Discussion Paper No. 1136

CONES OF COOPERATION FOR INDEFINITELY REPEATED, GENERALIZED PRISONER'S DILEMMA GAMES

by

Michael A. Jones

June 1995

CONES OF COOPERATION FOR INDEFINITELY REPEATED, GENERALIZED PRISONER'S DILEMMA GAMES*

MICHAEL A. JONES

June 1995

ABSTRACT. A continuation probability is introduced to develop a theory of indefinitely repeated games where the extreme cases of finitely and infinitely repeated games are specific cases. The set of publicly correlated strategies (vectors) that satisfy a matrix inequality equivalent to the one-stage-deviation principle forms a cone of cooperation. The geometry of these cones provides a means to verify intuition regarding the levels of cooperation attained when the discount parameter and continuation probability vary.

A bifurcation point is identified which indicates whether or not a cooperative subgame perfect publicly correlated outcome exists for the indefinitely repeated game. When a cooperative equilibrium exists, a recursive relationship is used to construct an equilibrium strategy. New cooperative behavior is demonstrated in an indefinitely repeated game with infrequent shocks (a subsequence of the continuation probability goes to zero).

1. INTRODUCTION

The disparity between the size of equilibrium sets for the finitely repeated Prisoner's Dilemma game and the infinitely repeated game has spawned a number of attempts and methods to determine where cooperation begins between finitely and infinitely repeated games. Past work has considered the discount parameter either as a discount on future payoffs or as the probability that the game continues to the next round [10]. By introducing a time-dependent continuation probability describing the conditional probability that round k occurs given that round $k-1$ occurs, the dual roles of the discount parameter are separated. The discount parameter is used only to represent a discount on future payoffs.

Key words and phrases. Repeated games, Prisoner's Dilemma, continuation probability, subgame perfect publicly correlated equilibria.

*This research stems from my dissertation work under D.G. Saari at Northwestern University; it was supported in part by NSF grant IST 9103180. This paper has benefitted from comments and suggestions by E. Kalai, M. Kilgour, J. Ratliff, and S. Zabell.

The continuation probability introduces a continuum of indefinitely repeated games which, in a sense, lie in between the two extreme cases of finitely and infinitely repeated games. From this class, I examine indefinitely repeated generalized Prisoner's Dilemma games with finite action spaces. Bernheim and Dasgupta [3] examine indefinitely repeated generalized Prisoner's Dilemma games with both finite and continuous action spaces. Their results are markedly different as they examine local behavior near a Nash equilibrium. Both Carroll [4] and Becker and Dudd [2] discuss indefinitely repeated games. As Becker and Dudd point out, Carroll considers indefinitely, but finitely repeated Prisoner's Dilemma games. Further, Becker and Dudd discuss the usefulness of indefinitely repeated games with and without a definite endpoint.

Past work has concentrated on the role of the discount parameter in determining the set of possible equilibrium payoffs for infinitely repeated games. For example, Stahl [15] provides a graph of possible payoffs of subgame perfect publicly correlated equilibria under different discount parameters. My approach introduces a linear operator, a matrix inequality equivalent to the one-stage-deviation principle, and a means for determining levels of cooperation. This is done by examining the set of "cooperation vectors" associated with subgame perfect publicly correlated equilibria satisfying the matrix inequality; the set forms a "cone of cooperative outcomes." The model and linear operator approach are discussed in Section 2. Important for this discussion is the observation that the matrix is defined solely by the discount parameter and the continuation probability. The cone of cooperation and the motivation of the continuation probability approach are presented in Section 3.

An advantage of this approach is that the intuition that "more" cooperation is possible, when payoffs are discounted less and when the probability that the game continues increases, is quantified. (The result is, in part, described by the monotonicity theorem which appears in Section 3 and is a direct consequence of the geometry of the cone of cooperation.) Also, a "bifurcation point of cooperation" is defined in Section 4 which identifies the border between indefinitely repeated games with cooperative subgame perfect publicly correlated equilibria and those with only noncooperative subgame perfect publicly correlated equilibria. These theorems determine when cooperative outcomes occur in between finitely and infinitely repeated games, as well as the size of the equilibrium set. Section 5 contains examples of finitely, infinitely, and indefinitely repeated Prisoner's Dilemma games and displays the usefulness of the bifurcation point. The indefinitely repeated game demonstrates cooperation for a game that previously was indescribable.

2. The Model and the Linear Operator Approach

I proceed by introducing notation for a stage game with standard information, *i.e.*, where there is only one state of nature and the players know all the past moves of each player. Let $N = \{1, 2, \ldots, n\}$ be the set of all players. Let $D = \times_{i=1}^{n} \Delta D_i$ be the set of actions of the stage game where $D_i = \{d_1^i, d_2^i, \ldots, d_{m_i}^i\}$ is the set of actions for player i at any round of the repeated game. Here ΔD_i is the space of probability distributions over D_i. The utility function for player i, $u_i : D \to \mathbb{R}$, describes the unweighted payoff to player i for any action taken by the players. The indefinite nature is described by the continuation probability.

Definition 2.1. The *time-dependent continuation probability* β is defined by $\beta = (\beta_1, \beta_2, \dots)$ where $\beta_k \in [0, 1]$. The number β_k is the probability that the game will continue to the k^{th} round given that the $(k-1)^{th}$ round occurs.

Notice that both finitely and infinitely repeated games can be represented by continuation probabilities. The continuation probability for a finitely repeated game of r rounds is represented by the sequence consisting of r 1's followed by a tail of zeroes. For an infinitely repeated game, it is the sequence of all 1's.

The continuation probability is assumed to be part of the common knowledge of the players, just as the duration of finitely and infinitely repeated games is common knowledge. (For a discussion of incomplete information regarding the duration of repeated Prisoner's Dilemmas, see Ratliff [13] or Samuelson [14].) The discount parameter $\delta \in [0, 1]$ denotes how the players' utilities decrease over time.

An indefinite round game is defined by its stage game, its discount parameter, and its continuation probability. Therefore, an *indefinitely repeated game* is defined by the quintuple $\Gamma = (N, D, u, \beta, \delta)$. The expected payoff for the indefinitely repeated game weights future payoffs by both the discount parameter and the continuation probability.

I examine stage games that generalize the Prisoner's Dilemma. Specifically, I consider stage games with a single inefficient Nash equilibrium and normalized payoffs so that the each player receives zero utiles under the Nash equilibrium. Denote the set of these games as $\mathcal{S}$. The normalization implies that there exists at least one efficient strategy whose payoffs are positive and, thereby, strictly Pareto dominate the payoffs of the Nash equilibrium.

Bernheim and Dasgupta [3] examine indefinitely repeated games with time dependent continuation probabilities. However, they primarily consider continuous action spaces but retain the assumptions of complete information and full rationality. Their results yield subgame perfect equilibria that are Pareto superior to the repeated Nash outcome. Their work depends, however, on continuous action spaces that provide a local flavor to the analysis. An example of this is where they consider a Pareto superior strategy in the strategy space of the stage game close to the inefficient Nash equilibrium. Bernheim and Dasgupta explore finite action spaces but restrict their analysis to continuation probabilities that are "asymptotically finite" (that is, $\lim_{k\to\infty} \beta_k = 0$).[1]

[1] My analysis of finite action space games was initially limited to the Prisoner's Dilemma. Bernheim and Dasgupta's work suggests the normalization of games in $\mathcal{S}$.

The normalizations of the payoffs and the existence of a single Nash equilibrium are achieved by the following definitions. Let Σ be the set of mixed strategies for the stage game. Therefore, $\Sigma = \{\sigma = (\sigma_1, \sigma_2, \dots \sigma_n) \mid \text{where } \sigma_i \in \Delta D_i\}$. Let (N, D, u) be a stage game in $\mathcal{S}$ and let the single Nash equilibrium in the stage game be denoted by σ^*. By definition, $u_i(\sigma^*) = 0$ for all players i. Define a function $f_i : \Sigma \to \mathbb{R}$ that gives the maximum gain to player i if her opponents stick to a fixed strategy σ and if she is allowed to deviate from σ. Therefore, $f_i(\sigma) = \max_{\tau_i \in \Delta D_i} [u_i(\tau_i, \sigma_{-i}) - u_i(\sigma)]$. Notice that $f_i(\sigma^*) = 0$ for every player i by the definition of a Nash equilibrium. Also, notice that $f_i(\sigma) > 0$ for at least one player i for all mixed strategies $\sigma \neq \sigma^*$ of the stage game.

Since σ^* is inefficient by assumption then there exists a $\tau^* \in \Sigma$ such that $u_i(\tau^*) > u_i(\sigma^*) = 0$ for all i. Let τ^* be a fixed strategy such that $u_i(\tau^*) > u_i(\sigma^*)$ for all i and

$$\tau^* \in \arg\min_{\tau} \left\{ \max_i \frac{f_i(\tau)}{u_i(\tau)} \right\}.$$

Assume that players publicly correlate their cooperative efforts. So, players essentially base their collective strategy on the public observance of the flipping of a λ-coin. Let $\lambda[\tau^*] + (1-\lambda)[\sigma^*]$ represent the correlated strategy where, with probability λ, the players all play τ^* and with probability $(1-\lambda)$ the players all play σ^*. The parameter λ represents the level of cooperation where $\lambda = 0$ is noncooperative and $\lambda = 1$ is purely cooperative. The pre-discounted expected utility for player i using this correlated strategy in the stage game is

$$u_i(\lambda[\tau^*] + (1-\lambda)[\sigma^*]) = \lambda u_i(\tau^*) + (1-\lambda) u_i(\sigma^*) = \lambda u_i(\tau^*).$$

Therefore, cooperation increases the expected payoff, but also gives a more lucrative incentive for a player to deviate from the strategy since $f_i(\tau^*) > 0$ for some i.

Let S be the strategy profile for the indefinitely repeated game where the players play the correlated strategy $\lambda_k[\tau^*] + (1-\lambda_k)[\sigma^*]$ at round k until a player deviates. Enforcement is by a grim trigger mechanism, where deviation from τ^* warrants punishment with σ^* in all subsequent rounds. Deviation from $\lambda_k[\tau^*] + (1-\lambda_k)[\sigma^*]$ is immediately obvious to the other players due to complete information. There is an important temporal aspect in this definition. I am requiring that the deviation is decided before the observation of the probabilistic event. One can think of this restriction as adhering to a mechanism and making the decision to deviate in a particular round at the beginning of the game. Similar results, however, hold if the deviation occurs after the event. Relying on grim trigger strategies is sufficient since any subgame-perfect outcome can be supported by a grim trigger mechanism [1].

Define $\Lambda = (\lambda_1, \lambda_2, \lambda_3, \dots)$ to be the cooperation vector associated with the strategy profile S. The cooperation vector Λ measures cooperation for the indefinitely repeated game just as λ_k measures cooperation for round k. The strategy profile

S, with an associated cooperation vector Λ, is a subgame perfect equilibrium if it satisfies the one-stage deviation principle.

Proposition 2.1. *The strategy profile S with the associated cooperation vector Λ is a subgame perfect equilibrium if and only if, for every round k and every player i,*

$$\beta_k\Big[\lambda_k f_i(\tau^*) + \lambda_k u_i(\tau^*)\Big] \leq \beta_k\lambda_k u_i(\tau^*) + \sum_{r=1}^{\infty} \delta^r \left(\prod_{j=1}^{r} \beta_{k+j-1}\right) \lambda_{k+r} u_i(\tau^*)$$

which is equivalent to

$$\lambda_k f_i(\tau^*) \leq \sum_{r=1}^{\infty} \delta^r \left(\prod_{j=1}^{r} \beta_{k+j}\right) \lambda_{k+r} u_i(\tau^*).$$

Proof. ($\Leftarrow$)The left hand side of the inequality is the expected payoff from deviating at round k. The right hand side of the inequality is the expected payoff from adhering to the strategy represented by the cooperation vector. Every player at every round receives a greater expected utility by adhering to the strategy. Therefore, the strategy induced by the cooperation vector satisfies the one-stage-deviation principle which implies that the strategy is subgame perfect. ($\Rightarrow$) The other direction is similar. $\square$

It is possible to arrive at an equivalent definition for S to be subgame perfect by considering the following matrix.

$$M(\delta,\beta) = \begin{pmatrix} 0 & \delta\beta_2 & \delta^2\beta_2\beta_3 & \delta^3\beta_2\beta_3\beta_4 & \delta^4\beta_2\beta_3\beta_4\beta_5 & \dots \\ 0 & 0 & \delta\beta_3 & \delta^2\beta_3\beta_4 & \delta^3\beta_3\beta_4\beta_5 & \dots \\ 0 & 0 & 0 & \delta\beta_4 & \delta^2\beta_4\beta_5 & \dots \\ \vdots & \vdots & \vdots & \ddots & \ddots & \ddots \end{pmatrix}. \tag{2.1}$$

The system of inequalities for i fixed in Proposition 2.1 is equivalent to the following matrix inequalities.

$$f_i(\tau^*)\Lambda \leq u_i(\tau^*) M(\delta,\beta)\Lambda$$

or

$$\frac{f_i(\tau^*)}{u_i(\tau^*)}\Lambda \leq M(\delta,\beta)\Lambda \tag{2.2}$$

since $u_i(\tau^*) > 0$. This matrix inequality must be satisfied for all players i. Due to the similar structure of the matrix inequalities for the players, only the matrix inequality with the largest $\frac{f_i(\tau^*)}{u_i(\tau^*)}$ need be considered due to transitivity.

Corollary 2.2. *The strategy profile S with the associated cooperation vector Λ is a subgame perfect publicly correlated equilibrium if and only if, for every player i,*

$$\frac{f_i(\tau^*)}{u_i(\tau^*)}\Lambda \leq M(\delta, \beta)\Lambda.$$

Notice that the left hand side of the inequality in Equation 2.2 depends solely on the stage game while the right hand side depends on the continuation probability and the discount parameter. This inequality separates the information from the stage game and the repeated process. The effects of changing a single parameter of the matrix inequality yields information about the set of subgame perfect equilibria and when cooperation occurs.

It is useful to consider $M(\delta, \beta)$ as a linear operator from ℓ^∞ to ℓ^∞. In doing so, recall that the spectral radius of $M(\delta, \beta)$ bounds the spectra of $M(\delta, \beta)$, including the absolute value of the eigenvalues. (See Dunford and Schwarz [5] for information on computing the spectral radius and its properties.) The spectral radius plays an integral part of establishing the bifurcation point of cooperation in Section 4. In general, the spectral radius is difficult to compute. However, since $M(\delta, \beta)$ is nonnegative and upper triangular, the spectral radius is easier to compute. This is discussed in Jones [7], as well as in Section 4.

It is possible to examine the properties of spectra of the operators associated with continuation probabilities and discount parameters. The spectral properties allow continuation probabilities to be classified by whether or not there exists a stage game in $\mathcal{S}$ where non-trivial cooperation vectors satisfy the matrix inequality, *i.e.*, the existence of non-trivial subgame perfect publicly correlated equilibria. This work appears in Jones [7].

3. The Cone of Cooperation and its Applications

In finitely repeated Prisoner's Dilemma games it is well known that the only subgame perfect equilibrium requires playing the Nash equilibrium in every round (*e.g.*, Luce and Raiffa [9]). However, when infinitely often with sufficiently large discount parameters, these games have an infinite number of outcomes preferable to defecting in every round (Myerson [11]).

Radner [12] discusses the "discontinuity" of the set of equilibria that arises when the number of rounds of a repeated game tends to infinity. He shows, under certain bounded rationality conditions, that equilibria in finitely repeated Prisoner's Dilemma games can be cooperative up to a number of rounds before the repeated game ends. In my framework, this intuition is represented by the continuation probability of the finitely repeated game of r rounds approaching the continuation probability of the infinitely repeated game as r tends to infinity. The difference in thought requires a few words.

Specifically, let p^r be the continuation probability for the finitely repeated game of r rounds. Let p be the continuation probability for the infinitely repeated game. As r tends to infinity, p^r and p match up in the first r terms.

$$p^r = (\overbrace{1, 1, 1, \dots, 1}^{r \text{ terms}}, 0, 0, \dots)$$
$$p = (1, 1, 1, \dots, 1, 1, 1, \dots).$$

Indeed, the sense that p^r is getting "closer" to p can be justified by using the weak topology of ℓ^∞ when ℓ^∞ is considered as the dual space of ℓ^1. Alternatively, capturing the sense of "bounded rationality" p^r converges to p under a metric that concentrates more weight on the beginning terms of the sequence than the tail of the sequence. (For example, if d^* defines a metric on continuation probabilities where $d^*(p^r, p) = \sum_{k=1}^{\infty} \epsilon^k \mid (p^r)_k - (p)_k \mid$ then p^r converges to p under d^* when ϵ is between 0 and 1.) Under the sup norm metric ($d(p^r, p) = \sup_{i \in \mathbb{Z}}\{\mid (p^r)_i - p_i \mid\} = 1$) used here, however, for all r, p^r does not converge to p under d. Thus, the approach described here introduces a conceptually different way to resolve this puzzle.

The disparity of the size of equilibrium sets of finitely and infinitely repeated games is considered a discontinuity since as the number of rounds of the finitely repeated game increases, the equilibrium sets do not converge to the equilibrium set of the infinitely repeated game. Under a sup norm metric, we don't expect to have the equilibrium sets of finitely repeated games approach the equilibrium set of the infinitely repeated game since the continuation probabilities do not converge.

This idea is the basis for results substantiating our intuition about cooperation increasing as the probability that the game continues increases. By concentrating on the geometry of the subgame perfect publicly correlated equilibria and not the expected payoffs of the strategies, I am able to justify the intuition that more cooperation is possible when the game is discounted less or has a greater probability of continuation.

The set of cooperation vectors satisfying the matrix inequality in Equation 2.2 forms a cone intersected with the closed unit sphere in ℓ^∞. This set is called the *cone of cooperation*. The cone of cooperation is convex and closed under multiplication by a positive scalar, as long as the restriction that the cooperation vectors all have entries between 0 and 1 inclusive holds.

The convexity allows for the construction of an infinite number of subgame perfect equilibria just by knowing two different cooperation vectors that satisfy the matrix inequality. But since the vector $(0, 0, 0, \dots)$ solves the matrix inequality, then any nonzero cooperation vector yields an infinite number of subgame perfect equilibria. The closure under multiplication by a positive scalar insures that when the cone is nonempty then there exists a subgame perfect equilibrium that achieves full cooperation for at least one round; actually, there exists a subgame perfect publicly correlated equilibrium with $\sup_{k\in\mathbb{N}} \lambda_k = 1$. The proof of the following theorem follows directly from the linearity of $M(\delta, \beta)$; it appears in Jones [6].

Theorem 3.1. *For a fixed indefinitely repeated game, the set of cooperation vectors that induce subgame perfect equilibria forms a cone intersected with the closed unit ball in ℓ^∞.*

Intuition suggests that more cooperation is possible for a repeated game with a greater probability of continuation. Further, as the discount parameter increases, we expect a greater range of cooperation vectors to be in the cone of subgame perfect equilibria of indefinitely repeated games. Not only does the cone of cooperation increase under the above changes, but the larger cone contains the smaller cone. This idea is captured in the following monotonicity theorem. Let $C(\delta, \beta)$ represent the cone of cooperation vectors that induce subgame perfect equilibria for the indefinitely repeated game with fixed stage game, discount parameter δ, and continuation probability β.

Theorem 3.2. *If α and β are continuation probabilities such that $\alpha_t \le \beta_t$ for all t, and δ_1 and δ_2 are discount parameters such that $\delta_1 \le \delta_2$, then*

$$C(\delta_1, \alpha) \subseteq C(\delta_2, \beta).$$

Proof. Suppose $\Lambda \in C(\delta_1, \alpha)$. It follows that for all players i and rounds k

$$\lambda_k f_i(\tau^*) \le \sum_{r=1}^{\infty} \delta_1^r \left(\prod_{j=1}^{r} \alpha_{k+j}\right) \lambda_{k+r} u_i(\tau^*).$$

But $\alpha_t \le \beta_t$ and $\delta_1 \le \delta_2$ imply

$$\begin{aligned}\sum_{r=1}^{\infty} \delta_1^r \left(\prod_{j=1}^{r} \alpha_{k+j}\right) \lambda_{k+r} u_i(\tau^*) &\le \sum_{r=1}^{\infty} \delta_1^r \left(\prod_{j=1}^{r} \beta_{k+j}\right) \lambda_{k+r} u_i(\tau^*)\\ &\le \sum_{r=1}^{\infty} \delta_2^r \left(\prod_{j=1}^{r} \beta_{k+j}\right) \lambda_{k+r} u_i(\tau^*).\end{aligned}$$

Transitivity implies that $\Lambda \in C(\delta_2, \beta)$; hence, $C(\delta_1, \alpha) \subseteq C(\delta_2, \beta)$. □

While this theorem insures weak containment, the inequality becomes strong when $C(\delta_1, \alpha)$ is nontrivial (that is, $C(\delta_1, \alpha) \neq (0, 0, 0, \dots)$) because it is easy to construct a vector in $C(\delta_2, \beta)$ that is not in $C(\delta_1, \alpha)$. When the discount parameter or continuation probability changes, the cones expand and collapse in response. How do we know when the cone of cooperation is trivial or nontrivial? This question is addressed in Section 4.

4. The Bifurcation Point of Cooperation

When does a nontrivial solution exist to an indefinitely repeated game? Further, how do we find a cooperative equilibrium? The bifurcation point of cooperation, μ^*, is a function of δ and $\mathcal{E}$, a number associated with the continuation probability. It divides the matrix inequality, $\mu\Lambda \leq M\Lambda$, into noncooperative and cooperative regions depending on μ. To see how the bifurcation point works, realize that a stage game gives the ratio $\frac{f_i(\tau^*)}{u_i(\tau^*)}$ for every player i. Essentially, this number represents how much deviation player i prefers to cooperation where she would receive $u_i(\tau^*)$ in future rounds. Let μ be the maximum of $\left\{\frac{f_i(\tau^*)}{u_i(\tau^*)}\right\}_{i\in N}$. If $\mu < \mu^*$, then there exists a cooperative outcome. If $\mu > \mu^*$, then there does not exist a cooperative equilibrium. When $\mu = \mu^*$, the existence of a cooperative equilibrium depends on the characteristics of the continuation probability (*e.g.*, is the continuation probability monotone?). Jones [7] explores this last consideration.

For a fixed continuation probability β and a fixed discount parameter δ, let $M = M(\delta, \beta)$. Further, denote the cone of cooperation vectors satisfying $\mu\Lambda \leq M\Lambda$ as $C(\mu)$. The proofs of the following theorems appear in the Appendix and are accompanied by lemmas and text which develops the necessary functional analysis. However, to give intuition to the techniques employed, remember that M is an operator from ℓ^∞ to ℓ^∞. The bifurcation point is actually the spectral radius of the operator M. The spectral radius is a nonnegative real number that bounds all of the eigenvalues with eigenvectors in ℓ^∞. In general, the spectral radius of an operator can be difficult to compute (see Dunford and Schwartz [5]); however, since M is upper triangular and nonnegative, it is easier to compute the spectral radius of M.

I limit the inspection to continuation probabilities with all positive terms, *i.e.*, $\beta_k > 0$ for all k. If $\beta_k = 0$ for some k, then the game is finite and there are no cooperative outcomes.

The following equation relates the spectral radius to the continuation probability and discount parameter. This equation is due in part to the spectral radius being a limit point of a sequence of eigenvalues and the spectral radius bounding all eigenvalues (with eigenvectors in ℓ^∞) of the operator. The proof of the following equation appears in Jones [8]. For a continuation probability with all positive terms, the spectral radius of M is $\dfrac{\mathcal{E}\delta}{1-\mathcal{E}\delta}$ where

$$\mathcal{E} = \lim_{n\to\infty}\left(\prod_{j=2}^{n+1}\beta_j\right)^{\frac{1}{n}}. \tag{4.1}$$

The assumption that $\beta_k > 0$ for all k does not insure that $\mathcal{E} > 0$. In fact, if $\lim_{k\to\infty}\beta_k = 0$ then $\mathcal{E} = 0$ and the spectral radius is zero. In this case, there does not exist *any* generalized Prisoner's Dilemma game where cooperation is a subgame perfect equilibrium. A classification of continuation probabilities as to when the spectral radius is zero appears in Jones [6] and [7]. In the following theorem, remember that $\mu > 0$ and $\mu^* \geq 0$.

Theorem 4.1. *a) When $\mu > \mu^*$, there does not exist a cooperative equilibrium satisfying $\mu\Lambda \leq M\Lambda$. b) When $\mu < \mu^*$, there exist cooperative equilibria satisfying $\mu\Lambda \leq M\Lambda$.*

We now know when there exists a nontrivial solution to $\mu\Lambda \leq M\Lambda$. One question remains: how do we find an equilibrium strategy? The following recursive formula yields one such way.

Theorem 4.2. *For any real number a, such that $\mu \leq a < \mu^*$, the following the recursive relationship*

$$\lambda^*_{k+1} = \left(\frac{a}{(a+1)\delta\beta_{k+1}}\right)\lambda^*_k \tag{4.2}$$

and $\lambda^*_1 = 1$ defines a vector solution to $\mu\Lambda \leq M\Lambda$. Upon rescaling, Λ^* defines an equilibrium strategy for the indefinitely repeated game.

The proof of the above theorem is contained in the Appendix. The idea stems from finding eigenvectors for the matrix, M, which have a as their eigenvalue. Rescaling may be necessary to insure that the entries of Λ^* define probabilities.

As μ increases, it follows that deviating becomes more profitable for at least one player. We expect that less cooperation should be supported by subgame perfect equilibria as μ increases. This intuition is substantiated below in Theorem 4.3. The proof of this theorem is in the Appendix and relies on constructing a vector in $C(\mu_2)$ that is not in $C(\mu_1)$.

Theorem 4.3. *If $\mu^* \neq 0$ and $0 \leq \mu_2 < \mu_1 < \mu^*$ then $C(\mu_1) \subsetneq C(\mu_2)$.*

5. Examples of the Bifurcation Point of Cooperation

The specific examples of finitely and infinitely repeated games demonstrate how this type of matrix analysis encompasses the standard results from game theory. Although the results are standard in these cases, the method is quite different and yields a different insight as to when cooperation is sustained by a subgame perfect publicly correlated equilibrium. Example 5.3 examines a game that cannot be analyzed by traditional game theory. An indefinitely repeated, Prisoner's Dilemma game is considered with a continuation probability that has a subsequence converging to zero. A representative cooperation vector of the cone of cooperation is examined.

Example 5.1. *Consider a finitely repeated game of length r. The associated continuation probability β satisfies $\beta_k = 1$ for $k \leq r$ and $\beta_k = 0$ for $k > r$. Notice that $M(\delta, \beta)$ is a nilpotent operator; specifically, $M^r(\delta, \beta)$ is the zero matrix. Therefore, the spectral radius of $M(\delta, \beta)$ is zero. (If μ^* is the spectral radius, then $M^r(\delta, \beta)\Lambda^* = (\mu^*)^r\Lambda^* = [0]\Lambda^*$. Therefore, $\mu^* = 0$.) By Theorem 4.1, there does not exist a nontrivial cooperation vector Λ that satisfies $\mu\Lambda \leq M(\delta, \beta)\Lambda$ for any μ. This solution coincides with traditional game theory, but utilizes the linear operator approach.*

The familiar result that cooperation in every round can be supported by a subgame perfect equilibrium in an infinitely repeated game as long as the discount parameter is large enough is examined in the next example. Recall that for an infinitely repeated game with stage game in $\mathcal{S}$, full cooperation can be supported in every round (*i.e.*, $\lambda_k = 1$ for all k) by a subgame perfect publicly correlated equilibrium as long as, for every player i,

$$\frac{f_i(\tau^*)}{[f_i(\tau^*) + u_i(\tau^*)]} \leq \delta. \tag{5.1}$$

For the infinitely repeated game, I show that the spectral radius is nonzero and that the cooperation vector $\Lambda_1 = (1, 1, 1, \dots)$ representing full cooperation in every round is an eigenvector for $M(\delta, \beta)$ as long as the discount parameter is large enough.

Example 5.2. *Let β be the continuation probability representing the infinitely repeated game. So, $\beta_k = 1$ for all k. By definition, $\mathcal{E} = 1$. Therefore, the spectral radius of $M(\beta, \delta)$ is $\mu^* = \dfrac{\delta}{1-\delta}$. By Theorem 4.1, there exists a nontrivial cooperation vector satisfying $\mu\Lambda \leq M(\delta, \beta)\Lambda$ for all $\mu < \mu^*$.*

Traditional game theory insures that full cooperation in every round can be supported by a subgame perfect publicly correlated equilibria as long as Equation 5.1 holds. Notice that $\Lambda_1 = (1, 1, 1, \dots)$ is an eigenvector of $M(\delta, \beta)$ for the eigenvalue $\mu^* = \frac{\delta}{1-\delta}$. For Λ_1 to yield a subgame perfect equilibrium, it follows for every player i that

$$\frac{f_i(\tau^*)}{u_i(\tau^*)}\Lambda_1 \leq M(\delta, \beta)\Lambda_1 = \left(\frac{\delta}{1-\delta}\right)\Lambda_1$$

must hold. But this matrix inequality is true, when for every i,

$$\frac{f_i(\tau^*)}{u_i(\tau^*)} \leq \frac{\delta}{1-\delta}.$$

However, this inequality can be rewritten as Equation 5.1.

Example 5.3. *Consider the Prisoner's Dilemma game with the bimatrix form given in Figure 5.1. This game is in $\mathcal{S}$. In particular, there exists a single Nash equilibrium, the pure strategy σ^* =(defect, defect) with $u_i(\sigma^*) = 0$ for $i = 1$ and 2. Let τ^* be the* *pure strategy (cooperate, cooperate). Then τ^* Pareto dominates σ^*. It follows that $f_i(\tau^*) = 1$ and $u_i(\tau^*) = 2$ for each player i.*

		Player 2: *defect*	Player 2: *cooperate*
Player 1	*defect*	*0,0*	*3,-1*
	cooperate	*-1,3*	*2,2*

Bimatrix form of the Prisoner's Dilemma
Figure 5.1

Let β be the continuation probability defined by

$$\beta_k = \begin{cases} \left(\frac{3}{4}\right)^{\sqrt{k}-1} & if \quad \sqrt{k} \in \mathbb{N} \\ 1 & otherwise. \end{cases}$$

For instance, $\beta_4, \beta_9, \beta_{25}, \ldots$ have a decreased likelihood of the game continuing, while the probability that the game continues is assured otherwise ($\beta_1, \beta_2, \beta_3, \beta_5, \ldots$ are all 1). Think of the continuation probability as adding severe shocks and dispersing the shocks further apart as time goes on. Even though the shocks become more severe, the shocks are spread further and further apart. Let the discount parameter be $\delta = \frac{2}{3}$.

By definition,

$$\mathcal{E} = \lim_{n\to\infty}\left(\prod_{j=2}^{n+1}\beta_j\right)^{\frac{1}{n}} = \lim_{n^2\to\infty}\left(\prod_{j=2}^{n^2+1}\beta_j\right)^{\frac{1}{n^2}} = \lim_{n^2\to\infty}\left(\prod_{j=2}^{n}\left(\frac{3}{4}\right)^{j}\right)^{\frac{1}{n^2}}$$

$$= \lim_{n^2\to\infty}\left[\left(\frac{3}{4}\right)^{\frac{n(n-1)}{2}}\right]^{\frac{1}{n^2}} = \left(\frac{3}{4}\right)^{\frac{1}{2}} = \frac{\sqrt{3}}{2}.$$

Substituting in the values for $\mathcal{E}$ and δ yields $\mu^* = \dfrac{1+\sqrt{3}}{2}$. Since the spectral radius is greater than $\frac{f_i(\tau^*)}{u_i(\tau^*)} = \frac{1}{2}$ (for both i), then there exists a nontrivial subgame perfect publicly correlated equilibrium for the indefinitely repeated game described in this example. Therefore, the cone of cooperation consists of nontrivial vectors including the eigenvector strategy associated with the eigenvalue 1 given below. This strategy de-

fines a subgame perfect publicly correlated strategy since $\frac{1}{2} \leq 1 < \mu^*$. (A legitimate strategy is defined by the eigenvector associated with μ^* also.)

Define a cooperation vector $\Lambda = (\lambda_1, \lambda_2, \lambda_3, \dots)$ by the recursive relationship of Equation 4.2 for the eigenvalue 1. The vector Λ is given by $\lambda_1 = 1$ and

$$\lambda_{k+1} = \left[\left(\frac{3}{4}\right)^k \Big/ \left(\prod_{j=2}^{n} \beta_j \right) \right].$$

To give a better idea of how Λ and β compare, consider the first 11 terms of each given below.

	1	*2*	*3*	*4*	*5*	*6*	*7*	*8*	*9*	*10*	*11*
Λ	*1*	$\left(\frac{3}{4}\right)$	$\left(\frac{3}{4}\right)^2$	$\left(\frac{3}{4}\right)^2$	$\left(\frac{3}{4}\right)^3$	$\left(\frac{3}{4}\right)^4$	$\left(\frac{3}{4}\right)^5$	$\left(\frac{3}{4}\right)^6$	$\left(\frac{3}{4}\right)^5$	$\left(\frac{3}{4}\right)^6$	$\left(\frac{3}{4}\right)^7$
β	*1*	*1*	*1*	$\left(\frac{3}{4}\right)$	*1*	*1*	*1*	*1*	$\left(\frac{3}{4}\right)^2$	*1*	*1*

The continuation probability offers shocks to the game. As the equilibrium strategy approaches a shock, the cooperation decreases since the game may end with high probability at round k^2. However, after surviving a shock, the amount of cooperation increases since the next shock does not occur until round $(k+1)^2$. As the number of rounds increase, the shocks dissipate but become more severe. The existence of a subgame perfect equilibrium strategy indicates that the players are forward looking and see beyond the high probability of the game ending, even though there is a subsequence of the continuation probability going to zero.

6. Concluding Remarks

By limiting my focus to simple publicly correlated strategies, I have been able to examine games in between finitely and infinitely repeated generalized Prisoner's Dilemma games. The techniques and ideas in this paper exploit the geometry of the set of subgame perfect publicly correlated equilibria, rather than considering the geometry of the payoffs. Considering the restricted type of strategy, I have been able to justify intuition of what happens in between finitely and infinitely repeated games. However, I concentrate on only a small portion of the continuum of repeated games.

Throughout, I assumed equal levels of cooperation, *i.e.*, where all players play the "cooperative" strategy with the same probability. It is possible to extend these results with publicly correlated strategies which utilize unequal levels of cooperation. However, the analysis yields a nonlinear operator which is difficult to manipulate. This nonlinear operator is "bounded" by an appropriate linear operator which yields results similar to those in this paper, see Jones [6].

Finally, it is possible to examine the spectral properties of matrices in relation to the continuation probabilities. Continuation probabilities can be classified with regards to the existence of a cooperative outcome for *any* generalized Prisoner's Dilemma game by the spectral properties of the associated operator. This appears as the Classification Theorem of Continuation Probabilities in Jones [7].

7. Appendix

As in Section 4, all the entries of the continuation probability are positive. Since the continuation probability and discount parameter are fixed, then let $M = M(\delta, \beta)$.

Lemma 7.1. *Every number a is an eigenvalue of $M : \mathbb{R}^\infty \to \mathbb{R}^\infty$ with an associated eigenvector unique up to multiples.*

Proof. The vector Λ^* is an eigenvector for the eigenvalue a if $M\Lambda^* = a\Lambda^*$, or equivalently,

$$a\lambda_k^* = \sum_{r=1}^{\infty} \left(\delta^r \left(\prod_{j=1}^{r} \beta_{k+j} \right) \lambda_{k+r}^* \right)$$

for all k.

Let $\lambda_1^* = 1$. Define the other terms of Λ^* by the equation

$$\lambda_{k+1}^* = \left(\frac{a}{a+1} \right)^k \left(\prod_{j=2}^{k+1} \beta_j \right)^{-1} \delta^{-k}.$$

Substitution of this equation below yields,

$$\sum_{r=1}^{\infty} \delta^r \left(\prod_{j=1}^{r} \beta_{k+j} \right) \lambda_{k+r}^* = \sum_{r=1}^{\infty} \left[\delta^r \left(\prod_{j=1}^{r} \beta_{k+r} \right) \left(\frac{a}{a+1} \right)^{k+r-1} \left(\prod_{j=2}^{k+r} \beta_j \right)^{-1} \left(\delta^{1-k-r} \right) \right]$$

$$= \delta^{1-k} \left(\prod_{j=2}^{k} \beta_j \right)^{-1} \left(\frac{a}{a+1} \right)^k \sum_{r=0}^{\infty} \left(\frac{a}{a+1} \right)^r = \lambda_k^* a = a\lambda_k^*$$

which proves that $\Lambda^* = (\lambda_1^*, \lambda_2^*, \dots)$ is an eigenvector with eigenvalue a.

Let $\Lambda \neq \Lambda^*$ be an eigenvector associated with a. Then $a\Lambda = M\Lambda$. It follows that, for all k,

$$\begin{aligned} a\lambda_k &= \sum_{r=1}^{\infty} \left[\delta^r \left(\prod_{j=1}^{r} \beta_{k+r} \right) \lambda_{k+j} \right] \\ &= \delta\beta_{k+1}\lambda_{k+1} + \sum_{r=2}^{\infty} \left[\delta^r \left(\prod_{j=1}^{r} \beta_{k+j} \right) \lambda_{k+r} \right] \\ &= \delta\beta_{k+1}\lambda_{k+1} + \delta\beta_{k+1} \sum_{r=1}^{\infty} \left[\delta^r \left(\prod_{j=1}^{r} \beta_{k+j+1} \right) \lambda_{k+r+1} \right] \\ &= \delta\beta_{k+1}\lambda_{k+1} + \delta\beta_{k+1} a\lambda_{k+1} = \delta\beta_{k+1}(1+a)\lambda_{k+1}, \end{aligned}$$

since Λ is an eigenvector. Isolating λ_{k+1} yields the recursive relationship in Equation 4.2. Therefore, Λ^* is unique up to multiples. □

The eigenvector associated with any number a may not define a strategy. Recall that all entries must be between 0 and 1 inclusive. If the eigenvector contains bounded entries, then a multiple of the eigenvector will define a strategy after rescaling. Notice that the sign of all of the entries are the same.

Lemma 7.2. *When $a < \mu^*$, then $\Lambda^* \in \ell^\infty$.*

Proof. Let Λ^* be the eigenvector associated with a. It follows that

$$\begin{aligned} \lim_{k\to\infty} (\lambda_k^*)^{\frac{1}{k}} &= \lim_{k\to\infty} \left(\left(\frac{a}{a+1} \right)^k \left(\prod_{j=2}^{k+1} \beta_j \right)^{-1} \delta^{-k} \right)^{\frac{1}{k}} \\ &= \frac{a}{(a+1)\mathcal{E}\delta} < 1 \end{aligned}$$

since $a < \frac{\mathcal{E}\delta}{1-\mathcal{E}\delta}$. The limit implies that the sequence $\{\lambda_k^*\}_{k=1}^{\infty}$ is bounded. Therefore, $\Lambda^* \in \ell^\infty$. □

Proof. (of Theorem 4.2) The result follows from Lemma 7.2. □

Proof. (of Theorem 4.1*a*)) This is a direct application of Lemma 7.1. □

Proof. (of Theorem 4.1*b*)) Assume that $\Lambda \in \ell^\infty$ and consists of nonnegative terms with at least one nonzero entry. Further, assume that Λ satisfies the matrix inequality, $M\Lambda \geq \mu\Lambda$. Define $\Lambda^* = (\Lambda / \sup_{k\in\mathbb{N}}\{\lambda_k\})$. The vector Λ^* is on the unit sphere in ℓ^∞. Due to the linearity of M, Λ^* is also a solution to $M\Lambda^* \geq \mu\Lambda^*$.

Since M contains nonnegative entries and M is a linear operator, then

$$M^2\Lambda^* = M(M\Lambda^*) \geq M(\mu\Lambda^*) \geq \mu^2\Lambda^*.$$

Upon iteration, it follows that $M^n\Lambda^* \geq \mu^n\Lambda^*$. The vector $\Lambda_1 = (1, 1, 1, \dots)$ is termwise greater than or equal to Λ^* since Λ^* is on the unit sphere. Since M^n consists of nonnegative entries then $M^n\Lambda_1 \geq M^n\Lambda^* \geq \mu^n\Lambda^*$. Under the sup norm in ℓ^∞, it follows that $\|M^n\Lambda_1\|_\infty \geq \|\mu^n\Lambda^*\|_\infty = \mu^n$. By the definition of the spectral radius of M and M having nonnegative entries,

$$\mu^* = \lim_{n\to\infty} \|M^n\Lambda_1\|_\infty^{\frac{1}{n}} \geq \lim_{n\to\infty} \|\mu^n\Lambda^*\|_\infty^{\frac{1}{n}} \geq \mu.$$

However, this is a contradiction since $\mu > \mu^*$. Therefore, $C(\mu)$ consists solely of the zero vector. $\square$

Proof. (of Theorem 4.3) Let $\Lambda \in C(\mu_1)$. Then Λ satisfies $\mu_1\Lambda \leq M(\delta, \beta)\Lambda$. Since $\mu_2\Lambda \leq \mu_1\Lambda$, transitivity implies that $\mu_2\Lambda \leq M(\delta, \beta)\Lambda$. The cone $C(\mu_1)$ is weakly contained in $C(\mu_2)$. It follows that the eigenvector Λ^* associated with μ^* is in both cones.

To show that $C(\mu_1)$ is contained in and not equal to $C(\mu_2)$ requires proving the existence of a vector that is in $C(\mu_2)$ but not in $C(\mu_1)$. The method is to construct such a vector by perturbing Λ^*.

Assume that Λ^* is not the vector consisting of all 1's. If Λ^* is the vector of all 1's, use the vector $(\Lambda^*/2)$ for the following construction. Let k be the first term where $\lambda_k^* < 1$. Since Λ^* is in the cones $C(\mu_1)$ and $C(\mu_2)$, the following inequalities hold

$$\mu_2\lambda_k^* < \mu_1\lambda_k^* \leq \sum_{r=1}^{\infty} \delta^r \left(\prod_{j=1}^{r} \beta_{k+j}\right) \lambda_{k+r}^*.$$

Let ϵ be defined to make the following equality hold

$$\mu_2(\lambda_k^* + \epsilon) = \sum_{r=1}^{\infty} \delta^r \left(\prod_{j=1}^{r} \beta_{k+j}\right) \lambda_{k+r}^*.$$

Clearly, $\epsilon > 0$. If $\lambda_k^* + \epsilon \leq 1$ then define

$$\Lambda' = (\underbrace{0, 0, \dots, 0}_{k-1 \text{ zeros}}, \lambda_k^* + \epsilon, \lambda_{k+1}^*, \lambda_{k+2}^*, \lambda_{k+3}^*, \dots).$$

This vector Λ' satisfies $\mu_2\Lambda' \leq M(\delta, \beta)\Lambda'$, but does not satisfy the same inequality with μ_1.

The inequality $\mu_1\Lambda' \leq M(\delta, \beta)\Lambda'$ does not hold since

$$\begin{aligned} \mu_1\lambda_k' &= \mu_1(\lambda_k^* + \epsilon) > \mu_2(\lambda_k^* + \epsilon) \\ &= \sum_{r=1}^{\infty} \delta^r \left(\prod_{j=1}^{r} \beta_{k+j}\right) \lambda_{k+r}^* = \sum_{r=1}^{\infty} \left(\prod_{j=1}^{r} \beta_{k+j}\right) \lambda_{k+r}'. \end{aligned}$$

If $\lambda_k^* + \epsilon > 1$ then define

$$\lambda' = \left(\underbrace{0, 0, \ldots, 0}_{k-1 \text{ zeros}}, 1, \frac{\lambda_{k+1}^*}{\lambda_k^* + \epsilon}, \frac{\lambda_{k+2}^*}{\lambda_k^* + \epsilon}, \frac{\lambda_{k+3}^*}{\lambda_k^* + \epsilon}, \ldots \right)$$

and proceed in the same fashion. It follows that $C(\mu_1)$ is contained in and not equal to $C(\mu_2)$. □

References

1. Abreu, Dilip. 1988. On the Theory of Infinitely Repeated Games with Discounting. *Econometrica* (2) 56: 383-196.
2. Becker, N.C. and A.E. Cudd. 1990. Indefinitely Repeated Games: A Response to Carroll. *Theory and Decision* 28: 189-195.
3. Bernheim, B.D. and A. Dasgupta. 1993. Repeated Games with Asymptotically Finite Horizons. *Journal of Economic Theory*, forthcoming.
4. Carroll, J.W. 1987. Indefinite Terminating Points and the Iterated Prisoner's Dilemma. *Theory and Decision* 22: 247-256.
5. Dunford, N. and J. Schwartz. 1957. *Linear Operators Part I: General Theory*. Somerset, NJ: Interscience Publishers, Inc.
6. Jones, M.A. 1994. Indefinitely Repeated Games and Cooperation. Doctoral Thesis. Northwestern University.
7. Jones, M.A. 1995. The Classification of Continuation Probabilities. Unpublished Manuscript.
8. Jones, M.A. 1995. Perron-Fröbenius Theory and the Repeated Prisoner's Dilemma. Unpublished Manuscript.
9. Luce, R.D. and H. Raiffa. 1957. *Games and Decisions: Introduction and Critical Survey*. Somerset, NJ: John Wiley and Sons, Inc.
10. Murnighan, J., and A. Roth. 1978. Equilibrium Behavior and Repeated Play of the Prisoner's Dilemma. *Journal Mathematical Psychology* 17: 189-198.
11. Myerson, R. 1991. *Game Theory: Analysis of Conflict*. Cambridge, MA: Harvard University Press.
12. Radner, R. 1986. Can Bounded Rationality Resolve the Prisoner's Dilemma? In *Contributions to Mathematical Economics in Honor of Gérard Debreu*, ed. W. Hildenbrand and A. Mas-Colell. New York, NY: Elsevier Science Publishers.
13. Ratliff, J. 1993. Repeated Prisoners' Dilemmas with Asymmetric Information about Duration. Unpublished Manuscript.
14. Samuelson, L. 1987. A Note on Uncertainty and Cooperation in a Finitely Repeated Prisoner's Dilemma Game. *International Journal of Game Theory*. (3) 16: 187-195.
15. Stahl, D. 1991. The Graph of Prisoners' Dilemma Supergame Payoffs as a Function of the Discount Factor. *Games and Economic Behavior* (3) 3: 368-384.

Department of Mathematical Sciences, United States Military Academy, West Point, NY 10996

E-mail address: am7426@euler.math.usma.edu