On finitely repeated games

Essay in the course Introduction to Game Theory

Albin Fredriksson

June 8, 2010

This essay is based on two papers by Benoît and Krishna [1, 2].

Introduction

In infinitely repeated games, there often exist subgame perfect equilibria in which the players' actions are not Nash equilibria of the stage game in all periods, since actions that are profitable in the short term might not coincide with the actions that are profitable in the long term. A player who takes an action that is profitable in the short term—and that reduces other players' payoffs—may be punished in the infinitely many following periods. If the players do not discount the future, each individually rational payoff vector of the stage game can be approximated by the payoff of an equilibrium of the repeated game.

However, games that are repeated a finite number of times do not always have more equilibria than the stage game. For example, a game with a unique Nash equilibrium has a unique subgame perfect equilibrium—playing the Nash equilibrium in all periods. This is seen by backwards induction: since each subgame must be played in an equilibrium, the last period must be played in the unique Nash equilibrium of the stage game, which implies that there is no credible threat in the penultimate period, whence also that period must be played in the Nash equilibrium, and so on. An example of such a game is the prisoner's dilemma.

Due to the backwards induction, some finitely repeated games are trivial, which is unsatisfactory since repeated games in reality typically are finite. It would be desirable that the inifitely repeated games be the limits of finitely repeated games as the number of repetitions goes to infinity. This would motivate the use of the former as approximations of the latter.

Nevertheless, a finitely repeated game for which the stage game has multiple Nash equilibria has multiple subgame perfect equilibria, which are not necessarily sequences of Nash equilibria of the stage game. Threatening to play a worse Nash equilibrium in the last period is a credible threat, which enables other outcome paths to be subgame perfect. However, it is not always possible to approximate every possible individually rational payoff vector by a subgame perfect equilibria, even if the game is repeated a large number of times. Benoît and Krishna states conditions under which the behavior of the finitely repeated games are similar to that of infinitely repeated games as the number of repetitions is large, i.e., conditions that enables approximation of every possible individually rational payoff vector by a subgame perfect equilibrium of the finitely repeated game.

Finitely repeated games

The set of subgame perfect equilibria can be characterized by the optimal punishments. An optimal punishment for player i is the subgame perfect outcome path that results in the worst payoff for player i. Before optimal punishments are formally defined, some notation must be introduced.

Consider an n player one-shot game $G = (A_1, ..., A_n; U_1, ..., U_n)$, where A_i is the strategy space for player i and $U_i : A \to R$ is the payoff function of player i, where $A = A_1 \times ... \times A_n$. An $a \in A$ is referred to as an outcome of G. Let $G(T)$ denote the game that results when G is repeated T times. A strategy for player i in the game $G(T)$ is a function σ_i which selects, for any history of play, an element of A_i. An n-tuple of strategies is denoted by σ and defines an outcome path $(a^1(\sigma), ..., a^T(\sigma))$, in which $a^t(\sigma)$ denotes the outcome of the game $G(T)$ in period t as a function of the strategy n-tuple σ. Let $P(T)$ denote the set of perfect equilibrium outcome paths, i.e., outcome paths that result from some perfect equilibrium strategy of $G(T)$. Optimal punishments can now be defined:

Definition 1 *The* optimal K-period punishment *for player i is any path in $P(K)$ that results in a payoff of $w_i(K)$ for player i, where $w_i(K)$ is player i's worst perfect equilibrium payoff, given by*

$$w_i(K) = \min \left\{ \sum_{t=1}^{K} U_i(a^t) : (a^1, ..., a^K) \in P(K) \right\}.$$

By considering these punishments for all players and all subgames, the set of subgame perfect equilibria can be completely characterized: an outcome path $(a^1, ..., a^T)$ is an element of $P(T)$ if and only if, for all players and for all $t < T$, the payoff resulting from a^t is at least as good as the payoff obtained by deviating from a^t and being optimally punished for the remainder of the game. It should be noted that a player i that is subject to an optimal punishment can come off worse than would be the result of playing all remaining periods in the Nash equilibrium of G that is worst for player i.

Optimal punishments are problematic in the sense that they are hard to determine. Therefore, Benoît and Krishna utilize three phase punishments to approximate optimal punishments. A three phase punishment is, as the name suggests, constructed by three phases: In the first phase, the player i to be penalized is severely punished. In this phase, the other players take their minmax actions against player i, which means that they take the actions $a_j \in A_j$ such that

$$a_j \in \underset{a_j \in A_j}{\operatorname{argmin}} \left(\max_{a_i \in A_i} U_i(a) \right)$$

for the extension of the phase. In the second phase, the players that participated in the punishment are compensated; it may be so costly for players $j \neq i$ to participate in the punishment that they might want to deviate, but the reward in the second phase is used to motivate the players to punish player i in the first phase. Care is taken so that the punished player does not recoup. In the third phase, a Nash equilibrium of the G is played long enough as to ensure compliance during the reward phase.

Three phase punishments are easier to determine than optimal punishments, and can under mild conditions approximate optimal punishments when the number of repetitions is large. Given an arbitrary n-tuple of punishments, it is then possible to construct a three phase punishment for a player that is more severe than the original punishment. This is utilized to construct a hierarchy of three phase punishments that converge to optimal punishments when the number of repetitions approaches infinity. Three phase punishments can then be used in the proof of a limit folk theorem for finitely repeated games. The theorem states that when certain weak conditions are satisfied, every possible individually rational outcome vector in the stage game can be approximated by the mean payoff of a subgame perfect equilibrium of the finitely repeated game, if only the game is repeated a sufficient number of times. The conditions that must be satisfied are that there for each player must exist a Nash equilibrium in the stage game in which the player's payoff is strictly better than in the worst Nash equilibrium of the player, and that F, the feasible region of the payoff vectors—the convex hull of the set of payoff vectors that can result from the choices of the players—has dimension n. The dimensionality condition can be dispensed with in games of two players. Formally, the theorem goes as follows:

Theorem 1 *Suppose that (i) for every j, there exists a Nash equilibrium of G such that player j's payoff is strictly better than $w_j(1)$, the worst one-shot equilibrium payoff for player j; and (ii) the dimension of the feasible region F of payoff vectors is n.*

Let u be any feasible and individually rational payoff vector. Then for all $\varepsilon > 0$ there exists a T_0 such that for all $T \geq T_0$ there exists an outcome path $(a^1, ..., a^T) \in P(T)$ such that

$$\left\| \frac{1}{T} \sum_{t=1}^{T} U(a^t) - u \right\| < \varepsilon.$$

Renegotiation in finitely repeated games

The result that all possible individually rational payoff vectors can be approximated in many games can be criticized, since some of the outcomes are enforced by threats that are Pareto dominated by other subgame perfect outcome paths. If the players were able to negotiate between periods, they would therefore possibly choose a dominating path, since all players would benefit from doing so. This critique implies that some equilibria are not probable to occur, since the enforcing threats can be considered as not being credible. Therefore, Benoît and Krishna extends the theory by allowing renegotiation between games and consider the equilibria that are not enforced by Pareto dominated threats. These equilibria are not the same as the Pareto efficient subgame perfect equilibria.

Players are assumed to be able to communicate with each other before each period, but they cannot form binding agreements. The subset of subgame perfect equilibria that is considered is called "renegotiation proof" and is constructed recursively. In a one-shot game, the set of acceptable equilibria are the ones that are not Pareto dominated by any other equilibria. In $G(2)$, the second period must be played in an efficient one-shot equilibrium. In the first period, only outcomes that are sustained by threats of efficient one-shot equilibria are considered, and only the efficient ones of these are played. The definition for $G(T)$ for arbitrary T can be constructed by continuing in this manner.

Formally, for any set $S \subset \mathbb{R}^2$, let Eff $S = \{x \in S : \neg\exists y \in S : y \succ x\}$ where for $x, y \in \mathbb{R}^2$, $x \succ y$ denotes that for $i = 1, 2$, $x_i > y_i$. Let $U(\sigma) = \sum_{t=1}^{T} U(a^t(\sigma))$. For $K < T$, let $h(K) = (a^1, ..., a^K)$ denote the K-period history and let $\sigma|_{h(K)}$ denote the strategy combination induced by σ on the subgame $G(T-K)$ following $h(K)$. Then $U(\sigma|_{h(K)})$ is called the continuation payoff prescribed by σ on $G(T-K)$ following $h(K)$. Finally, let $\tilde{P}(T) = \{U(\sigma) : \sigma \text{ is a perfect equilibrium of } G(T)\}$ (notice that $\tilde{P}(T)$ is not the same as $P(T)$ the previous paper). The renegotiation proof equilibria of $G(T)$ can then be defined as follows:

Definition 2 *A perfect equilibrium σ of $G(T)$ is said to be* renegotiation proof *if $U(\sigma) \in R(T)$, where the set $R(T)$ of renegotiation proof equilibria payoffs is defined by*

$$\begin{aligned} Q(1) &= \tilde{P}(1), \\ R(1) &= \mathit{Eff}\ Q(1), \end{aligned}$$

and for $T > 1$,

$$\begin{aligned} Q(T) &= \{U(\sigma) \in \tilde{P}(T) : \ \textit{all continuation payoffs prescribed by} \\ &\qquad\qquad \sigma \textit{ on } G(T-1) \textit{ lie in } R(T-1)\}, \\ R(T) &= \mathit{Eff}\ Q(T). \end{aligned}$$

Benoît and Krishna then studies the set of renegotiation proof equilibria as the number of repetitions of the game approaches infinity. They show that when this set exists—which is an open question but which they conjecture—it either is a singleton set or a connected subset of the Pareto frontier. Surprisingly, when the set is a singleton, it does not necessarily lie on the Pareto frontier, which implies that it corresponds to an inefficient renegotiation proof equilibrium. The theorem is stated as follows:

Theorem 2 *The set*

$$\lim_{T\to\infty} (1/T)R(T)$$

is a singleton, or it is a subset of Eff F.

Conclusion

Benoît and Krishna have given conditions under which finitely repeated games approach their infinitely repeated counterparts as the number of repetitions approaches infinity. This motivates the study of infinitely repeated games as an approximate model of reality, in which games typically are finite. It also implies that finitely repeated games generally are non-trivial, and motivates a further study of these games.

They have also shown how a subset of the set of subgame perfect equilibria, called the set of renegotiation proof equilibria, in which only threats with a certain measure of efficiency are considered credible, behaves as the number of repetitions of certain games approaches infinity. In the limit, the set of mean payoffs corresponding to these renegotiation proof equilibria can sometimes be reached by inefficient paths. If it can, it is a singleton set.

References

[1] Benoît, J.-P. och Krishna, V. (1985) Finitely Repeated Games. *Econometrica.* **53** 905-922.

[2] Benoît, J.-P. och Krishna, V. (1993) Renegotiation in Finitely Repeated Games. *Econometrica.* **61** 303-323.