

Anonymous?

Working Paper

Breaking the Barriers of the Feasible Set: On Repeated Games with Different Time Preference

Discussion Paper, No. 1072

Provided in Cooperation with:

Kellogg School of Management - Center for Mathematical Studies in Economics and Management Science, Northwestern University

Suggested Citation: Anonymous? (1993) : Breaking the Barriers of the Feasible Set: On Repeated Games with Different Time Preference, Discussion Paper, No. 1072, Northwestern University, Kellogg School of Management, Center for Mathematical Studies in Economics and Management Science, Evanston, IL

This Version is available at:

<https://hdl.handle.net/10419/221429>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

Discussion Paper No. 1072

**BREAKING THE BARRIERS OF THE FEASIBLE SET:
ON REPEATED GAMES WITH
DIFFERENT TIME PREFERENCES**

by

Ady Pauzner

and

Ehud Lehrer

Math Center web site:
<http://www.kellogg.nwu.edu/research/math>

2. Breaking the Barriers of the Feasible Set: on Repeated Games with Different Time Preferences

2.1 Introduction

Repeated games in which all players have identical time preferences have been extensively studied. For such games, folk theorems (See, for instance, Aumann and Shapley (1976), Rubinstein (1976), Aumann (1981), Fudenberg and Maskin (1986), Abreu, Dutta and Smith (1994)) assert that every feasible and individually rational vector of payoffs in the underlying one-shot game can be supported as an equilibrium outcome when players are sufficiently patient.

In the case where players have identical discount factors, the set of feasible payoffs of the repeated game coincides with that of the stage-game. Moreover, the set of equilibrium outcomes of the repeated game approaches the set of its feasible and individually rational payoffs as the discount factors tend to 1. When players have different discount factors, however, both statements are false. First, the set of feasible payoffs of the repeated game is typically larger than that of the stage game. Second, even when players become very patient, not all (repeated game's) feasible payoffs are supported by equilibria.

It is well known, that replicating the same game many times gives rise to two main issues. The first is the possibility to enforce certain types of behaviors on the basis of punishment threat; players are threatened by a punishment plan and as a result are refrained from deviation. This is the reason that the set of equilibrium payoffs is so large that it may reach the set of all feasible and individually rational payoffs. The second issue is the information acquisition (e.g., about the strategies played by the opponents, or about the opponents' types) during the course of the repeated game. Both these aspects play a role also in this paper.

When players have different time preferences, however, iteration of the same game gives rise to a third issue. In this case players may agree on playing different joint actions in different periods: actions that entail high payoffs to the impatient players first and actions that are patient players' favorites later. In doing so, stage payoffs are distributed over time in a way that drives players overall utility out of the one-shot feasible set. Therefore, the repeated game's feasible set is typically larger than that of the stage-game. For the same reason it is also natural to expect that in equilibrium, players can sustain payoffs that are located outside the one-shot game's feasible payoffs set. It turns out, however, that this is not always the case. That is, in some cases, most notably zero-sum games, there is no equilibrium payoff out of the one-shot game's feasible set.

The next section provides two illustrative examples. In both, the Pareto frontier of the feasible set is as in a zero-sum game. The two games differ only in their individually rational levels. In the first example there is no equilibrium outside the one-shot feasible set, while in the second game, Pareto-superior outcomes are sustainable in equilibrium.

Section 3 contains the formal model, and Section 4 is devoted to the characterization of the equilibrium payoffs set in a 2-player case. This is a folk theorem of games with two players having different time preferences. The intuition behind the characterization is simple. Consider, for example, a Pareto optimal equilibrium payoff of the repeated game. It is typically supported by a sequence of stage payoffs that are not all identical. Along the sequence, the patient player payoffs are increasing, while the impatient one's are decreasing. In equilibrium, all continuation payoffs, corresponding to all tails of the sequence, ought to be individually rational for both players. Since the patient player payoffs are increasing along the optimal path, if her initial payoff is individually rational, so are all her continuation payoffs.

As for the impatient player's payoffs, since his stage-payoffs are decreasing, so are the continuation payoffs. Consequently, only stage-payoffs that are impatient-player individually rational can participate in supporting a Pareto optimal equilibrium. In short, in a Pareto optimal equilibrium, all the patient player's constraints are satisfied as long as the initial present value is above her individually rational level. On the other hand, in order to prevent

the impatient player from deviating, only individually rational stage-payoffs are allowed. This observation gives rise to a simple geometric characterization of the 2-player folk theorem.

Different time preferences have appeared in a number of applications. Rubinstein (1982) discusses the alternating offers game model of bargaining between two players having different discount factors. He concludes that the more patient player will end up with a larger share of the pie. The subject of reputation building by a relatively more patient player was studied by Fudenberg and Levine (1989), Aoyagi (1993), Celentani et al. (1995), and others. Under various assumptions, they show that the more patient player reaps all the surplus of cooperation. In all these models, there exists a unique equilibrium in which a point on the Pareto frontier of the stage game is constantly played. By contrast, we focus on the inter-temporal trade in payoffs, made possible by differential time preferences, and characterize the set of all equilibrium payoffs.

In Section 5, we analyze the equilibrium set for n -player games. We characterize the class of games in which there exist equilibrium payoffs outside the one-shot game's feasible set. We were unable to describe the equilibrium set in the general case and we leave the question open.

Section 6 concludes with a few remarks. In particular, we note that the results can be extended to subgame perfect equilibrium, under the full dimensionality condition introduced by Fudenberg and Maskin (1986) for the

case of players with identical discount factors. We also discuss extensions to games with incomplete information.

2.2 Illustrative Examples

Consider the following 2-player 0-sum stage game:

1,-1	-1,1
-1,1	1,-1

At each stage k , both players choose a mixed strategy and receive stage payoff $(X(k), Y(k))$, where $X(k) \in [-1, 1]$ and $Y(k) = -X(k)$. By mixing $\frac{1}{2}$ on each action, each player can guarantee his or her individually rational (henceforth, IR) level of 0.

Assume that the players evaluate their infinite stream of stage payoffs using discounting factors $1 > \delta_P \geq \delta_I > 0$ (P and I stand for patient and impatient, respectively.) In other words, the repeated game payoffs of the impatient and patient players are, respectively,

$$U_I = (1 - \delta_I) \sum \delta_I^k X(k) \text{ and } U_P = (1 - \delta_P) \sum \delta_P^k Y(k).$$

In the case where $\delta_I = \delta_P$, $U_I + U_P = 0$, i.e., the repeated game is also 0-sum. Since the IR levels are 0, the only equilibrium outcome is $(U_I, U_P) = (0, 0)$. On the other hand, if $\delta_I < \delta_P$ there exist feasible payoffs (in the repeated game) that are Pareto-superior to $(0, 0)$. For example, the players may agree on receiving the payoff $(X(k), Y(k)) = (1, -1)$ up to a certain period, say, K , and $(X(k), Y(k)) = (-1, 1)$ thereafter. In other words, the patient player "lends"

payoff to the impatient one until period K , and gets refunded after period K . Both players are better off than by receiving a payoff of 0.

Unfortunately, this plan is not an equilibrium of the game. This is so because, at period K , the impatient player would deviate and fail to repay the debt. It turns out, as we show in the sequel, that the only equilibrium outcome of this repeated game is $(0,0)$. We therefore obtain the traditional folk theorem in this case: the repeated game's equilibrium set coincides with the set of the one-shot game's feasible and individually rational payoffs.

Consider now the following modification of the stage-game:

1,-1	-1,1	-2,-2
-1,1	1,-1	-2,-2
-2,-2	-2,-2	-2,-2

The additional "threat" actions reduce the IR levels from 0 to -2. Now, if δ_i is close enough to 1, the borrowing plan is sustainable at equilibrium; if the impatient player defects on his loan, he gets punished down to his IR level. That is, the patient player's threat enables the impatient one to make a credible promise to pay back a loan¹. The players can now carry out a self-enforcing plan that improve their utilities; both get more than 0.

¹ The reader should not be bothered, at this stage of the exposition, with the lack of subgame perfection. This can be rectified as we note in the last section.

While in the first example the repeated game's equilibrium outcomes set coincides with the set of all feasible and individually rational payoffs of the stage-game, in the second example an outside-of-the-feasible set payoff can be supported by an equilibrium. It is quite clear that what distinguishes, in this respect, between these two examples is the different IR levels.

It turns out that for a fixed Pareto frontier of the stage-game, the lower the IR level, the greater the equilibrium set. This is not surprising, though. However, as opposed to the case of the traditional folk theorem where the equilibrium set (of the repeated game) expands only in the direction of the reduced IR levels, here the expansion has a different feature. When the IR levels are reduced the Pareto frontier of the equilibrium set is pushed out. That is, there may be new equilibrium payoffs that strongly Pareto dominate former (with the high IR levels) equilibrium points. The reason is that more vulnerable players can trust each other more, and therefore can achieve a higher degree of cooperation. This phenomenon can be perceived as another instance of the advantage of vulnerability.

2.3 The Repeated Game and its Feasible Set

2.3.1 The Model

We consider an n -person stage game consisting of finite action sets $A_j, j = 1 \dots n$, and utility functions, $u_j : A \rightarrow \mathfrak{R}, j = 1 \dots n$. Without loss of generality, we assume that all payoffs are bounded between -1 and 1. Let V be the feasible set of the one-shot game. That is, V is the convex hull of all possible joint payoffs (i.e., the range of $u = (u_1 \dots u_n)$). Denote by ir_j player j 's individually rational level.

The stage game is repeated infinitely many times to form the repeated game². Player j evaluates an infinite stream of stage-payoffs according to a discount factor, δ_j .

Some of our results pertain to discount factors close to 1 (i.e., where players are patient). An equivalent way to say that the discount factor tends to 1 is to say that the time lap between two consecutive encounters diminishes to zero. Thus, let player j 's discount factor $0 < \delta_j < 1$ be fixed throughout, representing j 's evaluation of payoff delayed by one time unit.

Suppose that the interval between two consecutive repetitions of the stage game is Δ time units. Therefore, a unit of payoff received at the k -th

² We eliminate the full description assuming the reader is familiar with repeated games, where stage-actions are fully monitored.

repetition is worth $\delta_j^{k\Delta}$ (or $e^{k\Delta \log \delta_j}$) units of payoff at the outset. Properly normalized, the present value of the stream of player j 's payoffs, $X_j(0), X_j(1), \dots$, is

$$U_j^\Delta = (1 - \delta_j^\Delta) \sum_{k=0}^{\infty} \delta_j^{k\Delta} X_j(k)$$

The continuous counterpart, which will be used for characterization purposes is:

$$U_j^0 = (-\log \delta_j) \int_0^{\infty} \delta_j^t X_j(t) dt$$

Denote by F^Δ and F^0 the convex hulls of the ranges of $U^\Delta = (U_1^\Delta, \dots, U_n^\Delta)$ and $U^0 = (U_1^0, \dots, U_n^0)$, respectively. For the sake of simplicity we assume, like Fudenberg and Maskin (1986), the existence of a public randomizing device (e.g., a sunspot). Thus, players can agree on randomizing between various payoffs. With a public randomizing device at players' disposal, F^Δ is simply the feasible set of the repeated game corresponding to Δ . Obviously, F^Δ and F^0 are closed sets.

Note that, as Δ goes to zero, the public randomizing device becomes dispensable, because convexification can be attained by alternating between pure actions with the appropriate frequency.

2.3.2 The Feasible Set, F^Δ .

As mentioned above, F^Δ is a convex and closed set. Therefore it can be described by the points on its boundary maximizing linear functionals. Let $\alpha \in \Re^n$ be a direction representing a linear functional. For the sake of conciseness, denote $h_\alpha(k) = ((1 - \delta_1^\Delta) \delta_1^{k\Delta} \alpha_1, \dots, (1 - \delta_n^\Delta) \delta_n^{k\Delta} \alpha_n)$, for any integer k . Thus, $h_\alpha(k)$ is a vector in $\Re^n$. To find a point in the frontier corresponding to the direction α we solve the maximization problem:

$$(1) \quad \text{Max}_{\{X(k)\}_{k=0}^\infty} \sum_{k=0}^{\infty} h_\alpha(k) \cdot X(k) \quad \text{s.t. } \forall k, X(k) = (X_1(k) \dots X_n(k)) \in V$$

(A central dot, $\cdot$, denotes the product of two vectors.) This problem amounts to solving separately, for each k , the linear program:

$$(2) \quad \begin{aligned} & \text{Max}_{X(k)} h_\alpha(k) \cdot X(k) = \\ & \text{Max}_{X(k)} ((1 - \delta_1^\Delta) \delta_1^{k\Delta} \alpha_1, \dots, (1 - \delta_n^\Delta) \delta_n^{k\Delta} \alpha_n) \cdot (X_1(k) \dots X_n(k)) \quad \text{s.t. } X(k) \in V \end{aligned}$$

For every k we have the same feasible polygon, while the direction of ascent of the objective function changes gradually with k . The maximum of the objective function over the polygon is attained, for every k , at some vertex. So, for $k = 0$, we start at an initial vertex. Then, for each successive k , the direction of ascent of the objective function, is multiplied coordinate-wise by $\delta_1^\Delta, \dots, \delta_n^\Delta$ and the maximum is attained either still at the same vertex, or at another vertex. As k passes through all periods a path of vertices is followed. From a certain point on, when k is large enough, the maximum is attained at

the vertex that gives the highest payoff to the most patient players (since for $\delta_i > \delta_j$, the ratio $(1 - \delta_i^\Delta) \delta_i^{k\Delta} \alpha_i / (1 - \delta_j^\Delta) \delta_j^{k\Delta} \alpha_j$ goes to infinity).

For some hyperplanes α , there may exist some k at which the maximum is attained on a whole facet of the polygon (which can be of any dimension between 1 and $n - 1$). In this case, any point on the facet is optimal, and by choosing different points, utility can be transferred between the players at a fixed ratio of α_i / α_j . This implies that the Pareto frontier of F^Δ is tangent to the hyperplane corresponding to α over a whole facet of the same dimension.

For $\alpha = (\alpha_1, \dots, \alpha_n)$ define $g_\alpha(t) = ((-\log \delta_1) \delta_1^t \alpha_1, \dots, (-\log \delta_n) \delta_n^t \alpha_n)$, a vector in $\mathfrak{R}^n$. The set F^0 can also be described by those points that maximize linear functionals. That is, a point f is in the frontier of F^0 if and only if there is a direction $\alpha \in \mathfrak{R}^n$ s.t. f is the payoff associated with a solution to the maximization problem

$$(3) \quad \text{Max}_{X(t)} \int_0^\infty g_\alpha(t) \cdot X(t) \, dt \quad \text{s.t. } \forall t \in [0, \infty), X(t) \in V \text{ and } X(\cdot) \text{ is integrable.}$$

This global maximization problem (finding $X(t)$ for every t) can be split into a continuum of problems; for any t

$$(4) \quad \text{Max}_{X(t)} g_\alpha(t) \cdot X(t) \quad \text{s.t. } X(t) \in V$$

The solution is to be found in extreme points of V . Moreover, one can divide the time space $\mathfrak{R}_+$ to intervals $[0, t_1), [t_1, t_2), \dots$ such that at any time t the

solution to (3) is constant over each interval (up to a set of measure 0). In fact, there is a need for only finitely many intervals: $[0, t_1), [t_1, t_2), \dots, [t_l, \infty)$, because for t large enough there is only one extreme point that maximizes all the maximization problems (4) according to $g_a(t)$. This is the vertex that gives the highest payoff to the patient player. Since the solutions to (4), one for each t , form an integrable function $X(t)$, the condition in (3) requiring that $X(t)$ be integrable is in fact redundant.

Consider a 2-player game, and $\alpha = (\alpha_i, \alpha_p) \gg 0$. We can now give more structure to an optimal path. Notice that the Pareto optimal vertices of V can be simply ordered: a vertex that gives more to the impatient player necessarily gives less to the patient one. So we order them $v_1, \dots, v_n$ by the magnitude of the patient player's payoff or, equivalently, by ratio between both players' payoffs. An optimal path starts at some vertex v_s , and uses vertices with increasing indices until, at the tail, it reaches v_n . For F^Δ small enough (or in the continuous counterpart, F^0), all vertices between v_s and v_n are used consecutively (we never skip a vertex).

The following picture illustrates graphically the shape of the feasible sets. The gray area is the stage game's feasible set, the broken line is the Pareto frontier of F^Δ , while the arc is the Pareto frontier of F^0 .




Figure 2.1: The Pareto Frontier of the Feasible Set

The following theorem asserts that the feasible set of the n -player continuous limit "game" is a good approximation when F^Δ is small.

Theorem 2.3.1 F^Δ converges uniformly, from inside, to F^0 as $\Delta \rightarrow 0$.

Proof: Since F^Δ is convex, it is sufficient to show that:

- a. For any $\Delta > 0$, $F^\Delta \subset F^0$; and
- b. For any $\varepsilon > 0$ there exists $\underline{\Delta} > 0$, s.t. for any $f \in F^0$ and any $\Delta < \underline{\Delta}$, there exists $f_\Delta \in F^\Delta$ s.t. $|f - f_\Delta| < (\varepsilon, \dots, \varepsilon)$.

Part a: Since $\left(-\log \delta_j\right) \int_{k\Delta}^{(k+1)\Delta} \delta_j^t dt = (1 - \delta_j^\Delta) \delta_j^{k\Delta}$, every payoff in F^Δ can be achieved in (3) by setting $X(t)$ to be constant over intervals of the form $[k\Delta, (k+1)\Delta)$.

Let $\bar{X}(1), \bar{X}(2), \dots$ be a solution to (1), and thereby support a frontier point of F^Δ .

The continuous path: $X_j(t) = \bar{X}_j(k)$ whenever $k\Delta \leq t < (k+1)\Delta$, sustains the same point when evaluated by U^0 . (Notice that F^Δ is not necessarily a decreasing set in Δ . For fixed Δ_1 and Δ_2 , unless Δ_1 is an integer fraction of Δ_2 , F^{Δ_2} is not necessarily a subset of F^{Δ_1} .)

Part b: Since F^Δ is convex it is sufficient to show that the frontier F^0 can be uniformly approximated by points in F^Δ . As explained above, one can divide the time space into intervals $[0, t_1), [t_1, t_2), \dots, [t_l, \infty)$ s.t. a solution to (3), say, $X(\cdot)$, is constant over each interval.

For a given Δ define the discrete path $\bar{X}(k) = X(k\Delta)$. As in part (a), we extend $\bar{X}(\cdot)$ into a continuous path $\hat{X}(\cdot)$, by setting it constant over intervals of the form $[k\Delta, (k+1)\Delta)$. As before, $\bar{X}(\cdot)$, evaluated w.r.t. to U^0 , and $\hat{X}(\cdot)$, evaluated w.r.t. U^Δ , yield the same payoff.

The paths $\hat{X}(\cdot)$ and $X(\cdot)$ differ from each other on at most l intervals $[k\Delta, (k+1)\Delta)$, because there are l times where $X(\cdot)$ changes its value. Thus, the difference between the corresponding player i 's payoffs is bounded by $2l(1 - \delta_i^\Delta) \delta_i^\Delta$. (The factor 2 arises from the maximal one-stage payoffs differences which are bounded between -1 and 1; l arises from the number of

intervals where $\hat{X}(\cdot)$ and $X(\cdot)$ might be different: and $(1 - \delta_i^\Delta)\delta_i^\Delta$ from the weight of the first interval of length Δ , which is the heaviest.) Since l is fixed and Δ can be arbitrarily small, we conclude that any point in the frontier of F^0 can be approximated by points in F^Δ . The uniform approximation, as claimed by (b), is implied by compactness arguments. $\square$

2.4 Characterization of the Equilibrium Set in 2-Players Games

It can be quite easily seen that the equilibrium set is closed and convex. We therefore need only to find its frontier. Finding an extreme point of the set in the direction $\alpha \in \mathbb{R}^n$ differs somewhat from finding extreme points of the feasible set, since the path constructing an equilibrium point has to satisfy individual rationality for each player. That is, at each stage of the game, each player's evaluation of future payoffs should exceed her IR level.

The IR constraints break the symmetry between different paths that support extreme equilibrium outcomes corresponding to α 's that belong to different quadrants. We shall first characterize the Pareto frontier of the equilibrium set, i.e., the extreme points corresponding to α 's in the first (positive) quadrant. This is the most interesting part of the frontier from an economic point of view. We shall then explain the difference between this construction and the ones pertaining to the construction of the three other parts of the equilibrium payoffs frontier. A full characterization will follow immediately.

Given a bounded set A , discount factors $1 > \delta_P \geq \delta_I > 0$ and the time period Δ between two consecutive repetition of the stage-game, we introduce the following notation:

$F^\Delta(A)$: the set of feasible payoffs in the repeated game when A is the set of available payoffs. (Sometimes we will restrict V and, therefore, we introduce a notation applied to a general set A).

E^Δ : the set of equilibrium payoffs set of the repeated game.

IR_I, IR_P : the sets of the individually rational payoffs of the impatient and the patient players, respectively³:

$$IR_I = \{(x, y) \in V : x \geq ir_I\}, \quad IR_P = \{(x, y) \in V : y \geq ir_P\}$$

$IR_I^\epsilon, IR_P^\epsilon$: the sets of payoffs that satisfy "strong" individual rationality:

$$IR_I^\epsilon = \{(x, y) \in V : x \geq ir_I + \epsilon\}, \quad IR_P^\epsilon = \{(x, y) \in V : y \geq ir_P + \epsilon\}. \text{ Also denote}$$

$$IR = IR_I \cap IR_P \text{ and } IR^\epsilon = IR_I^\epsilon \cap IR_P^\epsilon.$$

Consider the path $\{X(k)\}_{k=0}^\infty$ that generates a point f on the Pareto frontier of $F^\Delta(A)$. The path consists of, at most, $l+1$ time intervals (the last one corresponds to the tail of the game), over which actions are constant.

A *tail of a path*, or a continuation path, is a sequence of the sort $\{X(k)\}_{k=K}^\infty$. In equilibrium, at any period K , players' rationality requires that the present value of the corresponding tail must be individually rational. As for the converse, if the value of any tail of the path is strongly individually rational for both players (i.e., at least ϵ above it), and if Δ is small enough,

³ We denote them in capitals to distinguish them from the corresponding IR levels of the previous section denoted in lower case.

then the path can be extended to an equilibrium, using the threat that any deviation will be punished to the IR level.

The value of the tail of any Pareto-optimal path, from period K on, is increasing with k for the patient player, and decreasing for the impatient one. Thus, if the individual rationality constraint for the patient player is satisfied for the whole path, it is satisfied for any tail. In contrast, individual rationality conditions for the impatient player require that all stage payoffs be in IR_I . These observations provide the intuition for the following theorem, which is the main result of the paper. The statement of the theorem requires following notation:

Definition 2.4.1: Let A and B be two sets in the Euclidean space. We say that $A \preceq B$ if for every $a \in A$ there exists $b \in B$ s.t. b weakly Pareto dominates a .

Theorem 2.4.1.

For any $\varepsilon > 0$ there exists $\underline{\Delta} > 0$, s.t. for any $\Delta < \underline{\Delta}$,

$$IR_P^\varepsilon \cap F^\Delta(V \cap IR_I^\varepsilon) \preceq E^\Delta \preceq IR_P \cap F^\Delta(V \cap IR_I)$$

Combined with Theorem 2.3.1, Theorem 2.4.1 can be given a simple geometric interpretation. To obtain the upper bound for the Pareto-frontier of the equilibrium set, we first intersect the stage game's feasible set V with the

impatient player's individually rational half-plane IR_I . We treat the resulting set, $V \cap IR_I$, as if it were a one-shot set of payoffs, and construct from it the Pareto frontier of the continuous time feasible set, $F^0(V \cap IR_I)$. We then intersect this set with the patient player's individually rational half-plane, IR_P . In order to get a lower bound we employ the same procedure (with F^Δ instead of with F^0) using the smaller sets of strongly individually rational payoffs. When Δ gets close to 0, the lower bound approaches the upper one uniformly.

The following picture illustrates the construction of the Pareto frontier of the equilibrium set, explained above. The gray area is the one-shot feasible set, and the axes denote the IR levels. The arc is the limit frontier, and is uniformly approached by the polygons of the Δ -discrete games when Δ tends to 0. The stripped area is the limit folk theorem as the lap, Δ , goes to zero.




Figure 2.2: The Equilibrium Set

We can also construct the Pareto frontier of the equilibrium sets for the two examples presented in Section 2:




Figure 2.3: The two examples

Proof of theorem 2.4.2:

To show the left inequality we need the following lemma, which says that the players' payoffs are monotone.

Lemma 1. Let $\{(X(k), Y(k))\}_{k=0}^{\infty}$ be an optimal path in the direction α , i.e., it maximizes $\sum_{k=0}^{\infty} \alpha_I (1 - \delta_I^{\Delta}) \delta_I^{k\Delta} X(k) + \alpha_P (1 - \delta_P^{\Delta}) \delta_P^{k\Delta} Y(k)$ s.t. $\forall k, (X(k), Y(k)) \in A$.

Then, $\alpha_P Y(k)$ is (weakly) increasing,
and $\alpha_I X(k)$ is (weakly) decreasing.

Proof: Assume $k_2 > k_1$ and consider the following modification of the path:

$$(\hat{X}(k), \hat{Y}(k)) = (X(k), Y(k)) \text{ for } k \neq k_1, k_2$$

$$(\hat{X}(k_2), \hat{Y}(k_2)) = (X(k_1), Y(k_1))$$

$$(\hat{X}(k_1), \hat{Y}(k_1)) = (1 - \delta_I^{(k_2-k_1)\Delta}) (X(k_1), Y(k_1)) + \delta_I^{(k_2-k_1)\Delta} (X(k_2), Y(k_2))$$

One can verify that the impatient player's valuation of the path is unchanged, while the patient's one changes by $d = \delta_P^{k\Delta} (\delta_P^{(k_2-k_1)\Delta} - \delta_I^{(k_2-k_1)\Delta}) (Y(k_1) - Y(k_2))$. If the path is already a solution to the maximization problem, then the proposed modification must not increase the optimal value. In particular, $\alpha_P d$ has to be non-positive. Since $\delta_P^{k\Delta} (\delta_P^{(k_2-k_1)\Delta} - \delta_I^{(k_2-k_1)\Delta}) > 0$, we must have

$$\alpha_P (Y(k_1) - Y(k_2)) \leq 0.$$

Using a slightly different modification of the path (the same changes except for one) one can verify that along optimal paths we also have,

$$\alpha_I (X(k_1) - X(k_2)) \geq 0.$$

□ *lemma 1*

We proceed with the first half of the proof. Let f be a Pareto optimal point in $IR_P^e \cap F^\Delta (V \cap IR_I^e)$. We show now that f is an equilibrium payoff. By Lemma 1, the patient player's payoffs are increasing (the Pareto frontier corresponds to $(\alpha_I, \alpha_P) \gg 0$). Thus, there is a path generating f s.t. the values, for each player, of any tail of the game, are strongly individually rational. For Δ small enough, this path can be extended to an equilibrium by designing punishing plans; any deviation from the equilibrium path will imply a

punishment of the deviator that will push down his/her payoff to the IR level. Thus, $f \in E^\Delta$. This grants us the left inequality.

The second half of the proof is a little more subtle than the first. We have to take an equilibrium point $f \in E^\Delta$, and construct a point in the right hand side that Pareto dominates it. Again, we need to use the fact that an optimal path uses monotone payoffs. Lemma 1, though, does not apply now, since an equilibrium path has a (potentially) weaker restriction: that all tail payoffs be IR (instead of all stage payoffs). Lemma 2 helps us overcome this obstacle.

Denote the players' valuations of future payoffs at time k by:

$$U_I^\Delta(k) = (1 - \delta_I^\Delta) \sum_{l=0}^{\infty} \delta_I^{\Delta l} X(k+l) \quad \text{and} \quad U_P^\Delta(k) = (1 - \delta_P^\Delta) \sum_{l=0}^{\infty} \delta_P^{\Delta l} Y(k+l)$$

Lemma 2. Let $(\alpha_I, \alpha_P) \gg 0$, and let $\{(X(k), Y(k))\}_{k=0}^{\infty} \in V^\infty$ maximize the function $\alpha_I U_I^\Delta(0) + \alpha_P U_P^\Delta(0)$ subject to the constraint: $\forall k, U_I^\Delta(k) \geq ir_I$. Then $\forall k, X_k \geq ir_I$.

Proof: Otherwise, there exists $\hat{k}$ such that $X(\hat{k}) < ir_I$. Let k be the first period after $\hat{k} - 1$ such that $X(k+1) > X(k)$ (such k must exist since $U_I^\Delta(\hat{k}) \geq ir_I$, and $U_I^\Delta(\hat{k})$ is a weighted average of $X(\hat{k}), X(\hat{k}+1), \dots$). By the definition of k , $X(k) < ir_I$, since $X(\hat{k})$ is. However, since $U_I^\Delta(k) = (1 - \delta_I^\Delta)X(k) + \delta_I^\Delta U_I^\Delta(k+1)$, we must have $U_I^\Delta(k+1) > ir_I$. Choose $\varepsilon > 0$ s.t. $U_I^\Delta(k+1) > ir_I + 2\varepsilon$.

Consider now the following modification of the path:

$$\begin{aligned}
(\hat{X}(k+1), \hat{Y}(k+1)) &= (1-\varepsilon)(X(k+1), Y(k+1)) + \varepsilon(X(k), Y(k)) \\
(\hat{X}(k), \hat{Y}(k)) &= (1-\varepsilon)(X(k), Y(k)) + \varepsilon[(1-\delta_i^\Delta)(X(k), Y(k)) + \delta_i^\Delta(X(k+1), Y(k+1))]
\end{aligned}$$

One can verify that the sequence $\{(\hat{X}(k), \hat{Y}(k))\}_{k=0}^\infty$ still satisfies the constraint.

That is, $\forall k, U_i^\Delta(k) \geq ir_i$. Moreover, $U_i^\Delta(0)$ is unchanged, while $U_p^\Delta(0)$ is increased by $d = \varepsilon\delta_p^\Delta(\delta_p^\Delta - \delta_i^\Delta)(Y(k) - Y(k+1))$. Since $\{(X(k), Y(k))\}_{k=0}^\infty$ is already optimal, and $X(k+1) > X(k)$, we must have $Y(k+1) < Y(k)$. Thus, $d > 0$, in contradiction with the assumption that the original path is optimal. $\square$ lemma 2

Corollary to Lemma 2.

$$\begin{aligned}
& \text{Max}_{\{(X(k), Y(k))\}_{k=0}^\infty \in \Xi} \alpha_i U_i^\Delta(0) + \alpha_p U_p^\Delta(0) \text{ s.t. } \forall k, U_i^\Delta(k) \geq ir_i = \\
& = \text{Max}_{\{(X(k), Y(k))\}_{k=0}^\infty \in \Xi} \alpha_i U_i^\Delta(0) + \alpha_p U_p^\Delta(0) \text{ s.t. } \forall k, X(k) \geq ir_i
\end{aligned}$$

We now complete the proof of the theorem and show the right hand side inequality. Let $f \in E^\Delta$, and let $\{(\hat{X}(k), \hat{Y}(k))\}_{k=0}^\infty$ be the equilibrium path generating f^4 . The path must satisfy:

$$(I) \quad \forall k, U_i^\Delta(k) \geq ir_i$$

$$(II) \quad \forall k, U_p^\Delta(k) \geq ir_p$$

Thus, $(\alpha_i, \alpha_p) \cdot f \leq \text{Max}_{\{(X(k), Y(k))\}_{k=0}^\infty \in \Xi} \alpha_i U_i^\Delta(0) + \alpha_p U_p^\Delta(0) \text{ s.t. } (I) \quad \forall k, U_i^\Delta(k) \geq ir_i$

$$(II) \quad U_p^\Delta(0) \geq ir_p.$$

⁴ If f results from mixed or behavior strategies, then $\{(\hat{X}(k), \hat{Y}(k))\}_{k=0}^\infty$ denotes the expected stage-payoffs.

By the corollary to lemma 2, there exists $\tilde{f} \in IR_p \cap F^\Delta(V \cap IR_I)$ which achieves the maximum of this maximization problem. Since $(\alpha_I, \alpha_p) \gg 0$, $\tilde{f}$ Pareto dominates f . $\square$

Corollary to Theorem 2.4.1

The Pareto frontier of the repeated game's equilibrium set E^Δ is strictly dominated by the Pareto frontier of the repeated game's feasible and individually rational payoffs set (even for small Δ), if and only if the highest payoff to the patient player corresponds to a payoff of the impatient player that is below ir_I .

This corollary states, in particular, that if there is a point in the Pareto frontier of V which is outside IR_I , then E^Δ , the repeated game's equilibrium set, is strictly smaller than the repeated game's feasible set, F^Δ , intersected with IR . Such a case is impossible when both players are identical in their time preferences.

We now complete the characterization of the equilibrium set. Recall that an optimal path corresponding to $\alpha \in \mathbb{R}_+^2$ is characterized by a present value of the patient player's tail payoffs which increase with time (the farther the tail starts the higher the payoff), while the present value of the tail's

payoffs, for the impatient player, are decreasing with time. For this reason, if the IR constraint for the patient player was satisfied at the beginning of the game, it is satisfied along the path. We could use stage-payoffs that are not IR for the patient player, as long as the present value of the whole path is IR. In contrast, since along an optimal path, the impatient player's payoffs are decreasing, we could never use a payoff that is not IR for her. Otherwise the present value of the rest of the path would violate her IR constraint. Thus, one-shot payoffs that are not IR for the impatient player had to be eliminated before the construction of a repeated game feasible set.

When we treat extreme points in other quadrants, the nature of payoffs used in an optimal path differs. The following table summarizes their behavior:

quadrant	α_i	α_p	impatient player's payoff	patient player's payoff
1	+	+	decreasing	increasing
2	-	+	increasing	increasing
3	+	-	decreasing	decreasing
4	-	-	increasing	decreasing

Table 2.1: The Inter-Temporal Structure of Optimal Payoffs

The following illustration of the quadrants and the direction of the payoffs along optimal paths may also be helpful:




Figure 2.4: The Four Quadrants

It is easy to see that the construction of the frontier for the first quadrant generalizes in the following way. First, intersect V with the IR half planes of the player(s) whose payoffs are decreasing along optimal paths in the given quadrant. Next, construct the feasible frontier for the repeated game, and then, intersect the resulting set with the IR half-planes of the players whose payoffs are increasing.

Stating the full characterization theorem requires additional notation.

Definition 2.4.2: Let A and B be two sets in the two dimensional plane. For a quadrant Q , we say that $A \leq_Q B$ if B dominates A in the direction associated with quadrant Q .

(I.e., for any $a \in A$ there exists $b \in B$ satisfying $\alpha \cdot a \leq \alpha \cdot b$, where $\alpha = (\alpha_I, \alpha_P)$ is a two dimensional vector consisting of either 1 or -1 depending upon Q according to the previous table. (For instance, if $Q=3$ then $\alpha = (1, -1)$.)

Theorem 2.4.3

For any $\varepsilon > 0$ there exists $\underline{\Delta} > 0$, s.t. for any $\Delta < \underline{\Delta}$,

$$\begin{aligned} IR_p^\varepsilon \cap F^\Delta(V \cap IR_l^\varepsilon) &\leq_1 E^\Delta \leq_1 IR_p \cap F^\Delta(V \cap IR_l) \\ IR^\varepsilon \cap F^\Delta(V) &\leq_2 E^\Delta \leq_2 IR \cap F^\Delta(V) \\ F^\Delta(V \cap IR^\varepsilon) &\leq_3 E^\Delta \leq_3 F^\Delta(V \cap IR) \\ IR_l^\varepsilon \cap F^\Delta(V \cap IR_p^\varepsilon) &\leq_4 E^\Delta \leq_4 IR_l \cap F^\Delta(V \cap IR_p) \end{aligned}$$

The proof is only a slight modification of the proof of Theorem 2.4.2. Notice that lemma 1 is already stated for the general case. Lemma 2, however, can be easily generalized to show that exactly for those players whose payoffs are decreasing along the optimal path (according to the table), when all tail-payoffs are individually rational, so are all the stage-payoffs. The detailed proof is left to the reader.

Notice that since V must have at least one point (weakly) dominated by (ir_l, ir_p) , the statement of the theorem regarding the fourth quadrant reduces to: $IR^\varepsilon \cap V \leq_4 E^\Delta \leq_4 IR \cap V$, as in the traditional folk theorem. In other words, the south-western frontier of the equilibrium set is not affected by the fact that the players have different discounting factors.

2.5 Breaking the Barriers of the One-Shot Feasible Set

In this section we characterize the cases where there are equilibrium points of the repeated game outside V .

Denote $R^\Delta = \{(x_1, \dots, x_n) \in \mathfrak{R}^n : \forall j, x_j > ir_j + 2(1 - \delta_j^\Delta) / \delta_j^\Delta\}$, and let P^Δ be the Pareto frontier of $R^\Delta \cap V$.

Theorem 2.4

Assume that all the discount factors are different.

There is an equilibrium point in E^Δ that Pareto dominates a point in P^0 for sufficiently small Δ if and only if there are at least two points in P^0 .

The theorem states a necessary and sufficient condition for the equilibrium set of the repeated game to break out of the feasible set of the stage game. Intuitively, if there are two points in P^Δ , then one can do better than some point on the line segment connecting them, by playing first the point that is best for an impatient player, and then the point that is best for a patient player. On the other hand, in the case where there is only one point in P^0 , then for any Δ , the only Pareto optimal outcome is attained by always choosing that point. In the other case, where the intersection of R^0 and V is

empty, we show that it is impossible to give all players strictly individually rational payoffs in equilibrium. We now provide the formal proof.

Proof:

If: Suppose that there are two points in P^0 . Then, for small enough Δ , there are two points in P^Δ . Let $\alpha = (\alpha_1, \dots, \alpha_n) \in \mathfrak{R}^n$ represent a hyperplane and choose two distinct point $x, y \in P^\Delta$ such that $\alpha \cdot x = \alpha \cdot y \geq \alpha \cdot v \ \forall v \in V$ (this is possible because V is a polygon). In words, α is a hyperplane tangent to P^Δ , and touches the points x and y . Let $\{Z^*(k)\}_{k=0}^\infty$ be a solution of the problem:

$$(5) \quad \text{Max}_{\{Z(k)\}_{k=0}^\infty} \sum_{k=0}^{\infty} h_\alpha(k) \cdot Z(k) \text{ s.t. } \forall k, Z(k) \text{ is either } x \text{ or } y$$

Notice that if all $Z(k)$ are equal to x , then the value attained is $\alpha \cdot x$, and the same is true for y . For some k large enough, $h_\alpha(k) \cdot x \neq h_\alpha(k) \cdot y$ (because as k goes to infinity the coefficient pertaining to the patient players, the one with the large δ , becomes dominant since $\alpha_j(1 - \delta_j^\Delta)\delta_j^{k\Delta} / \alpha_i(1 - \delta_i^\Delta)\delta_i^{k\Delta} \xrightarrow{k \rightarrow \infty} \infty$ when $\delta_j > \delta_i$). Assume that $h_\alpha(k_0) \cdot x > h_\alpha(k_0) \cdot y$.

Consider now a (possibly suboptimal) sequence for (5), where $\tilde{Z}(k) = y$ for $k \neq k_0$ and $\tilde{Z}(k_0) = x$. The value attained for this particular sequence is greater than $\alpha \cdot y$ because $\alpha \cdot y$ is supported by a path $Z(k)$ which

always coincides with y , and $h_\alpha(k_0) \cdot x > h_\alpha(k_0) \cdot y$. Therefore, this value lies outside V . We will show that it is sustained by an equilibrium.

Notice that the strategy that prescribes $\bar{Z}(k)$ at the k -th repetition and in case of deviation a punishment of the deviator down to the individually rational level is an equilibrium. This is so because the payoffs are bounded by 1 and a one-shot gain due to a deviation can be at most $2(1 - \delta_j^\Delta)$, while a punishment would reduce future payoffs by at least $\delta_j^\Delta 2(1 - \delta_j^\Delta) / \delta_j^\Delta$ (the first δ_j^Δ is for the present value of a stream starting at the repetition right after the deviation took place, and the second δ_j^Δ , in the denominator, is due to the definition of R^Δ).

Only if: We divide the proof into two cases. In the first, there is exactly one point, say, p , in P^0 , while in the second P^0 is empty.

Case I In this case, the point p must be the only Pareto optimal point in V . Otherwise let y be a vertex of V , adjacent to p . The line connecting between p and y is a part of the Pareto optimal frontier. Hence, there are many points between p and y in P^0 . This contradicts the assumption regarding the uniqueness of the Pareto optimal point in P^0 .

Since there is only one Pareto optimal point in V , for every $\alpha \in \mathfrak{R}_+^n$, the function $h_\alpha(k) \cdot x$ is maximized at that point, for every k . Provided Δ is

sufficiently small, the only Pareto optimal point of the repeated game's feasible and individually rational set must be p . Therefore, this is true also for the equilibrium payoff set of the repeated game. In other words, p is the only equilibrium point.

Case II. Suppose now that P^0 is empty. Let $\tilde{V} = \text{Conv}(V \cup \{(ir_1 \dots ir_n)\})$. It is easy to see that the convex sets R^0 and $\tilde{V}$ can be separated by a hyperplane. Let $\alpha \in \mathbb{R}_+^n$ be the vector representing the hyperplane separating them. There is, therefore, a constant c s.t.

$$\alpha \cdot v \leq c < \alpha \cdot r \quad \forall v \in V, r \in R^0.$$

Since $(ir_1 \dots ir_n)$ is both in $\tilde{V}$ and in the boundary of R^0 , one obtains,

$$\alpha \cdot (ir_1 \dots ir_n) = c.$$

Let $X_i(k)$ be the stage k expected payoff of player i in an equilibrium. Denote

$$(6) \quad U_{j,i}^\Delta(k) = (1 - \delta_j^\Delta) \sum_{s=k}^{\infty} \delta_j^{(s-k)\Delta} X_i(s).$$

The term $U_{j,i}^\Delta(k)$ is the evaluation of player i 's payoffs with respect to the discount factor of player j . Notice that $U_{i,i}^\Delta(k) = U_i^\Delta(k)$.

When $\delta_j > \delta_i$ it is easy to check that

$$(7) \quad U_{j,i}^\Delta(0) = \frac{1 - \delta_j^\Delta}{1 - \delta_i^\Delta} U_i^\Delta(0) + \sum_{k=1}^{\infty} \delta_j^{(k-1)\Delta} (\delta_j^\Delta - \delta_i^\Delta) \frac{1 - \delta_j^\Delta}{1 - \delta_i^\Delta} U_i^\Delta(k).$$

Moreover, the sum of all the coefficients of $U_i^\Delta(k)$ is 1. That is,

$$(8) \quad \frac{1 - \delta_j^\Delta}{1 - \delta_i^\Delta} + \sum_{k=1}^{\infty} \delta_j^{(k-1)\Delta} (\delta_j^\Delta - \delta_i^\Delta) \frac{1 - \delta_j^\Delta}{1 - \delta_i^\Delta} = 1.$$

Now suppose, without loss of generality, that $\alpha_1 > 0$ and that for all i 's with $\alpha_i > 0$, $\delta_1 > \delta_i$. Since $X(k) \in V$, $\alpha \cdot X(k) \leq c$, we obtain,

$$\begin{aligned} \alpha_1 U_1^\Delta(0) &= (1 - \delta_1^\Delta) \sum_{k=0}^{\infty} \delta_1^{k\Delta} X_1(k) \alpha_1 \leq (1 - \delta_1^\Delta) \sum_{k=0}^{\infty} \delta_1^{k\Delta} (c - X_2(k) \alpha_2 - \dots - X_n(k) \alpha_n) \\ &= c - \alpha_1 U_{1,2}^\Delta(0) - \dots - \alpha_n U_{1,n}^\Delta(0). \end{aligned}$$

By (7), denoting $d_i^k = \delta_1^{(k-1)\Delta} (\delta_1^\Delta - \delta_i^\Delta) \frac{1 - \delta_1^\Delta}{1 - \delta_i^\Delta}$ for $k \geq 1$, and $d_i^0 = \frac{1 - \delta_1^\Delta}{1 - \delta_i^\Delta}$, we have,

$$\alpha_1 U_1^\Delta(0) \leq c - \alpha_2 \sum_{k=0}^{\infty} d_2^k U_2^\Delta(k) - \dots - \alpha_n \sum_{k=0}^{\infty} d_n^k U_n^\Delta(k).$$

In equilibrium each $U_i^\Delta(k)$, player i 's expected payoff from period k on, is at least ir_i . Thus,

$$\alpha_1 U_1^\Delta(0) \leq c - \alpha_2 \sum_{k=0}^{\infty} d_2^k ir_2 - \dots - \alpha_n \sum_{k=0}^{\infty} d_n^k ir_n.$$

By (8), $\sum_k d_i^k = 1$ for every i . Recall also that $\alpha \cdot (ir_1, \dots, ir_n) = c$. We obtain,

$$\alpha_1 U_1^\Delta(0) \leq c - \alpha_2 ir_2 - \dots - \alpha_n ir_n = \alpha_1 ir_1.$$

Since in equilibrium $\alpha_1 U_1^\Delta(0) \geq \alpha_1 ir_1$, we conclude that all the inequalities in the process are actually equalities. Thus,

$$c - \alpha_2 \sum_{k=0}^{\infty} d_2^k U_2^\Delta(k) - \dots - \alpha_n \sum_{k=0}^{\infty} d_n^k U_n^\Delta(k) = c - \alpha_2 \sum_{k=0}^{\infty} d_2^k ir_2 - \dots - \alpha_n \sum_{k=0}^{\infty} d_n^k ir_n.$$

Since all d_i^k are positive, and since, in equilibrium all $U_i^\Delta(k) \geq ir_i$, we have, in particular, $U_i^\Delta(0) = ir_i$ for all the i 's having positive α -coefficients. This implies that $\alpha \cdot (U_1^\Delta(0), \dots, U_n^\Delta(0)) = \alpha \cdot (ir_1, \dots, ir_n) = c$, which means that $(U_1^\Delta(0), \dots, U_n^\Delta(0))$ is not in the interior of R^0 . $\square$

Notice that the proof of the 'if' part does not depend on Δ . In other words, the statement holds for any Δ . In fact, we proved a stronger version of the theorem:

Theorem 2.5.4'. Suppose that all discount factors are different.

- a. There is an equilibrium payoff of the repeated game that Pareto dominates a point in P^Δ if there are at least two points in P^Δ .
- b. If there is only one point in P^0 , then there is no equilibrium payoff in P^Δ that dominates it, for any time interval $\Delta > 0$.

Remark: One may get a general result concerning equilibrium points which are out of the feasible set of one-shot games (in other directions than the Pareto optimum one). Such a result states that a sufficient condition for the existence of equilibrium points (of the repeated game) beyond the one-shot feasible set is the existence of at least one point in the interior of R^0 . This point ensures the existence of at least two individually rational payoffs in the

boundary of V . Using these two points one can construct a repeated game equilibrium which is not in V . The technique is similar to the one in Theorem 2.5.4, and therefore omitted.

Theorem 2.5.4 deals with equilibrium points in the interior of the individually rational region (i.e., in R^0). The question arises as to what can be said about equilibrium points where some players cannot receive more than their individually rational level. Is it true, for instance, that two Pareto optimal points in V on the IR boundary are sufficient to ensure an equilibrium point which is not feasible in the one-shot game? The following two examples show that the answer can be either yes or no.

Consider any game that satisfies the hypothesis of Theorem 2.5.4. Thus, there are two Pareto optimal points in $R^0 \cap V$, and therefore one can find an equilibrium payoff which is not feasible in the one-shot game. By adding a dummy player whose payoffs are flat, say, zero, the set $R^0 \cap V$ becomes void. There are, therefore, no two Pareto optimal points of V in R^0 . Nevertheless, there exists an equilibrium point which is not in V . This example shows that there may be no Pareto points which are strictly greater than the IR levels while there still exists an equilibrium payoff which is not feasible in the one-shot game.

On the other hand, consider the following 4-player game, where each player has two actions: a_j and b_j . All the payoffs are zeros in twelve of the actions combinations. The four exceptions are: $u(a_1, a_2, a_3, a_4) = (1, 2, 0, 0)$, $u(a_1, b_2, a_3, a_4) = (2, 1, 0, 0)$ and $u(a_1, a_2, a_3, b_4) = u(a_1, b_2, a_3, b_4) = (0, 0, -10, 1)$. It is clear that the individually rational levels are all 0 (i.e., $ir_j = 0, j = 1, \dots, 4$). Moreover, the payoffs $(1, 2, 0, 0)$ and $(2, 1, 0, 0)$ form two individually rational and Pareto optimal points in V which are not in R^0 .

Any attempt to play (a_1, a_2, a_3, a_4) or (a_1, b_2, a_3, a_4) will drive player 4 to deviate from a_4 to b_4 and thereby gain 1. Thus, in this example there are two Pareto optimal points in V which are individually rational but no equilibrium point out of V exists. In fact, the only equilibrium point is 0. This kind of example is plausible only when the number of players is four or more as stated in the following proposition.

Proposition 2.5.1: Suppose that there are three players or less, and that all the discount factors are different. If there exist two distinct Pareto optimal points in V which are individually rational, then there exists an equilibrium point in the repeated game outside of V , when Δ is sufficiently small.

Proof: The case of two players is already covered by Theorem 2.5.1, since in

that case the part of the Pareto frontier of V connecting these two points is in P^0 . Thus assume that there are three players and that the two Pareto optimal points not in R^0 .

Whenever there exist two individually rational and Pareto optimal points in V , one can find two points, say $x = (x_1, x_2, x_3)$ and $y = (y_1, y_2, y_3)$, in which at least two players receive more than their IR levels, say, $x_1, y_1 > ir_1$ and $x_2, y_2 > ir_2$. Therefore, no problem might be created by these two players when playing x or y ; any deviation of players 1 or 2 will be followed by a punishment. A problem may arise when the player who receives exactly her IR level, player 3, has a profitable deviation (no punishment prevents her from deviating because she already receives her IR level. I.e., $x_3, y_3 = ir_3$).

Denote by $z = (z_1, z_2, z_3)$ the payoff generated by a possible deviation of player 3. Since $z_3 > ir_3$, there exists a convex combination of x and z whose coordinates are all greater than the respective IR levels. Thus, there is a joint point of V and R^0 . This immediately implies that there are at least two points in P^0 , which contradicts the assumption. Therefore, player 3 has no profitable deviation.

We conclude that one can sustain a combination of x and y without any player having a profitable deviation in the repeated game. A combination of x and y played in different stages can sustain a payoff outside V . $\square$

2.6 Concluding Remarks

Subgame Perfect Equilibrium

The folk theorem we provide is for the Nash equilibrium solution concept. The extension to subgame perfect equilibrium does not differ from the extension of folk theorem for players with the same discount factor. For example, the method used by Fudenberg and Maskin (1986) for games where the set V of feasible payoffs has full dimension, applies also here. The only difference is that in our case, along the optimal equilibrium paths, payoffs are changing. However, for the punishment phases, we need not worry about optimality, and thus the constant punishments used by Fudenberg and Levine will do the job.

Notice that in the 2-player case, unless the matrix of payoffs is constant, there are exactly two possibilities: either V is 2-dimensional, or the players payoffs lie on a line with an increasing slope, in which case the game is strategically equivalent to one where the players have the same payoffs (this immediately follow from the fact that V must have at least one point (weakly) dominating (ir_I, ir_P) , and at least one point (weakly) dominated by it). In the first case, the folk theorem holds also with the subgame perfection refinement, while in the second, the only subgame perfect equilibrium is the

Pareto optimal one⁵.

Games with Incomplete Information

In the analysis above, we confine ourselves to games with complete information, where a known stage-game is played repeatedly. When incomplete information is involved, matters become much more complicated. As opposed to the complete information case, even zero-sum games may have equilibrium points which are not zero-sum. This is exemplified by the following game.

Consider a game where the impatient player can be one of two types: t_1 or t_2 . Assume that his type is chosen with probability $\frac{1}{2}$, and that once chosen, the type becomes known only to him. The patient player knows only the distribution according to which the impatient player's type is selected. Furthermore, assume that player I is extremely impatient relative to player P. The game played is the zero-sum game:

1	0
0	0

if the type chosen is t_1 , and

0	0
0	1

⁵ See Abruë, Dutta and Smith (1994) for a more detailed discussion of the "non-equivalent utilities" condition.

if the type chosen is t_2 . The impatient player, who knows the matrix played will play top in the first stage if the type is t_1 and bottom if the type is t_2 . At this point, he reveals his type and the patient player knows to respond in later stages by playing right against t_1 and left against t_2 . I's payoff (recall, he is extremely impatient) is nearly $\frac{1}{2}$ (this is his expected payoff at the first stage) while P's payoff is close to 0 (which is the stage-payoff from stage 2 on). Thus, the payoffs in the repeated game are approximately $\frac{1}{2}$ and 0 for the impatient and patient players, respectively. This is certainly not a zero-sum payoff.

The reason why it may happen in incomplete information games is that once information is used, as a result of Bayesian updating, the game is not the same anymore; players are endowed with different knowledge than they started the game with. In this respect the situation is not stationary; anytime the game is played the game is actually different due to information acquisition by player P. The lack of stationarity is the source of the difference between games with and without complete information.

Convergence of Discount Factors to 1

In order to establish a folk theorem, one needs to have discount factors close to 1. Here, we are concerned with players who have different discount factors. Therefore, we need to retain the difference between the players while the discount factors converge to 1. There are many converging paths of the n -

vector of discount factors to the vector $(1, \dots, 1)$. For instance, the ratios $(1 - \delta_i) / (1 - \delta_j)$ could be kept fixed.

A path of discount factors can be interpreted in different ways. One may think of a fixed game, being played by different n -tuples of players, that become increasingly patient. This interpretation makes the choice of the convergence path somewhat arbitrary. Another interpretation, is to consider specific players, with fixed time preferences, while shortening the time lap between two consecutive stages.

The approach we adopted above is the second. The discount factors δ_j , representing the present value of payoff delayed by 1 time unit, are fixed throughout. The stage discount factors, i.e., the factors that represent the difference between one payoff unit in two consecutive stages, are δ_j^Δ . When Δ goes to zero, all stage discount factors go to 1.

Another way to express the same idea is to take players' discount factors that converge to 1, while keeping the ratio between their logarithms, $\log \delta_i / \log \delta_j$, constant. In other words, gradually shortening the time lap between stages to 0 is equivalent to choosing a specific convergence path for the discount factors.

The Distinction Between the 2-Player and the n-Player Cases

In Section 4, we could provide a characterization of the equilibrium payoffs set only for the 2-player case. In the 2-player case for Δ small enough and for the hyperplane represented by $\alpha \in \mathbb{R}^2$, in order to sustain a solution of the maximization problem (1), a path of adjacent extreme point in V yielding monotone payoffs to each player should be followed. For instance, in the Pareto direction one should start at a point that yields greater payoffs for the impatient player first and to increase gradually the payoffs of patient player at the expense of the impatient one. In other words, what makes the characterization possible is the fact that for a given $\alpha \in \mathbb{R}^2$ the path generating an extreme point follows a simple one dimensional curve (the Pareto frontier of V).

In the n-player case, in contrast, for a given direction $\alpha \in \mathbb{R}^n$, the sequence of extreme points of V followed does not have any monotonicity property. At any repetition the optimal path (of (2)) may yield tail payoffs below the IR level of any player. We could not even explicitly define a sequence whose continuation values (starting at any k) are above the IR levels of all players. This is the reason we leave the general folk Theorem characterization open.

For some particular cases, however, we can be more specific:

Proposition 2. In the case where all the Pareto optimal points of V are strongly individually rational, any Pareto optimal point of the repeated game's feasible set is also an equilibrium point, for Δ small enough.

A more detailed result, concerning the general n -player case, would be of great interest.

The case of extremely different players

To make the point consider two players whose time preferences are extremely different. That is, the ratio $r = (1 - \delta_p)/(1 - \delta_l)$ is close to 0. Furthermore, suppose the following game is played:

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

In this example the IR levels are 0. By Theorem 2.4.1, when r is close to 0, one can support an equilibrium payoff close to (1,1). In other words, if r goes to zero, and if the time lap between iterations of the stage game, Δ , goes to zero, then (1,1) can be approximated by equilibrium of the repeated game.

More generally, when r goes to 0, the following pair of payoffs:

(the maximal payoff in V of the impatient player,

the maximal payoff in $V \cap IR_i$ of the patient player),

can be approached by equilibria of the repeated game.