On the Folk Theorem in Infinitely Repeated Games with Discounting

July 15, 2024

Jiana Raj Shroff

1 Abstract

Assuming a basic knowledge of game theory and graph theory, this paper discusses different versions of the Folk theorem in infinitely repeated games, applied when players are very patient and consider future interactions heavily; essentially when the discount factor approaches 1. This paper aims to provide a concise understanding of how equilibria in the folk theorem can deliver virtually any average payoff outcome that is achievable in a single-stage game.

2 Table of Contents

Contents

3 Introduction

The Folk Theorem gets its name from being a part of game theory's oral tradition or "folk wisdom" long before it was recorded in print. In game theory, folk theorems state that any outcome can be a Nash Equilibrium in an infinitely repeated game as long as it provides a per-period payoff that is, on average, greater than the minimum average payoff each player can secure by unilaterally deviating from the cooperative agreement. In other words, The Folk Theorem suggests that if the players are patient enough and far sighted (i.e. if the discount factor $\delta \to 1$, then repeated interaction can result in a feasible and individually rational perfect equilibrium This paper proves the existence of such equilibria one step at a time, initially proving the existence of a nash equilibrium in a one-time game, and building upon it by proving Freidman's Theorem on the existence of Nash equilibria in a repeated game, to provide readers with a basic understanding. Then we prove the existence of perfect equilibria in two player games, using the prisoner's dilemma as an example throughout.

4 Acknowledgements

I'd like to thank Dr. Simon Rubenstein-Salzedo and Mr. Lucas Fagan for their assistance with writing this paper.

5 The Model

I will begin this discussion by introducing the model of an infinitely repeated game, stating and defining key terms used to understand the Folk Theorem. I will do this using 'The Prisoner's Dilemma' as an example of a strategic form game. We will use this model to define what an infinitely repeated game is.

		Player 2: $cooperate(C)$	Player 2: $defect(D)$
Player 1	$cooperate$	$1, 1$	$-1, 2$
Player 1	$defect$	$2, -1$	$0, 0$

Fig 6.1 - The Prisoner's Dilemma

The building block of an infinitely repeated game is called the stage game (which is also known as the sub game), which is the game that is repeated. Let G represent the game as a whole, and g denote the stage game. We define g as:

Definition 5.1. The stage game, g, in a strategic form consists of a triplet $g = (I, (S_i)_{i \in I}, (U_i)_{i \in I})$, where

1. I is the set of players ($I = 1, \ldots, n$)
2. For each $i \in I, S_i$ is a Player i's set of strategies, where $S = \prod_{i \in I} S_i$ is the set of strategy profiles.
3. For each $i \in I, U_i : S \to \mathbb{R}$ is the payoff (von-Neumann-Morgenstern utility) function on the set of strategy profiles. Implicitly, it is as if we had an outcome function $g : S \to A$ and a utility function $g : S \to \mathbb{R}$. Then, for every $s \in S$, set $U_i(s) = U_I(g(s))$. (In simpler terms, a payoff function maps the n-tuple strategies s to real number values)

Remark 5.2. A denotes the set of action profiles, where $a_i \in A_i$ for player i.

It is important to note that:

1. We assume that the strategy spaces are compact sets and that the payoff functions are continuous.
2. For all player(s) other than i, their strategy space(s) are denoted by S_{-i} (this is true for strategies, mixed strategies, payoff vectors, and other components, which we will define later).

Definition 5.3. A mixed strategy σ is a strategy consisting of possible moves and a probability distribution which corresponds to how frequently each move is to be played.

This completes the knowledge we require about the components of a stage game. We discuss the components, and ultimately, define an infinitely repeated (IR) game below.

In an IR game, for all S, $a \in S$ can be referred to as an outcome of G. $G(T)$ denotes the game that results when G is successively played T times (T is a positive integer). For $t = 1, 2, \ldots, T$, if $a^t \in A$ denotes the outcome of the game $G(T)$ at time t, player i's payoff in $G(T)$ is given by:

$$\frac{1}{T} \sum_{t=1}^{T} U_i \left(a^t\right) \tag{5.1}$$

In each period $t = 0, 1, 2, \ldots$, players players simultaneously choose an action $a_i \in A_i$ and the chosen action profile $(a_1, a_2, a_3, \ldots, a_n)$ is observed by all players. Additionally, offs. That is, if $(a_1(t), \ldots, a_n(t))$ is the vector of actions played in period t. Then the player moves to period $t+1$ and the game continues in the same manner. Each information set of each player i associated with a finitely repeated game corresponded to a history of action profiles chosen in the past. We can represent each information set of player i by a history:

$$h^0 = (\emptyset), h^1 = (a^0 := (a_1^0, \ldots, a_n^0), \ldots, h^t = (a^0, a^1, \ldots, a^{t-1})$$

We denote the set of all histories at time t as H^t. For example, if the stage game is the prisoner's dilemma, at period 1, there are 4 possible histories;

$$H^1 = \{(C_1^0, C_2^0), (C_1^0, D_2^0), (D_1^0, C_2^0), (D_1^0, D_2^0)\}.$$

Therefore, for time t, H^t consists of 4^t possible histories.

We define strategies in IR games as:

Definition 5.4. A strategy for player i in the game $G(\infty)$ is a function s which selects, for any history of play, an element of A_i. Formally, $s_i = (s_i^1, s_i^2, \ldots, s_i^T)$, where $s_i^t \in A_i$ and for $t > 1, s_i^t : A^{t-1} \to A_i$.

An n-tuple of strategies, s, inductively defines an outcome path $(a^1(s), a^2(s), \ldots, a^T(s))$ of the game $G(T)$ as follows: $a^1(s) = s^1$ and for $t > 1$, $a^t(s) = s^t\left(a^1(s), \ldots, a^{t-1}(s)\right)$

It is denoted by the function:

$$s_i : \bigcup_{t \geq 0} H^t \to A_i \tag{5.2}$$

A strategy of incredible importance in this paper is the 'Grim Trigger Strategy' (a strategy where you always cooperate until someone defects), which is represented as the following:

$$s_i\left(h^t\right) = \begin{cases} C_i & \text{if } t = 0 \text{ or } h^t = (C, C, \ldots, C) \\ D_i & \text{otherwise} \end{cases} \tag{5.3}$$

To determine the payoffs, we suppose the strategies $s_1, \ldots, s_n$ are played which lead to an infinite sequence of action profiles $a^0, a^1, \ldots, a^t, a^{t+1}, \ldots$, therefore making the payoff of player i in a repeated game:

$$\sum_{t=0}^{\infty} \delta^t u_i\left(a^t\right) \tag{5.4}$$

where 'δ' is used to represent the discount factor, which is defined below.

Definition 5.5. In a game, a discount factor is a value $0 \leq \delta \leq 1$ used to represent a player's pure time preference.

Its interpretation corresponds to

$$\delta = e^{-r\Delta} \tag{5.5}$$

In the above equation, r denotes the rate of time preference and Δ denotes the length of the period. The discount factor can represent the probability of the game terminating at the end of each period.

To compute payoffs in a repeated game, consider is the strategy profile $s_i(h^t) = C_i$ for all $i = 1, 2$ for all h^t. In this case, player 1's repeated game is given by:

$$\sum_{t=0}^{\infty} \delta^t = \frac{1}{1-\delta} \tag{5.6}$$

Using the above information, we can finally define an infinitely repeated game.

Definition 5.6. Given a stage game G, let (∞, δ) denote the infinitely repeated game (IR)in which G is repeated infinitely and the players share a discount factor. For each t, the outcomes of the $t-1$ preceding plays of the stage game are observed before the t^{th} stage begins. Each player's payoff in $G(\infty, \delta)$ is the present value of the player's payoffs from the infinite sequence of stage games.

(in simpler terms, an IR game is a situation where players always believe that the game extends one more period with high probability)

The notation and definitions presented above are adapted from Rubenstein (1979) [8].

6 Key Concepts in Folk Theorems

Definition 6.1. In this paper, a Nash Equilibrium in a subgame is an action profile a_i^* ($a_i^* \in A_i$) for player i, where, for all $i \in N$ and $s_i \neq s_i^*$,

$$u_i \left(s_i^*, s_{-i}^* \right) \geq u_i \left(s_i, s_{-i}^* \right)$$

(In other words, A Nash equilibrium is a strategy profile s_i^* with the property that no player i can do better by choosing a strategy different from s_i^*, given that every other player $-i$ adheres to s_{-i}^*)

(For reference, in the Prisoner's Dilemma, the Nash Equilibrium is (D, D).

We now prove Nash's equilibrium using definition 6.2, and after this, we will explain how it applied to pure strategies [4].

Theorem 6.2. *(Nash 1950) There exists a mixed Nash Equilibrium in all strategic form games.*

To prove this, we will apply Kakutani's fixed point theorem to the players' *reaction correspondences*, where the Nash equilibrium is the said fixed point.

Definition 6.3. Player i's reaction correspondence 'r' maps each strategy profile s to mixed strategies that maximise i's payoff when $-i$ plays s_{-i}. In particular,

$$r(s_{-i}) = \text{argmax}_{s_i \in S_i} u_i \left(s_i, s_{-i} \right)$$

Theorem 6.4. *(Kakutani, 1938) In a non - empty, compact, and convex space, under certain conditions, there will be a point within this space that gets mapped by a special function to a set containing itself.*

A fixed point of r is a s such that $s \in r(s)$, so that, for player i, $s_i \in r_i(s)$. A fixed point of r is the Nash Equilibrium.

Therefore, to satisfy Kakutani's fixed - point theorem for $r : S \to S$ in this paper, we need to satisfy 4 conditions:

1. S is a compact, convex nonempty subset of a Euclidean space.
2. $r(s)$ is non-empty for all s.
3. $r(s)$ is convex for all s.
4. it has a closed graph (this is the same as upper - hemicontinuity)

Before we prove these conditions, consider the following definitions:

Definition 6.5. Convexity A set $C \subset \mathbb{R}^m$ is convex if for every $x, y \in C$ and $\lambda \in [0,1]$, $\lambda x + (1-\lambda)y \in C$. For vectors $x^0, \ldots, x^n$ and nonnegative scalars $\lambda_0, \ldots, \lambda_n$ satisfying $\sum_{i=0}^{n} \lambda_i = 1$, the vector $\sum_{i=0}^{n} \lambda_i x^i$ is called a convex combination of $x^0, \ldots, x^n$. For example, a cube is a convex set in $\mathbb{R}^3$; a bowl is not.

Definition 6.6. n-simplex An n-simplex, denoted $x^0 \cdots x^n$, is the set of all convex combinations of the affinely independent set of vectors $\left\{x^0, \ldots, x^n\right\}$, i.e.

$$x^0 \cdots x^n = \left\{ \sum_{i=0}^{n} \lambda_i x^i : \forall i \in \{0, \ldots, n\}, \lambda_i \geq 0; \text{ and } \sum_{i=0}^{n} \lambda_i = 1 \right\}.$$

Each x^i is called a vertex of the simplex $x^0 \cdots x^n$ and each k-simplex $x^{i_0} \cdots x^{i_k}$ is called a k-face of $x^0 \cdots x^n$, where $i_0, \ldots, i_k \in \{0, \ldots, n\}$.

We will now prove the existence of a Nash Equilibrium,

Proof. **Condition 1:** Each S_i is a simplex of dimensions (the number of $S_i - 1$). Each player's payoff function is linear, and is therefore continuous in his own mixed strategy, therefore, condition 1 is satisfied.

Condition 2: Continuous functions on compact sets attain maxima, hence satisfying condition 2.

Condition 3: We prove this by contradiction. If $r(s)$ were not convex, there would be a $s' \in r(s)$, a $s'' \in r(s)$ and a $\lambda \in (0,1)$ such that $\lambda s' + (1-\lambda)s'' \notin r(s)$. But for each player i,

$$u_i(\lambda s_i' + (1-\lambda)s_i'', s_{-i}) = \lambda u_i(s_i', s_{-i} + (1-\lambda)u_i(s_i'', s_{-i}$$

therefore, if s_i' and s_i'' are the best responses to s_{-i}, then so is their average. This contradicts the statement we made, therefore confirming statement 3.

Condition 4: We also prove this by contradiction. If violated, there is a sequence

$$(s^n, \varrho^n) \to (s, \varrho), \varrho^n \in r(s^n), \varrho \notin r(s)$$

Then, $\varrho_i \notin r_i(s)$ for player i. Thus, there is an $\varepsilon > 0$ and a s_i' such that $u_i(s_i', s_{-i}) > u_i(\hat{s}_i, s_{-i}) + 3\varepsilon$. Considering that u_i is continuous and $(s^n, \varrho^n) \to (s, \varrho)$, we have,

$$u_i\left(s_i', s_{-i}^n\right) > u_i\left(s_i', s_{-i}\right) - \varepsilon > u_i\left(\hat{s}_i, s_{-i}\right) + 2\varepsilon > u_i\left(\hat{s}_i^n, s_{-i}^n\right) + \varepsilon.$$

Therefore, s_i' does strictly better against s_i^n than ϱ_i^n does, which contradicts the claim $\varrho_i^n \in r_i(s^n)$. Thus, condition 4 is verified.

Definition 6.7. A **minimax** strategy is a decision rule used to minimize the worst-case potential loss, where each player minimizes the maximum payoff possible for the other. For player i, it is denoted by:

$$\left(M_1^i, \ldots, M_{i-1}^s, M_{i+1}^i, \ldots, M_n^i\right) \in \arg\min_{s_{-j}} \max_{s_i} G_i\left(s_i, s_{-i}\right).$$

[9]

Definition 6.8. v is player i's reservation value, where we refer to $(v_1^*, \ldots, v_n^*)$ as the minimax point, and v_i^* is the smallest payoff that the players can keep player player i below. (In any equilibrium of G, whether or not it is repeated, player i's expected average payoff must be at least v_i^*. v_j^* can be denoted by:

$$v_j^* \equiv \max_{a_j} g_j\left(a_j, M_{-j}^j\right) = g_j\left(M^j\right).$$

Definition 6.9. In a repeated game, a strategy profile is a subgame perfect equilibrium if it represents a Nash equilibrium of every subgame of the original game. This is also called a perfect equilibrium [2].

Remark 6.10. It is important to note that the strategy profiles of perfect equilibria may differ from those in Nash Equilibrium. We demonstrate this in the following proposition.

Proposition 6.11. *In the infinitely repeated prisoners' dilemma, if $\delta \geq \frac{1}{2}$ there is an equilibrium in which (C, C) is played in every g*

Proof. Considering the 'Grim Trigger' strategy profile. To prove that there is no profitable single deviation, suppose D has already been played. Then, player i has two choices:

1. Play C for a payoff of $-1 + \delta \times 0 + \delta^2 \times 0 + \ldots = -1$ (referring to equation 6.4)
2. Play D for a payoff of $0 + \delta \times 0 + \delta^2 \times 0 + \ldots = 0$

Assuming that player i will want to maximise his payoff, player i should play D. Now, assume D has not been played, then, player i will have 2 choices:

1. Play C for a payoff of $1 + \delta + \delta^2 + \ldots = 1/(1-\delta)$
2. Play D for a payoff of $2 + \delta 0 + \delta^2 \times 0 + \ldots = 2$.

Therefore, we prove that, for $\delta \geq \frac{1}{2}$ it is best to play C. So we have a perfect equilibrium.

Remark 6.12. We cannot say that people always cooperate when they interact. Cooperation is only one possible SPE outcome. There are many others.

1. For any δ, there is a SPE in which players play D in every period this is because (D, D) is the Nash Equilibrium.
2. For $\delta \geq \frac{1}{2}$, there is a SPE in which the players play D in the first period and C in every following period.
3. For $\delta \geq \frac{1}{\sqrt{2}}$, there is a SPE in which the players alternate between (C, C) and (D, D).

In an infinitely repeated game, one uses the single-deviation principle in order to check whether a strategy profile is a subgame-perfect Nash equilibrium. In such a game, single-deviation principle takes a simple form and is applied through augmented stage games.

Definition 6.13. Augmented stage game for s^* and g is the same game as the stage game in the repeated game except that the payoff of each player i from each terminal history h of the stage game is:

$$U_i\left(h \mid s^*, g\right) = u_i(h) + \delta PV_{i,t+1}\left(g, h, s^*\right)$$

where $PV_{i,t+1}(h, g, s^*)$ is the present value of player i at $t+1$ from the payoff stream that results when all players follow s^* starting with the history $(g, h) = (a_0, \ldots, h)$, which is a history at the beginning of date $t+1$.

Definition 6.14. An outcome of a game is **Pareto dominated** if some other outcome would make at least one player better off without making another player worse off.

7 Existence of Nash Equilibria in IR games with discounting

In this section, we will be proving the existence of a Nash Equilibrium in infinitely repeated games with discounting [5] [7] .

We normalise the payoffs of game g so that

$$(v_1^*, \ldots, v_n^*) = (0, \ldots, 0)$$

. Let
$U = \{ \{(v_1, \ldots, v_n) \mid \exists\,(a_1, \ldots, a_n) \in A_1 \times \cdots \times A_n$
with $g\,(a_1, \ldots, a_n) = (v_1, \ldots, v_n)\}\,,$

where $v \in V$, and V is a convex hull of U, consisting of feasible payoffs, these are strictly individually rational payoffs.

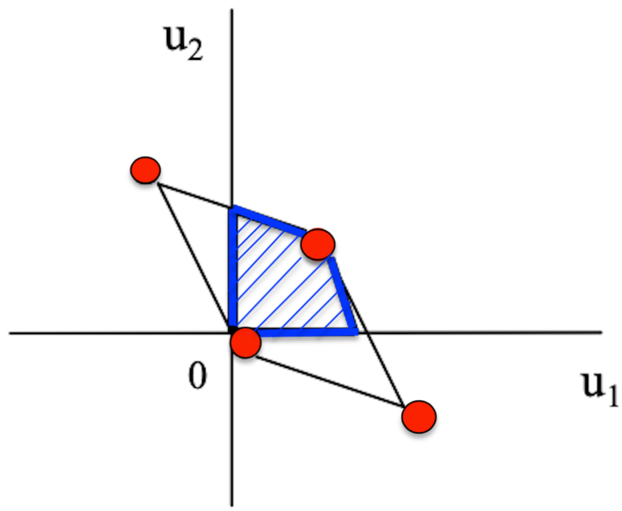

Fig 7.1 - Individually rational payoffs in the Prisoner's Dilemma

$$V^* = (v_1, \ldots, v_n) \in V | v_i > 0 \forall i$$

V^* consists of feasible payoffs that pareto dominate the minimax point.

7.1 The Classical Folk Theorem

We now introduce the primary folk theorem, which is more commonly known as the classical folk theorem.

Theorem 7.1. *(Friedman, 1971) For any $(v_1, \ldots, v_n) \in V^*$, if players discount the future sufficiently little, there exists a Nash equilibrium of the infinitely repeated game, where, for all i, player i's average payoff is v_i [3].*

The following one is short and basic, but will be built upon in the rest of the paper, providing more information.

Proof. Let $(s_1, \ldots, s_n \in \prod_{i \in I}^n A_i$ be a vector of strategies, or, if necessary, correlated strategies, such that $g(s_i \ldots, s_n) = (v_1, \ldots, v_n)$. Using the concept of the grim trigger strategy (6.3), so when player i plays s_i until some player $-i$

deviates from s_{-i} (if more than one player deviate simultaneously, we can suppose that the deviations are ignored). Then, we assume player i plays M_i^{-i} to minimax his opponents payoff. We can say that these strategies form a Nash Equilibrium if there isn't a lot of discounting (in other words, players' pure - time preferences remain the same). This is because any momentary gain may be accrue to player $-i$ if she deviates from s_{-i} and is swamped by the prospect of being minimaxed forever after. Therefore, both players will be playing their best responses (for example, in the prisoner's dilemma, defecting is the best response for each player, therefore the nash equilibrium is (D, D), however, in the following proposition, we show how equilibria differ from nash equilibria in a game played once).

[6]

For the above proof, Forges, Mertens and Neyman hypothesised that ensuring that the v_i's are positive is important.

7.2 The Folk Theorem in the context of Pareto Dominance

Theorem 7.2. *If $(v_n, \ldots, v_n) \in V^*$ Pareto dominates the payoffs $(u_1, \ldots, u_n)$ of a (one-shot) Nash equilibrium $(e_1, \ldots, e_n)$ of g, then, if players discount the future sufficiently little discounting, there exists a perfect equilibrium of infinitely repeated game where, for all i, player i's average payoff is v_i.*

Proof. Suppose that players play actions that sustain $(v1, \ldots, vj)$ until someone deviates, after which they play $(e_1, \ldots, e_n)$ forever. With sufficiently little discounting, this behavior constitutes a perfect equilibrium.

7.3 Aumann-Shapely/Rubinstein strategies

We now look into Aumann-Shapely/Rubinstein [1] [9] strategies, as in the theorem below.

Theorem 7.3. *(Aumannmann-Shapley/Rubinstein): For any $(v_1, \ldots, v_n) \in V^*$, there exists a perfect equilibrium in the infinitely repeated game with no discounting, where for all i, player i's expected payoff each period is v_i.*

Here, it is important to note that If there is no discounting, the sum of single-period payoffs cannot serve as a player's repeated game payoff since the sum may not be defined. Aumann and Shapley use the average payoff, and Rubinstein considers both this and the overtaking criterion, and the sketch of the proof we offer corresponds to this latter rule.

Remark 7.4. Aumann-Shapely and Rubenstein arguements suggest that only pure strategies have been played, which leads to a smaller equilibrium set.

Proof. We use the idea of Grim-Trigger strategies (refer to 6.3).

Consider the following events:

1. As long as previously, everyone has cooperated, player i will continue to play their strategies s_i, leading to a payoff vector of v_i.
2. some player $-i$ defects
3. they are minimaxed long enough to eliminate possible gains obtained from their defection.
4. Once the punishment is completed, players continue to play their s_is.

These events lead to 2 outcomes:

1. Punishers are compelled to carry out their minimaxing by the possibility that, should one of them stray from their plan, the others will minimax her for a length of time that will render the deviation unnecessary.
2. Her punishers will be punished if any one of them deviates.

Therefore, it is suitable to conclude that considering these punishments, a player's best response (assuming everyone has previously cooperated) is cooperation. And considering that every player is cooperating, the game is in a Nash Equilibrium.

Proposition 7.5. *We cannot use the Aumann-Shapley/ Rubinstein (AS/R) strategies once there is discounting.*

Remark 7.6. If there was discounting, strategies would not be able to sustain all individually rational points.

Proof. Using the following version of the Prisoner's Dilemma,

		Player 2	
		$cooperate(C)$	$defect(D)$
Player 1	$cooperate$	$1, 1$	$0, -2$
	$defect$	$-2, 0$	$-1, -1$

Fig 7.3.1 - Another version of The Prisoner's Dilemma

the minimax point is (0,0), and therefore, a "folk theorem" would require us to sustain strategies that choose (C, C) for $(\varepsilon + 1)/2$ (where $0 < \varepsilon < 1$) and (D, D) for the remaining amount of time (refer to proposition 6.11). Keeping the proof for proposition 6.11 in mind, note that for δ near 1, these strategies yield average payoffs of approximately $(\varepsilon, \varepsilon)$, which are individually rational.

However, Rubenstein's theorem states that such behavior cannot be part of an AS/R equilibrium. Suppose player i played C in a period where she was supposed to play D, in an AS/R equilibrium, player $-i$ would punish i by playing D for a time long enough to deem the advantage player i's deviation brought forward redundant. i's gain from the deviation is 1, and i's best response to D is C, with a payoff of 0. So, if the punishment lasts t_1 periods, t_1 must satisfy:

$$\frac{\delta\varepsilon(1-\delta^{t_1})}{1-\delta} > 1 + t_1 \times 0 = 1$$

that is,

$$t_1 > \frac{\log\left(\frac{\varepsilon\delta-1+\delta}{\delta\varepsilon}\right)}{\log\delta}. \tag{7.1}$$

equation 7.1 is satisfied as long as

$$\delta > \frac{1}{1+\varepsilon} \tag{7.2}$$

However, in order to punish player i, $-i$ would need to suffer a payoff of -2 for t_1 periods. ds. To induce him to submit to such self-laceration, he must be threatened with a t_2-period punishment, where

$$-2\frac{(1-\delta^{t_1})}{1-\delta} + \frac{\delta^{t_1}\varepsilon\left(1-\delta^{t_2-t_1+1}\right)}{1-\delta} > 1$$

when

$$t_2 > -1 + \log\frac{\delta^{t_1}\varepsilon - 3 + 2\delta^{t_1} + \delta}{\varepsilon} / \log\delta \tag{7.3}$$

t_2 exists when

$$\delta^{t_1}\varepsilon - 3 + 2\delta^{t_1} + \delta > 0$$

which requires a δ such that

$$\delta > \left(\frac{2}{2+\varepsilon}\right)^{1/t_1} \tag{7.4}$$

7.4 is stricter than 7.2 since

$$(2/(2+\varepsilon))^{1/t_1} > \frac{1}{1+\varepsilon}$$

Continuing iteratively, we find that, for successively higher order punishments, δ is bounded below by a sequence of numbers converging to 1 . Since $\delta < 1$, however, this is impossible, and so an AS/R equilibrium is impossi-

ble.

7.4 The existence of Perfect Equilibria

In theorem 7.1, we introduce the existence of Nash equilibria in repeated games. Now, we prove the existence of perfect equilibria.

Theorem 7.7. *Suppose that the set of feasible payoffs of G is I-dimensional. Then, for any feasible and strictly individually rational payoff vector v, $\exists\$\underline{\delta}$ such that, for any $\delta \geq \underline{\delta}$, there is a perfect equilibrium s^* of G such that the average payoff vector to s^* is v_i for player i.*

Remark 7.8. The Folk Theorem says that anything that is individually rational is possible.

Proof. When players play an action profile with payoff v, if some player i deviates, he will be punished by getting minimaxed by the others, as they play some s_{-i} for T periods. From this, the maximum he can get is $\underline{v_i}$. After the punishment, the players other than i get rewarded for carrying out their punishments. In order for this to happen, switch to an action profile that allows players $-i$ a

payoff v_{-i}, where $_{-i} > \underline{v_{-i}}$

Keep in mind, the above proof was a basic version and doesn't take into account several anomalies. We now introduce a slight refinement of Theorem 7.9 by Fudenberg and Maskin and a proof that follows it, which, hopefully, provides a firm understanding of the existence of perfect equilibria.

Theorem 7.9. *For any $(v_1, v_2) \in V^*, \exists \underline{\delta} \in (0,1)$ such that, $\forall \delta \in (\underline{\delta}, 1)$, there exists a subgame perfect equilibrium of the infinitely repeated game in which player i's payoff is v_i when players have a discount factors δ.*

Remark 7.10. Since the proof of this theorem for 3 or more players is pretty involved, we prove this for two players.

Proof. Let M_1 be player ones's minimax strategy against two, and let M_2 be player two's minimax strategy against one. Take

$$\bar{v}_i = \max_{a_1, a_2} g_i\left(a_1, a_2\right)$$

.

For $(v_1, v_2) \in V^*$, choose $\underline{v}$ and $\underline{\delta}$ such that, for $i = 1, 2$,

$$v_i > \bar{v}_i(1 - \underline{\delta}) + \underline{\delta} v_i^{**}, \tag{7.5}$$

where

$$v_i^{**} = \left(1 - \underline{\delta}^{\underline{\nu}}\right) g_i\left(M_1, M_2\right) + \underline{\delta}^{\underline{\nu}} v_i, \tag{7.6}$$

and

$$v^{**} > 0 \tag{7.7}$$

Finally, to confirm that $\underline{v}$ and $\underline{\delta}$ exist, choose $\underline{\delta}$ close enough to 1 so that

$$v_i > \bar{v}_i(1 - \underline{\delta}) \tag{7.8}$$

and (7.7) holds for when

$$\underline{v} = 1$$

If, with $\underline{v} = 1$, (7.5) is violated, considering that we are raising $\underline{v}$. From 7.8, 7.5 will eventually be satisfied. For $\underline{\delta}$ close to 1, (7.6) declines as $\underline{v}$ increases, by taking $\underline{\delta}$ near enough to 1 we can ensure that (7.7) will be satisfied for the first $\underline{v}$ for which (7.5) holds.

(7.5) guarantees that player i prefers receiving v_i forever to receiving his maximum possible payoff $(\bar{v}_i)$ once, then receiving $g_i\left(M_1, M_2\right)$ for $\underline{\nu}$ periods, and receiving v_i thereafter. (7.7) ensures that being punished for deviating is still better than receiving the reservation value, zero, forever. Clearly, for any $\delta > \underline{\delta}$ there is a corresponding $\nu(\delta)$ such that (7.5) and (7.7) hold for $(\delta, \nu(\delta)\ )$.

Let (s_1, s_2) be correlated one-shot strategies corresponding to $(v_1, v_2) :\ g_i\left(s_1, s_2\right) = v_i$. Consider the following repeated game strategies for player i : (A) Play s_i each period as long as (s_1, s_2) was played last period.

After any deviation from (A): (B) Play $M_i\nu(\delta)$ times and then start again with (A). If there are any deviations while in phase (B), then begin phase (B) again.

These strategies form a subgame-perfect equilibrium. This is because (7.5) guarantees that deviation is not profitable in phase (A). In phase (B), player i receives an average payoff of at least v_i^{**} by not deviating. If he deviates, he can obtain at most 0 in the first period (because his opponent, $-i$, is playing M_{-i}), and thereafter can average at most v_i^{**}. Hence deviation is not profitable in phase (B). 

7.5 The lower - hemicontinuity of perfect equilibria payoffs

Since we have proven the existence of perfect equilibria, we now investigate the lower hemi-continuity of the perfect equilibrium average payoff correspondence. This is where the discount factor is the independent variable. Using this knowledge, we investigate whether Theorem 7.1 holds for perfect equilibrium, instead of Nash Equilibrium.

Theorem 7.11. *Let $V(\delta) = \{(v_1, \ldots, v_n) \in V^* \mid (v_1, \ldots, v_n)$ are the average payoffs of a perfect equilibrium of the infinitely repeated game where players have discount factor $\delta\}$. The correspondence $V(\cdot)$ is upper hemicontinuous at any $\delta < 1$.*

Remark 7.12. We find it easier to show that $V(\cdot)$ cannot be lower hemicontinuous at $\delta < 1$.

Proof. We prove this by contradiction. Using the version of the prisoner's dilemma in Fig. 6.1 Using information from remark 7.3, For $\delta < 1/2$ there are no equilibria of the repeated game other than players' choosing $\cdot D$ every period. However at $\delta = 1/2$ many additional equilibria appear, including playing C each period until someone deviates and thereafter playing D. Thus $V(\cdot)$ is not lower hemicontinuous at $\delta = 1/2$.

This corresponds to the initial proof of the Nash Equilibrium (see Theorem 7.4) and the fact that it was a closed graph (upper hemi-continuity). With this, we conclude the information that we had to share in this paper.

8 Axelrod's tournament

Considering that you've made it so far, you are clearly interested in the theory of repeated games! Let me leave you with one last bit of information; about an experiment carried out by Robert Axelrod (University of Michigan) in 1979, where game theorists submitted strategies run by computers to play a repeated version of the prisoner's dilemma [**?**].

Some of the strategy submissions included:

1. Always cooperate
2. Always Defect
3. Random, where players cooperated 50% of the time
4. Tit for tat, where the strategy cooperates on the first move, and then does whatever its opponent has done on the previous move.
5. Tit for two tats

The winner of Axelrod's tournament was the 'Tit for Tat' strategy, this is becauase, when matched against the all-defect strategy, TIT FOR TAT strategy always defects after the first move. When matched against the all-cooperate strategy, TIT FOR TAT always cooperates. This strategy has the benefit of both cooperating with a friendly opponent, getting the full benefits of cooperation, and of defecting when matched against an opponent who defects. When matched against itself, the TIT FOR TAT strategy always cooperates.

Interestingly, beyond the realm of theoretical abstraction, tit for tat strategy found practical utility in diverse fields, ranging from conflict resolution to social psychology. Its simplicity and effectiveness rendered it a valuable tool for mitigating conflicts, with research indicating its efficacy in fostering cooperation and defusing tensions. By leveraging the principles of reciprocity and behavioural assimilation, tit for tat offered a pragmatic framework for navigating complex social dynamics, facilitating trust-building and conflict mitigation.

References

[1] Robert Aumann and Lloyd Shapley. Long term competition-a game theoretic analysis. UCLA Economics Working Papers 676, UCLA Department of Economics, 1992.

[2] Jean-Pierre Benoit and Vijay Krishna. Finitely repeated games. *Econometrica*, 53(4):905–922, 1985.

[3] James W. Friedman. A non-cooperative equilibrium for supergames. *The Review of Economic Studies*, 38(1):1, Jan 1971.

[4] Fudenberg and Tirole. *Game Theory*.

[5] Drew Fudenberg and David K. Levine. *A Long-run Collaboration on Long-run Games*. World Scientific, 2009.

[6] Drew Fudenberg and Eric Maskin. The folk theorem in repeated games with discounting or with incomplete information. *Econometrica*, 54(3):533–554, 1986.

[7] Sergiu Hart. Sergiu hart / papers / lecture notes - special topics in game theory.

[8] A. Rubinstein. Strong perfect equilibrium in supergames. *International Journal of Game Theory*, 9(1):1–12, Mar 1980.

[9] Ariel Rubinstein. Equilibrium in supergames. *Springer eBooks*, page 17–27, Jan 1994.