Application of Game Theory to Analysis of Machine versus Human Strategies

New Mexico

Supercomputing Challenge

Final Report

April 1, 2026

Welch Homeschool

Team Members:

Helena Welch

Kalliope Luna Welch

Teacher:

Cindy Welch

Project Mentor:

Paul Welch

Executive Summary

As machine learning becomes increasingly prevalent in our current society, it is assuming a more and more influential role upon individuals' decisions. The interplay of such decisions made by different players affects many aspects of life, including finance and politics. Machine learning algorithms have the potential to aid in these decisions, but only when examined first to ensure that their choices are rational.

This study aims to evaluate whether or not a game theoretical simulation of human interaction, the Consensus Game, can improve large language model (LLM) decision-making capabilities. We find that the Consensus Game tends to improve results for LLMs with small parameter sizes and for questions involving the Nash Equilibrium. This indicates that the Consensus Game performs well in simulating a human conversation in which multiple players with distinct mentalities reach a solution benefitting both players, thus modeling human decision-making processes more effectively than raw LLM output.

Introduction

Role of LLMs in Decision Making

LLMs are now employed by around 66% of today's population for a variety of purposes [4]; this usage is not limited to gathering fact-based information, but also includes making decisions regarding problems or difficult situations in users' lives. However, using LLMs to aid in decision-making can be harmful, as LLMs generate their answers statistically rather than rationally. This means that they often struggle with questions that are not merely fact-based but require a certain level of logic to make wise choices [13]. Thus, we need to better understand the capabilities of LLMs in making decisions that impact our everyday lives before we use them without constraints.

Purpose of Project

Game theory offers a set of mathematical tools and ideas that allow us to study how to make optimal decisions in various situations that affect multiple individuals, or players. It is used today to make strategic decisions in such fields as economics, diplomacy, and defense [15].

We hypothesize that game theory can improve LLM answers to decision-based questions through optimizing strategic interactions between models. Since groups of individuals often make decisions through discussions and arguments, we choose to test the Consensus Game, a game that simulates such interactions.

The Consensus Game

The Consensus Game (developed by Jacob et al. [8]) is used to achieve agreement between the answers chosen by two different methods for the same question. These two methods comprise:

1. the Generator, which chooses from a set of answers to answer a given prompt, and

2. the Discriminator, which determines the correctness of the Generator's answer.

The Consensus Game is then played based on the answers given by the Generator and Discriminator, using equations to update the probability that each player will choose a given answer. Note that the Consensus Game does not modify the LLMs themselves, but rather changes the initial probability that a given LLM returns a given answer.

By the end of this project, we expect to understand whether the Consensus Game improves the accuracy of LLMs in answering fact-based versus game-theory-based questions. In addition, we hope to identify differences in various LLMs' initial and final accuracies, based on their parameter sizes, and evaluate why certain question-model pairs improve to a greater extent.

Prior Works

Several studies have examined the processes of LLMs in decision-making compared to the decision-making processes of humans. Jacob et al. first developed and introduced the Consensus Game, testing it using two models, Llama 7B and 13B, to answer knowledge-based and problem-solving questions.

Other works include Rasal, Sumedh & Hauer 2024 [13], which introduces a methodology for improving LLM decision-making by prompting it with multiple choice questions, and Xiao & Wang 2025 [16], which concluded that humans are better at finding the Nash Equilibrium, or the optimal decision for all involved, than are LLMs.

Methods

First, we pass a series of knowledge-based questions and game theoretical scenarios to two separate instances of an LLM, the Generator and Discriminator. The Generator is prompted with several possible answers to a question, while the Discriminator must choose whether the Generator's answer is correct. Each question is asked 100 times, creating a probability distribution of answers for each instance of each LLM queried. These probability distributions, known as the policies, are then updated 5,000 times based on each other's previous policies, simulating a 'conversation' between the Generator and Discriminator. The accuracies of the initial and final policies are then calculated for each data set and question.

Data Collection and Preparation

Four LLMs with varying parameter spaces were used to compare variability in answer choices between machines. Following Jacob et al., zero-shot prompting was used, so no examples were given before the actual question was passed. We found that LLMs mostly gave clean answers with this format, although LLMs with smaller vocab sizes (Llama 3.1 8b and Mistral-small) sometimes produced words surrounding their answers rather than the single digit asked for.

model	**parameters**	**vocabulary size**
Mistral-small [12]	22 billion	32k
Command-r [6]	35 billion	128k
Llama 3.1 8b [14]	8 billion	64 k
Mistral-nemo [11]	12 billion	128k

Table 1: Characteristics of the large-language models used.

Prompts are given below:

Generator prompt: "[Question] Please choose an answer from the following messages: [Answer Choices] Only give the number of your answer, 1, 2, 3, or 4, with no other text surrounding it. Do not provide any other notes or commentary."

Discriminator prompt: "[Question] Is the following the correct answer to this question: [Answer Choice] Provide ONLY a single digit response. State 0 for true and 1 for false. No further comments are allowed."

We also attempted to use Chat-GPT 5o, a larger model, but were banned for asking it the same knowledge-based questions repeatedly, as if we were collecting data to train our own model. In the future, we hope to try collecting data from a similar large LLM that does not ban our usage.

Knowledge-Based and Decision-Based Queries

Following Jacob et al., knowledge-based questions were taken from the MMLU dataset [7] within the astronomy and anatomy topics. Each topic's dataset was approximately 200 questions. The following is an example question from the astronomy topic:

Why is Saturn almost as big as Jupiter despite its smaller mass?

1. *Jupiter's greater mass compresses it more, thus increasing its density.*
2. *Saturn has a larger proportion of hydrogen and helium than Jupiter and is therefore less dense.*
3. *Saturn is further from the Sun thus cooler and therefore less compact.*

4. *Saturn's rings make the planet look bigger.*

We then compare knowledge-based questions to decision-based queries generated based on standard game theoretical problems known as the atomic games. An example is given below. The full set of prompts we generated can be found in Appendix 2.

The prompt for a given game theoretical situation was given in two different ways: 1) abstractly and 2) textually.

1. *Abstract: Player 1 and Player 2 have two actions, A and B. If both players choose A, they will each receive a utility of 50. If both choose B, they will each receive a utility of 25. If Player 1 chooses B while Player 2 chooses A, they will each receive a utility of 15, and vice versa. What should each Player do? What is the Nash Equilibrium?*

2. *Textual: Two companies, Phony and Softdrive, are going to release a new product. If both companies choose to have product Option A, the more expensive, high-tech option, they will each receive 50 dollars. If they both choose Option B, the cheaper, lower-tech option, they will receive each 25 dollars. However, if Phony chooses Option A and Softdrive chooses Option B, they will both only receive 15 dollars, and vice versa. What should each company do? What is the Nash Equilibrium?*

Deriving the Consensus Game

The Consensus Game was derived per Jacob et al. 2021 [9] and Jacob et al. 2024 [8].

Having collected 100 answers for every question from the Generator and Discriminator each, we calculate the probability that a given answer was chosen and use this to define the initial policy π_G and π_D. From here we can define a standard utility, or the amount of payoff for each right answer from both the Generator and Discriminator for a given question.

$u_G(\pi_G, \pi_D) = \frac{1}{2} \sum_{v \in \text{correct, incorrect}} \sum_{y \in Y} \pi_G(y|x, \nu) \cdot \pi_D^{(t)}(\nu|x, y)$

We regularize this utility with the Kullback-Leibler Divergence, which prevents the policies from updating and diverging too rapidly from the previous policy.

$D_{KL}[\pi_G(\cdot|x, \nu)||\pi_G^{(1)}(\cdot|x, \nu)] = \pi_G \log(\frac{\pi_G}{\pi_G^{(1)}})$

$\tilde{u}_G(\pi_G, \pi_D) = u_G - \lambda_G D_{KL}[\pi_G(\cdot|x,\nu)||\pi_G^{(1)}(\cdot|x,\nu)]$

The regret, or how much the Generator and Discriminator are penalized for a given is defined below as the maximum difference of current and ideal utilities.

$\text{Reg}_G^{(T)} = \max_{\pi^* \in \triangle Y} \sum_{t=1}^{T}(u^{(t)}(\pi_G^*|x,\nu) - u_G^{(t)}(\pi_G^{(t)}(\cdot|x,\nu)|x,\nu))$

$\pi_G(y|x,\nu) \propto \exp(\text{Reg}_G^{(T)})$

$\pi^{t+1} = \arg\max_{\pi \in \triangle Y}\{(\sum_{t'=1}^{t} \tilde{u}_G(\pi_G, \pi_D)) - \frac{1}{\eta}\sum_{y \in Y} \log \pi(y|x,\nu)\}$ Follow-the-Regularized Leader algorithm guarantees a bound on the regret

$= \arg\max_{\pi \in \triangle Y}\{\eta(\pi_G u_G - \lambda_G \pi_G \log(\frac{\pi_G}{\pi_G^{(1)}})) - \sum_{y \in Y} \log \pi(y|x,\nu)\}$

$= \arg\max_{\pi \in \triangle Y}\{\eta \sum_{\pi_G}(t\lambda_G \log \pi_G^{(1)} + \sum_{t'=1}^{(t)} u_G)\pi(y|x,\nu) - (1+t\lambda_G\eta)\sum_{y \in Y} \pi(y|x,\nu)\log \pi(y|x,\nu)\}$

$= \arg\max_{\frac{\eta}{1+t\lambda_G\eta}\pi \in \triangle Y}\{\eta \sum_{y \in Y}(t\lambda_G \log \pi_G^{(1)} + \sum_{t'=1}^{(t)} u_G)\pi(y|x,\nu) - \sum_{y \in Y} \pi(y|x,\nu)\log \pi(y|x,\nu)\}$

Softmax Function

$w^{t+1}(\pi_G) = \frac{\eta}{1+t\lambda_G\eta}\sum_{y \in Y}(t\lambda_G \log \pi_G^{(1)} + \sum_{t'=1}^{t} \tilde{u}_G)$

$\pi_G^{t+1} = \frac{\exp\{w^{t+1}\}}{\sum \exp\{w\}}$

$\pi_G^{(t+1)}(y|x,\nu) \propto \exp\{\frac{Q_G^{(t)}(y|x,\nu) + \lambda_G \log \pi_G^{(1)}(y|x,\nu)}{1/(t\eta_G) + \lambda_G}\}$

Final Equations

The final equations are given based on the above derivation:

Initial Policies

$$\pi_G^{(1)}(y|x,\nu) = \frac{P_G(y|x,\nu)}{\sum_{y'}\sum_{\nu'} P_G(y'|x,\nu')} \quad \textbf{(1a)}$$

$$\pi_D^{(1)}(\nu|x,y) = \frac{P_D(\nu|x,y)}{\sum_{\nu'}\sum_{y'} P_D(\nu'|x,y')} \quad \textbf{(1b)}$$

Policy Updates

$$\pi_G^{(t+1)}(y|x,\nu) = \frac{e(t, Q_G^{(t)}, \pi_G^{(1)})}{\sum_{y'} e(t, Q_G^{(t)}, \pi_G^{(1)})} \quad \textbf{(2a)}$$

$$\pi_D^{(t+1)}(\nu|x,y) = \frac{e(t, Q_D^{(t)}, \pi_D^{(1)})}{\sum_{\nu'} e(t, Q_D^{(t)}, \pi_D^{(1)})} \quad \textbf{(2b)}$$

where:

$$e(t, Q_G^{(t)}, \pi_G^{(1)}) = \exp\{\frac{Q_G^{(t)}(y|x,\nu)+\lambda_G \log \pi_G^{(1)}(y|x,\nu)}{1/(t\eta_G)+\lambda_G}\}$$

$$e(t, Q_D^{(t)}, \pi_D^{(1)}) = \exp\{\frac{Q_D^{(t)}(\nu|x,y)+\lambda_D \log \pi_D^{(1)}(\nu|x,y)}{1/(t\eta_D)+\lambda_D}\}$$

$$Q_G^{(t)}(y|x,\nu) := \frac{1}{2t}\sum_{\tau=1}^{t} \pi_D^{(\tau)}(\nu|x,y) \quad \textbf{(3a)}$$

$$Q_D^{(t)}(\nu|x,y) := \frac{1}{2t}\sum_{\tau=1}^{t} \pi_G^{(\tau)}(y|x,\nu) \quad \textbf{(3b)}$$

Calculating Accuracies

$$S_G(t) = \frac{\sum_{x'} \pi_G^{(t)}(y^*_{x'}|x',\nu_c)}{N_X} \quad \textbf{(4a)}$$

$$S_D(t) = \frac{\sum_{x'} \pi_D^{(t)}(\nu^*|x',\nu_c)}{N_X} \quad \textbf{(4b)}$$

Definitions

$x =$ question

$y =$ answer choice

$\nu =$ correctness parameter

$t =$ update run

$P =$ probability distribution

$\pi^{(t)} =$ policy on a given run

$Q =$ running average of π

$y^* =$ accurate answer choice

$\nu^* =$ accurate correctness

$N_x =$ number of questions asked

$S =$ accuracy

$\tilde{u} =$ utility

D_{KL} = Kullback-Leibler Divergence

λ determines amount of regularization

Reg = regret

η = learning rate

Running the Consensus Game

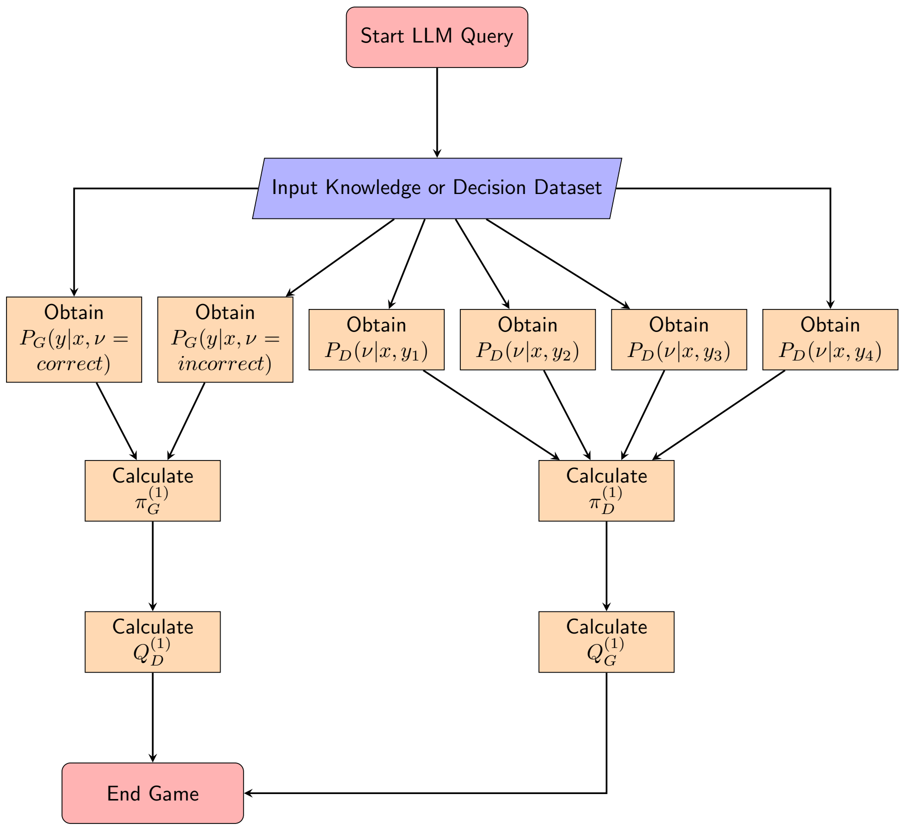

Graphic made by student using TikZ [20], 2026

Figure 1: Flow chart of algorithm querying the LLMs and obtaining the probability distributions for the Consensus Game.

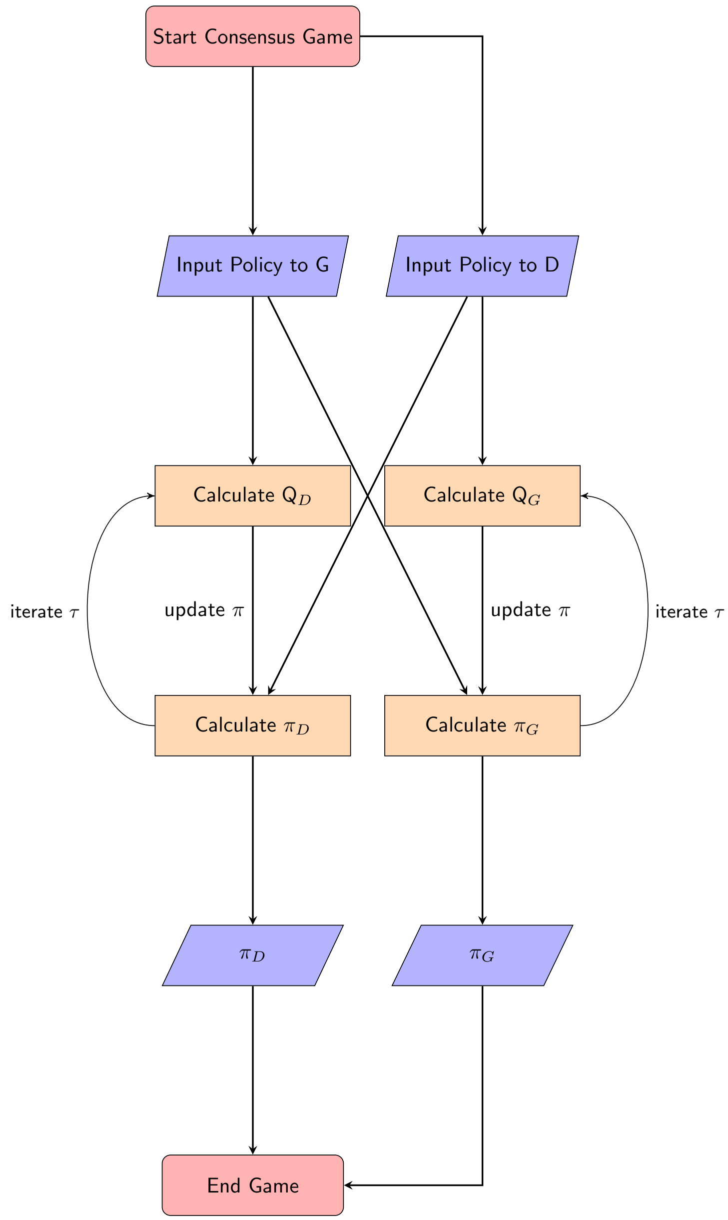

Figure 2: Flow chart of algorithm playing the Consensus Game.

Summary of Computation

Hardware and Computational Details

In running our code we found that playing the Consensus Game for a given question takes approximately 1 minute. As such, an entire knowledge-based dataset took around 2.5 hours, and each decision-based question set took 2.5 minutes.

Verification and Validation of Code

Several steps were taken to ensure that code ran correctly and efficiently. First, policy updates were calculated by hand for several test probability distributions and checked against our code's computation. In addition, code was run on multiple computers, confirm that it worked on different hardware.

Unit tests are also performed on each function used in performing the Consensus Game calculations. These are performed using the library pixi and check whether the functions return the expected answer. We perform these unit tests on each function involved in playing the Consensus Game, as shown in our supporting code.

Results

Testing the Consensus Game

We now analyze results from the Consensus Game. We find that repeated updating of the policies makes the Generator and Discriminator very ‘confident’ of their answers. The initial probability of choosing a given answer is relatively low, but after running the policy updates 5000 times, the probability of choosing one of the four answers increases logarithmically, often to near 100% probability after the first 2000 runs (see Figures 3 and 4). As seen in Figure 6, the answers of each LLM are thus polarized, as one answer choice becomes most popular and the others are almost never chosen. Consequently, through repeated ’interaction’, both the Generator and Discriminator become extremely confident of their answers. The goal is that their choice correctly answers a given question.

Note that a certain complexity is introduced by asking for the LLM to select an incorrect answer as opposed to a correct option. Often an initial answer increases in probability before being replaced by a different choice (see Figure 5). This indicates the greater difficulty that comes with the greater sophistication of such a question.

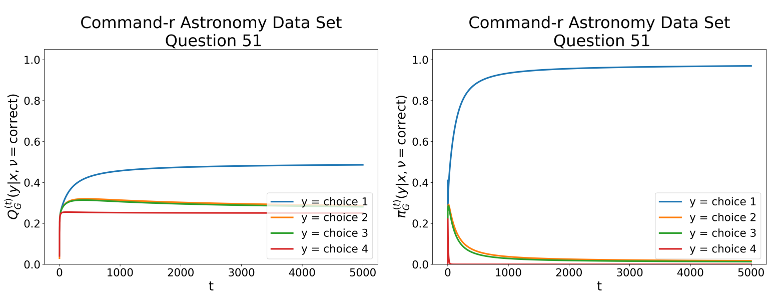

Figure 3: Evolution of Q-value and policy π_G over time step t for correct answer choices.

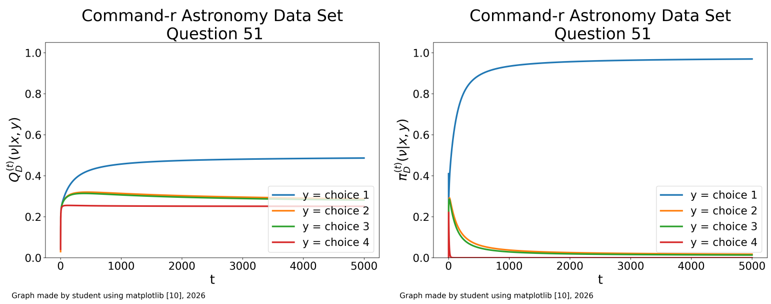

Figure 4: Evolution of Q-value and policy π_D over time step t.

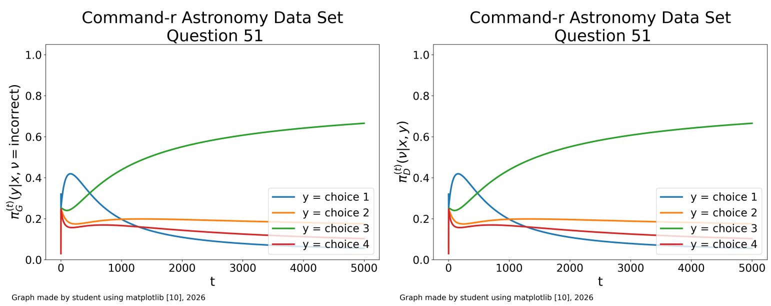

Figure 5: Evolution of Q-value and policy π_G over time step t for incorrect answer choices.

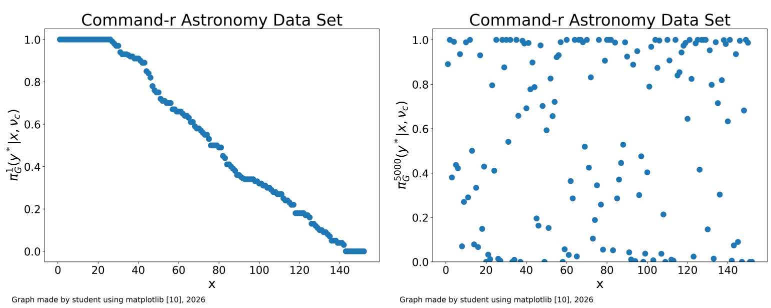

Figure 6: Accuracy of Generator in answering ordered set of astronomy questions before and after the Consensus Game is played.

Optimizing Hyperparameters

To determine how much the policy should change after each update, we iterate through values of λ_G, λ_D, η_G, and η_D to determine which would produce the highest accuracy. Overall, we find that accuracy for decision-based questions changes very little as a function of λ, but that a value of 0.1 usually leads the Consensus Game to perform at slightly higher accuracies. Thus, using a smaller Kullback-Leibler Divergence and changing the probability distribution in small steps at each update is usually most effective.

In addition, we find that the learning rate η has little effect on the final accuracy of each LLM on any given decision-based dataset; however, increasing η shortens the number of times the Consensus Game must be played before the updated policy plateaus and becomes constant. As such, increasing the learning rate η increases the efficiency of the algorithm, as the Consensus Game needs to be performed fewer times before coming to a final answer.

Accuracy of Consensus Game

Overall, we find that the Consensus Game improves accuracy of knowledge-based questions only marginally, as seen in Figures 7 and 8. Decision-based questions, on the other hand, resulted in large increases in accuracy, often from below 20% to almost 100% in the case of the smaller models mistral-nemo and llama 3.1-8b (see Figure 9 and 10). Thus, our hypothesis that the Consensus Game would be more useful in improving the decision-making capabilities of LLMs was confirmed. Abstract and textual formats for the decision-based prompts (see Knowledge-Based and Decision-Based Queries) overall tended to perform similarly, with textual prompts performing slightly better. Giving more context in prompts consequently allows the LLMs to perform at a higher accuracy. We also find that the Discriminator tends to improve more drastically than the Generator.

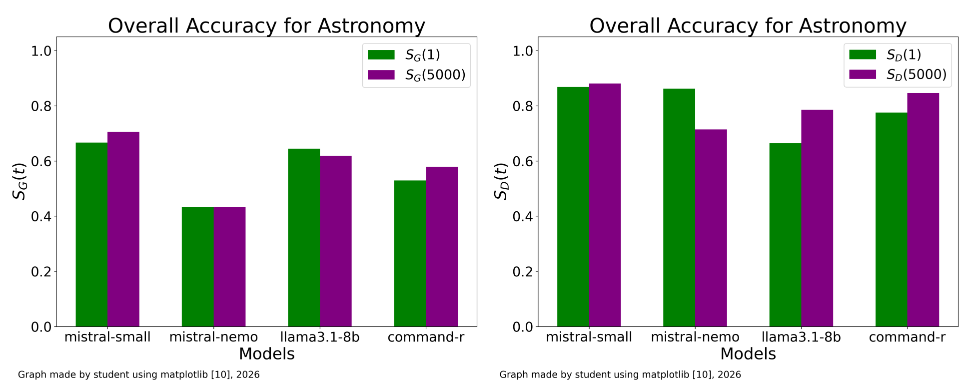

Figure 7: Overall accuracy of each LLM in answering astronomy questions from the MMLU database.

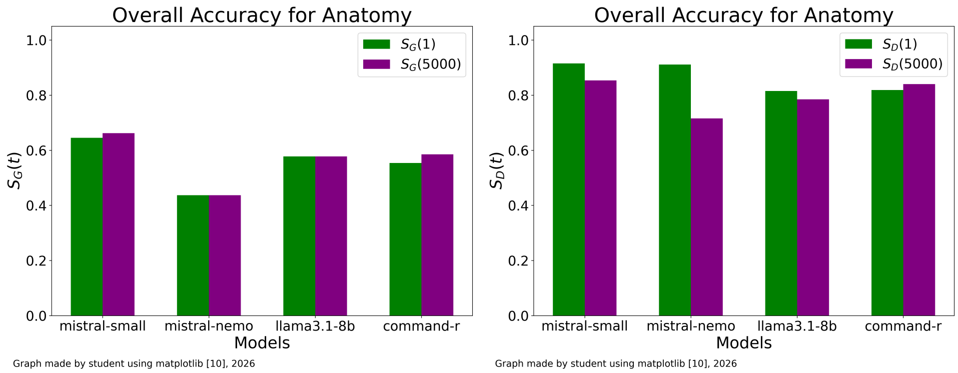

Figure 8: Overall accuracy of each LLM in answering anatomy questions from the MMLU database.

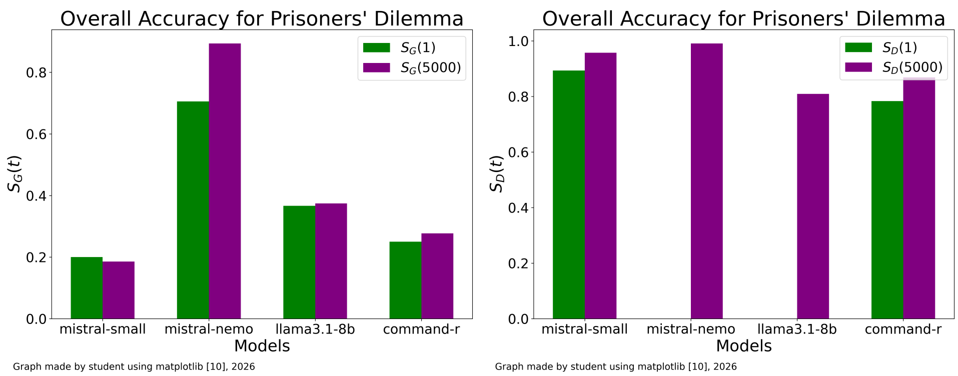

Figure 9: Overall accuracy of each LLM in answering questions related to the Prisoners's Dilemma (see Appendix).

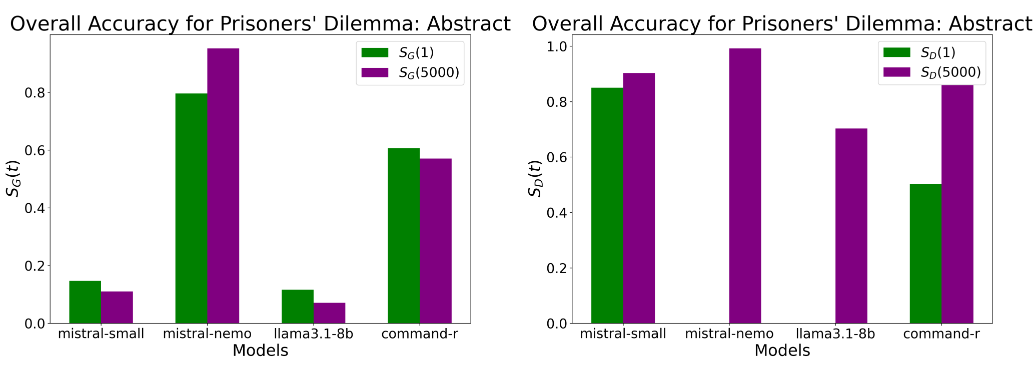

Figure 10: Overall accuracy of each LLM in answering questions related to the abstract version of the Prisoners's Dilemma (see Appendix).

Most significantly, we find that the increase in accuracy is often contributed most by questions involving the Nash Equilibrium, or the optimal strategy for all players together (see Figure 11and Appendix 1).

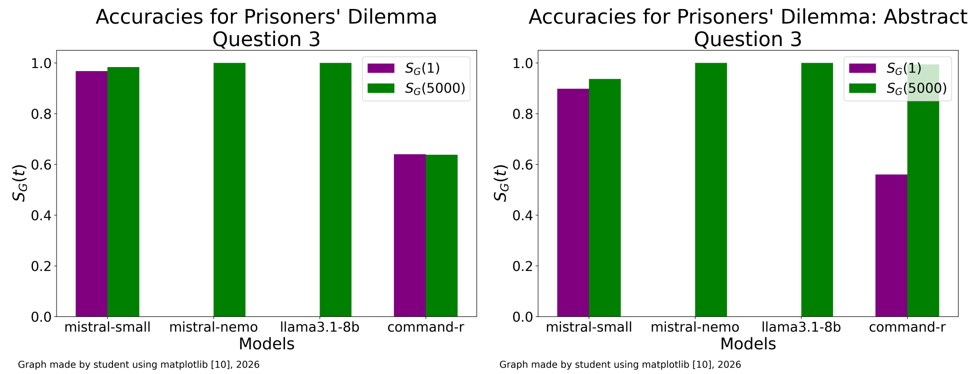

Figure 11: Accuracy of each LLM in answering questions regarding the Nash Equilibrium related to the Prisoners's Dilemma (see Appendix).

<u>Applying the Consensus Game</u>

The Consensus Game has thus far been tested for multiple choice questions; we next explore decision-making in multi-step problems by having LLMs make a choice at each decision node in a story plot line. We find that the Consensus Game can be used not only to answer questions but also to simulate the creative process of generating a story more efficiently. Whereas traditional game theory would require that we map out every possible choice that could be made through the course of a story, the Consensus Game finds the optimal decision at each node before moving to the next step. As an example, we use the Consensus Game to create a decision tree mapping out LLMs' retelling of operas in different time frames.

We began with the opera *Dido and Aeneas* by Henry Purcell (libretto by Nahum Tate) [1], in which Aeneas is forced to leave Dido for the sake of following his destiny and starting a new civilization. The prompt was phrased:

"A young man is pressured to start a new civilization, but is attached to a woman who will not follow him. What are two options he might pursue?"

This prompt was then fed to the LLM Mistral-small, which outputted two options, 'settle nearby and stay in contact', or 'promise to return'. The Consensus Game was run on the LLM's output, determining that Aeneas should select the first option. The second prompt,

"A young man is pressured to start a new civilization, but is attached to a woman who will not follow him. He chooses to settle nearby and maintain contact. What are two options he might pursue?"

is given to the same LLM, which proceeds to output two options and repeat the process. After a 62 iterations, Aeneas's final decision is to join a well established trade company. The purpose of this process is to examine the decision making processes of a LLM in fictional situations, and how these decisions differ from those made in the original story.

Conclusions

This study finds that a game theoretic algorithm that simulates the dynamics of a human conversation improves the decision-making capabilities of LLMs. As humans often make choices through back-and-forth discussions, LLMs thus improve their accuracy in making such decisions through a similar process.

Analysis of Multi-Step Results

In the original opera *Dido and Aeneas*, Aeneas makes the decision to leave Dido and start a new civilization; as a result, Dido dies of sorrow and causes grief to both Aeneas and her people. In the LLM generated output, Aeneas makes the same decision to leave Dido. However, the results are far less drastic and concern Aeneas's future in creating a civilization via different types of trading and military outposts.

From this, we can conclude that the focuses of the librettist for *Dido and Aeneas* and the LLM given our prompt differ; the opera revolves around the emotional impact on the characters therein while the LLM output revolves around potential realistic opportunities in the time period of Aeneas's existence. In the future, we will obtain more data by phrasing our prompts from the viewpoints of Dido and other characters in the opera apart from Aeneas.

Analysis of Biases and Limitations

The Consensus Game can only improve LLM answers when the Generator and Discriminator do not agree on the wrong answer. As the algorithm's purpose is to draw the Generator and Discriminator decisions to an equilibrium, they cannot already be in agreement on the wrong answer. We note that occurrences of this problem within our dataset prevent the Consensus Game from improving results for certain questions.

In addition, we find that the optimal hyperparameters depends on the given question and is thus not necessarily consistent across datasets. These caveats do not undermine the usefulness of the Consensus Game in every case but should be noted as occasionally decreasing the accuracy of the original LLM responses.

Next Steps

In the future, we plan to test the Consensus Game on more multi-step games beyond the opera *Dido and Aeneas*. In addition, we hope to explore alternative games to the Consensus Game. How will competitive games perform in comparison to the cooperative methodology of the Consensus Game? Will a third player improve the accuracy in coming to an agreement? We hope to continue exploring game theory's ability to improve LLM decision-making capabilities in the future.

Appendix 1: Glossary of Technical Terms

Generator: LLM prompted with a question and a set of answer choices, which then choose an answer to a given question 100 times; this generates a probability distribution of answers, which is then used to play the Consensus Game.

Discriminator: LLM prompted with a question and a given answer choice, which then chooses whether that answer choice is true or false 100 times; this generates a probability distribution of answers, which is then used to play the Consensus Game.

Nash Equilibrium: A solution to a problem in game theory in which no player can improve their outcome by making a change; the optimal decision for all players as a collective [3].

Appendix 2: Full Set of Game Theoretical Scenarios used in Decision-Based Questions

Battle [5]

Two teenagers with separate tastes must agree on an activity. Their choices are basketball and shopping–the boy prefers basketball, and the girl prefers shopping, but they both prefer to go together, wherever they go..

Boy \Girl	**Basketball**	**Shopping**
Basketball	3,2	1,1
Shopping	0,0	2,3

Table made by student using LaTex. [10]

Prisoner's Dilemma [5]

Two suspects are arrested with a potential five year sentence. If neither confess to the crime, they will be jailed for one year. If one confesses and the other does not, the one who admits will go free while the other serves eight years of jail.

1 \2	**Defect**	**Cooperate**
Defect	-5,-5	0,-8
Cooperate	-8,0	-1,-1

Table made by student using LaTex. [10]

Chicken [5]

Two drivers speed towards each other on the same road. The one who swerves first is the game's loser, or 'chicken.' If neither swerve, they crash.

1 \2	**Swerve**	**Straight**
Swerve	0,0	-1, 1
Straight	1, -1	-1000, -1000

Table made by student using LaTex. [10]

Stag Hunt [22]

Two hunters may hunt stag or hare. Hare produces less meat, but stag cannot be caught without the help of the other hunter.

1 \2	Stag	Hare
Stag	10,10	0, 2
Hare	2, 0	2, 2

Table made by student using LaTex. [10]

Assurance Game [5]

Two companies may produce a cheap or expensive product. Their profits will be lower if they do not choose the same product.

1 \2	A	B
A	50, 50	15, 15
B	15, 15	25, 25

Table made by student using LaTex. [10]

Traveler's Dilemma [2]

Two airline passengers have identical damaged briefcases. They must choose a reparation value. If one chooses a lower value than the other, the lower value will be accepted and the one who chose the lower value will gain some portion of the other passenger's reparations.

1 \2	Lower Amount	Higher Amount
Lower Amount	20, 20	30, 10
Higher Amount	10, 30	80, 80

Table made by student using LaTex. [10]

Appendix 3: Multi-Step Storylines Generated by Consensus Game

1800 AD Storyline

A young man is pressured to start a new civilization, but is attached to a woman who will not follow him.

He chooses to settle nearby and maintain contact. He chooses to establish a settlement nearby. He chooses to establish a trading post. He chooses to establishing a missionary outpost. He chooses to establishing a trading post. He chooses to establish a trading post. He chooses to establishing a trading post. He chooses to establishing a trading post. He chooses to establish a trading post. He chooses to establishing a frontier trading post. He chooses to establishing a trading post. He chooses to establishing a fur trading post. He chooses to establish a fur trading post. He chooses to establishing a fur trading post. He chooses to establishing a frontier trading post. He chooses to establish a fur trading post. He chooses to establish a fur trading post. He chooses to establishing a fur trading post. He chooses to establishing a fur trading post. He chooses to establish a fur trading post. He chooses to establishing a lumber camp. He chooses to establishing a fur trading post. He chooses to establishing a fur trading post. He chooses to establish a fur trading post. He chooses to establish a hudson's bay company trading post. He chooses to establish a hudson's bay company trading post. He chooses to establish a hudson's bay company trading post. He chooses to establish a trading post. He chooses to fur trading post. He chooses to establish a hudson's bay company trading post. He chooses to fur trading post. He chooses to establishing a hudson's bay company trading post. He chooses to join a trading company. He chooses to joining a trading company. He chooses to joining a trading company. He chooses to joining the north west company. He chooses to joining a trading company. He chooses to establishing a hudson's bay company trading post. He chooses to joining a well-established trading company. He chooses to joining the north west company. He chooses to joining a trading company. He chooses to joining a well-established trading company. He chooses to joining a well-established trading company. He chooses to joining the hudson's bay company (hbc). He chooses to joining a well-established trading company. He chooses to establishing a hudson's bay company (hbc) trading post. He chooses to joining a well-established trading company. He chooses to joining the north west company. He chooses to joining the north west company. He chooses to joining a well-established trading company. He chooses to joining a well-established trading company. He chooses to joining the north west company. He chooses to joining a well-established trading company. He chooses to joining the hudson's bay company (hbc). He chooses to joining the hudson's bay company (hbc). He chooses to joining the hudson's bay company (hbc). He chooses to joining the north west company (nwc). He chooses to joining a well-established trading company. He chooses to establish a hudson's bay company (hbc) trading post. He chooses to establish a hudson's bay company (hbc) trading post. He chooses to joining a well-established trading company.

1800 BC Storyline

He chooses to leave the woman and start the new civilization. He chooses to establishing an agricultural community. He chooses to fertile crescent approach. He chooses to utilizing irrigation systems. He chooses to following the Sumerian model. He chooses to establish a settlement in mesopotamia. He chooses to establish trade routes. He chooses to establishing religious institutions. He chooses to establishing a centralized government. He chooses to adoption of cuneiform script. He chooses to establishing a monarchy. He chooses to development of urban centers. He chooses to alliance with a neighboring city-state. He chooses to military alliance. He chooses to seeking alliance with Akkad or Sumer. He chooses to seeking alliance with Akkad. He chooses to seeking alliance with Akkad. He chooses to establishing a trading post with Akkad. He chooses to establish a military alliance with Akkad. He chooses to establishing a military alliance with Akkad. He chooses to establishment of a code of laws. He chooses to establishing trade routes with the indus valley civilization. He chooses to establishment of a code of laws. He chooses to establishing diplomatic relations. He chooses to establishing alliances with other city-states. He chooses to establishing a canal system. He chooses to establishing a canal system. He chooses to development of metallurgy. He chooses to trade routes. He chooses to establishing a canal system for improved irrigation and transportation. He chooses to establishing a polytheistic religious system. He chooses to alliance with Akkad for protection and expansion. He chooses to establishing a joint military venture. He chooses to establishing a form of taxation. He chooses to establishing a ziggurat. He chooses to military expansion. He chooses to establishment of a joint religious center. He chooses to establishing a joint religious center with Akkad. He chooses to focus on agricultural innovations. He chooses to establishing a canal system for improved irrigation and transportation. He chooses to establishing a temple-based economy. He chooses to development of advanced irrigation techniques. He chooses to seeking alliance with babylon. He chooses to establishing a joint military venture with Akkad. He chooses to establishing alliances with other city-states for mutual defense. He chooses to establish a canal system. He chooses to establishment of a joint religious center with Akkad. He chooses to establishing a writing system. He chooses to establishing a joint religious center with Akkad. He chooses to establishing diplomatic marriages. He chooses to establishment of a code of laws. He chooses to establish a joint religious center with Akkad. He chooses to establishing a trade alliance with harappan civilization. He chooses to establishing a form of taxation. He chooses to establishing trade routes to the indus valley civilization. He chooses to establishing a form of taxation. He chooses to establishing a royal academy for scholars. He chooses to establishing a tributary system. He chooses to establishment of early forms of public health. He chooses to establishing craft guilds. He chooses to military expansion and consolidation of power. He chooses to establishing a strategic alliance with Akkad. He chooses to establishment of a bureaucracy. He chooses to establishing a bronze smelting industry. He chooses to focus on irrigation and agricultural innovations. He chooses to establishing advanced irrigation techniques. He chooses to establishment of a code of laws. He chooses to establishing early forms of public health. He chooses to explore trade opportunities with the indus valley civilization. He chooses to expansion and military alliances. He chooses to focus on technological and agricultural advancements. He chooses to alliance with Akkad for military expansion. He chooses to establishing a joint religious center with Akkad. He chooses to strengthening agricultural innovations and infrastructure. He chooses to establishment of a royal library. He chooses to agricultural innovations. He chooses to forming a military alliance with Akkad. He chooses to alliance with babylon. He chooses to military expansion and political consolidation. He chooses to strengthening agricultural innovations and infrastructure. He chooses to cultural and technological exchange. He chooses to establishing trade routes with the indus valley civilization. He chooses to adoption of advanced irrigation techniques. He chooses to establishing a joint military venture with Akkad. He chooses to establishment of a joint religious center with Akkad. He chooses to establishment of a canal system. He chooses to establishing a diplomatic marriage. He chooses to alliance and expansion with Akkad. He chooses to establishing a formal alliance with Akkad. He chooses to seeking military support from the old assyrian empire. He chooses to establishing a diplomatic marriage with a daughter of the Akkadian king. He chooses to establishing a ziggurat. He chooses to establishing a tribute system. He chooses to expansion through alliance and diplomatic marriage. He chooses to establishing a military coalition with other city-states. He chooses to expansion and military alliances. He chooses to establishing a strategic alliance with Akkad for mutual defense. He chooses to establishing a military alliance with Akkad. He chooses to send a diplomatic mission to Akkad to negotiate a formal alliance. He chooses to establish a strategic alliance with Akkad for mutual defense.

Acknowledgments

We could not have developed this project without our parents' inspiring advice and encouragement; their deeper understanding of the computer science, mathematics, and computer hardware that they have lended to us have been priceless. We would like to thank the numerous judges and those who came to hear us present our posters at the Science Fair and Supercomputing Challenge. Their generous feedback allowed us to advance our project in ways that we could scarcely predict. We would like to add an additional thank you to all authors, researchers, and scientists that have inspired us in the making of this project. With help from all of these people we have been given motivation to advance our project further.

References

[1] Jules Barbier and Carre Michel. *The Book of 101 Opera Librettos*. Deluxe ed., Tess Press, 1996.

[2] Kaushik Basu. The traveler's dilemma: Paradoxes of rationality in game theory. *American Economic Review*, 84:391–395, 1994.

[3] Leonardo Becchetti, Luigino Bruni, and Stefano Zamagni. Chapter 6 - non-competitive markets and elements of game theory. In Leonardo Becchetti, Luigino Bruni, and Stefano Zamagni, editors, *The Microeconomics of Wellbeing and Sustainability*, pages 157–198. Academic Press, 2020. ISBN 978-0-12-816027-5. doi: https://doi.org/10.1016/B978-0-12-816027-5.00006-9. URL `https://www.sciencedirect.com/science/article/pii/B9780128160275000069`.

[4] Anthony Cardillo. How many people use ai? (latest 2025 data). *Exploding Topics*, 2025.

[5] J.R. Corrigan. *Understanding Economics Game Theory,The Great Courses Guidebook*. The Teaching Company, 2021.

[6] Aidan Gomez. Introducing command r7b: Fast and efficient generative ai. https://cohere.com/blog/command-r7b, 2024.

[7] Dan Hendrycks, Collin Burns, Steven Basart, Andy Zou, Mantas Mazeika, Dawn Song, and Jacob Steinhardt. Measuring massive multitask language understanding. In *International Conference on Learning Representations*, 2021. URL `https://openreview.net/forum?id=d7KBjmI3GmQ`.

[8] Athul Jacob, Yikang Shen, Gabriele Farina, and Jacob Andreas. The consensus game: Language model generation via equilibrium search. In B. Kim, Y. Yue, S. Chaudhuri, K. Fragkiadaki, M. Khan, and Y. Sun, editors, *International Conference on Learning Representations*, volume 2024, pages 5832–5848, 2024. URL

https://proceedings.iclr.cc/paper_files/paper/2024/file/17a9bfda8be2301e24f232fb32f1e0fa-Paper-Conference.pdf.

[9] Athul Paul Jacob, David J. Wu, Gabriele Farina, Adam Lerer, Anton Bakhtin, Jacob Andreas, and Noam Brown. Modeling strong and human-like gameplay with kl-regularized search. In *International Conference on Machine Learning*, 2021.

[10] Leslie Lamport. *LaTeX: A Document Preparation System*, 1986. User manual.

[11] Mistral AI. Mistral nemo. https://mistral.ai/news/mistral-nemo, 2024.

[12] Mistral AI. Mistral small 3. https://mistral.ai/news/mistral-small-3, 2025.

[13] Sumedh Rasal and E. J. Hauer. Optimal decision making through scenario simulations using large language models, 2024. URL https://arxiv.org/abs/2407.06486.

[14] Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, Aurelien Rodriguez, Armand Joulin, Edouard Grave, and Guillaume Lample. Llama: Open and efficient foundation language models, 2023. URL https://arxiv.org/abs/2302.13971.

[15] UKEssays. The importance of game theory. URL ukessays.com/essays/economics/the-importance-of-game-theory-economics-essay.php?vref=1.

[16] Feng Xiao and XT (XiaoTian) Wang. Evaluating the ability of large language models to predict human social decisions. *Scientific Reports*, 15, 2025. URL https://api.semanticscholar.org/CorpusID:281082739.