

Forgiving-Proof Equilibrium*

Miguel Aramendia and Luis Ruiz

Economia Aplicada IV, Universidad del Pais Vasco, Bilbao 48015, Spain.

Abstract

In repeated games, when a player deviates from the agreed strategy this may result in losses for other players, but it may also result in gains. If the profits of the betrayed players drop they will react to prevent such actions from recurring. But if their profits increase, where is the sense in punishing? We explore a new equilibrium concept that rules out strategies in which punishments are carried out for deviations that have not harmed any of the other players. The goal is to prevent potential renegotiations that could end up in an agreement to forgive. We show that this definition significantly reduces the set of equilibrium outcomes in many games.

Keywords: Infinitely repeated games; Renegotiation.

JEL classification: C72.

*Financial support from Ministerio de Educación y Cultura (project BEC2003-08182) is gratefully acknowledged.

1. Introduction

It is well-known that repeating a game may lead to the emergence of new equilibria which can sustain many cooperative outcomes. Folk theorems assert that under simple conditions any feasible and individually rational payoff of the one-shot game can be sustained as a subgame perfect equilibrium in the infinitely repeated game with discounting provided that the players are sufficiently patient (i.e. that the discount factor is sufficiently close to 1).

Since the set of equilibrium payoffs is so large, and each payoff vector can also be sustained by a wide variety of strategies, before the game starts the players must agree on what strategies are to be followed. This means that there must be communication, at least at the beginning of the game. And it seems natural that the players keep the channels of communication open throughout all periods. This opens the door to renegotiation. So, as many authors have suggested, at the beginning of each period the players should have the option of renegotiating their prescribed strategies.

With the aim of avoiding subgame perfect equilibria which are likely to be renegotiated, various concepts of renegotiation-proofness have been introduced. The idea underlying these concepts is to set up a kind of filter at the beginning of the game to rule out strategies that may be subgame perfect equilibria but run the risk of not being carried out for some reason. Here we propose a different way of setting up such a filter, and one which, as we shall see, significantly restricts the set of strategies that can be used, and in many cases also restricts the set of payoffs that can be sustained in equilibrium.

The idea underlying this new definition can be illustrated by this example. Consider the Cournot model with two firms and complete information. Suppose that one of the firms suffers a strike by its workers or needs to stop work to refurbish its machinery. The result is a drop in that firm's output for the period in question. The strategies that appear in the literature on repeated games recommend punishing the firm for failing to meet its agreed output quota, but the reduction in output drives the price up, and that is beneficial for the other firm. Does it make sense for the second firm to punish the first for an action from which it has benefited? We think that, at least in this case, we would all agree that it does not. It would be more logical to consider only those strategies that recommend continuing to cooperate as if nothing had happened.

Now let us look at "stick and carrot" strategies as introduced by Abreu (1988). These are fully symmetrical strategies that use a single punishment, regardless of

the type of deviation that takes place. That punishment has two stages: first comes the "stick", which comprises a period of negative profits. Then comes the "carrot", which is less severe. Let us take the case of two players with identical cost functions. Observe that if one of the two deviates during the stick in search of better results, it has to reduce its production, because for profits to be negative the price must be lower than the marginal cost. But, as in the previous case, this deviation is also beneficial for the other player. In these circumstances the betrayed player is likely to decide to forgive rather than punish. It can perhaps be argued that a player cannot forgive on its own account, because by doing so it is contravening the strategy and must itself be punished in the following period. This is true, but we must not forget that we are operating in a context where everything is negotiable. In this case it is perfectly possible for the betrayer to attempt to renegotiate with its opponent in terms such as the following:

- Look, you are much better off than expected, and if you punish me we will both lose out. I know you are supposed to punish me, because if you don't I am supposed to punish you in the next period, but I promise that if you don't punish me I won't punish you for not punishing me. It would be absurd for me to do so, as we would both have to go through the stick again.

The only difference between the first example and this one is that in the former situation the player that cuts production gains less than agreed, while in the latter the deviating player gains more. We believe that it is not sensible for a player to decide whether to forgive or punish on the basis of what its opponent gains or loses. From a rational viewpoint it is more coherent to make its decision on the basis of how the action of the other player affects our own profits. If its profits drop (compared to the agreed level) it will react to prevent such actions from recurring. But if its profits increase, where is the sense in punishing?

Our new concept rules out strategies such as the stick and carrot, in which punishment stages are carried out for deviations that have not harmed any of the other players. Intuitively, the definition of the forgiving proof equilibrium is quite simple. We say that a strategy profile is a forgiving-proof equilibrium (FPE) if it is a subgame perfect equilibrium (SPE) and it does not punish deviations which are harmless to the other players.

This definition significantly reduces the set of equilibrium outcomes in many games. For instance in the Cournot model with linear demand and two players the FPE-payoff area is a quarter of the SPE-payoff area. We will also look at some

matrix games in which the FPE-payoff set is reduced to a single point.

The rest of the paper is organized as follows. Section 2 presents some preliminaries and then introduces the definition and some necessary and sufficient conditions. Section 3 studies some examples. Section 4 concludes and shows the difference between FPE and the weakly renegotiation proof equilibrium of Farrell and Maskin (1989).

2. Forgiving-Proof Equilibrium

Let $G = (A_1, \dots, A_n; \Pi_1, \dots, \Pi_n)$ be an n -player game where $N = \{1, \dots, n\}$ is the set of players, A_i is the set of (pure or mixed) actions a_i of player i and $\Pi_i : A_1 \times \dots \times A_n \rightarrow R$ is player i 's payoff function. We assume that the action sets A_i are compact and the payoff functions Π_i are continuous. For the sake of simplicity we consider that the game G has at least one Nash equilibrium. The associated infinitely repeated game with discounting is denoted by $G^\infty(\delta)$ where $\delta \in (0, 1)$ is the discount factor. If $a(t) = (a_1(t), \dots, a_n(t))$ is the vector of actions played in period t , then $\{a(1), \dots, a(t)\}$ is a history h of length t . A strategy σ_i of player i in $G^\infty(\delta)$ is a sequence of functions σ_i^t from the set of all histories of length $t - 1$ to A_i , so $\sigma_i^1 \in A_i$ is the initial action of player i . We assume perfect monitoring, that is, every player can observe the past actions of the rest of the players.

A stream of action profiles $\{a(t)\}_{t=1}^\infty$ is referred to as an outcome path and is denoted by S . Any strategy profile $\sigma = (\sigma_1, \dots, \sigma_n)$ generates an outcome path $S(\sigma) = \{a(\sigma)(t)\}_{t=1}^\infty$ defined inductively by:

$$\begin{aligned} a(\sigma)(1) &= \sigma^1 \\ a(\sigma)(t) &= \sigma^t(a(\sigma)(1), \dots, a(\sigma)(t-1)), \text{ if } t > 1. \end{aligned}$$

The value $\Pi_i(a(t))$ denotes the payoff of player i in period t when the outcome in that period is $a(t)$. And $\Pi_i^\delta(S)$ denotes the average discounted payoff of player i for the outcome path $S = \{a(t)\}_{t=1}^\infty$:

$$\Pi_i^\delta(S) = (1 - \delta) \sum_{t=1}^{\infty} \delta^{t-1} \Pi_i(a(t)).$$

Then, the average discounted payoff of player i in $G^\infty(\delta)$ obtained with the strategy profile $\sigma = (\sigma_1, \dots, \sigma_n)$ is:

$$\Pi_i^\delta(\sigma) = \Pi_i^\delta(S(\sigma)).$$

A strategy profile $\sigma = (\sigma_1, \dots, \sigma_n)$ is a Nash equilibrium in $G^\infty(\delta)$ if for all $i \in N$, σ_i is the best response to $(\sigma_1, \dots, \sigma_{i-1}, \sigma_{i+1}, \dots, \sigma_n)$. And it is a subgame perfect

equilibrium (SPE) if after every history h , σ_h (i.e. the continuation of σ after h) is a Nash equilibrium in the corresponding subgame.

As many authors have argued, when renegotiation is possible subgame perfection is a necessary condition to avoid deviations, but it is not sufficient. In recent years different concepts of renegotiation-proofness have been introduced. Each of these concepts reduces the equilibrium strategies by eliminating those strategies that, under various criteria, run the risk of not being met (in spite of their being SPE).

The concept that we define here rules out those strategies that recommend punishing deviations that have not harmed any of the other players. The goal is to prevent potential renegotiations that could end up in an agreement to forgive.

Definition 1. A strategy profile σ is a Forgiving-Proof Equilibrium if σ is a subgame perfect equilibrium and does not punish deviations which are harmless to the other players. If an equilibrium σ is FPE, then we also say that the payoffs $(\Pi_1^\delta(\sigma), \dots, \Pi_n^\delta(\sigma))$ are FPE.

Now let us see that FPE strategies can only recommend actions of one type, which we call “individually efficient actions”. We decompose any action profile $a = (a_i, a_{-i})$.

Definition 2. We say that an action profile a is *individually efficient* if there is no a single deviation of a player which is profitable for him and harmless for the others. That is, if there is no $\bar{a}_i \in A_i$ for a player i such that

- i) $\Pi_i(\bar{a}_i, a_{-i}) > \Pi_i(a_i, a_{-i})$,
- ii) $\Pi_j(\bar{a}_i, a_{-i}) \geq \Pi_j(a_i, a_{-i})$ for all $j \neq i$.

We denote by I the set of all the individually efficient actions of the game G . It is easy to see that if the action profile a is a Nash equilibrium then it is individually efficient. As, by assumption, the game G has a Nash equilibrium, the set I is not empty. It can also be proven that the set I includes the set of efficient actions (in the Pareto sense).

Theorem 3. If the strategy profile σ is FPE for any $\delta < 1$, then the actions which are recommended by σ after any history are individually efficient.

Proof. Assume that σ does not punish deviations which are harmless to the other players and there exists a non individually efficient action profile recommended after a history. Then there exists a player \hat{i} who could deviate profitably for one period without harming the others and therefore not be punished by σ . If he does deviate in this way and reverts to complying with the strategy immediately afterwards, then this deviation is profitable for him in the associated subgame, and consequently σ is

not SPE. ■

Note that the converse is not true. For instance, a strategy may recommend playing individually efficient actions in all subgames and not be FPE on the grounds of setting punishments for deviations that are not harmless.

Note also that the infinite repetition of the one-shot Nash equilibrium is FPE.

Below, we establish necessary and sufficient conditions for a payment vector to be FPE. We denote by $\Pi_i^D(a)$ the best response payoff of player i from the action a , that is, $\Pi_i^D(a) = \max\{\Pi_i(a_1, \dots, a_{i-1}, \bar{a}_i, a_{i+1}, \dots, a_n) \mid \bar{a}_i \in A_i\}$.

Theorem 4. If the payoff profile $v = (v_1, \dots, v_n)$ is FPE for some $\delta < 1$, then there exist individually efficient action profiles a^i such that $\Pi_i^D(a^i) \leq v_i$, for $i = 1, \dots, n$.

Proof. We prove the proposition by contradiction. Assume there exists an FPE σ such that $\Pi_i^\delta(\sigma) = v_i$ for $i = 1, \dots, n$, and there exists a player \hat{i} such that

$$\Pi_{\hat{i}}^D(a) > v_{\hat{i}} \text{ for all } a \in I. \quad (1)$$

From Theorem 3, all actions recommended by σ after any history are individually efficient. Now, suppose that player \hat{i} deviates optimally right from the beginning and forever. Then, by (1), he will obtain an average discounted payoff greater than $v_{\hat{i}}$ and, therefore, strategy σ is not SPE, which contradicts the initial assumption. ■

The condition in Theorem 4 itself is not sufficient. A sufficiency theorem is presented below that is easily applicable to the examples in Section 3. This theorem has the virtue of having a very simple constructive proof making use of fair simple strategy profiles. We define these fair simple strategy profiles similarly to the simple strategy profiles introduced by Abreu (1988), but bearing in mind that these strategies do not punish deviations which are harmless to the other players, and that all actions are individually efficient.

Formally, a fair simple strategy profile is determined by $n + 1$ outcome paths S^0, S^1, \dots, S^n of individually efficient actions that induce the following strategy profile $\sigma_F(S^0, S^1, \dots, S^n)$:

- i) Play S^0 until a player deviates singly from S^0 with harm for another player.
- ii) Play S^j if the j th player deviates singly from S^i with harm for another player, where S^i is an ongoing previously specified path, $i = 0, 1, \dots, n$. Continue with S^i if no deviations occur, or single deviations are harmless, or two or more players deviate simultaneously.

It is clear that $\Pi_i^\delta(\sigma_F(S^0, S^1, \dots, S^n)) = \Pi_i^\delta(S^0)$. According to Abreu (1988, p. 391), only one-shot deviations must be checked to ensure that a simple strategy profile

is SPE, where a one-shot deviation from a strategy in $G^\infty(\delta)$ consists of deviating from the strategy for a single period, and sticking to it subsequently. If we follow the same steps as Abreu, it is easy to see that a fair simple strategy profile is FPE if and only if one-shot harming deviations are not profitable. Note that if a player one-shot deviates from any path without harm for the rest of the players then, since actions are individually efficient, the deviation is not profitable for him in the short-run (i.e. just that period), or indeed in the long-run, since the continuation is the same.

Theorem 5. Let $\Pi(a) = v$ for an individually efficient action a . If there exist n individually efficient actions a^1, \dots, a^n such that $\Pi_i^D(a^i) < v_i$ and $\Pi_i(a^i) < \Pi_i(a^j)$ for $i, j = 1, \dots, n, i \neq j$, then it is possible to construct an FPE strategy with average payoff v for δ close enough to 1.

Proof. We consider the fair simple strategy profile $\sigma_F(S^0, S^1, \dots, S^n)$ where the cooperative path $S^0 = \{a, a, \dots\}$ and the n punishing paths are defined as follows

$$S^i = \{a^i, \dots, a^i, a, a, \dots\},$$

T times

which consists of playing a^i for T periods and then reestablishing the path S^0 . For the sake of simplicity, we consider the same T for the n punishing paths. By construction all the recommended actions are individually efficient. Next, we see which values of T are appropriate to guarantee that this fair simple strategy will be FPE for δ close enough to 1. Note that, in this particular situation, only the following one-shot deviations must be checked.

i) If player i deviates from the cooperative path S^0 with harm for another player, then his payoff is at most $(1 - \delta)M + \delta\Pi_i^\delta(S^i)$, where M is a finite upper bound for the payoffs of any player in G , and $\Pi_i^\delta(S^i) = (1 - \delta^T)\Pi_i(a^i) + \delta^T v_i$. Set $R(\delta) = (1 - \delta)M + \delta((1 - \delta^T)\Pi_i(a^i) + \delta^T v_i) - v_i$. We have that $\lim_{\delta \rightarrow 1} R(\delta) = 0$ and $\lim_{\delta \rightarrow 1} R'(\delta) = -M - T\Pi_i(a^i) + (T + 1)v_i$. Thus, taking

$$T > \frac{M - v_i}{v_i - \Pi_i(a^i)},$$

$\lim_{\delta \rightarrow 1} R'(\delta) > 0$, and $R(\delta) \leq 0$ for δ close enough to 1. Therefore there is no incentive for any player to deviate from the cooperative path when δ is close enough to 1.

ii) Next, we show that there is no incentive for player i to deviate with harm for another player from his punishing phase in any period. As $\Pi_i(a^i) \leq \Pi_i^D(a^i) < v_i$, it follows that $\Pi_i^D(a^i) < (1 - \delta^T)\Pi_i(a^i) + \delta^T v_i$ for δ close enough to 1. Therefore $(1 - \delta)\Pi_i^D(a^i) + \delta((1 - \delta^T)\Pi_i(a^i) + \delta^T v_i) < (1 - \delta^T)\Pi_i(a^i) + \delta^T v_i \leq (1 - \delta^t)\Pi_i(a^i) + \delta^t v_i$ for $t = 1, \dots, T$ and for δ close enough to 1.

iii) Finally, we see that there is no incentive for player i to deviate with harm for another player from the punishment phase of player j ($j \neq i$). If player i does not deviate then his payoff is $(1-\delta^t)\Pi_i(a^j)+\delta^t v_i$ for $t = 1, \dots, T$. If player i deviates, then his payoff is at most $(1-\delta)M+\delta((1-\delta^T)\Pi_i(a^i)+\delta^T v_i)$. We have two different cases. If $\Pi_i(a^j) \geq v_i$, then, from above $(1-\delta)M+\delta((1-\delta^T)\Pi_i(a^i)+\delta^T v_i) \leq v_i \leq (1-\delta^t)\Pi_i(a^j)+\delta^t v_i$. If $\Pi_i(a^j) < v_i$, set $R(\delta) = (1-\delta)M+\delta((1-\delta^T)\Pi_i(a^i)+\delta^T v_i)-(1-\delta^t)\Pi_i(a^j)-\delta^t v_i$. It holds that $\lim_{\delta \rightarrow 1} R(\delta) = 0$ and $\lim_{\delta \rightarrow 1} R'(\delta) = -M+v_i-T\Pi_i(a^i)+Tv_i+t(\Pi_i(a^j)-v_i) \geq -M+v_i-T\Pi_i(a^i)+Tv_i+T(\Pi_i(a^j)-v_i) = -M+v_i+T(\Pi_i(a^j)-\Pi_i(a^i))$. Thus, taking

$$T > \frac{M - v_i}{\Pi_i(a^j) - \Pi_i(a^i)},$$

$\lim_{\delta \rightarrow 1} R'(\delta) > 0$ and $(1-\delta)M+\delta((1-\delta^T)\Pi_i(a^i)+\delta^T v_i) < (1-\delta^t)\Pi_j(a^i)+\delta^t v_j$ for $t = 1, \dots, T$ and for δ close enough to 1.

Summing up, the strategy profile $\sigma_F(S^0, S^1, \dots, S^n)$ is FPE for δ sufficiently close to 1 when the number of periods T satisfies $T > \max\{\frac{M-v_i}{v_i-\Pi_i(a^i)}, \frac{M-v_i}{\Pi_i(a^j)-\Pi_i(a^i)} \mid i, j = 1, \dots, n, i \neq j\}$. ■

3. Some Examples

Example 1. Consider the Cournot duopoly linear model in which marginal cost is c and inverse demand is given by $p(q) = D - bq$ (if $q \leq \frac{D}{b}$ and 0 otherwise), where $b > 0$ and $D > c$. The actions for both firms are quantities $q_i \in [0, K]$ (K is big enough) and the payoffs are $\Pi_i(q_1, q_2) = ((p(q_1 + q_2) - c)q_i$. In this case, we have that the set of individually efficient actions is

$$\left\{ (q_1, q_2) \in [0, K] \times [0, K] \mid 0 < q_1 \leq \frac{D-c-bq_2}{2b}, 0 < q_2 \leq \frac{D-c-bq_1}{2b} \right\} \cup \left\{ \left(0, \frac{D-c}{2b}\right), \left(\frac{D-c}{2b}, 0\right) \right\}$$

where $\frac{D-c-bq_2}{2b}$ is the (unique and decreasing) best response of firm 1 to action q_2 of firm 2. To see this, note that if a firm's quantity is bigger than its best response, then the firm profits by reducing to its best response and this does not harm its opponent since the price does not decrease. Similarly, the action profile $(q_1, 0)$ with $q_1 \neq \frac{D-c}{2b}$ is not individually efficient, since firm 1 profits and firm 2 remains with zero payoff when firm 1 deviates to its best response. By contrast, $(\frac{D-c}{2b}, 0)$ is individually efficient since firm 1 loses when either of the two firms deviates. Last, if both firms' quantities are positive and not bigger than their best responses, then a downward

deviation of a firm produces a reduction in its own profits, and an upward deviation produces a loss for the opponent since the price decreases. So these action profiles are individually efficient.

As $\Pi_1^D(q)$ is decreasing on q_2 , the minimum best response payoff of firm 1 from individually efficient actions is obtained when the quantity of firm 2 is as big as possible, that is,

$$\begin{aligned} \min\{\Pi_1^D(q_1, q_2) \mid (q_1, q_2) \text{ is individually efficient}\} &= \Pi_1^D\left(0, \frac{D-c}{2b}\right) \\ &= \Pi_1\left(\frac{D-c}{4b}, \frac{D-c}{2b}\right) = \frac{(D-c)^2}{16b} = \frac{\Pi^M}{2}, \end{aligned}$$

where $\Pi^M = \frac{(D-c)^2}{8b}$ is the symmetric monopoly outcome for both firms.

Now, Theorem 4, Theorem 5 with $a^1 = (0, \frac{D-c}{2b})$ and $a^2 = (\frac{D-c}{2b}, 0)$, and some easy calculations for the border payoffs show that the set of FPE payoffs is

$$\{(v_1, v_2) \mid v_1 + v_2 \leq 2\Pi^M, v_1 > \frac{\Pi^M}{2}, v_2 > \frac{\Pi^M}{2}\}.$$

It is well known that the set of SPE payoffs is $\{(v_1, v_2) \mid v_1 + v_2 \leq 2\Pi^M, v_1 > 0, v_2 > 0\}$, so the FPE-payoff area is 25% of the SPE-payoffs area. Note that the lower bound $\frac{(D-c)^2}{16b}$ of FPE payoffs is smaller than $\frac{(D-c)^2}{9b}$ the one-shot Nash equilibrium payoff.

If we consider the nonlinear inverse demand $p(q) = 110 - \sqrt{q}$ (if $\sqrt{q} \leq 110$ and 0 otherwise) and $c = 10$, then the best response functions are also decreasing. Following the same steps as in the above example, the set of FPE payoffs is

$$\{(v_1, v_2) \mid v_1 + v_2 \leq 148148.2, v_1 > 41794.5, v_2 > 41794.5\}.$$

Now the FPE-payoff area is 18.98% of the SPE-payoffs area.

If Cournot duopolists face the nonlinear industry demand $q = \frac{1}{p^2}$ of constant elasticity 2, and $c = 0.5$, then, the best response functions are not decreasing and the set of individually efficient actions is the enclosed area below (Figure 1)




Figure 1




Figure 2

The minimum best response payoff of firm 1 from individually efficient actions is

$$\begin{aligned} \min\{\Pi_1^D(q_1, q_2) \mid (q_1, q_2) \text{ is individually efficient}\} &= \Pi_1^D(0.59259, 1.18519) \\ &= \Pi_1(1.11251, 1.18519) = 0.17768. \end{aligned}$$

And the set of FPE payoffs is

$$\{(v_1, v_2) \mid v_1 + v_2 \leq 0.5, v_1 > 0.17768, v_2 > 0.17768\}.$$

In this case the FPE-payoff area is 8.36% of the SPE-payoffs area (see Figure 2 above).

Example 2. Consider the following version of the prisoner's dilemma

	C	F
C	2,2	-1,3
F	3,-1	0,0

In this case all (mixed) actions are individually efficient. And all feasible and individually rational payoffs can be sustained as FPE. For instance, using the "fair-trigger" strategy: play an individually efficient action cooperatively until one player deviates with harm to the other player, then both players play forever a Nash equilibrium of the matrix game.

Example 3. Consider the following matrix game (Aramendia 2006)

	L	R
U	2,3	-1,-2
D	3,0	0,-1

Let α, β be the probabilities that players I and II play their first pure action, respectively. The payoffs are $\Pi_1(\alpha, \beta) = 3\beta - \alpha$ and $\Pi_2(\alpha, \beta) = 4\alpha\beta - \alpha + \beta - 1$. We have that $\Pi_2(\alpha, \bar{\beta}) > \Pi_2(\alpha, \beta)$ iff $\bar{\beta} > \beta$, and $\Pi_1(\alpha, \bar{\beta}) \geq \Pi_1(\alpha, \beta)$ iff $\bar{\beta} \geq \beta$, so there is no profitable single deviation of player 2 which is harmless for player 1 iff $\beta = 1$. Moreover, $\Pi_1(\bar{\alpha}, 1) > \Pi_1(\alpha, 1)$ iff $\bar{\alpha} < \alpha$, and $\Pi_2(\bar{\alpha}, 1) \geq \Pi_2(\alpha, 1)$ iff $\bar{\alpha} \geq \alpha$. Therefore the set of individually efficient actions is $\{(\alpha, 1) \mid 0 \leq \alpha \leq 1\}$. Since $\Pi_1^D(\alpha, 1) = 3$, the only FPE payoff is $(3, 0)$ which is the Nash equilibrium payoff of the matrix game. This contrasts with the SPE-payoff set $\{(v_1, v_2) \mid 3v_1 + v_2 \leq 9, 5v_1 - 3v_2 \geq 1, v_1 > 0, v_2 > 0\}$.

Example 4. Consider the following symmetric matrix game (Abreu 1988, p. 386)

	L	M	H
L	10, 10	3, 15	0, 7
M	15, 3	7, 7	-4, 5
H	7, 0	5, -4	-15, -15

Bearing in mind that if a player deviates from H to M both players become better off, it is left to the reader to see that the set of individually efficient actions is

$$\{(\alpha, \beta, 0), (\bar{\alpha}, \bar{\beta}, 0) \mid \alpha + \beta \leq 1, \alpha \geq 0, \beta \geq 0, \bar{\alpha} + \bar{\beta} \leq 1, \bar{\alpha} \geq 0, \bar{\beta} \geq 0\},$$

where α, β are the probabilities that player I plays his first and second pure action, respectively, and similarly $\bar{\alpha}, \bar{\beta}$ for player II. The payoffs that can be sustained as FPE (for instance, using the fair trigger strategy) are

$$\{(v_1, v_2) \mid 7v_1 + 5v_2 \leq 120, 5v_1 + 7v_2 \leq 120, v_1 \geq 7, v_2 \geq 7\}.$$

The FPE-payoff area is now 10.14% of the SPE-payoffs area.

4. Conclusions

If a strategy is not a subgame perfect equilibrium that is because there is a player that can gain more by betraying than by acting as agreed. Subgame perfection is therefore a necessary condition to avoid betrayal, but according to most authors, it is not a sufficient condition. In any period all players can reach an agreement to renegotiate and abandon the initial strategy if this works out best for them. In this paper we have introduced a renegotiation proof concept that rules out strategies in which punishment are carried out after deviations that have not harmed any of the other players. In our opinion this is not a sufficient condition to ensure that there is no renegotiation either: rather it is a necessary condition. For instance, note that the strategy that we call "fair-trigger" is FPE, but it nevertheless recommends that there be no further cooperation following a betrayal. It seems likely that the players will renegotiate at some time in the punishment stage to resume the cooperation stage. Indeed, other equilibrium concepts such as the Weakly Renegotiation Proof of Farrell and Maskin (1989), which enjoys the approval of recent literature, eliminate it.

Remember that a subgame-perfect σ equilibrium is Weakly Renegotiation Proof (WRP in short) if there are no continuation equilibria σ^1, σ^2 of σ such that σ^1 strictly Pareto dominates σ^2 . When it comes to filtering strategies, this definition coincides with ours in some points. For instance in the Cournot model both definition eliminate strategies that propose negative payments in any period. But in general the differences are greater than the similarities. Some strategies, such as the fair-trigger strategy, are FPE but not WRP. And conversely some strategies are WRP but not FPE. In the Cournot model with two players and linear demand the area of WRP payoffs is 87.65% of the area of SPE payoffs. So in this case at least our FPE concept

is considerably more restrictive than WRP (remember that the FPE payoff area is 25% of the area of SPE payoffs), which means that many WRP strategies do not get through the FPE filter.

Section 3 presents several examples in which the FPE concept gives a substantially smaller set of equilibrium payoffs than SPE payoffs. For instance, in the Cournot model, payments that are too low and too asymmetric are ruled out. To some extent FPE eliminates the "bad" equilibria which are repeatedly mentioned (but never precisely defined) in the relevant literature.

As far as the number n of players is concerned, it must be stressed that although the examples in Section 3 refer to the simplest case ($n=2$), the FPE concept is defined for any n . However, when $n>2$ coalitions may come into play. SPE strategies are designed to prevent only single player deviations, and when two or more players deviate they recommend ignoring the deviation. Group deviations therefore go unpunished. Horniaček (1996) and Larrea and Ruiz (2004) argue in different ways that group deviations cannot be ignored. What they prove, roughly speaking, is that if all members of a coalition can become better off by deviating then that deviation can be sustained with a strategy from which no subcoalitions will deviate. The conclusion is that group deviations must not be ignored, they must be punished. We believe that when there are more than two players the concept of FPE would have to be refined to take this possibility into account.

In any event, to be able to state truthfully that a strategy will not be renegotiated we would need an even more comprehensive definition. Apart from solving this problem, we would have to take into account the possibility that all players may decide to renegotiate to switch to another strategy that proposes better payments for them all. And perhaps this comes closer to what can really be understood by renegotiation. In the words of Bergin and MacLeod (1993) (p.43):

- *An equilibrium is renegotiation proof if agents as a group cannot find a better alternative.*

To prevent players from seeking improvements through other strategies we would have to propose "optimum" payments at both the cooperative and punishment stages, and that is not an easy thing to do, since in general the bigger the payments received by players in punishment stages are the lower the sustainable cooperative payments are.

To quote Bergin and MacLeod again (p. 43):

- *A difficulty arises when agents can negotiate away from bad outcomes: the*

coercive power of threats is weakened, and this in turn reduces the scope for sustaining good outcomes.

This is a complex problem which, like the previous one, deserves further research.

References

Abreu, D., 1988. On the Theory of Infinitely Repeated Games with Discounting. *Econometrica* 56, 383-396.

Aramendia, M., 2006. Asymmetric Finite Punishments in Repeated Games. *Economic Letters*, forthcoming.

Bergin, J. and MacLeod, W. B., 1993. Efficiency and Renegotiation in Repeated Games. *Journal of Economic Theory* 61, 42-73.

Farrell, J. and Maskin, E., 1989. Renegotiation in Repeated Games. *Games and Economic Behavior* 1, 327-360.

Horniaček, M., 1996. The Approximation of a Strong Perfect Equilibrium in a Discounted Supergame. *Journal of Mathematical Economics* 25, 85-107.

Larrea, C. and Ruiz, L., 2004. Simple Coalitional Strategy Profiles. Discussion Paper 43, Dto Economía Aplicada IV, Basque Country University, Bilbao, Spain.