


UNIVERSITY OF
ILLINOIS LIBRARY
AT URBANA-CHAMPAIGN
BOOKSTACKS



Digitized by the Internet Archive
in 2011 with funding from
University of Illinois Urbana-Champaign

<http://www.archive.org/details/equilibriumperfe825cave>

1283
NO 825

cop. 2

OPEN SHELF



BEBR

FACULTY WORKING
PAPER NO. 825

Equilibrium and Perfection in Discounted
Supergames. I: Public Lotteries

Jonathan A. K. Cave

LIBRARY U. OF I. URBANA-CHAMPAIGN

The Trustees of the
JAN 21 1982
University of Illinois
at Urbana-Champaign

College of Commerce and Business Administration
Bureau of Economic and Business Research
University of Illinois, Urbana-Champaign

~~5110-3~~
~~6315-2~~
Cap. 2

330
B385
no. 825

BEBR

FACULTY WORKING PAPER NO. 825

College of Commerce and Business Administration

University of Illinois at Urbana-Champaign

November 1981

Equilibrium and Perfection in Discounted
Supergames, I: Public Lotteries

Jonathan A. K. Cave, Assistant Professor
Department of Economics

In this paper, we discuss equilibrium and perfect equilibrium in a simplified model of the supergame. We assume that players can observe the mixed moves employed by all players at each previous stage. For this model, we obtain a complete characterization of the set of equilibrium outcomes, and a fairly weak sufficient condition for this set to coincide with the set of perfect equilibrium outcomes. Inter alia, simple proofs of the Folk Theorem and the result that the requirement of perfection does not eliminate any equilibrium outcomes for the undiscounted supergame are presented.

Equilibrium and Perfection in Discounted Supergames, I: Public Lotteries

I. Introduction

This paper deals with some results on the characterisation of payoffs sustainable by equilibria or perfect equilibria of infinitely-repeated games with discounting. This work extends previous work on supergames without discounting by Aumann (1976), Aumann and Shapley, and Rubenstein (1977). In this paper, we work with a simplified model of the supergame used by Roth and Rubenstein (1977), in which players can observe the behavioral strategies used by their opponents at the conclusion of each play. In a subsequent paper, we show that passage to the more general model in which only the realisations of these strategies can be observed does not materially affect the results.

In the undiscounted case, where players evaluate the infinite streams of payoffs accruing to them by the limit of means, should it exist, it has been demonstrated that any outcome that is feasible and individually rational in the stage game can be sustained by a Nash Equilibrium of the supergame. In this context, an outcome is feasible if it belongs to the convex hull of the pure-strategy payoffs and is individually rational if it yields each player a payoff at least as great as his minmax payoff in the game. Moreover, it has been demonstrated that the additional requirement of perfection does not affect the set of outcomes.

In the supergame with discounting, we shall find that neither of these results goes through. In the first place, not all outcomes in the convex hull of the pure-strategy payoffs are feasible. Of those outcomes which are feasible, not all the individually rational ones can be obtained as equilibrium outcomes for the discounted supergame, since

myopic players will not be deterred by the promise of eventual punishment. Finally, not all equilibrium outcomes can be supported by perfect equilibrium outcomes, although there are several sufficient conditions that include many games of theoretical and economic interest.

The organisation of the paper is as follows: in the first section we describe the model of the supergame we are using. In the second section, we describe the set of attainable outcomes in the discounted supergame. The third section characterises the set of equilibrium outcomes, while the fourth section describes certain sufficient conditions for this set to co-incide with the set of perfect equilibrium outcomes, and provides a simple economic example of the possibility that these sets may differ.

II. The model

The verbal description of the model is as follows: we begin with a stage game in normal (strategic) form, with a finite number of players, each of whom selects one of a finite number of pure strategies. This game is to be played a denumerable infinity of times, and at each play, the choices of the players are allowed to depend on the entire previous history of the game. In particular, this means that each player is allowed to observe the mixed strategy in the stage game used by each of opponents at each previous stage. This is a strong assumption, and requires some justification. One justification is that one might think of this as a situation in which only pure strategies are allowed, and where the stage game has a continuous payoff function defined over convex, compact and finite-dimensional pure strategy sets for each player. Another interpretation is that the players meet on

successive days, but that play during each day consists of a sufficiently large number of repetitions of the specified mixed strategies for that day that each player can observe the mixtures used by his opponents with probability arbitrarily close to unity. To this, we add the further condition that discounting is done on a daily, rather than a continuous basis, and that players' strategies are fixed during the course of each day's play. A stronger justification will be provided in the sequel, where we demonstrate that the only effect of relaxing the assumption is to shrink the set of attainable outcomes, and that the restrictions of the sets of equilibrium and perfect equilibrium outcomes found in this paper to the new set of attainable outcomes form the new sets of equilibrium and perfect equilibrium outcomes. The result of an n-tuple of supergame strategies is an infinite sequence of n-tuples of mixed strategies in the stage game; to this, we associate a corresponding infinite random sequence of payoffs. There are various ways for the players to evaluate these sequences, but we shall concentrate on the discounted sum, normalised to lie within the convex hull of the payoffs in the stage game.

2.1 Definition: The stage game is a triple $[N, S, h]$, where N is a finite set of players; $S = \times_{i \in N} S_i$ is the set of n-tuples of pure strategies (also finite), and $h: S \rightarrow R^n$ is the payoff function. We also define the mixed extension of the stage game to be the triple $[N, M, H]$, where N is as above; $M = \times_{i \in N} M_i$ is the set of mixed strategies, with generic member

$$m = (m_1, \dots, m_n) \text{ where, for each } i$$

$m_i \in \Delta |S_i|$ is a probability distribution on the members of S_i .

Thus, m is a probability distribution on the members of S , although not all such probability distributions (called correlated strategies) can be represented as members of M . If $m(s)$ is the probability that the n -tuple of pure strategies s , will be played:

$$m(s_1, \dots, s_n) = \prod_{i \in N} m_i(s_i)$$

we define the expected payoff $H(m)$ by

$$H(m) = \sum_s m(s)h(s)$$

It is clearly a continuous function, linear in each of the numbers $m_i(s_i)$.

2.2 Definition: Let $[N, S, h] = G$ be a stage game; $G^* = [N, F, P]$ is a supergame of G if the following are satisfied: $F = \times_{i \in N} F_i$ is the space of pure supergame strategies. For each member $f_i \in F_i$, we write $f_i = (f_i^1, \dots, f_i^t, \dots)$, where

$$f_i^1 \in S_i$$

$$f_i^t: [S]^{t-1} \rightarrow S_i$$

and $P = [P_1, \dots, P_n]$ where each P_i is a partial ordering on the space R^∞ of infinite sequences of real numbers. P_i represents the preferences of player i over the infinite streams of payoffs resulting from plays of the supergame.

2.3 Definition: Let G^* be a supergame, and f an n -tuple of pure strategies for the supergame. We can calculate a sequence $s(f) = (s^1(f), \dots, s^t(f), \dots)$ of outcomes in the stage game as follows:

$$s^1(f) = r^1(f) = (f_1^1, \dots, f_n^1) = f^1$$

$$s^t(f) = (f_1^t(r^{t-1}(f)), \dots, f_n^t(r^{t-1}(f))) = f^t(r^{t-1}(f))$$

$$r^t(f) = (r^{t-1}(f), s^t(f))$$

Thus, $s^t(f)$ is the action specified by f for the t^{th} play of the game (if all previous plays have been according to f), and $r^t(f)$ is the cumulative record of play up to and including the play on date t . To the sequence $s(f)$ we can associate a sequence of payoffs $g(f) = (g_1^1(f), \dots, g_n^1(f), \dots, g_1^t(f), \dots)$ in the obvious way, using the pure strategy payoff function h from the stage game:

$$g_i^t(f) = h_i(s^t(f))$$

and for convenience, we shall write $g(f) = (g^1(f), \dots, g^t(f), \dots)$ and $g(f) = (g_1(f), \dots, g_n(f))$.

2.4 Definition: A discounted supergame is a supergame where each player i is characterised by a discount rate $\delta_i \in [0, 1]$, and has the preference relation P_i defined by $(x^1, \dots, x^t, \dots) P_i (y^1, \dots, y^t, \dots)$ (where $x, y \in R^\infty$) iff

$$\bar{h}_i^\delta(x) = (1 - \delta_i) \sum_{t=1}^{\infty} \delta_i^{t-1} x^t \geq (1 - \delta_i) \sum_{t=1}^{\infty} \delta_i^{t-1} y^t = \bar{h}_i^\delta(y)$$

By an obvious abuse of notation, we can define a payoff function $h^\delta: F \rightarrow R^n$ for the discounted supergame:

$$h_i^\delta(f) = \bar{h}_i^\delta(g_i(f))$$

One of the nicest features of the discounted supergame is that with this payoff function, the set of payoffs $h^\delta(F)$ is a compact subset of $CH(h(S))$ and also that $h^\delta(f)$ is continuous in f . There are other preference relations that have been used, including the first Caesaro mean of the payoffs and the overtaking relation. However, the first of these will not give an answer for payoff sequences that are not Caesaro summable, while the second cannot be represented by a payoff function. The next problem is that of mixed strategies. Since these supergames are games of perfect recall, it follows from Aumann's (1964) extension of Kuhn's (1953) theorem that it is sufficient to confine our attention to behavioral strategies; in this case, a behavioral strategy is a device which selects a mixed strategy in each stage game.

2.5 Definition: Let G^* be a supergame, and define $\bar{F} = \prod_{i \in N} \bar{F}_i$ to be the space of m -tuples of behavioral strategies for the supergame. The generic member $\bar{f}_i \in \bar{F}_i$ is defined by:

$$\bar{f}_i^1 \in M_i$$

$$\bar{f}_i^t: [M]^{t-1} \rightarrow M_i$$

For any n -tuple \bar{f} of behavioral strategies, we can calculate a sequence $m(\bar{f})$ of mixed stage game outcomes as follows:

$$m^1(\bar{f}) = (\bar{f}_1^1, \dots, \bar{f}_n^1) = \bar{f}^1 = r^1(\bar{f})$$

$$m^t(\bar{f}) = (\bar{f}_1^t(r^{t-1}(\bar{f})), \dots, \bar{f}_n^t(r^{t-1}(\bar{f}))) = \bar{f}^t(r^{t-1}(\bar{f}))$$

$$r^t(\bar{f}) = [r^{t-1}(\bar{f}), m^t(\bar{f})]$$

Since the choices are independent at each stage, the expected payoff sequence is well-defined by $G(\bar{f}) = [G_i^t(\bar{f}) : i \in N, t = 1, \dots]$, where $G_i^t(\bar{f}) = H_i(m^t(\bar{f}))$. The discounted-supergame "payoff function" resulting from this definition is $H^\delta : \bar{F} \rightarrow R^n$, where $H_i^\delta(\bar{f}) = \sum_{t=1}^{\infty} \delta^{t-1} G_i^t(\bar{f})$.

It only remains to define equilibrium and perfect equilibrium for the discounted supergame.

2.6 Definition: Let $\bar{f} \in \bar{F}$ be an n-tuple of behavioural strategies for the discounted supergame G^* . We say that \bar{f} is a Nash Equilibrium iff, for each player i , and each behavioral strategy $\bar{f}'_i \in \bar{F}_i$, we have

$$H_i^\delta(\bar{f}) \geq H_i^\delta(\bar{f}_{(i)}, \bar{f}'_i)$$

where $\bar{f}_{(i)}$ denotes the n-1 tuple $(\bar{f}_1, \dots, \bar{f}_{i-1}, \bar{f}_{i+1}, \dots, \bar{f}_n)$ of behavioral strategies used by the other players. \bar{f} is a perfect equilibrium iff, for every t , and for every member m' of $[M]^t$, the "continuation strategy" $\bar{f}'_i(\cdot; m')$ defined by:

$$\bar{f}'_i{}^{\tau}(m^1, \dots, m^{\tau-1}; m') = \bar{f}_i{}^{t+\tau}(m', m^1, \dots, m^{\tau-1})$$

is an equilibrium in G^* . As a further matter of notation, we shall denote the set of members of $CH(h(S))$ that can be achieved as $H^\delta(\bar{f})$ for some Nash Equilibrium \bar{f} by e.p., while the subset of e.p. that can be achieved by perfect equilibria will be denoted p.e.p.

III. The Set of Attainable Outcomes

This section concerns the observation that, for sufficiently small values of the individual discount rates, it may happen that not all members of $CH(h(S))$ can be achieved in behavioral strategies, let alone in pure strategies. This is in sharp contrast to the situation for the undiscounted game, where any point in $CH(h(S))$ can be achieved in pure strategies, by playing the relevant pure strategy n-tuples of the stage game with frequencies that correspond to the weights used in the convex combinations forming $CH(h(S))$.

For simplicity, we work with the case where $\delta_i = \delta$, for all i . In any game $[N, S, h]$, we can isolate three subsets of $CH(h(S))$, corresponding to the outcomes that can be achieved using pure, mixed and correlated strategies in each stage. Let $C = CH(S)$ be the set of correlated strategies for the stage game, we define

3.1 Definition: $D_p^\delta = \{x \in CH(h(S)) \text{ s.t. there exists an infinite sequence } (s^1, \dots, s^t, \dots) \text{ of members of } S \text{ with the property that}$

$$x = (1 - \delta) \sum_{t=1}^{\infty} \delta^{t-1} h(s^t) \}$$

$D_m^\delta = \{x \in CH(h(S)) \text{ s.t. there exists an infinite sequence } (m^1, \dots, m^t, \dots) \text{ of members of } M \text{ with the property that}$

$$x = (1 - \delta) \sum_{t=1}^{\infty} \delta^{t-1} H(m^t)$$

$D_c^\delta = \{x \in CH(h(S)) \text{ s.t. there exists an infinite sequence } (c^1, \dots, c^t, \dots) \text{ of members of } C \text{ with the property that}$

$$x = (1 - \delta) \sum_{t=1}^{\infty} \delta^{t-1} H(c^t)$$

Clearly, since $H(C) = CH(h(S))$, we have $D_C^\delta = CH(h(S))$ for all δ . Moreover, in cases where $H(M) = CH(h(S))$, as with Prisoner's Dilemma, we have $D_m^\delta = CH(h(S))$. In general, the definition of the set of attainable points will hinge on the set of weights that can be obtained through the use of the relevant strategies. For example, if the discount rate is 0.1, the first pure (or mixed) payoff will have a weight of 0.9; this means that the eventual discounted sum must be close to one of the original pure or mixed strategy payoffs for the stage game. To see what this means in terms of D_p^δ , let us observe that the weight given to a particular pure-strategy outcome must be of the form

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

where $a = (a_1, \dots, a_t, \dots)$ is some infinite sequence of 0's and 1's. If we are taking convex combinations of m such pure-strategy outcomes, we need a characterisation of the feasible convex combinations.

3.2 Definition: Let $\lambda \in (0,1)$, and let Δ^m be the standard m -simplex.

Define

$$\Delta_\delta^m = \{\lambda \in \Delta^m: \text{there exist sequences } a_1, \dots, a_m \text{ s.t.}$$

- i) $a_i = (a_i^1, \dots, a_i^t, \dots), a_i^t \in \{0,1\}$ each i,t
- ii) for all $t, \sum_{i=1}^m a_i^t = 1$, and
- iii) $\lambda_i = (1 - \delta) \sum_{t=1}^{\infty} a_i^t \delta^{t-1}$

This is the set of weights on pure-strategy payoffs available via the use of pure strategies in the discounted supergame. We have the following result.

It is clear that for all $t, i, \lambda_i^t \geq 0$, and that for all i , $\lim_{t \rightarrow \infty} \lambda_i^t = 0$. Moreover, from the construction it follows that for all T ,

$$\lambda_i^0 = \lambda_i^{T+1} + (1 - \delta) \sum_{t=1}^T a_i^t \delta^{t-1}$$

which completes the proof.

QED

When we turn to mixed strategies, we should expect some relaxation of this condition. Indeed, in many cases we can achieve the entire set of outcomes. We shall not continue with our characterisation of the set of attainable outcomes, since there is insufficient generality to warrant it. However, we shall observe that there is a natural upper bound for m that in most cases is less than the number of pure-strategy combinations. $CH(h(S)) \subset R^n$, so that we need use no more than $N+1$ pure-strategy combinations. On the other hand, if $N \geq 2$ and $|S_i| \geq 2$ for each i , then $|S| > N+1$ so we obtain:

Theorem 2.4: If $\delta \geq 1 - \frac{1}{N}$, then the set of outcomes obtainable via pure strategies coincides with $CH(h(S))$.

We close this section with an example of a well-known game where

$D_p^\delta \neq D_m^\delta \neq D_c^\delta$: the Battle of the Sexes. This is a two-player game, where each player has two pure strategies:

	L	R
T	(2,1)	(0,0)
B	(0,0)	(1,2)

To begin with, the following three figures show the stage-game payoffs to pure, mixed and correlated strategies, respectively.

3.3 Proposition: $\delta \geq \frac{m-1}{m}$, iff $\Delta_\delta^m = \Delta^m$.

Proof: if $a_i^1 = 1$, then $\lambda_i \geq (1 - \delta)$, so a necessary condition for the conclusion of the Proposition is that $\max_i \lambda_i \geq (1 - \delta)$ for every $\lambda \in \Delta^m$. But $\lambda \in \Delta^m$ implies that $\max_i \lambda_i \geq \frac{1}{m}$, and this bound is tight, so we know the necessary condition is only satisfied if

$$\frac{1}{m} \geq 1 - \delta; \text{ or } \delta \geq \frac{m-1}{m}$$

It remains to be shown that no further restrictions result from the choice of subsequent weights. To do this, we shall exhibit a procedure by which the a_i^t can be calculated explicitly. To begin with, fix $\lambda^0 \in \Delta^m$. Let $i_1 \in \operatorname{argmax}_i \lambda_i^0$. We know that $\lambda_{i_1}^0 \geq \frac{1}{m}$, so that by hypothesis, $\lambda_{i_1}^0 \geq 1 - \delta$. Therefore, let us set $a_{i_1}^1 = 1$, and $a_j^1 = 0$ for all $j \neq i_1$, and form a new vector λ^1 by $\lambda_j^1 = \lambda_j^0$ for $j \neq i_1$, and $\lambda_{i_1}^1 = \lambda_{i_1}^0 - (1 - \delta)$. This new vector belongs to


$$\Delta_1^m = \{ \lambda \in \mathbb{R}_+^m : \sum_{j=1}^m \lambda_j = 1 - (1 - \delta) = \delta \}$$

Once again, we know that $\max_i \lambda_i^1 \geq \frac{\delta}{m}$. The condition for us to be able to choose a^2 appropriately is that $\max_i \lambda_i^1 \geq (1 - \delta)\delta$, so that this condition is satisfied if


$$\frac{\delta}{m} \geq (1 - \delta)\delta; \text{ or } \delta \geq \frac{m-1}{m}$$

so that successive levels of choice of a^t introduce no new conditions.


We have now demonstrated the truth of the Proposition, since the procedure begun above can be iterated by choosing $i_t \in \operatorname{argmax}_i \lambda_i^{t-1}$, setting $a_{i_t}^t = 1 = 1 - a_j^t$, all $j \neq i_t$, and letting $\lambda_j^{t+1} = \lambda_j^t$ for $j \neq i_t$, while $\lambda_{i_t}^{t+1} = \lambda_{i_t}^t - \delta^t(1 - \delta)$.



pure




mixed




correlated


To construct the set of payoffs obtainable via pure strategies, given a small discount rate, we first add a shrunken replica of the convex hull of the original payoffs near each of the pure-strategy payoffs. The final payoff must lie within one of these convex hulls; which one is determined by the choice of the first-period strategy combination. Next to this we have repeated the process within each of the next convex hulls, adding shrunken replicas to the vertices representing the pure choices at the first and second stages. The process continues inductively, resulting in a sparse nondenumerable set of payoffs.



Stage 1





Stage 2



Stage 3



To construct D_m^δ , the process is only slightly more complicated. For each point in the set of payoffs obtainable via mixed strategies, there will be a shrunken replica of that set, and we must take the envelope of these attached replicas. In the following figures, we show a few of these attached replicas, and the envelope of those replicas.



First stage: some replicas

the envelope

The next step is to repeat the process for each point added at the second step: this means that we must remain within the envelope of the convex hulls of the shrunken replicas added at the first stage: this envelope, shown in the left-hand figure below, represents the maximum area we can hope for. In the figure on the right, we have illustrated the result of this second step.



The outer limit

the second-stage envelope

These pictures make it clear that $D_p^\delta \neq D_m^\delta \neq D_c^\delta$.

IV. Equilibrium in the discounted supergame

Any n-tuple of supergame strategies can be divided into two parts: the specified sequence $s(f)$ that results from adherence to the strategies f , and various contingent sequences resulting from deviations from the specified sequence. Since each player knows the strategy descriptions of the other players in the equilibrium, each player can predict the future course of play for any choice of his own actions.

One immediate consequence is that any n-tuple of supergame strategies whose specified sequence consists of equilibria of the stage game, and which makes the same prescription for any history is an equilibrium. Thus, any sequence consisting of members of the set of stage-game equilibria can be sustained as the outcome of such an "open loop" equilibrium for any monotonic evaluation relation.

In general we shall be concerned with outcomes that cannot be achieved in this manner, so that we shall need to consider the concept of punishment. In general, the worst punishment that can be inflicted on a player in any single play of the game is that which holds him to his minmax security level. There are two reasons for using this security level rather than the (lower) maxmin level. The first is that the punishment to be used against a defector forms part of the declared strategies of the other players, so that the defecting player can adapt his "defense" to the specific punishment. The second is that in this game all lotteries are public, so that there is no way that the other players can use a correlated punishment against the defector. Of course, if they

were able to use correlated punishments, there would be no difference between the minmax and maxmin levels; both would co-incide with the value of the two-person zero-sum game played between the defector (the maximizing player) and the others (the minmising player), over the defector's payoffs.

From this observation, it follows that the strongest punishment that can be inflicted on a defector in the supergame is to hold him to his minmax level in all plays following the detection of a deviation. Such a punishment is called a grim punishment. We observe that a player can be deterred from defecting from a particular specified sequence if and only if the threat of grim punishment is sufficient to deter him.

4.1 Definition: Let $[N, M, H]$ be the mixed extension of a normal-form game. For each $i \in N$, define $p^i \in M$, as follows:

$$p_{(i)}^i \in \arg \min_{m_{(i)}} [\max_{m_i} H_i(m_i, m_{(i)})]$$

$$p_i^i \in \arg \max_{m_i} H_i(m_i, p_{(i)}^i)$$

This is the minmax punishment and defense to be used when player i defects. Let $v_i = H_i(p^i)$ be player i 's minmax security level.

Let $[\bar{m}^1, \dots, \bar{m}^t, \dots]$ be an infinite sequence of members of M . The grim strategy supporting $\bar{m} = [\bar{m}^1, \dots]$ is the n -tuple f of supergame strategies defined by

$$f_i^1 = \bar{m}_i^1 \quad \text{for all } i, \text{ and}$$

$$f_i^t(m^1, \dots, m^{t-1}) = \begin{cases} p_i^j & \text{iff there exists } T < t \text{ s.t.} \\ & \text{i) } m^{t'} = \bar{m}^{t'} \text{ for all } t' < T \\ & \text{ii) } m_k^T = \bar{m}_k^T \text{ for all } k \neq j \\ & \text{iii) } m_j^T \neq \bar{m}_j^T \\ \bar{m}_i^t & \text{if not} \end{cases}$$

In other words, the grim strategy plays according to the cooperative sequence until the first date when a defection occurs. If a single player is responsible for that defection, then that player is punished forever.

In the undiscounted game, a player contemplating defection from a grim strategy supporting a specified sequence \bar{m} has the choice of two outcomes: $\lim_{t \rightarrow \infty} H_i(\bar{m}^t)$ if he does not defect, and v_i if he does. If the first limit exists, i.e., if the sequence $H_i(\bar{m}^t)$ is Caesaro-convergent, player i will adhere to the grim strategy iff the first of these numbers exceeds the second. This is the "first Folk Theorem" of the undiscounted supergame:

4.2 Theorem: the set of limiting-average payoffs to equilibria of the undiscounted supergame is $\{y \in CH(h(S)) : y_i \geq v_i, \text{ all } i \in N\}$.

One striking feature of this result is that the outcomes can be characterised purely by their payoffs; no strategic considerations enter in. In particular, the immediate profit earned by the defector plays no role. Unfortunately, this is not true in the discounted game, so that the characterisation of equilibrium outcomes involves explicitly the strategic aspects of the specified sequence. However, it is still the case that the sine qua non of equilibrium is the existence of a grim-strategy equilibrium supporting the outcome, so that we obtain:

4.3 Theorem: $y \in D_m$ is the outcome of an equilibrium of the discounted supergame iff there exists an infinite sequence $(\bar{m}^1, \dots, \bar{m}^t, \dots) = \bar{m}$ of members of M with the following properties:

i) for all i , $y_i = (1 - \delta_i) \sum_{t=1}^{\infty} \delta_i^{t-1} H_i(\bar{m}^t)$

ii) for all i, T

$$(1) \quad \sum_{t \geq T} \delta_i^{t-T} H_i(\bar{m}^t) \geq \max_{m_i} H_i(m_i, \bar{m}^t_{(i)}) + (\delta_i / (1 - \delta_i)) v_i$$

or

$$(2) \quad y_i \geq (1 - \delta_i) [\max_{m_i} H_i(m_i, \bar{m}^t_{(i)}) + \sum_{t=1}^{T-1} \delta_i^{t-1} H_i(\bar{m}^t)] + \delta_i v_i$$

In particular, if $y \in H(M)$, y is the outcome of a stationary equilibrium of the discounted supergame iff there exists $m^* \in M$ s.t.

iii) $H(m^*) = y$

iv) for all $i \in N$

$$(3) \quad y_i \geq (1 - \delta_i) \max_{m_i} H_i(m_i, m^*_{(i)}) + \delta_i v_i.$$

Proof:

Stationary Equilibrium: Let us suppose that player i wishes to defect from a stationary grim strategy supporting the sequence m^*, m^*, \dots . If he does not defect, his payoff will be y_i ; if he does defect, his payoff will be at most $\max_{m_i} H_i(m_i, m^*_{(i)})$ in the first period and at most v_i in all subsequent periods. The normalised discounted payoff to his best defection is therefore the RHS of condition (3). This shows the sufficiency of the condition. Necessity follows from the following

observation: in any supergame strategy combination with stationary specified sequence m^*, m^*, \dots , the payoff to the best defection will be greater than or equal to the payoff to the best defection against the grim strategy supporting this outcome.

In general, this will be true: if \bar{m} is any infinite sequence, and f is any supergame strategy supporting this sequence (i.e., $m^t(f) = \bar{m}^t$), then

$$\max_{f'_i} H_i^\delta(f'_i, f_{(i)}) \geq \max_{f'_i} H_i^\delta(f'_i, g_{(i)})$$

where g is the grim strategy supporting \bar{m} .

Nonstationary equilibrium: In general, it may be the case that D_m^δ strictly contains $H(M)$. If player i chooses to defect from a grim strategy g supporting the sequence \bar{m} at time T , he exchanges an expected payoff sequence worth

$$\sum_{t>T} \delta_i^{t-T} H_i(\bar{m}^t)$$

for one which pays at most $\max_{m_i} H_i(m_i, \bar{m}^T)$ on day T , and v_i in all subsequent periods, for a total expected payoff of

$$\max_{m_i} H_i(m_i, \bar{m}_{(i)}^t) + \delta_i v_i / (1 - \delta_i)$$

as of day T , which give us the LHS and the RHS of (1), respectively.

Finally, we can obtain condition (2) by applying condition (i) to equation (1). QED

We shall present an example which uses this theorem to characterise the equilibrium points of Prisoner's Dilemma. Before we do so, there

are several consequences of this theorem that are worth noting. In the first place, by letting all the discount rates go to 1 we obtain precisely the "first Folk Theorem": any feasible and individually-rational payoff can be supported by an equilibrium of the undiscounted game. Strictly-speaking this gives us a version of the Folk Theorem where players use the Abel limit, rather than the first Caesaro sum to evaluate payoff streams. However, this poses no problems, since Caesaro convergence implies Abel convergence.

Another interesting feature of this result can be noted by letting the discount rate shrink to 0, condition (3) becomes the usual condition for Nash Equilibrium, while condition (1) limits us to precisely those sequences with which we began this section; sequences calling for a stage-game Nash equilibrium at every stage.

Finally, it will be noted that condition (1) can be written:

$$(4) \quad \min_{T} [\sum_{t \geq T} \delta_i^{t-T} H_i(\bar{m}^t) - \max_{m_i} H_i(m_i, \bar{m}_{(i)}^T)] \geq \delta_i v_i / (1 - \delta_i)$$

and it is clear that this is monotonic in δ_i : if y is an equilibrium outcome at $\delta = (\delta_1, \dots, \delta_n)$ and $\delta'_i \geq \delta_i$ for each $i \in N$, then y is an equilibrium outcome at δ' .

We conclude this section by characterising the outcomes of equilibria of a simple version of the Prisoner's Dilemma. The mixed extension of this game is as follows: let m_i be the probability that player i uses the "Greedy" Strategy, and $1 - m_i$ be the probability that player i uses the "Helpful" strategy. The payoff functions are:

$$H_1(m_1, m_2) = 3 + m_1 - 3m_2$$

$$H_2(m_1, m_2) = 3 + m_2 - 3m_1$$

It follows that, for any pair (m_1, m_2) , the best defection is $m_i = 1$, so that

$$\max_{m_i} H_1(m_1, m_2) = 4 - 3m_2; \quad \max_{m_2} H_2(m_1, m_2) = 4 - 3m_1; \text{ and}$$

$$v_1 = 1 = v_2$$

Thus, we can determine that (\bar{m}_1, \bar{m}_2) is the outcome of a stationary supergame equilibrium iff:


$$(5) \quad 3 + \bar{m}_1 - 3\bar{m}_2 \geq (1 - \delta_1)(4 - 3\bar{m}_2) + \delta_1; \text{ and}$$

$$(6) \quad 3 + \bar{m}_2 - 3\bar{m}_1 \geq (1 - \delta_2)(4 - 3\bar{m}_1) + \delta_2$$

We can rearrange these linear inequalities to give a unified condition on \bar{m} :

$$(7) \quad 1 - \frac{1 - m_2}{3\delta_2} \geq m_1 \geq 1 - 3\delta_1(1 - \bar{m}_2)$$

In the following figure, we show the image of this set of strategies, for discount rates δ_i between 1 and $\frac{1}{3}$.



We have labelled the boundaries of this region, to facilitate translating these strategy pairs into pairs of payoffs. In region I, we have $\bar{m}_2 = 0$, $\bar{m}_1 \leq 1 - \frac{1}{3\delta_2}$; in Region II, we have $\bar{m}_1 = 0$, and $\bar{m}_2 \leq \frac{1}{3\delta_1}$; in region III, we have $\bar{m}_2 = 1 - \frac{1}{3\delta_1}(1 - \bar{m}_1)$; while in region IV, we have $\bar{m}_2 = 1 - 3\delta_2(1 - \bar{m}_1)$. In region III, \bar{m}_2 ranges between $1 - \frac{1}{3\delta_1}$ and 1, and in region IV between 0 and 1. Inserting these boundary values into the payoff function, we get the corresponding regions in payoff space:


$$\text{I: } H = 3 + \bar{m}_1, 3 - 3\bar{m}_1; \bar{m}_1 \in [0, 1 - \frac{1}{3\delta_2}]$$

$$\text{II: } H = 3 - 3\bar{m}_2, 3 + \bar{m}_2; \bar{m}_2 \in [0, 1 - \frac{1}{3\delta_1}]$$

$$\text{III: } H = [(\bar{m}_1 + \frac{1}{\delta_1}(1 - \bar{m}_1)), (4 - \frac{1}{3\delta_1} - \bar{m}_1(3 - \frac{1}{3\delta_1}))]$$

$$\text{IV: } H = [(\bar{m}_1 + 9\delta_2(1 - \bar{m}_1)), (4 - 3(\bar{m}_1 + \delta_2(1 - \bar{m}_1)))]$$

We have displayed these regions below.



Turning now to the non-stationary equilibrium outcomes, we observe that since any outcome in D_m^δ is also in $H(M) = CH(h(S))$, we have only to see whether there are any outcomes which are more stable when the

specified sequence is nonstationary. The crucial element in this is the incentive to defect at any stage. Suppose that we are trying to support an outcome paying (y_1, y_2) : the stationary strategies giving this outcome specify a repetition of (m_1^*, m_2^*) where

$$m_1^* = \frac{3}{2} - \frac{1}{8}[y_i + 3y_j] \quad i = 1, 2; j = 2, 1, j \neq i$$

On the other hand, if \bar{m} is an infinite sequence with the same payoff, and if \bar{m} is the outcome of an equilibrium of the stage game, we have four equations to satisfy.

$$(8) \quad \sum_{t=1}^{\infty} m_1^t \delta_1^{t-1} - 3 \sum_{t=1}^{\infty} m_2^t \delta_1^{t-1} = \frac{1}{1 - \delta_1} (y_1 - 3)$$

$$(9) \quad \sum_{t=1}^{\infty} m_2^t \delta_2^{t-1} - 3 \sum_{t=1}^{\infty} m_1^t \delta_2^{t-1} = \frac{1}{1 - \delta_2} (y_2 - 3)$$

$$(10) \quad \min_T \left[\sum_{t \geq T} m_1^t \delta_1^{t-T} - 3 \sum_{t \geq T} m_2^t \delta_1^{t-T} \right] \geq \frac{1 - 3\delta_1}{1 - \delta_1}$$

$$(11) \quad \min_T \left[\sum_{t \geq T} m_2^t \delta_2^{t-T} - 3 \sum_{t \geq T} m_1^t \delta_2^{t-T} \right] \geq \frac{1 - 3\delta_2}{1 - \delta_2}$$

The first two being feasibility conditions, and the last two being equilibrium conditions derived from conditions i) and ii) of Theorem 4.3.

Since we are interested in the set of outcomes, and not the sequences that give rise to them, we may assume that there is no stationary equilibrium paying (y_1, y_2) . Thus either

$$(12) \quad y_1 > \frac{4(1 - 3\delta_1) - 3(1 - \delta_1)y_2}{(1 - 9\delta_1)} \quad \text{or}$$

$$(13) \quad y_2 > \frac{4(1 - 3\delta_2) - 3(1 - \delta_2)y_1}{(1 - 9\delta_2)}$$

In terms of the strategies used by the players, from the definition of m^* we know that

$$(14) \quad \sum_{t=1} ([3 + m_1^t - 3m_2^t] - [3 + m_1^* - 3m_2^*]) \delta_1^{t-1} = 0 = \\ \sum_{t=1} ([3 + m_2^t - 3m_1^t] - [3 + m_2^* - 3m_1^*]) \delta_2^{t-1}$$

Let us suppose that condition (12) is satisfied; it is player 1 who will defect from the stationary equilibrium. This means that

$$(15) \quad 3 + m_1^* - 3m_2^* < (1 - \delta_1)(4 - 3m_2^*) + \delta_1; \quad \text{or}$$

$$(16) \quad m_1^* < 1 - 3\delta_1 + 3\delta_1 m_2^*$$

Condition (14) can be rearranged (using only the left-hand equation) to give:

$$(17) \quad \sum_{t=1} m_1^t \delta_1^{t-1} - 3 \sum_{t=1} m_2^t \delta_1^{t-1} = \frac{m_1^* - 3m_2^*}{1 - \delta_1}$$

Inserting (16) into this gives

$$(18) \quad \sum_{t=1} m_1^t \delta_1^{t-1} - 3 \sum_{t=1} m_2^t \delta_1^{t-1} < \frac{1 - 3\delta_1}{1 - \delta_1} - 3m_2^*$$

which contradicts condition (10). Thus, any equilibrium outcome in the prisoner's dilemma can be supported by a stationary equilibrium, so the sets defined above provide the entire set of equilibrium outcomes.

V. Perfect Equilibrium in the discounted supergame

The grim strategies usually fail to provide perfect equilibrium outcomes, as the following simple example shows: Let the mixed extension of a game with two players be given by:

$$H(m_1, m_2) = (m_1, 2m_1m_2 - m_2 + m_1 - 1)$$

The grim strategy secures an equilibrium of the game with the stationary result $m_1^* = 1$, $m_2^* = 0$, as long as the discount rate for player 2 is at least $\frac{1}{2}$, and certainly for the undiscounted game. However, once player 2 has defected, it is incredible that player 1, who is in no way injured by player 2's defection, should actually carry out the grim punishment which costs him his entire remaining profit in the game.

In the undiscounted game, the requirement of perfection does not actually affect the set of payoffs sustained by equilibrium behavior, a result discovered independently by Rubenstein and by Aumann and Shapley. We shall call this the "perfectness Folk Theorem" and include a simple proof for the present model.

5.1 Theorem: The set of limiting average payoffs to perfect equilibria of the undiscounted supergame coincides with the set of payoffs to equilibria of the undiscounted supergame.

Proof: Let $y \in CH(h(S))$ be a feasible and individually-rational payoff: $y_i \geq v_i$ for each player i . We know by Theorem 4.2 that there exists an equilibrium of the undiscounted game with limiting average payoff exactly y . Let us denote the specified sequence of this equilibrium by $m(y) = [m^1(y), \dots, m^t(y), \dots]$. By the properties of the Caesaro mean, for any finite T , we have

$$(19) \quad \lim_{t \rightarrow \infty} \frac{1}{t} \sum_{\tau=T}^t H(m^\tau(y)) = y$$

Thus, nothing that happens in finite time affects the limiting average payoff. Now, let f be any n -tuple of supergame strategies, and m^1, \dots, m^t an arbitrary history of length t . We define the set of last defectors from f according to $m = (m^1, \dots, m^t)$; $LD(f, m)$, and the time of last defection $t^*(f, m)$ as follows:

$$t^*(f, m) = \begin{cases} \max \{t' \leq t: m^{t'} \neq f^{t'}(m^1, \dots, m^{t'-1})\} & \text{if it exists} \\ t + 1 & \text{otherwise} \end{cases}$$

$$LD(f, m) = \begin{cases} \{i \in N: m_i^{t^*}(f, m) \neq f_i^{t^*}(m^1, \dots, m^{t^*-1})\} & \text{if } t^* = t^*(f, m) \leq t \\ \emptyset & \text{otherwise} \end{cases}$$

Now let e_t be an infinite sequence of positive numbers $e_t = o(t)$. We shall define the notion of a "debt to society," by stipulating that a player who defects at time t is to be punished until his cumulative average payoff is within e_t of his minmax payoff, at which point play returns to the specified sequence, or until another player or players defects. If another single player j defects at time t' subsequent to t , then j is to be punished to within e_t of his minmax payoff; if more than one player defects simultaneously, play returns to the cooperative sequence. To implement this idea, we must remove from $LD(f, m)$ those players who have paid their debt to society. The remaining criminals as of time t form a set $C^t(f, m)$ defined by:

$$C^t(f, m) = \{i \in LD(f, m): \frac{1}{t} \sum_{\tau \leq t} H_i(m^\tau) > v_i + e_{t^*(f, m)}\}$$

We can now define a perfect equilibrium strategy with outcome y :

$$f_i^1 = m_i^1(y) \quad \text{for all } i$$

$$f_i^t(m^1, \dots, m^{t-1}) = f_i^t(m) = \begin{cases} p_i^j & \text{iff } C^t(f, m) = \{j\} \\ m_i^t & \text{otherwise} \end{cases}$$

The following are consequences of this definition: $m(f) = m(y)$, so that the specified sequence does indeed have the outcome y . If player i contemplates defecting for at most a finite number of periods, his limiting average payoff will be y_i by equation (19) and the fact that $C^t(f, m)$ is always empty a finite number of periods after any last defection, if the players adhere to f : in other words, by adding at each stage the amount v_i to player i 's cumulative payoff, his cumulative average payoff reaches the trigger level $e_{t^*(f, m)} + v_i$ within finite time after $t^*(f, m)$. If player i defects an infinite number of times, then the strategy calls for him to be punished forever, since the trigger level approaches v_i . Now, consider any subgame in which players are punishing player i . If player j decides to defect by not playing p_j^i , then player j is punished. If j defects forever, he is held to v_j ; if he defects a finite number of times, play returns to $m(y)$ and his payoff is y_j ; if he does not defect, the punishment of player i ends in finite time, so his payoff is y_j .

Therefore, any feasible and individually-rational y can be sustained as a perfect equilibrium outcome by such a strategy. Since every perfect equilibrium outcome is a fortiori an equilibrium outcome, it follows that the set of perfect equilibrium outcomes co-incides with the set of equilibrium outcomes.

QED

This happy state of affairs does not persist in the discounted game, as we can show through analysis of a simple example related to the problem of strategic control of externalities.

5.2 Example: There are two players. In the stage game, player 1 can take a level of precaution $s_1 \in [0,1]$. It costs him nothing to take this precaution, but the result is a social cost of $C(1 - s_1)$ where C is a large positive number. Player 2 cannot take any precaution, but can compensate player 1 by paying him an amount s_2 from her initial wealth of 1. Player 1 is liable for a constant share, $L \in (0,1)$, of the social cost, and player 2 pays the balance. The payoffs to the two players are therefore:

$$h_1(x_1, s_2) = s_2 - L(1 - s_1)C$$

$$h_2(s_1, s_2) = 1 - s_2 - (1 - L)(1 - s_1)C$$

In the one shot game there is a unique equilibrium at $s_1 = 1 = 1 - s_2$.

In the undiscounted supergame, we can define the set of strong equilibrium outcomes to be the set of Pareto Optimal equilibrium outcomes; in general, a strong equilibrium is a situation from which no coalition can defect, making all of its members better off. Here, the only non-singleton coalition is the pair $\{1,2\}$. A strong equilibrium outcome of the undiscounted supergame is a pair of net wealths (a_1, a_2) s.t.

- i) $a_1 + a_2 = 1$ (Pareto Optimality)
- ii) $a_1 \geq 0$ (individual rationality for player 1)
- iii) $a_2 \geq 1 - (1 - L)C$ (individual rationality for player 2)

The set of equilibrium outcomes can be found by replacing condition i) by

$$i') \quad a_1 + a_2 \leq 1 \quad (\text{feasibility})$$

It is clear that player 1 can use the threat of diminishing his precaution to extract some money from player 2. Moreover, by Theorem 5.1 we know that this threat is credible in the undiscounted game, so conditions i-iii also give us the set of strong perfect equilibrium outcomes for the undiscounted game. Now let us move to the discounted game, with both players using the same discount rate, $d \in (0,1)$. From Theorem 4.3, we know that (a_1, a_2) is a Pareto Optimal outcome of an equilibrium of the discounted game iff it satisfies i), ii), and

$$iv) \quad 1 - a_2 = a_1 \geq d(1 - L)C$$

Now suppose that player 1 wishes to punish player 2 for some defection by playing the punishment sequence $(s_1^1, \dots, s_1^t, \dots)$. The ratio of the cost to player 1 of this sequence divided by the punishment inflicted on player 2 is:

$$\frac{\sum_{t=1}^{\infty} d^{t-1} [(1 - s_1^t)LC]}{\sum_{t=1}^{\infty} d^{t-1} [(1 - s_1^t)(1 - L)C]} = \frac{L}{1 - L}$$

So that player 1 may as well react immediately to defection with a punishment sufficient to have deterred defection in the first place. How strong must this punishment be?

If the outcome of the perfect equilibrium is to be the pair $(a, 1-a)$, it is easy to see that the cheapest punishment sequence sufficient to prevent defection is $(s_1, 1, 1, \dots)$ where

$$(20) \quad (1 - s_1)(1 - L)C \geq \frac{a}{d}$$


We must now see whether player 1 will be willing to execute this punishment. By the argument given above, player 1 has two alternatives, use the specified punishment at once, or hold off forever, for a payoff of 0 in each period. The condition for carrying out the punishment is therefore

$$(21) \quad \frac{da}{(1-d)LC} \geq 1 - s_1$$

Combining this with (20) gives us the condition for the equilibrium outcome $(a, 1-a)$ to be sustainable as the outcome of a stationary Pareto Optimal perfect equilibrium of the discounted supergame:

$$(22) \quad \frac{da}{(1-d)LC} \geq \frac{a}{(1-L)dC} \text{ or } \frac{d^2}{1-d} \geq \frac{L}{1-L} \text{ for } a \neq 0, L \in (0,1), \\ d \in (0,1)$$

Combining this condition with i, ii, and iv gives us the set of (strong) perfect equilibrium outcomes of the discounted supergame. The set of equilibrium outcomes is as shown below:



The set of perfect equilibrium outcomes is equal to the set of equilibrium outcomes if condition (22) is satisfied, and is equal to the single outcome (0,1), otherwise. We should remark that an argument similar to that used in analysing the set of equilibrium outcomes for the prisoner's dilemma lets us confine our attention to stationary outcomes of equilibria in this game.

Since there is no hope for a general result such as Theorem 5.1 for the discounted case, we close by presenting a sufficient condition for the set of perfect equilibrium outcomes to co-incide with the set of equilibrium outcomes. This condition turns out to be satisfied by quite a few games of economic interest.

5.3 Theorem: Let $[N, M, H]$ be the mixed extension of a stage game. Suppose that for each player i , there exists an n -tuple $\bar{m}^i \in M$ of mixed strategies for the stage game with the following properties:

- i) $\max_{m_i} H_i(m_i, \bar{m}_{(-i)}^i) = H_i(\bar{m}^i) = v_i$
- ii) for all $j \neq i$, $H_j(\bar{m}^i) \geq \delta_j v_j + (1 - \delta_j) \max_{m_j} H_j(m_j, \bar{m}_{(-j)}^i)$

Then the set of outcomes sustainable by perfect equilibria of the discounted supergame co-incides with the set of outcomes sustainable by equilibria of the discounted supergame.

Proof: Let y be an outcome sustainable by a grim-strategy equilibrium of the discounted supergame, and $m(y)$ the associated specified sequence. Recalling the definition of the last defector from a strategy f given

a partial history m used in the proof of Theorem 5.1, we define a perfect equilibrium strategy f by

$$f_i^1 = m_i^1(y)$$

$$f_i^t(m^1, \dots, m^{t-1}) = f_i^t(m) = \begin{cases} m_i^j & \text{iff } LD(f, m) = \{j\} \\ m_i^t(y) & \text{otherwise} \end{cases}$$

To see that this is indeed a perfect equilibrium strategy combination, we first observe that $m(f) = m(y)$, so that adherence to this strategy results in a payoff of y . Secondly, notice that this strategy calls for a grim punishment to be inflicted on any defector, regardless of whether that defection occurred while playing the specified sequence or a punishment sequence. It follows that no player will wish to unilaterally defect from the specified sequence. Now suppose that we are playing a punishment sequence. The condition of the theorem states that for each player i there is a way of implementing the grim punishment via a stationary equilibrium of the discounted supergame: by condition i) the punished player cannot improve his payoff in any stage; and by condition ii) no other player will find it in his interest to defect from the new stationary equilibrium. QED

As a special case, we remark that the condition is clearly satisfied if there is an equilibrium of the stage game that gives each player his minmax payoff. This is clearly the case with Prisoner's Dilemma, and also with many economic exchange games. In the latter, a player's

security level v_i is almost always the same as his payoff at the no-trade point, so if there is a no-trade equilibrium, the Theorem applies. Examples include: Kurz' "Altruism games"; the Shapley-Shubik-Dubey family of exchange games; and Wilson's Competitive bidding model.

REFERENCES

- Aumann, R. J. (1976); Lecture Notes on Game Theory, mimeo, Stanford University.
- Aumann, R. J. (1964); "Mixed and Behavior Strategies in Infinite Extensive Games," in Advances in Game Theory, Dresher, Shapley and Tucker, eds., Princeton University Press, Princeton, N.J, pp. 627-650.
- Aumann, R. J. and Shapley, L. S.; "Long-term Competition: a game-theoretic analysis," unpublished MS.
- Cave, J. A. K. (1980); "Equilibrium Behavior in Infinitely-Repeated Games," unpublished Ph.D. Dissertation, Stanford University.
- Kuhn, H. (1953); "Extensive Games and the Problem of Information," in Contributions to the Theory of Games, Vol. II, Kuhn, Tucker, eds., Princeton University Press.
- Kurz, M. (1978); "Altruism as an Outcome of Social Interaction," Amer. Econ. Rev. 68, pp. 216-222.
- Roth, A. E.; "A Self-supporting Equilibrium in the Supergame," unpublished MS.
- Rubenstein, A. (1977); "Equilibrium in Supergames," Res. Mem. 25, May, Institute for Advanced Studies, Hebrew University, Jerusalem.
- Shapley, L., Shubik, M. and Dubey, P. (1972-1980); A Theory of Credit and Financial Institutions, Cowles Foundation, Yale University.
- Wilson, R. (1978); "A Competitive Model of Exchange," Tech. Rept. 221, I.M.S.S.S., Stanford University.

HECKMAN
INDRY INC.



JUN 95

MANCHESTER

UNIVERSITY OF ILLINOIS-URBANA



3 0112 060296198