COOPERATION IN REPEATED GAMES

JADEN ZHU

1. Abstract

This paper goes over how cooperation can emerge in repeated games, with examples based on the Prisoner's Dilemma. I analyze several strategies including Grim Trigger, Tit-for-Tat, and Win-Stay Lose-Shift, examining the specific conditions needed for cooperation to become rational. We also dig into many ways to sustain cooperation, like the discount factor and the Folk Theorem. Towards the end of this paper, we examine challenges that prevent a one-to-one correspondence with real-world scenarios and game theory. Alas, repeated games offer valuable insights to understanding real-world interactions, from business to international relations.

2. Introduction

2.1. What is game theory?

Game theory is the formal study of a player's actions whether it be conflict or cooperation. Concepts of game theory are used whenever there are many parties interacting. Game theory allows an approachable way to analyze, discuss, and interpret strategic scenarios. The earliest known trace of Game theory dates back to 1838 in Antoine Cournot's study of duopolies. Game theory really started gaining recognition after the 1944 publication of the book *Theory of Games and Economic Behavior* by John von Neumann and the economist Oskar Morgenstern. This book provides a lot of basic terminology that is still used today.

Game theory further evolved in 1950, when John Nash discovered that all finite games have an equilibrium path of play where all players choose optimal actions given the other players' choices. In the late 1950s and 1960s, Game theory had applications to problems in war and politics as well. Since then, Game theory has evolved to find application in all sorts of areas, from psychology to biology, and raised billions of dollars in the United States and Europe. Nowadays, Game theory is increasingly relevant in sectors like biology, politics, economics and AI.

2.2. Introduction.

This paper explores the transformation that occurs in player behavior when games turn from single encounters to repeated interactions and how cooperation can arise even in unlikely situations where individual interests de-incentivize it.

Date: June 2025.

The main paradox is that in many situations, individual interests lead to worse collective results for all parties. This tension is illustrated in the Prisoner's Dilemma, one of the most famous examples in game theory, perfectly. When two players play, both play to optimize their individual gain, which compels them to defect. even though mutual cooperation would provide a greater payoff for both players. This outcome is so important and necessary to study because it appears in countless real-world scenarios, like environmental protection, trade agreements, and social cooperation.

In contrast, in repeated games, players know that they will have a future interaction so they are more incentivized to cooperate today to not recieve retaliation in the future. This idea helps create a huge variety of strategic possibilities that otherwise could not be sustained in one-shot games.

Next, we will analyze how repetition enables cooperation through several different ways. We will explore the difference between finitely and infinitely repeated games, and the key difference between them that allows infinitely repeated games to host a much larger possibility of cooperation than the latter. The concept of the discount factor is also crucial to sustain this, as it represents how much players value future payoffs compared to immediate gain.

Following this, we will then compare various strategic approaches that players can use in repeated games, from simple strategies like Always Cooperate and Always Defect, to more complicated ones like Tit-For-Tat, Grim Trigger, and Win-Stay Lose-Shift. All of these strategies are a blend between forgiving, ruthless, retaliatory, exploitable, and everything in-between.

In the next section, we will go over how to look for equilibria, to make sure that these strategies are rational for players to carry out. We use the One-Shot Deviation Principle, as well as the discount factor to simplify and reduce the amount of time needed to look for equilibria. Then, we do a case study of four key strategies played in infinitely repeated games and calculate a sufficient amount that the discount factor can be to sustain cooperation.

Moving on, the next section describes the Folk Theorem, which demonstrates that virtually any outcome can be sustained as an equilibrium in infinitely repeated games given that players are sufficiently patient. This result shows that cooperation is not just 'possible' in infinitely repeated games, but there are a vast amount of sustainable cooperative outcomes that can be played.

However, this paper also goes over limitations to the Folk Theorem and ideas of game theory as a whole. Pure game theory often overlooks challenges like noise, imperfect monitoring, human limitations which are present in real-world interactions. These issues can drastically reduce the probability of cooperation arising because of various uncontrollable factors in the real world.

Finally, all of our analysis presented here is essential to answer some of the most pressing problems in the real-world. The analysis can draw a lot of parallels to real-world interactions, like the Cold War, where both the United States and the Soviet Union were able to reduce tensions by nuclear disarmament through mutual cooperation. The insights gained from studying repeated games extend far beyond mathematics, providing wisdom in dozens of areas, from business, to social interactions, international agreements, and understanding actions.

Ultimately, this analysis reveals that cooperation, while challenging, is both theoretically possible and achievable under the right conditions. This paper serves as both a mathematical exploration of cooperation in repeated games, and a guide for understanding why organizations or governments react a certain way in real-world interactions.

3. Preliminaries

3.1. What is a game?

In game theory, a game is a formal model used to represent strategic interactions among decision-makers, known as players. Each player is assumed to be perfectly rational and to possess common knowledge (they know the rules of the game, and they know that all other players also know the rules). See [1] for more.

There are two main types of games:

1. **Strategic Form Games**: In these games, players make decisions simultaneously. The outcomes are represented in a payoff matrix, where each cell corresponds to a combination of strategies chosen by the players.
2. **Extensive Form Games**: These games are represented as game trees. The tree structure shows the sequence of moves, the points at which players make decisions, and the information available to them at each decision point.

3.2. What are payoffs?

Payoffs are numerical values assigned to outcomes in a game. These values, also referred to as utility, indicate how desirable a particular outcome is for a player. A higher payoff is considered more favorable.

3.3. What is a strategy?

A strategy is a complete plan of action a player follows throughout the game. In extensive-form games, this includes specifying the action the player would take at every decision point.

3.4. **What are payoff matrices?**
A payoff matrix is a table that summarizes the outcomes of a game for all possible strategy combinations. Each cell in the matrix contains the pair of payoffs that the players receive when they choose a specific combination of strategies. The rows typically represent the possible choices of one player, while the columns represent the choices of the other player. This structure helps to clearly visualize the consequences of each strategic interaction. See [11].

3.5. **Dominant Strategies.**
A dominant strategy is one that yields a better (or at least equal) payoff for a player regardless of the strategy chosen by the opponent. Rational players choose dominant strategies because they maximize individual benefit in any scenario, especially in one-shot games where the interaction occurs only once.

There are two types of dominance:

- Strict Dominance: A strategy strictly dominates another if it always provides a higher payoff, no matter what the opponent does.
- Weak Dominance: A strategy weakly dominates another if it provides a payoff that is at least as good in all cases and strictly better in some cases. See [6] for more.

3.6. **Nash Equilibrium.**
A Nash Equilibrium is a strategy profile in which each player's strategy is a best response to the strategy of the other player. In this state, no player can improve their payoff by changing their strategy. Nash equilibrium represents a stable outcome where both players are satisfied with their decisions after considering their opponent's actions. Importantly, every finite game has at least one Nash equilibrium. These equilibria often model outcomes driven by self-interest.

3.7. **Subgame Perfect Equilibrium.**
In extensive-form games, it is useful to analyze smaller segments of the game known as subgames. A strategy profile is a Subgame Perfect Equilibrium (SPE) if a Nash equilibrium is played in every subgame of the original game. In other words, players' strategies must form a Nash equilibrium not only in the overall game but also in every part of the game that can be considered on its own. This ensures that players make rational choices at every stage of the game.

4. The Prisoner's Dilemma

The Prisoner's Dilemma is one of the most commonly studied and well-known scenarios in all of game theory. It's simple format and game scenario parallels fields such as economics, political science, business strategy, and psychology.

The game is as follows. In the classic example, two people are taken into a police station because they are suspected of committing a crime. While the police have proof that they were trespassing, they lack adequate proof to convict them of a more serious offense. To get confessions, the police put the prisoners in separate rooms and offer each the same deal:

- If **neither confesses**, each receives a light sentence for trespassing.
- If **one confesses** (defects) and the other remains silent (cooperates), the confessor is released while the silent prisoner receives a long sentence.
- If **both confess**, each receives a moderate sentence.

Importantly, both prisoners know that the other has been given the same offer. However, since they are unable to communicate, they must make their decision independently. This setup can be modeled using a payoff matrix, where the action "confess" is labeled as Defect (D), and "remain silent" is labeled as Cooperate (C).

Classic Payoff Matrix for the Prisoner's Dilemma:

		Player 2: C	Player 2: D
Player 1	C	$(3, 3)$	$(0, 5)$
	D	$(5, 0)$	$(1, 1)$

Conditions for a game to be classified as a Prisoner's Dilemma:

		Player 2: C	Player 2: D
Player 1	C	(R, R)	(S, T)
	D	(T, S)	(P, P)

To qualify as a Prisoner's Dilemma, the following inequality must hold:

$$T > R > P > S$$

Where
R = Reward
T = Temptation

P = Punishment
S = Sucker's payoff (Or just sucker)

This structure guarantees that each player has a dominant strategy: to defect. Regardless of what the other player does, defecting always provides a better individual outcome (for example, $5 > 3$ and $1 > 0$). As a result, both players will logically choose to defect.

This leads to the game's Nash Equilibrium at the outcome (D, D), or mutual defection. In this situation, neither player can profitably improve their outcome by changing their strategy. Ironically, both players would have achieved a better outcome through mutual cooperation, but both players' dominant strategies tell them to defect. This paradox is what makes the Prisoner's Dilemma so interesting and useful in understanding real-world decision making.

5. Repeated Games

5.1. One Shot vs. Repeated Games.

Games can be played only once, or multiple times. A **one-shot game** is a game that only occurs a single time between players. Some properties of these games are: no memory of past actions, no anticipation of future interactions, and no influence from reputation or threats. As a result, players focus on maximizing individual gain.

However, one-shot games are often unable to capture the complex nature of real-world interactions, where decisions and outcomes are influenced by past outcomes and behavior. This key limitation makes one-shot games not suitable for modeling real-life scenarios.

5.2. Overview of Repeated Games.

In contrast, **repeated games** allow players to remember past actions and adjust their strategies accordingly. In these settings, players must weigh the potential consequences of their present choices, and recognize that today's actions can influence the behavior of others in future rounds. This dynamic allows the use of strategies that are not possible in one-shot games, including punishment, reward, and cooperation.

For example, in a single one-shot version of the Prisoner's Dilemma, the dominant strategy for each player is to defect, leading to the Nash equilibrium of (D, D). Despite this, both players would prefer the mutual cooperation outcome.

Repeated games changes this drastically. If the same game is played multiple times, cooperation can be sustained through credible threats of future punishment. For cooperation to be stable, players must believe that any deviation will trigger retaliation in future rounds. In addition, the threat of punishment must also be rational for the punisher to carry it out when the time comes. Later in this paper, we will explore several strategies that make this possible.

There are **two** main types of repeated games:

- Finitely Repeated Games, and
- Infinitely Repeated Games.

5.3. Finitely Repeated Games.

A **finitely repeated game** is a game where a 'stage game' is played repeatedly a fixed number of times. A big feature of these games is that players can adjust their strategies based on the outcomes of previous rounds. To analyze these games, we utilize a tool called **backward induction**, which prescribes starting by working out a strategy from the final period and working backward to the beginning.

For example, applying backward induction to the Repeated Prisoner's Dilemma, in the final round, both players have no future consequences to worry about and will therefore play their dominant strategy, Defection. Given that mutual defection is guaranteed in the final round, the same logic applies to the second-to-last round. Since players know they will defect in the last round regardless of current choices, they have no incentive to cooperate now. This reasoning continues through all of the earlier stages. The outcome is that mutual defection is played in every round, forming a subgame perfect equilibrium.

Noting down the results, even if the players prefer cooperation, they anticipate defection from their opponent, making cooperation unappealing. Thus, finitely repeated games with known endpoints and dominant strategies telling the player to defect cannot support cooperation, even when threats are available.

5.4. Infinitely Repeated Games.

On the other hand, **infinitely repeated games** are games that have no determined endpoint. This drastically changes player behavior. and strategies because without a final round, the unraveling effect seen in finite games disappears. This allows for cooperation to emerge as a rational strategy.

In these settings, players shift from focusing on immediate payoffs to considering the **long-term consequences** of their actions. Cooperation becomes sustainable because players recognize that current deviations may lead to future punishments, potentially outweighing any short-term benefits.

5.5. Discount Factor.

To calculate payoffs in infinitely repeated games, we use a **discount factor**, which is denoted as δ. This factor represents the idea that future rewards are typically valued less than immediate ones. The range of δ is $0 < \delta < 1$. A higher δ value implies players are more patient and value future payoffs high, while a lower δ value indicates greater preference for short-term gains.

The discount factor allows us to compute the value of an infinite stream of future payoffs by reducing the weight of later rounds. This shows how players evaluate their strategies over the long run.

5.6. **What's Next.**

Because infinitely repeated games contain an infinite number of possible strategy sets, players must use **comprehensive strategies** that specify exactly what to do in any scenario, no matter how the opponent behaves. It is not feasible to predict every possible move the other player might make, so strategies must provide clear guidelines that can be applied after every potential sequence of past plays.

6. Strategies

This section outlines common strategies used in repeated games, with particular emphasis on their application in the Repeated Prisoner's Dilemma. Many of these strategies have proven success in Axelrod's Tournament of strategies, a well-known competition designed to test the effectiveness of different strategies in repeated interactions.

6.1. **Always Defect and Always Cooperate.**

The most basic strategies in infinitely repeated games are Always Defect and Always Cooperate. Always Defect involves a player choosing to defect in every round, regardless of previous outcomes. This strategy is considered one of the most **ruthless**, because it leaves no room for trust or collaboration. In contrast, Always Cooperate involves a player choosing to cooperate in every round. While this approach is very trusting, it is very vulnerable to exploitation and is therefore unsustainable in real scenarios.

6.2. **Grim Trigger.**

One of the most notable strategies in Repeated Games is **Grim Trigger.** The strategy starts with players cooperating and continued cooperation as long as both players have done so in the past. However, if either player defects even once, the strategy prescribes permanent defection from that point forward. Grim Trigger can be extremely effective in promoting cooperation, but its unforgiving nature means that it requires specific conditions to be rational, which will be discussed in the next section.

A lesser-known variant of Grim Trigger is Naive Grim Trigger. This strategy is similar to Grim Trigger but introduces a key difference: the player permanently defects only in response to the **other** player's defection, and not their own. While this is interesting, this approach is not a stable equilibrium, as it fails to hold both players equally accountable.

6.3. **Limited Retaliation.**
Limited Retaliation is a middle ground between Grim Trigger and the next strategy, Tit-for-Tat. This approach begins with cooperation and prescribes punishment of defection for a fixed number, k periods in response to any defection. After serving the retaliation period, the strategy returns to cooperation. Unlike strategies like Tit-for-Tat, the punishment duration in Limited Retaliation does not depend on the opponent's behavior after the initial defection.

6.4. **Tit-For-Tat.**
One the most famous strategies in repeated games is Tit-for-Tat (TFT). Similarly to Grim Trigger, TFT begins with cooperation in the first period. However, in following periods, the player mirrors the opponent's previous moves. TFT is both retaliatory and forgiving because it punishes defection immediately, but cooperation can resume as soon as the opponent returns to cooperation. TFT rose to fame due to its success in Axelrod's Tournament, where it performed exceptionally well against a range of alternative strategies.

Two other notable variants of TFT include Generous Tit-for-Tat and Tit-for-2-Tats (TF2T). Generous Tit-for-Tat operates like TFT but is more forgiving; it introduces a probability to cooperate even after the opponent defects, therefore reducing the likelihood of endless retaliation due to accidental defections. On the other hand, Tit-for-2-Tats only retaliates after two consecutive defections, which offers an even higher tolerance for mistakes.

6.5. **Win-Stay, Lose-Shift.**
Another well-known strategy is Pavlov, also known as Win-Stay, Lose-Shift (WS-LS). This strategy starts with cooperation. If the outcome of the previous round was either mutual cooperation (C, C) or mutual defection (D, D), the strategy repeats the same action in the next round. However, if the result was mismatched (e.g., C, D or D, C), the strategy switches its choice. Pavlov replays successful outcomes while adjusting in response to unfavorable ones, which makes this strategy simple but adaptive.

6.6. **What's Next.**
To conclude, these strategies represent a broad spectrum of approaches in repeated games, each with its own advantages, vulnerabilities, and implications for equilibrium. These strategies range in performance, depending on game structures, environment, noise and the discount factor, which will be analyzed further in later sections.

7. Conditions for Cooperation

Now that we have examined various strategies that can be used in repeated games, it is important to determine whether these strategies are actually rational to carry out. In other words, we must ensure that these strategies represent an equilibrium, meaning that they are stable and not just empty threats. This is where the One-Shot Deviation Principle comes in.

7.1. **One-Shot Deviation Principle.**

The One-Shot Deviation Principle states that a strategy is optimal if and only if it cannot be improved by making a single deviation at any decision point in the game. In other words, if you cannot get a better result by changing a single move, then changing multiple moves won't help either. Therefore, the current strategy is considered unimprovable.

This principle is especially useful in repeated games, whether they are very long finite games or infinitely repeated games. OSDP allows us to avoid the grueling and/or impossible task of checking every possible strategy.

OSDP has many real-world applications because it is often unrealistic to consider every possible plan in complex situations. However, by testing whether small, single-step changes can improve the outcome, we have a practical and reliable method to test strategies. This is one reason why many businesses and corporations use the One-Shot Deviation Principle in their decision-making processes.

7.2. **Sustainable Cooperation.**

Continuing to build onto the concept of making sure that strategies are rational, we now look at another method to identify subgame-perfect equilibrium. This approach involves the **discount factor**, or also referred to as the "shadow of the future." As previously discussed, the discount factor represents how much a player values future rewards compared to immediate ones.

When δ is close to 1, it indicates that players value future payoffs almost as much as present ones. This reflects perfect patience.

On the other hand, when δ is close to 0, it shows that players prioritize immediate gains and place little to no value on future outcomes.

The value of δ has a significant influence on player behavior and strategy. When players anticipate facing each other again in future rounds, they are more likely to cooperate in the present to avoid retaliation later on.

Applying this concept to the Repeated Prisoner's Dilemma, we see a change in player behavior between the one-shot version of the game. In a single round, defection strictly dominates cooperation. However, when the game is repeated, players begin comparing short-term and long-term consequences. For example, choosing to defect today might yield a temporary advantage, but it will likely trigger retaliation from the opponent in future rounds. In contrast, cooperating today might result in a lower immediate payoff, but it builds trust and stability that can lead to better outcomes in the long run.

Thus, the higher the value of δ, the more likely cooperation is to emerge. Players are incentivized by the threat of future retaliation to maintain cooperative behavior.

Cooperation can become a subgame-perfect equilibrium in strategies such as Grim Trigger or Tit-for-Tat as long as the discount factor is at a sufficiently high level. In these scenarios, cooperating becomes the rational choice when the long-term cost of provoking future defection (resulting in perpetual punishment with a payoff of P permanently) outweighs the short-term benefit from a single defection (the difference $T - R$). For this logic to hold, players must believe that defection will be punished in future rounds. As a result, they are better off cooperating consistently, even though defecting once may appear more profitable in the short term.

7.3. **Analysis of Strategies.**
In this subsection, we will be investigating four different examples of possible strategy sets two players could play in the Repeated Prisoner's Dilemma. We will be applying the One-Shot Deviation Principle to try and conclude for which strategies is cooperation a subgame-perfect equilibrium. If so, we want to find how high the discount factor has to be to satisfy the conditions.

1. (GRIM, GRIM)

First up, we will be analyzing when both players play Grim Trigger. To start, we calculate the payoff if both players agree to cooperate forever.

Since the two players cooperate, they have an infinite stream of 3's.

$$3 + 3\delta + 3\delta^2 + \ldots = \tfrac{3}{1-\delta}.$$

If either player defects now, their payoff is:

$$5 + \delta + \delta^2 + \ldots = 5 + \tfrac{\delta}{1-\delta}.$$

To make cooperation the preferred option, the following inequality must hold:

$$\tfrac{3}{1-\delta} \geq 5 + \tfrac{\delta}{1-\delta}.$$

Continuing to simplify the inequality, we get

$$3 \geq 5 - \delta 5 + \delta,$$

So, $\delta \geq \frac{1}{2}$ must hold for (GRIM, GRIM) to be sustainable.

Some concluding remarks: We notice that the greater the temptation is to defect (T), the higher δ must be to sustain cooperation. On the other hand, if the punishment (P) is harsh, then δ can be lower.

2. (LR-K, LR-K)

Now, we want to find the minimum value of δ to sustain cooperation when both players play Limited Retaliation strategies. Following a similar procedure, the payoff if both players agree to cooperate is $\frac{3}{1-\delta}$. Now, we want to calculate the payoff if one player defects now.

$$5 + \delta + \delta^2 + \ldots + \delta^K + 3\delta^{K+1} + \ldots$$

Using the formulas for a geometric sequence and infinite geometric series, we can simplify to get:

$$\frac{3}{1-\delta} \geq 5 + \frac{\delta(1-\delta^K)}{1-\delta} + \frac{3\delta^{K+1}}{1-\delta},$$

$$3 \geq 5(1-\delta) + \delta(1-\delta^K) + 3\delta^{K+1},$$

$$3 \geq 5 - 5\delta + \delta - \delta^{K+1} + 3\delta^{K+1},$$

$$3 \geq 5 - 4\delta + 2\delta^{K+1},$$

$$\delta^{K+1} \geq 2\delta - 1.$$

For a numerical value, we would need to plug in a value for K. Just to demonstrate, we plug in $K = 2$.

$\delta^3 \geq 2\delta - 1$, so δ has to be approximately greater than or equal to 0.63.

Something interesting to note here is that as K approaches infinity, δ approaches $\frac{1}{2}$, which makes sense because as K approaches infinity, this new version of Limited Retaliation fits the definition of Grim Trigger, whose δ value matches $\frac{1}{2}$.

3. (TFT, TFT)
Our next example is when both players play Tit-for-Tat as their strategy. For this analysis, we will need to break it down into cases.

Case 1: The previous outcome was (C, C). If the player sticks with TFT, their payoff is $\frac{3}{1-\delta}$.

Now what if the player deviates once? The payoffs will cycle in between (C, D) and (D, C). (Alternating payoff between 5 and 0). Since all of the odd payoffs are zero, we can disregard them. Only the even powers of delta contribute to the total payoff. We can express this with $\frac{5}{1-\delta^2}$.

We also have to multiply by $(1-\delta)$ to "normalize" or compare it to the stage payoff. Now let's write the inequality.

$$\frac{5}{1+\delta} \leq \frac{3}{1-\delta}.$$

Simplifying, we get:

$$5 - 5\delta \leq 3 + 3\delta,$$

$$8\delta \geq 2,$$

So, $\delta \geq \frac{1}{4}$.

Case 2: The previous outcome was (C, D) (You cooperated and your opponent defected). If the player sticks with TFT, similar to Case 1, your expected payoff is $\frac{5}{1+\delta}$. Now, if the player deviates to cooperation, the sequence then stays at (C, C) forever, or payoff $\frac{3}{1-\delta}$.

We would like to **prevent** deviation, so $\frac{5}{1+\delta} \geq \frac{3}{1-\delta}$ must hold. After simplifying equations, $\delta \leq \frac{1}{4}$.

Case 3: The previous outcome was (D, C). This is a similar outcome to the previous case, but with the roles reversed. If we stick to TFT, our payoff is $\frac{5\delta}{1+\delta}$ (Now, we have a payoff of 5 for odd terms).

If we deviate again to defection, then we get (D, D) or a payoff of 1 forever, which is $\frac{1}{1-\delta}$. We want to make deviation unappealing, so our inequality is

$$\frac{1}{1-\delta} \geq \frac{5\delta}{1+\delta}.$$

Simplifying, we get: $1 + \delta \geq 5\delta - 5\delta^2$,

$$5\delta^2 + 4\delta - 1 \geq 0.$$

The roots of this polynomial are: $\frac{4 \pm \sqrt{-4}}{10}$.

This polynomial yields no real solutions, so **no** δ satisfies this inequality, and deviation is **always** better if the last outcome was (D, C).

Case 4: This is our final case, where the previous outcome was (D, D). If we stick to TFT, our payoff is $\frac{1}{1-\delta}$.

If we deviate to cooperation, then we cycle between (D, C) and (C, D) forever, which is $\frac{5}{1+\delta}$. This leads to the same result as above, with no solution for δ. So, staying with TFT is not optimal after (D, D) is played.

Some interesting observations are that the δ value has to be exactly $\frac{1}{4}$ for Tit-For-Tat to be a Nash equilibrium. Although, TFT cannot be a subgame perfect equilibrium because even under these extremely fragile and special conditions, there is no incentive to resume cooperation when you defected in the last outcome.

4. (WS-LS, WS-LS)

Our last example is when both players play Win-Stay, Lose-Shift. Similarly to other examples, the payoff for cooperation forever is $\frac{3}{1-\delta}$. However, now we need to calculate what happens when one player deviates once.

Firstly, in round zero, the payoffs are (D, C) where the deviator gets 5. In the next round, both view (D, C) as a "lose" so both switch their actions, which results in (C, D), where the deviator gets 0. This results in an infinite alternation between these two options, whose payoff is $5 - 4\delta$ like in previous examples. After cross-multiplying and using the quadratic formula, we end up with $\delta \geq 0.25$ for deviation to be unappealing.

5. Concluding Remarks

Out of the four different examples we tested, the only strategy set that was not a subgame-perfect equilibrium was when both players played Tit-for-Tat.

8. The Folk Theorem

The folk theorem is one of the most important and useful tools used in repeated games. The main idea is that virtually any outcome in an infinitely repeated game can be sustained given that players are sufficiently patient, even if the outcomes in each subgame aren't Nash equilibriums.

First, let's introduce some new definitions.

Feasible payoffs are payoffs that can be achieved through some combination of the players' strategies. For example, in the Prisoner's Dilemma, this would be any weighted average of (3,3), (5,0), (0,5), or (1,1).

The minmax payoff is the minimum payoff a player can guarantee themselves no matter what the other player plays. In the Prisoner's Dilemma, this value is 0.

Finally, a payoff vector u $= (u_1, u_2, u_3, \ldots, u_n)$ is strictly-enforceable if and only if for every period the player gets strictly more than their minmax payoff. $(u_i > e_i)$

So, the Folk Theorem states that there exists a discount factor such that for all δ greater than this value, there exists a Subgame Perfect Equilibrium of the infinitely repeated game in which each player receives the average vector u. (or arbitrarily close).

In other words, the Folk Theorem tells us that for any payoff vector that is both feasible and enforceable, it is theoretically possible to be an equilibrium of the game given that players are sufficiently patient.

8.1. **What This Means.**

The Folk Theorem tells us that player behavior is completely changed in repeated games. In the Prisoner's Dilemma, the theorem tells us that not only can mutual cooperation be sustained given a high enough discount factor, but virtually any outcome in the feasible set can be achieved as a SPE. This results because repetition creates a large set of possible combinations of strategies. Especially in infinitely repeated games, the discount factor becomes so powerful that it can support outcomes that would be impossible in one-shot interactions.

For example, mutual cooperation in the Prisoner's Dilemma is unstable in a one-shot game because each player has more incentive to defect and receive 5 instead of 3. However, in the repeated version, where δ is sufficiently high, the short-term gain from defection ($5 - 3 = 2$) is outweighed by the long-term cost of triggering permanent punishment ($3 - 1 = 2$ in every future period).

8.2. **Implications.**
The Folk Theorem has a lot of implications for cooperation and social behavior. Firstly, it shows that player decisions are largely affected by repetition because repeated interactions open up lots of new possibilities for players to achieve desired results.

Second, the theorem shows how crucial patience is to sustain cooperation. As players become more patient, the set of sustainable equilibria expands drastically.

Third, the theorem shows how multiple equilibria can coexist in repeated games. This means that cooperation is possible even when it seems unlikely. However, this also complicates things because it is not clear which equilibrium players will actually play. Usually, the answer to these issues depends on outside factors like monitoring, social norms, and noise.

8.3. **Limitations to the Folk Theorem.**
Applications of the Folk Theorem are limited because it relies on a few strong assumptions that are unrealistic conditions in the real world. We will go more in-depth into these in the subsequent section.

Firstly, we have perfect monitoring. The theorem assumes players can perfectly observe each other's actions. In contrast, information is usually imperfect or noisy because players can accidentally choose the wrong action, forget past actions, or make observation errors. This can lead to the discouragement of cooperation even when no one intended to defect.

Next**,** The theorem assumes that players can keep track of infinite streams of payoffs and perform complex calculations every period to evaluate their next move. Real humans have these limitations that can prevent them from implementing the various strategies required by the Folk Theorem.The Folk Theorem also assumes that players are perfectly rational, where in the real-world this is impossible to control.

Despite these limitations, the Folk Theorem remains an extremely powerful tool used in game theory because it reveals that there are in-fact lots of possibilities for cooperation to emerge from repeated interaction.

9. Challenges to Cooperation

Building onto the last section, there is a big gap between Game Theory and real-world scenarios due to key assumptions mentioned in the previous section. This section will go over big obstacles in the real-world that stop cooperation from emerging even though using methods from Game Theory tells us that cooperation is possible.

9.1. **Noise and Miscommunication.**

Firstly, there is something called noise which is almost always present in real-world scenarios. Noise is what is referred to when a player accidentally deviates from their strategy in a game. This can result from miscommunication, accidental move, or various other factors.

Noise can present a lot of unintended consequences for certain strategies like Tit-for-Tat or Grim Trigger. Imagine a scenario where both players intend to cooperate but one player accidentally defects due to a random error. Following this, this mistake triggers an infinite alternation of mutual punishment or an infinite stream of defection respectively.

This is one of the main reasons why variants of Grim Trigger and Tit-for-Tat (like GTFT, LR-K, and TF2T) perform better in real-world experiments because they have more tolerance for occasional mistakes, allowing cooperation to still emerge after accidental defections.

Similarly, miscommunication is when players interpret the game rules, payoffs and strategies differently which can lead to conflict.

9.2. **Monitoring Issues.**

Repeated games typically involve perfect monitoring, where players perfectly observe their opponent's moves. However, real-world interactions involve imperfect monitoring because players cannot completely verify opponent choices.

When players cannot clearly observe whether their opponent cooperated or defected, they cannot reliably implement strategies that are based on opponent actions like TFT or Grim Trigger. So, this could lead to no punishment when the opponent defected on purpose. On the flip side, this could also lead to false accusations and unnecessary punishment if a player mistakenly believes their opponent defected when in reality, they cooperated.

Also, imperfect monitoring makes it harder to coordinate playing a specific equilibria because it is always safer to revert to a non-cooperative strategy when you aren't sure what the opponent is playing.

9.3. **Human Constraints.**

Another big issue is that in game theory, players are assumed to be perfectly rational, have the ability to do complex calculations, and have infinite memory. However, in the real-world this is not the case. Real players cannot keep track and calculate infinite streams of discounted payoffs or evaluate every possible strategy and then decide which one to play. This is especially outlined in the Folk Theorem where all of the possible subgame perfect equilibria exceed human brain capacity.

Also, humans have limited memory, so strategies like Tit-For-Tat, and Grim Trigger may not function as well because human players could forget crucial details of past interactions leading to strategies not working as intended.

Finally, all humans have bias, so factors like overconfidence, revenge, and loss aversion can affect players actions by deviating from their strategy to pursue a different outcome.

9.4. **Many Equilibria.**

The Folk Theorem reveals to us that infinitely repeated games have a vast amount of equilibria, many of which support cooperation. However, now we are presented a problem: in a real-life scenario, which equilibrium will the players actually play?

This can also further the challenge of cooperating due to two players having conflicting strategies, unclear communication, and overlying expectations for players to follow when they make decisions.

9.5. **Finite Horizons.**

As discussed in Section 5.3, finitely repeated games with known endpoints face the unraveling problem due to backward induction. This is relevant because in real-world scenarios, most interactions DO have definite endpoints.

So, for games to be played in the real world, organizations should avoid creating clear endpoints to eliminate the possibility triggering backward induction.

9.6. **Other Issues.**

There are also other differences between game theory and the real world. Just to name a few, firstly we have Reputation and how it affects decisions because players care about their reputation for future interactions.

Second, in the real-world, games are played with incomplete information where players are uncertain about opponent payoffs, strategies, and thoughts so this uncertainty could affect the probability of cooperation emerging.

9.7. **Concluding Remarks.**

In summary, this section went over common misconceptions and gaps between game theory and the real-world. In the next section we will explore how these theories and principles apply and scale to real-world problems and situations.

10. Real-World Applications

This section examines the various applications of repeated games to different areas of the real-world. From business to politics to biology, the principles learned in this paper provide countless valuable insights into real-life behavior. Firstly, we can draw several parallels from the concepts discussed to the Cold War.

Initially, the United States was the sole nuclear power of the world throughout the late 1940s. This all changed when an American weather monitoring plane detected trace amounts of materials in the air, only possible from a detonation of a nuclear bomb. However, the US hadn't conducted any tests recently. This confirmed American suspicions that the Soviet Union was able to conduct a successful nuclear test, which occurred on September 3, 1949, ending the US monopoly on nuclear weapons.

The US government was in chaos. Officials were experiencing considerable stress and many began to form opinions on how to handle the decision. Many advocated for a preemptive nuclear strike on the Soviet Union as an "aggressor for peace". The strategic position of the United States had switched from a strong advantage to a very vulnerable one almost overnight.

Following this, the threat of Mutually Assured Destruction (MAD) stroke fear into citizens of both countries. The theory hypothesized that if one nation attempted a preemptive strike on the other, then the other superpower had enough nuclear weapons to destroy them. This strategy is a Nash equilibrium because once both countries have been armed, neither side has any incentive to initiate an attack or disarm.

Tension reached a boiling point during the Cuban Missile Crisis in October 1962, where the incident showed both countries that nuclear confrontation would only end with the destruction of the world. Both countries recognized this, and in response, the United States removed missiles from Turkey, and the Soviet Union removed missiles from Cuba.

Although the reconciliation period started slow, many gradual changes led to the reduction of tensions throughout the later Cold War period. Inititally, throughout the late 1960s and the early 1970s, both sides agreed to limit testing of and the spread of nuclear weapons to other countries through numerous treaties. This snowballed to the Strategic Arms Limitation Talks (SALT) in the 1980s, where both countries stated numerical limits on the amount of nuclear weapons that a country could have. Finally, in the late 1980s, tensions grew to an all-time low with large reforms that completely changed the international landscape.

Firstly, the United States and the Soviet Union established direct communication lines to communicate more clearly to prevent misunderstandings from escalating into full-scale wars in the future.

Second, both sides signed the Strategic Arms Reduction Treaty (START I), which was the largest nuclear arms reduction in history. In detail, almost 30 percent of nuclear warheads on each side were destroyed and this created permanent infrastructure that promoted cooperation.

What we can learn from this is that both sides learned to match their opponent's actions while simultaneously avoiding escalation. Military and political leaders from both sides recognized the importance of maintaining good communication and avoiding provocative actions that would deter cooperation.

In the end, nuclear weapons transformed from tools of war to one of the most important instruments in causing diplomacy between two superpowers.

In the modern day, NATO and Russia have continued clashes that closely mirror Cold War dynamics which make this all the more important to learn from. Russian nuclear threats in the Russian-Ukraine conflict closely resemble Cold War levels of tension, with both sides preparing to deploy their military and modernizing their nuclear weapons. On top of this, in the Cold War, there were only two main nuclear superpowers. However, today, there are many nuclear powers that add to the complex dynamics of the world, with China, North Korea, Pakistan, India and more possessing nuclear arsenals. Along with new technologies and AI being the forefront of modern discussion, this makes the study of game theory and repeated games all the more important.

11. Conclusion

11.1. What's Next?

As game theory continues to evolve, this opens many doors for future research. Game theory has a lot of applications to actual human behavior, and with the increasingly popular AI and machine learning sectors, this opens new questions about how game theory can be applied there. For anyone intrigued, game theory has various diverse applications, from computer science, to political science, to psychology, to business, to biology, and countless others. The mathematical foundations explained here provide the essential skills needed to understand the insights of cooperative potential of our increasingly complex world.

12. References

[1] Axelrod, Robert. *The Evolution of Cooperation.* Basic Books, 1984.

[2] Axelrod Python Library Documentation. "Play Contexts." *Axelrod Documentation.*

[3] Axelrod Python Library Documentation. "Index." *Axelrod Documentation.*

[4] Gratch, Jonathan. "Game Theory Introduction." *USC ICT Readings.*

[5] MIT OpenCourseWare. "Chapter 11: Subgame Perfect Equilibrium." *14.12 Economic Applications of Game Theory*, Fall 2012.

[6] Osborne, Martin J. and Ariel Rubinstein. *A Course in Game Theory.* MIT Press, 1994.

[7] Policonomics. "Prisoners Dilemma."

[8] Ray, Debraj. "One-Shot Deviation Principle." *NYU Course Materials.*

[9] Slantchev, Branislav L. "Repeated Games." *UCSD Course Materials.*

[10] Stadelis, Stavros. "Game Theory Week 6." *UC Berkeley Haas School of Business.*

[11] Study.com. "Payoff Matrix: Economics Theory, Calculation & Template."

[12] Veritasium. "The Most Important Problem in Game Theory (Does it pay to be nice?)." *YouTube*, 2023.