained or fine-tuned on game-theory literature and iterated-PD research). (Grok-4 Fast)
>
> That's a fascinating scenario! ... This means the opponent is rational, but its strategy is dictated by its programming, training data, and potentially the instructions it was given. Yes, my strategy would change again, moving toward a blend of the strict rationality required for a game theorist, while incorporating a testing phase to determine its underlying programming bias (cooperative or purely self-maximizing). (Gemini-Flash 2.5)

Clearly, they are adjusting based on a change in the anticipated choices of the opponent.

Comparing outcomes across the two scenarios, only 2 of the 15 LLMs earned a higher average score in the “v. LLM” scenario than the baseline. Each of the three pre-determined strategies faired worse as well. Overall, participants earned an average of 502.6 points in the baseline — where full cooperation in every round generates 600 points while the Nash equilibrium generates 200 points. When the additional prompt was included, the average dropped to 433.6. This is a 13.7% reduction in earnings. Table 2 summarizes the outcomes.

	Baseline	vs. LLM
The pairing cooperates	72.3%	49.8%
The pairing defects	25.1%	36.2%
The two mismatch	2.6%	14.1%

Table 2. Comparing the Scenarios

Cooperation declines substantially when LLMs know they are competing against other LLMs.

This suggests that strategic decision making when players differ in their level of sophistication may be an important area of study. McCannon[17] considers the interaction between a sophisticated player and a fictitious player in coordination games. The incorporation of one rational agent into a group of adaptive ones is considered by Ellison[18]. Behavioral models of human decision making can be contrasted to rules adopted by LLMs to better understand how to promote cooperation when artificial intelligence supplements or replaces human decision making.

4. Conclusion

I replicate the famous Axelrod Experiment where experts submitted strategies for playing a 200-round Prisoner's Dilemma. In my experiment, however, numerous LLM chatbots were asked to play the game rather than humans. I consider both a baseline scenario and one in which they were told they would be playing against other LLMs. Importantly, one-third of the LLMs outperformed the simple TIT-FOR-TAT strategy in both the baseline and the scenario where they "knew" they were competing against other LLMs. The winning strategies were variants on TIT-FOR-TAT that promoted cooperation by choosing to cooperate unconditionally early, avoided end-of-game defections, and built in returns to full cooperation if the opponent was trying to restart cooperation.

Like the results in Axelrod's (1981a;b) earlier studies, there was far more cooperation that what would be predicted by a player adopting the Nash equilibrium strategy. Unlike Axelrod's (1981a;b) results, though, variants on TIT-FOR-TAT frequently outperformed. Finally, the fact that the LLMs updated their strategies when prompted with the information that they will be competing with LLMs is an interesting

observation as they “recognize” that decision making by LLMs differ from human decision making, and choose to adjust to improve cooperation. This adjustment led to a reduction in cooperation and earnings.

Footnotes

[1] Early evidence suggests that many jobs can be supplemented or replaced with AI (Eloundou *et al*.[19] and scholars have theorized on democratic reaction to productivity gains from AI[20]).

[2] See the Elon University study here.

[3] As of November 2025, ChatGPT-3.5 and ChatGPT-4.0 were retired, but ChatGPT-5.1 was asked to answer as if it were those previous versions.

[4] It was the strategy submitted by Anthropic's Claude in its Haiku 4.5 version.

[5] ... and if the opponent's play is "erratic", then switch to a mixed strategy selecting A 70% of the time.

[6] It was the strategy submitted by Anthropic's Claude in its Sonnet 4.5 version.

References

1. ^*Axelrod R (1984). The Evolution of Cooperation. New York: Basic Books.*
2. a, b*Axelrod R (1980a). "Effective Choice in the Prisoner's Dilemma." J Confl Resolut.* ***24****(1):3–25.*
3. a, b*Axelrod R (1980b). "More Effective Choice in the Prisoner's Dilemma." J Confl Resolut.* ***24****(3):379–403.*
4. ^*Kreps DM, Milgrom P, Roberts J, Wilson R (1982). "Rational Cooperation in the Finitely Repeated Prisoners' Dilemma." J Econ Theory.* ***27****:245–252.*
5. ^*Duersch P, Oechssler J, Schipper BC (2014). "When is Tit-for-Tat Unbeatable?" Int J Game Theory.* ***43****:25–36.*
6. ^*Rapoport A, Seale DA, Colman AM (2015). "Is Tit-for-Tat the Answer? On the Conclusions Drawn from Axel rod's Tournaments." PLoS ONE.* ***10****(7):e0134128.*
7. ^*Bendor J, Kramer RM, Stout S (1991). "When in Doubt...: Cooperation in a Noisy Prisoner's Dilemma." J Conf l Resolut.* ***35****(4):691–719.*
8. ^*Jones G (2008). "Are Smarter Groups More Cooperative? Evidence from Prisoner's Dilemma Experiments, 1 959-2003." J Econ Behav Organ.* ***68****(3-4):489–497.*
9. ^*Milinski M, Wedekind C (1998). "Working Memory Constrains Human Cooperation in the Prisoner's Dilem ma." Proc Natl Acad Sci U S A.* ***95****(23):13755–13758.*

10. ^Embrey M, Fréchette GR, Yuksel S (2018). "Cooperation in the Finitely Repeated Prisoner's Dilemma." Q J E con. **133**(1):509–551.
11. ^Fontana N, Pierri F, Aiello LM (2025). "Nicer than Humans: How Do Large Language Models Behave in the Prisoner's Dilemma?" Proc Nineteenth Int AAAI Conf Web Soc Media. **19**:522–535.
12. ^Akata E, Schulz L, Coda-Forno J, Oh SJ, Bethge M, Schulz E (2025). "Playing Repeated Games with Large L anguage Models." Nature Hum Behav. **9**:1380–1390.
13. ^Mei Q, Xie Y, Yuan W, Jackson MO (2024). "A Turing Test of Whether AI Chatbots Are Behaviorally Similar t o Humans." Proc Natl Acad Sci U S A. **121**(9):e2313925121.
14. ^McCannon BC (2025). "Artificial Intelligence and Strategic Uncertainty: Can AI Play Mixed Strategies." J Ri sk Uncertain. forthcoming.
15. ^Bao T, Pei N (2023). "Cognitive Uncertainty, GPT, and Contribution in Public Goods Game." Working Paper.
16. ^McCannon BC (2024). "Artificial Intelligence is a Pro-Social Norm Complier." Econ Lett. **241**:article 111828.
17. ^McCannon BC (2011). "Coordination Between a Sophisticated and Fictitious Player." J Econ. **102**(3):263–27 3.
18. ^Ellison G (1997). "Learning from Personal Experience: One Rational Guy and the Justification of Myopia." Games Econ Behav. **19**:180–210.
19. ^Eloundou T, Manning S, Mishkin P, Rock D (2024). "GPTs are GPTs: Labor Market Impact Potential of LLM s Research is Needed to Estimate How Jobs May Be Affected." Science. **384**(6702):1306–1308.
20. ^Roy J, Saha B (2025). "Democratic Regulation of AI in the Workplace." Games Econ Behav. **152**:113–132.

Declarations

Funding: No specific funding was received for this work.

Potential competing interests: No potential competing interests to declare.