

Repeated Games

A repeated game (say, infinitely repeated prisoner's dilemma) is a special case of an extensive game.

The additional structure of the same game being repeated allows for new results. "Folk theorems" show that any payoffs that are feasible and enforcable (individually rational) can be achieved in equilibrium.

Repeated interaction allows for socially beneficial outcomes, essentially substituting for the ability to make binding agreements. One interpretation: if interaction is repeated, then socially beneficial outcomes that cannot be sustained by players with short-term objectives can be sustained by players with long-term objectives.

Proofs are constructive. We can interpret the equilibrium path as a “social norm,” which is supported by the threat of punishment.

But does Game Theory lose all predictive power?

There are folk theorems for finitely repeated games as well, but for many games the infinitely repeated game is very different from the finitely repeated game. If the prisoner’s dilemma is repeated 1,000,000 times, the only Nash equilibrium outcome is to defect in every period. The game unravels.

Which is the more appropriate model of human behavior, if people have finite lifetimes but do not perceive the distant future to be relevant?

Assume throughout a compact action space and continuous preferences.

Definition 137.1: Let $G = \langle N, (A_i), (u_i) \rangle$ be a strategic game, and let $A = \times_{i \in N} A_i$. An **infinitely repeated game** of G is an extensive game with perfect information and simultaneous moves $\langle N, H, P, (\succeq_i^*) \rangle$ in which we have

1. $H = \{\emptyset\} \cup \left(\bigcup_{t=1}^{\infty} A^t \right) \cup A^\infty$ (where A^∞ is the set of infinite sequences of action profiles),

2. $P(h) = N$ for every non-terminal history, $h \in H$,

3. \succeq_i^* is a preference relation on A^∞ that satisfies the following notion of weak separability: if $(a^t) \in A^\infty$, $a \in A$, $a' \in A$, and $u_i(a) > u_i(a')$, then we have for all t ,

$$(a^1, \dots, a^{t-1}, a, a^{t+1}, \dots) \succeq_i^* (a^1, \dots, a^{t-1}, a', a^{t+1}, \dots).$$

Note that a **strategy** for player i assigns an action in A_i for every finite sequence of outcomes in G .

What additional structure should we impose on \succeq_i^* ? For infinite games, you cannot simply add up the payoffs received at each stage.

1. Discounting. There is a discount factor $\delta \in (0, 1)$ such that $(a^t) \succeq_i^* (b^t)$ if and only if

$$\sum_{t=1}^{\infty} \delta^{t-1} u_i(a^t) \geq \sum_{t=1}^{\infty} \delta^{t-1} u_i(b^t).$$

With discounting, we normalize payoffs so that a payoff of v each period yields an overall payoff of v .

$$(1 - \delta) \sum_{t=1}^{\infty} \delta^{t-1} u_i(a^t)$$

2. Limit of Means. (a^t) is strictly preferred to (b^t) if and only if there exists $\varepsilon > 0$ such that

$$\frac{\sum_{t=1}^T [u_i(a^t) - u_i(b^t)]}{T} > \varepsilon$$

holds for all but a finite number of periods T .

This criterion treats periods symmetrically, and any one period has a negligible effect on the overall payoff. When the limiting average payoff exists, the payoff is

$$\lim_{T \rightarrow \infty} \left(\frac{\sum_{t=1}^T u_i(a^t)}{T} \right).$$

If the average payoffs cycle, so there are multiple limit points, then two sequences might not be comparable.

Fact: If (a^t) is strictly preferred to (b^t) according to the limit of means criterion, then there is a discount factor close enough to 1 such that (a^t) is strictly preferred to (b^t) according to the discounting criterion.

3. Overtaking. (a^t) is strictly preferred to (b^t) if and only if there exists $\varepsilon > 0$ such that

$$\sum_{t=1}^T [u_i(a^t) - u_i(b^t)] > \varepsilon$$

holds for all but a finite number of periods T .

This criterion treats periods symmetrically, and emphasizes the long run, but a single period can affect the overall preference.

Examples: $(1, -1, 0, 0, \dots)$ is strictly preferred to $(0, 0, \dots)$ according to discounting (for any δ), but the sequences are indifferent according to limit of means or overtaking.

$(-1, 2, 0, 0, \dots)$ is strictly preferred to $(0, 0, \dots)$ according to overtaking, but the sequences are indifferent according to limit of means. According to discounting, the preference depends on the discount factor.

Definition: A vector, $v \in \mathbb{R}^N$ is a **payoff profile** of $G = \langle N, (A_i), (u_i) \rangle$ if there is an outcome $a \in A$ for which we have $v = u(a)$. A vector, $v \in \mathbb{R}^N$ is a **feasible payoff profile** of $G = \langle N, (A_i), (u_i) \rangle$ if we have

$$v = \sum_{a \in A} \alpha_a u(a)$$

for some collection of nonnegative rational numbers $(\alpha_a)_{a \in A}$ with $\sum_{a \in A} \alpha_a = 1$.

Definition: Player i 's **minmax payoff** in G , denoted by v_i , is the lowest payoff that the other players can force upon i ,

$$v_i = \min_{a_{-i} \in A_{-i}} \max_{a_i \in A_i} u_i(a_{-i}, a_i).$$

A payoff profile w for which $w_i \geq v_i$ holds for all $i \in N$ is called **enforceable**, and **strictly enforceable** if the inequality is strict.

For each $i \in N$, choose $M^i = (M_1^i, \dots, M_{i-1}^i, M_{i+1}^i, \dots, M_N^i)$ such that $M^i \in \arg \min_{a_{-i} \in A_{-i}} \max_{a_i \in A_i} u_i(a_{-i}, a_i)$.

Folk Theorems for the Limit of Means Infinitely Repeated Game

Propositions 14.1 and 14.3: The set of feasible, enforceable payoff profiles of G is the set of **Nash** equilibrium payoff profiles of the limit of means infinitely repeated game of G .

Proof: First, note that every Nash equilibrium payoff profile of the limit of means infinitely repeated game of G must be feasible and enforceable. If a Nash equilibrium payoff profile is not enforceable for player i , then he/she has a profitable deviation to guarantee at least v_i in every period. (Obviously it must be feasible, Nash or not.)

Let $w = \sum_{a \in A} (\beta_a / \gamma) u(a)$ be a feasible, enforceable payoff profile, where each β_a is an integer and γ normalizes the coefficients to sum to one. Let (a^t) be a sequence of action profiles which cycle every γ periods, and each $a \in A$ is played β_a times. Let s_i be the strategy in the repeated game that chooses a_i^t unless there was a previous period t' in which a single player $j \neq i$ deviated from $a_j^{t'}$, in which case player i chooses M_i^j . This is a Nash equilibrium, because a deviating player receives at most v_i , but w is preferred since it is enforceable.

Proposition 146.2: Every feasible, strictly enforceable payoff profile of G is a **subgame perfect equilibrium** profile of the limit of means infinitely repeated game of G .

Proof Sketch: The “equilibrium path,” as before, consists of a cycle of actions of length γ . If some player j deviates, then once the cycle is finished, the other players play M^j long enough so that player j does not benefit from the deviation. After the punishment phase, all players return to the cycle. This is subgame perfect, because the payoff profile is w after every history. (A deviation during the punishment phase only increases one’s payoff during a finite number of periods.)

Note: The argument must be modified for the overtaking criterion. If player i refuses to play M_i^j during a punishment phase, then the players other than i (including j) minmax i long enough so that i does not benefit from the deviation. If one of the players punishing i (for not punishing j) deviates, then that player is punished, and so on.

Folk Theorems for the Discounted Infinitely Repeated Game

Proposition 145.2: Let w be a **strictly** enforceable, feasible payoff profile of G . For all $\varepsilon > 0$, there exists $\underline{\delta} < 1$, such that $\delta > \underline{\delta}$ implies the δ -discounted infinitely repeated game of G has a **Nash equilibrium** whose payoff profile w' satisfies $|w' - w| < \varepsilon$.

Proof uses the same trigger strategy argument, but no discount factor below 1 can hold a player to v_i if there is a beneficial short run deviation. Must be strictly enforceable. Cannot achieve the boundary of the convex hull by cycling, because of discounting.

For subgame perfect equilibrium with discounting, the arguments for limit of means and overtaking do not work. For overtaking, there could be a sequence of longer and longer punishments required, and no fixed discount factor can handle all the deviations. The problem arises when a punisher is hurt more severely than the player being punished.

Proposition (James Friedman): Let w be a strictly enforceable, feasible payoff of G that Pareto dominates the payoffs of a (one-shot) Nash equilibrium of G . Then, if δ is sufficiently close to one, there exists a subgame perfect equilibrium of the δ -discounted infinitely repeated game with payoff profile arbitrarily close to w .

Proof constructs strategies where a deviation triggers the one-shot Nash actions to be played afterwards. This argument works for all three payoff criteria, and involves simple strategies.

Theorem 1 (Fudenberg and Maskin 1986): Let w be a strictly enforceable, feasible payoff profile of a **two-player** game, G . For all $\varepsilon > 0$, there exists $\underline{\delta} < 1$, such that $\delta > \underline{\delta}$ implies the δ -discounted infinitely repeated game of G has a **subgame perfect equilibrium** whose payoff profile w' satisfies $|w' - w| < \varepsilon$.

Proof Sketch: After a deviation by either player, each player minmaxes the other for a certain number of periods, after which they return to the original path. If a deviation occurs during punishment, the punishment phase is begun again.

With three players, the action player 1 chooses to minmax player 2 might be different from the action player 1 chooses to minmax player 3. The proof of Theorem 1 does not extend.

1,1,1	0,0,0	0,0,0	0,0,0
0,0,0	0,0,0	0,0,0	1,1,1

For this example, to minmax player 1, we must have (\cdot, r, L) or (\cdot, l, R)

to minmax player 2, we must have (b, \cdot, L) or (t, \cdot, R)

to minmax player 3, we must have (t, r, \cdot) or (b, l, \cdot)

Not only does the proof not work, but the theorem is false without more assumptions. For this example, for any $\delta < 1$, there is no subgame perfect equilibrium in which the payoffs are less than $\frac{1}{4}$, even though any positive payoff for all players is strictly enforceable and feasible.

The problem is that you cannot differentially punish the players. This is a knife-edge example. In general, we can reward one player for punishing another.

Theorem 2 (Fudenberg and Maskin): Assume that the dimensionality of the set of strictly enforceable, feasible payoff profiles equals the number of players. Let w be a strictly enforceable, feasible payoff profile of G . For all $\varepsilon > 0$, there exists $\underline{\delta} < 1$, such that $\delta > \underline{\delta}$ implies the δ -discounted infinitely repeated game of G has a **subgame perfect equilibrium** whose payoff profile w' satisfies $|w' - w| < \varepsilon$.

What about mixed strategy equilibrium?

Previous results apply to games in which the action space is the space of probability distributions. However, for G to be a game of perfect information, players would have to observe the history of action profiles, which would mean the probability distributions themselves, and not just the outcomes of the mixing.

However, Fudenberg and Maskin (1986, Theorem 5) show that Theorem 2 continues to hold when players observe only the past actions rather than their mixed strategies.

When mixing is allowed, the minmax payoffs are often much lower, so more payoff profiles would be consistent with subgame perfect equilibrium.

Finitely Repeated Games

Definition: A T -period finitely repeated game of $G = \langle N, (A_i), (u_i) \rangle$ is an extensive game of perfect information satisfying the conditions of Definition 137.1, with the symbol ∞ replaced with T . Assume that preferences are represented by the mean payoff,

$$\frac{\sum_{t=1}^T u_i(a^t)}{T}.$$

For finitely repeated games where the Nash equilibrium payoffs of the one-shot game G coincide with the minmax payoffs, then each (a^t) must be a Nash equilibrium of G . (Unraveling in the finitely repeated prisoner's dilemma)

Proposition 156.1: Suppose $G = \langle N, (A_i), (u_i) \rangle$ has a Nash equilibrium \hat{a} in which we have $u_i(\hat{a}) > v_i$ for each player $i \in N$. Then for any strictly enforceable outcome a^* of G and any $\varepsilon > 0$, there exists an integer T^* such that $T > T^*$ implies the T -period repeated game of G has a **Nash equilibrium** in which the payoff to player i is within ε of $u_i(a^*)$.

Proof Sketch: On the equilibrium path, the action profile is a^* for the first $T - L$ periods and \hat{a} for the last L periods. If there is a deviation, the players minmax each other. Because we have $u_i(\hat{a}) > v_i$ for each player, we can calculate L such that the one-period gain from deviating from a^* is outweighed by the loss of reverting to the minmax payoff. Since L is independent of the length of the game, the result follows.

Obviously, the equilibrium is not subgame perfect. If G has a unique Nash equilibrium, then the finitely repeated game has a unique subgame perfect equilibrium in which the Nash profiles are played each period.

If G has multiple Nash equilibria, one of which dominates the other, then the threat to play the inferior NE at the end of the game is credible.

For large enough T , any strictly enforceable payoff profile can be achieved (within ε) as the average payoff profile in a subgame perfect equilibrium of the repeated game.