

Reputation and Perfection in Repeated Common Interest Games*

Martin W. Cripps and Jonathan P. Thomas

Department of Economics,

University of Warwick, Coventry CV4 7AL

March 1996

* We would like to express our gratitude to Klaus Schmidt, an associate editor and an anonymous referee for useful comments. Our thanks are also due to seminar participants at Bonn, Edinburgh, Erasmus, Exeter and Tilburg for comments. All remaining errors are our responsibility.

PROPOSED RUNNING HEAD: Reputation and Perfection

ABSTRACT: We consider a wide class of repeated common interest games perturbed with one-sided incomplete information: one player (the informed player) might be a commitment type playing the Pareto dominant action. As discounting, which is assumed to be symmetric, and the prior probability of the commitment type go to zero, it is shown that the informed player can be held close to her minmax payoff even when perfection is imposed on the equilibrium.

J.E.L. CLASSIFICATION NUMBERS: C73, D83

MAILING ADDRESS: Martin W. Cripps
Department of Economics
University of Warwick
Coventry CV4 7AL
England

1. INTRODUCTION

Two-person common interest games are defined as games with a strongly Pareto dominant payoff vector (Aumann & Sorin, 1989). If the game is repeated infinitely often, and if the players are patient, it might be expected that they would be able to coordinate and receive average payoffs close to the dominant payoff vector. It is, however, an implication of the Folk theorem for repeated games that there exist equilibria in which patient players receive payoffs substantially below the dominant payoff vector. Moreover, even imposing subgame perfection does not alter this general result (Fudenberg and Maskin, 1986). That such inefficient equilibria can survive in the long run when players are very patient seems counter-intuitive, and in this paper we shall investigate whether perfection, when applied to a simple "reputation" model, can lead to such undesirable equilibria being eliminated.

Specifically, we shall consider perturbing a common interest game with (only) the possibility that one of the players, say player 1, might be a type committed to playing a cooperative action, that is, the action corresponding to the dominant payoff vector. The other player, player 2, is unsure of player 1's type. This will allow the possibility of a reputation effect, where player 1 can mimic the commitment strategy (of always playing the cooperative action) in the hope of convincing player 2 of her cooperative intentions. The question we address is: will this form of incomplete information allow us to rule out at least the most undesirable equilibria as the players become very patient? If the equilibrium concept is that of Nash equilibrium, the answer is negative¹. We shall consider whether reputation arguments might nevertheless have a degree of power when the equilibrium concept is refined to incorporate some notion of perfectness. Our results will show that the answer to the question is still negative, in that a small perturbation of the original common interest game has little effect on the attainable equilibria, and extremely undesirable equilibria still exist.

In fact, if attention is restricted to *pure strategy* equilibria, imposing perfection *does* lead to payoffs close to the Pareto dominant pair. If a pure-strategy equilibrium leads to payoffs substantially below the Pareto-dominant ("cooperative") payoff pair, then there must be periods in which one or

¹Even the assumption of *two-sided* uncertainty of the type we assume does not force cooperation in Nash equilibrium. Aumann & Sorin (1989) construct a mixed-strategy counter example (to their main result, which assumes pure strategies) in which cooperation is not approximated as the players become patient.

both players do not play the cooperative actions. If, in the first period this occurs, player 1 is supposed to play noncooperatively, then by cooperating instead she will establish a reputation for being the commitment type, and cooperative payoffs are guaranteed thereafter; hence not cooperating cannot be an equilibrium strategy. It must therefore be player 2 who is first supposed to play noncooperatively, and in a Nash equilibrium this can be enforced by severe off-the-equilibrium-path punishments by player 1. But suppose that we impose *perfection* on the equilibrium. To punish player 2, player 1 must, at some point, play non-cooperatively. By not doing so, however, she will establish a reputation for being the commitment type and hence guarantee herself the cooperative payoff thereafter. Roughly speaking, player 1 cannot credibly punish because she can always "hide behind" the possibility of being the commitment type (and has every incentive to do so).² Our reason for studying common interest games is that in this class of games this argument seems to be most powerful, and so reputation has the best opportunity to work effectively.

We show, however, that if *mixed strategies* are permitted, then credible and severe punishments are still possible. Our main result, Proposition 3, establishes that in a wide class of repeated common interest games, as discounting goes to zero and as the prior probability attached to the commitment type goes to zero, the normal type of player 1 can be driven close to her minmax payoff. Hence this is a continuity result with the complete information game as the probability of the perturbation goes to zero. Mixed strategies play an important role in the construction because a randomization by the normal type of player 1 between the cooperative and some other action, can allow her to credibly punish player 2 if the latter deviates. Specifically, if she randomises and player 2 deviates, there is a probability that player 2 deviates simultaneously with player 1 revealing herself to be the normal type; if this happens, the continuation game is a complete information game where severe punishments are credible. Moreover, an equilibrium in which player 1 puts positive probability on an action other than the cooperative one, need not imply that she receives a continuation payoff equal to the cooperation payoff, should she play the cooperative action. In that case, player 2 will

² Formally, it is easy to establish that for a fixed probability of the commitment type, and for a given ϵ , there is a threshold discount factor above which all pure-strategy perfect Bayesian equilibrium payoffs are within ϵ of the dominant payoff pair. It should be noted, however, that this depends on the assumption that the perturbation *only* involves the above described automaton.

revise upwards the probability he attaches to facing the commitment type, but not to one, so the continuation payoff need not equal the cooperation payoff. Consequently punishment can be threatened by player 1 in a way which does not imply cooperation payoffs thereafter.

Although this result is, in the context of the reputation literature, a negative one, we see it additionally as a first step towards investigating perfect equilibria in general incomplete information games. This is of interest because, to our knowledge, nothing is known about general properties of the equilibrium payoff set (Nash or perfect) of general discounted incomplete information games as discounting becomes small.³ This is in contrast to the undiscounted case where complete characterizations exist, although only for Nash equilibria (see Forges, 1992).

2. THE MODEL AND RESULTS

We begin by describing a broad class of common interest games. When these games are infinitely repeated, with both players discounting the future with the same factor δ ($0 < \delta < 1$), there is a large set of possible equilibrium outcomes. In particular, given any pair (g_1, g_2) of feasible strictly individually rational payoffs, there is a discount factor $\underline{\delta}$ such that for all $\delta > \underline{\delta}$ there exists a subgame perfect equilibrium with the payoffs (g_1, g_2) . The repeated common interest games we consider are then perturbed so that player 1 is either a "normal" type, or a commitment type, hereafter "automaton", that always plays the Pareto optimal action. Player 2 has prior beliefs that attach a probability μ to player 1 being the automaton and a probability $1 - \mu$ to her being a normal type. In our main result, Proposition 3, we show that given any $\omega > 0$, there exists a $\underline{\delta}$ ($0 < \underline{\delta} < 1$) and $\underline{\mu} > 0$ such that for any $\delta > \underline{\delta}$ and $\mu < \underline{\mu}$ there is a perfect Bayesian equilibrium in which the normal type receives a payoff within ω of her minmax payoff.

2.1. A Class of Common Interest Games

³ Bergin (1989) shows that sequential equilibria have a Markov property; unfortunately this result does not directly have bearing upon the set of payoffs which can be realised in equilibrium. The same can be said for the results of Kalai and Lehrer (1993) and Jordan (1995) who have studied the long-run properties of equilibrium play in contexts more general than the current one.

In this subsection we shall describe the class of common interest games that are studied in this paper. First we define some notation. A finite 2-player game in strategic form is denoted by

$$g = (g_1, g_2) : A_1 \times A_2 \rightarrow \mathbb{R}^2,$$

where A_i is player i 's finite action space (we assume $\#A_i = 2$, $i=1, 2$) and g_i is player i 's payoff function, $i=1, 2$. Let $a := (a_1, a_2)$ denote an action profile for the two players and $A := A_1 \times A_2$ be the set of all action profiles. The convex hull of all payoffs is the set $G := \text{co}\{ (g_1(a), g_2(a)) \mid a \in A \}$.⁴ Let M be a positive number that bounds the payoffs of the players: $M = \max\{|g_i(a)| \mid a \in A, i=1, 2\}$. Also let the pair (\bar{g}_1, \bar{g}_2) denote the players' minmax payoffs:

$$\bar{g}_i := \min_{\alpha_j} \max_{\alpha_i} E_{\alpha_i, \alpha_j} g_i(a_1, a_2), \quad j \neq i, i = 1, 2,$$

where α_i is a mixed action for player i . Define the set of feasible and strictly individually rational payoffs to be $G^* := G \cap \{ (g_1, g_2) \in \mathbb{R}^2 \mid g_1 > \bar{g}_1, g_2 > \bar{g}_2 \}$.

We consider a class of common interest games, that is a class of games with a strongly Pareto dominant *payoff* pair (Aumann and Sorin (1989)), although we shall restrict attention to games in which the payoff vector to one pair of *actions* strictly Pareto dominates all others. Let $(a_1^*, a_2^*) \in A$ denote the action pair corresponding to the Pareto dominant pair, that is $g_1^* := g_1(a_1^*, a_2^*) > g_1(a)$ and $g_2^* := g_2(a_1^*, a_2^*) > g_2(a)$ for all $a \in A$ where $a \neq (a_1^*, a_2^*)$. We make three assumptions about the structure of the payoffs. These assumptions place some limits on the generality of our results but simplify the arguments considerably.⁵

(i) Let $\hat{A}_2 \subset A_2$ be the set of actions for player 2 that give player 1 no more than \bar{g}_1 if she plays her Pareto optimal action: $\hat{A}_2 := \{a_2 \in A_2 \mid g_1(a_1^*, a_2) = \bar{g}_1\}$. (By the definition of \bar{g}_1 the set \hat{A}_2 is non-empty.) The set \hat{A}_2 could be interpreted as the set of possible punishments for player 1 if she is playing a_1^* . The first assumption we make is that action a_1^* is not always the unique best response to an action in \hat{A}_2 ; that is, for some $\hat{a}_2 \in \hat{A}_2$, there exists $\hat{a}_1 \neq a_1^*$ such that

⁴ $\text{co}(X)$ denotes the convex hull of the set X .

⁵ Assumption (ii) can be relaxed at the cost of some additional complications and an appropriate reformulation of Proposition 3. Assumption (iii) is also not essential; the case where the feasible set is one-dimensional was treated in an earlier version, although some of the constructions needed differ. We conjecture that assumption (i) is likewise inessential, although we have not proved this.

$$(1) \quad g_1(\hat{a}_1, \hat{a}_2) = g_1(a_1^*, \hat{a}_2), \quad \text{for all } a_1 \in A_1.$$

A sufficient (though by no means necessary) condition for this is if action a_1^* does not ensure player 1 her minmax payoff in the game, that is if $\min_{a_2 \in A_2} g_1(a_1^*, a_2) < \bar{g}_1$. Henceforth (\hat{a}_1, \hat{a}_2) will refer to a fixed action pair, with $\hat{a}_1 \neq a_1^*$ and $\hat{a}_2 \in \hat{A}_2$, which satisfies (1). The payoffs when actions (\hat{a}_1, \hat{a}_2) are taken will be denoted $\hat{g}_1 := g_1(\hat{a}_1, \hat{a}_2)$, $\hat{g}_2 := g_2(\hat{a}_1, \hat{a}_2)$.

(ii) Our second assumption is that there exists feasible and individually rational payoffs that hold both players down to their minmax levels. That is,

$$(2) \quad \exists (\tilde{g}_1, \tilde{g}_2) \in G \text{ such that (a) } \tilde{g}_i = \bar{g}_i, \quad (b) \tilde{g}_j = \bar{g}_j, \quad i=1,2, j \neq i$$

(iii) Our third assumption is that the set G^* has a non-empty interior.

[Figure 1 about here]

Given the second and third assumptions above, the set G has the form shown in Figure 1. The dashed line between $(\tilde{g}_1, \tilde{g}_2)$ (see (2)) and (g_1^*, g_2^*) will be used in the construction of an equilibrium. This line will be described by the equation $g_2 = \alpha + \beta g_1$, where $\beta > 0$.

2.2. The Repeated Game of Complete Information

The game in strategic form described above is played in the periods $t=0,1,2,\dots$. In each period, players are aware of all (pure) actions taken in previous periods. Player i 's payoff in this infinitely repeated game is given by the expected discounted sum of its normalized stage-game payoffs, $E(1 - \delta) \sum_{t=0}^{\infty} \delta^t g_i(a_1^t, a_2^t)$ ($i=1,2$), where $a^t := (a_1^t, a_2^t) \in A$ is the players' action profile in period t , δ is their common discount factor ($0 < \delta < 1$), and E denotes expectations. We will let $G(\delta)$ denote the infinitely repeated game of complete information. Given our assumptions on the structure of payoffs in the stage game, the Perfect Folk Theorem applies to $G(\delta)$ provided δ is sufficiently close to one. By Fudenberg & Maskin (1991) the following result holds for the repeated game of complete information $G(\delta)$ when only pure strategies are observed and there is no public randomisation.

Result 1 (Fudenberg & Maskin (1991)) : *For any $(g_1, g_2) \in G^*$ there exists $\underline{\delta} < 1$ such that for all $1 > \delta > \underline{\delta}$ there is a subgame perfect equilibrium of $G(\delta)$ in which player i 's average payoff is g_i .*

In general, the lower bound $\underline{\delta}$ in Result 1 will vary with the point $(g_1, g_2) \in G^*$ that is being sustained as the equilibrium payoff vector. This is because the threshold $\underline{\delta}$ varies with the threat point $(g_1', g_2') < (g_1, g_2)$ used in the proof. By considering those payoff pairs $(g_1, g_2) \in G^*$ that can be supported as equilibrium payoffs using a fixed threat point (g_1', g_2') , the following corollary to Result 1 is immediate.

Corollary : *Let $\epsilon > 0$ be given, and define $G_{\epsilon}^* := G \leftrightarrow \{ (g_1, g_2) \in \mathbb{R}^2 \mid g_1 = \bar{g}_1 + \epsilon, g_2 = \bar{g}_2 + \beta \epsilon \}$; then provided G_{ϵ}^* is non-empty, there is a $\underline{\delta}_{\epsilon} < 1$ such that for all $\underline{\delta}_{\epsilon} < \delta < 1$ and any $(g_1, g_2) \in G_{\epsilon}^*$ there is a subgame perfect equilibrium of $G(\delta)$ in which player i 's average payoff is g_i .*

(Recall that the parameter $\beta > 0$ is the slope of the dashed line in Figure 1.)

2.3. The Perturbed Repeated Game

We now introduce a perturbation of the repeated game of common interests $G(\delta)$ described above. Before the play commences there is a move of nature, the outcome of which is not observed by player 2. With probability $1 - \mu$, nature selects player 1 to be a type with payoffs as described above, and with probability μ , nature selects a player 1 to be a type that always plays action a_1^* independently of history. We will call the first type of player 1 "the normal type" and the second type of player 1 "the automaton". As player 2 does not observe nature's move, this gives a repeated game of one-sided incomplete information which we will denote $G(\mu, \delta)$ and we will study the perfect Bayesian equilibria (PBE's) of this game.⁶

⁶ The automaton can also be thought of as a type with a standard payoff matrix in which the payoffs in the row corresponding to a_1^* are all equal and strictly greater than all other payoffs. At a PBE this type will play a_1^* after every history, including those off the equilibrium path.

We adopt the definition of perfect Bayesian equilibrium given by Fudenberg & Tirole (1991a), which in this context amounts to the following. If h^t is any history of actions taken by both players up to and including period t , then given player 2's beliefs about facing the automaton, say $\mu(h^t)$, at the start of period $t+1$, strategies must yield a Bayesian Nash equilibrium for the continuation game.⁷ Moreover Bayes' rule is used to update beliefs whenever possible, that is, $\mu(h^{t+1})$ is derived from $\mu(h^t)$ by Bayes' rule whenever player 1 plays an action at period t which player 2 had expected to be played with positive probability.

Proposition 1 exploits the natural recursive structure of the repeated games of incomplete information $G(\mu, \delta)$ to determine a relationship between a PBE of $G(\mu, \delta)$ and a PBE of $G(\pi\mu, \delta)$ where $\pi < 1$. The principal idea of the proof is very simple. It takes as given a PBE of $G(\mu, \delta)$ with payoffs (γ_1, γ_2) to the normal type and player 2 respectively, and uses this PBE to construct a PBE of $G(\pi\mu, \delta)$. In the first period of play in $G(\pi\mu, \delta)$ the normal type of player 1 randomises, playing a_1^* with probability $q = \pi(1-\mu)/(1-\pi\mu)$ and \hat{a}_1 with probability $1-q$ (where \hat{a}_1 is defined below (1)). Player 2 plays \hat{a}_2 in the first period. Conditional upon observing a_1^* in the first period, player 2 will revise his priors (about player 1 being an automaton) by Bayes's Theorem to precisely π (given our choice of q). Thus if (a_1^*, \hat{a}_2) is played in the first period, we specify that the PBE of $G(\mu, \delta)$ is then played out subsequently, with payoffs (γ_1, γ_2) . In this case the expected payoff to the normal type is $(1-\delta)g_1(a_1^*, \hat{a}_2) + \delta\gamma_1$. In order for randomization for the normal type to be optimal in the first period, she must be indifferent between this payoff and what she would receive from playing \hat{a}_1 in the first period. After the first period history (\hat{a}_1, \hat{a}_2) , however, she reveals herself to be the normal type, so the players are in the complete information game $G(\delta)$. Thus it is necessary that an equilibrium of $G(\delta)$ can be chosen which makes player 1 indifferent (this equilibrium can also be used as the continuation after all actions of player 1 other than a_1^* since \hat{a}_1 is a best response to \hat{a}_2 , so the normal type will not wish to deviate). In addition, the continuation equilibrium must be selected so that it is

⁷ The reader is referred to Fudenberg and Tirole (1991a, 1991b) for formal definitions of all equilibrium concepts used here. For the purpose of the definitions, the automaton should be interpreted as a payoff-matrix type as described in footnote 6; the strategy of such a type in a PBE must be identical to the automaton strategy. This is in contrast to a Nash equilibrium where the payoff-matrix type need only follow the commitment strategy on the equilibrium path. In the equilibria we construct, beliefs off-the-equilibrium path put probability zero on the automaton if player 1 has deviated from a_1^* in the past, which is consistent with the idea of an automaton which cannot deviate.

optimal for player 2 to choose \hat{a}_2 in the first period; this requires that another equilibrium of $G(\delta)$ can be chosen as the continuation after (\hat{a}_1, a_2^0) where $a_2^0 \succ \hat{a}_2$, which is sufficiently severe to prevent player 2 from deviating. Provided these two equilibria can be constructed, a PBE of $G(\pi\mu, \delta)$ has been found with the payoff $(1-\delta)g_1(a_1^*, \hat{a}_2) + \delta\gamma_1$ to the normal type. Since by construction $g_1(a_1^*, \hat{a}_2) = \bar{g}_1$, it follows that the equilibrium of $G(\pi\mu, \delta)$ has a lower payoff for the normal type than the equilibrium of $G(\mu, \delta)$, a property that will, by repeated application of Proposition 1, permit the construction of a PBE with payoffs for the normal type arbitrary close to her minmax payoff.

Define $\underline{\varepsilon} > 0$ to be such that $G_{\underline{\varepsilon}}^*$ is nonempty. Then we can state:

Proposition 1: *Let $\varepsilon, 0 < \varepsilon < \underline{\varepsilon}$, and $\delta > \underline{\delta}_\varepsilon$ be given. Also let (γ_1, γ_2) be the expected payoffs to the normal type of player 1 and to player 2 at a PBE for $G(\mu, \delta)$. Then $G(\pi\mu, \delta)$ has a PBE where the normal type of player 1 receives the expected payoff $(1-\delta)g_1(a_1^*, \hat{a}_2) + \delta\gamma_1$ provided π (with $0 < \pi < 1$) and γ_1 satisfy*

$$(3) \quad \gamma_1 = \bar{g}_1 + (1-\delta)\delta^{-1}(\hat{g}_1 - g_1(a_1^*, \hat{a}_2) + 2M) + \varepsilon,$$

$$(4) \quad \frac{\pi}{1-\pi} \leq \frac{\delta(\alpha + \beta\gamma_1 - \bar{g}_2 - \beta\varepsilon) - (1-\delta)(g_2^* - \hat{g}_2 + \beta(\hat{g}_1 - g_1(a_1^*, \hat{a}_2)))}{(1-\delta)(g_2^* - g_2(a_1^*, \hat{a}_2))}.$$

Proof: See Appendix.

The above proposition allows us to generate an equilibrium for $G(\pi\mu, \delta)$ using an equilibrium of $G(\mu, \delta)$, provided the bounds (3) and (4) on γ_1 and π are satisfied. Proposition 2 below repeatedly applies Proposition 1. The first step is to describe a PBE for $G(1, \delta)$, which is the complete information game between player 2 and the automaton. Player 2's best response to the automaton is to play the action a_1^* in every period, so in $G(1, \delta)$ there is a PBE where the players play (a_1^*, a_2^*) in every period. This equilibrium is used as a starting point for repeated applications of Proposition 1. The next step is to apply Proposition 1 to this PBE to find a PBE for $G(\pi^1, \delta)$ (where $\pi^1 < 1$); at this PBE player 2 plays \hat{a}_2 for one period and then $G(1, \delta)$ is played if $a_1^0 = a_1^*$. The whole process can be repeated by applying Proposition 1 to the PBE of $G(\pi^1, \delta)$ to find a new PBE for $G(\pi^1\pi^2, \delta)$ (where $\pi^2 < 1$); at this PBE \hat{a}_2 is played for two periods before play settles on (a_1^*, a_2^*) in

$G(1, \delta)$. Proposition 2 repeatedly applies Proposition 1 in this fashion until some step $N(\delta)+1$ where the constraint (3) is finally violated. At this last equilibrium \hat{a}_2 is played $N(\delta)+1$ periods against the automaton, and in period $N(\delta)+2$ play finally settles on (a_1^*, a_2^*) . The process described above thus generates a finite family of equilibria for the sequence of games $G(\mu, \delta)$ with priors $\mu=1 \supseteq \pi^1 \supseteq \pi^2 \dots \pi^{N(\delta)+1} \supseteq \pi^n$, $n=0, 1, \dots, N(\delta)+1$.

The PBE described in Proposition 2 is parameterised by three sequences: $\{\gamma^n\}_{n=0}^{N(\delta)}$, $\{\pi^n\}_{n=0}^{N(\delta)}$ and $\{\mu^n\}_{n=0}^{N(\delta)}$. The terms of all of these sequences depend on δ , although this dependence is suppressed in the notation. The sequences will be defined inductively, because Proposition 1 describes a relationship between their adjacent terms. Suppose we have found an equilibrium for the game $G(\mu^n, \delta)$, where player 1's normal type has an expected payoff of γ^n ; then Proposition 1 determines an equilibrium for the game $G(\mu^{n+1}, \delta)$ (where $\mu^{n+1} = \mu^n \pi^{n+1}$) where player 1's normal type gets the payoff $\gamma^{n+1} = (1-\delta)g_1(a_1^*, \hat{a}_2) + \delta\gamma^n$. Thus, given the pair (γ^n, μ^n) , Proposition 1 determines the parameters $(\gamma^{n+1}, \mu^{n+1}, \pi^{n+1})$. The initial values of these sequences are determined so γ^0 is player 1's equilibrium payoff at the equilibrium of $G(1, \delta)$ described above: $\mu^0=1$, $\gamma^0=g_1^*$. The following recursion describes how the successive terms $(\gamma^{n+1}, \mu^{n+1}, \pi^{n+1})$ are generated:

$$(5) \quad \gamma^{n+1} = (1-\delta)g_1(a_1^*, \hat{a}_2) + \delta\gamma^n,$$

$$(6) \quad \mu^{n+1} = \mu^n \pi^{n+1},$$

$$(7) \quad \frac{\pi^{n+1}}{1-\pi^{n+1}} = \frac{\delta(\alpha + \beta\gamma^n - \bar{g}_2 - \beta\varepsilon) - (1-\delta)(g_2^* - \hat{g}_2 + \beta(\hat{g}_1 - g_1(a_1^*, \hat{a}_2)))}{(1-\delta)(g_2^* - g_2(a_1^*, \hat{a}_2))}.$$

Proposition 2: *Let ε , $0 < \varepsilon < \underline{\varepsilon}$, and $\delta > \underline{\delta}_\varepsilon$ be given and let $N(\delta)$ be the largest positive integer (if one exists) such that*

$$(8) \quad (1-\delta^{N(\delta)})g_1(a_1^*, \hat{a}_2) + \delta^{N(\delta)}g_1^* = \bar{g}_1 + (1-\delta)\delta^{-1}(\hat{g}_1 - g_1(a_1^*, \hat{a}_2)) + \varepsilon.$$

Then for $n=1, 2, \dots, N(\delta)+1$, if $\mu=\mu^n$ there exists a PBE of $G(\mu, \delta)$ where the normal type of player 1's payoff is γ^n .

Proof: We have shown that $\gamma^0 = g_1^*$ is a PBE payoff for the normal type of player 1 in $G(1, \delta)$. It is also true that g_1^* is a PBE payoff for the normal type of player 1 in $G(\mu, \delta)$ for all μ . Thus the proposition is true when $n=0$. Now suppose the proposition is true for $n=n'$ where $n'=N(\delta)$. As the proposition is true for $n=n'$, the game $G(\mu^{n'}, \delta)$ has a PBE where the normal type of player 1 receives the payoff $\gamma^{n'} = (1-\delta^{n'})g_1(a_1^*, \hat{a}_2) + \delta^{n'}g_1^*$. Apply Proposition 1 to this equilibrium; $\gamma^{n'}$ satisfies (3) because $n'=N(\delta)$. Set $\gamma_1 = \gamma^{n'}$ in (4). The largest value for π that satisfies (4) will solve (4) with equality. This defines $\pi^{n'+1}$ as in (7). Hence from Proposition 1, if $\mu = \mu^{n'+1} = \pi^{n'+1}\mu^{n'}$ the game $G(\mu, \delta)$ has a PBE where the normal type's payoff is $\gamma^{n'+1} = (1-\delta)g_1(a_1^*, \hat{a}_2) + \delta\gamma^{n'}$. *Q.E.D.*

Proposition 2 goes a long way towards achieving the result described in the introduction, because it shows that for any δ we can find a μ such that the game $G(\mu, \delta)$ has a PBE where the normal type gets approximately $\bar{g}_1 + (1-\delta)\delta^{-1}(\hat{g}_1 - g_1(a_1^*, \hat{a}_2)) + \varepsilon$. As δ becomes close to unity, therefore, the normal type's payoff can be made within ε of her minmax level \bar{g}_1 . We want a stronger result, however, so that given ε , there are threshold values for δ and for μ such that *for all* δ bigger than its threshold value and *all* μ less than its (strictly positive) threshold value, equilibrium payoffs within ε of \bar{g}_1 exist. It is therefore necessary that we consider how the family of equilibria described in Proposition 2 varies as δ approaches unity for a given value $\varepsilon > 0$. The equilibrium payoffs in Proposition 2 define a piecewise continuous function $\hat{\gamma}_\delta(\mu)$ where $\hat{\gamma}_\delta(\mu) = \gamma^n$ for $\mu^{n+1} < \mu \leq \mu^n$. The function $\hat{\gamma}_\delta(\mu)$ describes how the payoffs at the PBE we construct are related to the priors. The sequence $\{\mu^n, \gamma^n\}_{n=0}^{N(\delta)}$ determines the properties of the function $\hat{\gamma}_\delta(\mu)$ and these are both shown in Figure 2. (The line labelled γ_δ is referred to in the proof of Lemma 1.) We are particularly interested in how $\hat{\gamma}_\delta(\mu)$ behaves as $\delta \nearrow 1$, and as this happens the figure changes in two ways. First, $N(\delta)$ becomes arbitrarily large and each individual line segment becomes arbitrarily short. Secondly, the points (μ^n, γ^n) become closer together, with $\|(\mu^n, \gamma^n) - (\mu^{n+1}, \gamma^{n+1})\| \searrow 0$, and so the step sizes shrink. The following technical lemma shows that on the interval $(0, 1]$ the step function in the picture converges uniformly to a continuous function $\gamma^*(\mu)$. Moreover, this limiting function $\gamma^*(\mu)$ is continuously differentiable for all but one value of μ .

[Figure 2 about here]

Lemma 1: *If $\varepsilon < (\alpha/\beta) + g_1(a_1^*, \hat{a}_2) - (\bar{g}_2/\beta)$, then as $\delta \searrow 1$ the function $\hat{\gamma}_\delta(\mu)$ converges uniformly to a continuous decreasing function $\gamma^*(\mu)$ on $(0, 1]$, where*

$$(9) \quad \gamma^*(\mu) = \max \left\{ \bar{g}_1 + \varepsilon, \quad g_1(a_1^*, \hat{a}_2) + \left(\frac{A\mu^\kappa}{1 - A\beta\mu^\kappa} \right) \kappa (g_2^* - g_2(a_1^*, \hat{a}_2)) \right\},$$

$$\kappa = \left(\frac{\alpha + \beta g_1(a_1^*, \hat{a}_2) - \bar{g}_2 - \beta \varepsilon}{g_2^* - g_2(a_1^*, \hat{a}_2)} \right) \quad \text{and} \quad A = \left(\frac{g_1^* - g_1(a_1^*, \hat{a}_2)}{g_2^* - \bar{g}_2 - \beta \varepsilon} \right).$$

Proof: See Appendix.

Lemma 1 describes the properties of the equilibria as $\delta \searrow 1$. We have shown that as $\delta \searrow 1$ the function $\gamma^*(\mu)$ is a good approximation for the payoffs at a PBE of $G(\mu, \delta)$. Moreover, the function $\gamma^*(\mu)$ can be made arbitrarily close to \bar{g}_1 (by varying ε) at some strictly positive value of μ . This is now used to prove the main result: there exists an equilibrium where the normal type gets a payoff arbitrarily close to \bar{g}_1 for all games where the players are sufficiently patient and the probability of the automaton is small.

Proposition 3: *For any $\omega > 0$ there exists a $\underline{\delta}_\omega > 0$ and a $\mu_\omega > 0$ such that for any δ, μ satisfying $1 > \delta > \underline{\delta}_\omega$ and $\mu_\omega > \mu > 0$, the game $G(\mu, \delta)$ has a PBE where the normal type of player 1 receives an expected payoff γ_1 within ω of her minmax payoff, that is, satisfying $\bar{g}_1 < \gamma_1 < \bar{g}_1 + \omega$.*

Proof: Let $\omega > 0$ be given and choose $\varepsilon = \omega/2$ (without loss of generality assume $\beta \varepsilon < \alpha + \beta g_1(a_1^*, \hat{a}_2) - \bar{g}_2$); then, since $\omega > \varepsilon$ and using Lemma 1, $\gamma^*(\mu) = \bar{g}_1 + \omega$ is equivalent to

$$\bar{g}_1 + 2\varepsilon = g_1(a_1^*, \hat{a}_2) + \left(\frac{A\mu^\kappa}{1 - A\beta\mu^\kappa} \right) \kappa (g_2^* - g_2(a_1^*, \hat{a}_2)).$$

After some rearrangement this implies the unique solution for μ to $\gamma^*(\mu) = \bar{g}_1 + \omega$ satisfies

⁸ This condition is merely to rule out $\kappa = 0$, which would change the method of solving the differential equation studied below (but not the conclusion). Note that κ can be negative.

$$A\beta\mu^\kappa = 1 - \frac{\alpha + \beta g_1(a^*, \hat{a}_2) - \bar{g}_2 - \beta\epsilon}{\alpha + \beta \bar{g}_1 - \bar{g}_2 + \beta\epsilon}.$$

The quotient is less than unity if and only if $\kappa > 0$, so for all $\kappa > 0$ there is $0 < \mu^* < 1$ that satisfies this equation with equality and $\mu < \mu^*$ if and only if $\gamma^*(\mu) < \bar{g}_1 + \omega$. The functions $\hat{\gamma}_\delta(\mu)$ are non-decreasing and converge uniformly to $\gamma^*(\mu)$ so there exists $\mu_\omega < \mu^*$ and a δ_ω such that provided $\mu < \mu_\omega$ and $\delta > \delta_\omega$, $\hat{\gamma}_\delta(\mu) < \bar{g}_1 + \omega$. *Q.E.D.*

Remark: Although Proposition 3 establishes that player 1 can be held close to her worst payoff, it is easy to show under the same assumptions that equilibria can be constructed in which she receives (approximately) any payoffs between \bar{g}_1 and g_1^* : in Proposition 1, in addition to constructing an equilibrium with payoffs less than γ_1 , an equilibrium with payoffs *equal* to γ_1 can be constructed. Using this repeatedly, as δ goes to one in Proposition 3, all points to the "left" (see Figure 2) of the limiting function, $\gamma^*(\mu)$, can be approximated by equilibria.

3. CONCLUDING COMMENTS

We have shown that small perturbations of a large class of common interest games, in which one of the players might be a type committed to playing in a cooperative fashion, do not rule out low payoffs, even when sequential rationality is imposed on the equilibrium concept. In a broader context, these results also have implications for the reputation literature following Fudenberg and Levine (1989), which considers games between a long-run and a sequence of short-run players, perturbed with the possibility that the long-run player might be committed to some fixed action. Their results were extended to games with two long-run players by Schmidt (1993) for "conflicting interest games", and, for general stage games, by Cripps *et al.* (1996). The latter paper develops a lower bound on the Nash equilibrium payoffs of the informed player which is applicable to the class of games studied here, but the result applies only *if the informed player is arbitrarily patient relative to the uninformed player*. It is certainly the case that in some common interest games satisfying our conditions, this lower bound is above the informed player's minmax payoff. Hence our results imply that *with*

symmetric discounting no such lower bound exists in this class of common interest games, even when perfection is imposed.⁹

REFERENCES

AUMANN, R. J., AND SORIN, S. (1989). "Cooperation and Bounded Recall," *Games and Economic Behavior* **1**, 5-39.

BERGIN, J. (1989). "A Characterization of Sequential Equilibrium Strategies in Infinitely Repeated Incomplete Information Games," *Journal of Economic Theory* **47**, 51-65.

CELENTANI, M., FUDENBERG, D., LEVINE D. K., AND PESENDORFER, W. (1994). "Maintaining a Reputation Against a Long-Lived Opponent," *Econometrica* **64**, 691-704.

CRIPPS, M. W., K. M. SCHMIDT, AND J. P. THOMAS (1996). "Reputation in Perturbed Repeated Games," *Journal of Economic Theory* **69**, 387-410.

FORGES, F. (1992). "Non-Zero-Sum Repeated Games of Incomplete Information," in *Handbook of Game Theory, Vol 1* (R.J.Aumann, and S.Hart, Eds.). Amsterdam: North Holland.

FUDENBERG, D., AND LEVINE, D. K. (1989). "Reputation and Equilibrium Selection in Games with a Patient Player," *Econometrica* **57**, 759-778.

FUDENBERG, D., AND MASKIN, E. (1986). "The Folk Theorem in Repeated Games with Discounting and with Incomplete Information," *Econometrica* **54**, 533-554.

FUDENBERG, D., AND MASKIN, E. (1991). "On the Dispensability of Public Randomization in Discounted Repeated Games," *Journal of Economic Theory* **53**, 428-438.

FUDENBERG, D., AND TIROLE, J. (1991a). "Perfect Bayesian Equilibrium and Sequential Equilibrium," *Journal of Economic Theory* **53**, 236-260.

FUDENBERG, D., AND TIROLE, J. (1991b). *Game Theory*. Cambridge, MA: MIT Press.

JORDAN, J. S. (1995). "Bayesian Learning in Repeated Games," *Games and Economic Behavior*, forthcoming.

⁹ An example is constructed in Celentani *et al.* (1996) which is similar to the type of construction we use, and which establishes that payoffs below the Stackelberg payoff can be sustained in a PBE. For their game, however, the Cripps *et al.* (1996) bound is just the minmax payoff.

KALAI, E., AND LEHRER, E. (1993). "Rational Learning Leads to Nash Equilibrium," *Econometrica* **61**, 1019-1045.

SCHMIDT, K.M. (1993). "Reputation and Equilibrium Characterization in Repeated Games of Conflicting Interests," *Econometrica* **61**, 325-352.

APPENDIX

Proof of Proposition 1: The players' equilibrium strategies for the game $G(\mu, \delta)$ are described below.

- Player 1:** Play a_1^* with probability q and \hat{a}_1 with probability $(1-q)$ in period zero, where $q = \pi(1-\mu)/(1-\pi\mu)$.
 If player 1's period zero action (a_1^0) was a_1^* , then from period one play the PBE that gives the payoffs (γ_1, γ_2) (i.e. deviations by player 2 in period zero are ignored when $a_1^0 = a_1^*$).
 If (\hat{a}_1, \hat{a}_2) is played in period zero then from period one play a subgame perfect equilibrium for $G(\delta)$ to achieve the payoffs $(x, y) = G_\epsilon^*$ (where (x, y) is described below).
 If $a_1^0 = \hat{a}_1$ and $a_2^0 \neq \hat{a}_2$, then play out a subgame perfect equilibrium for $G(\delta)$ that gives player 2 a payoff $\bar{g}_2 + \beta\epsilon$.
- Player 2:** Play \hat{a}_2 in period zero.
 If $a_1^0 = a_1^*$, then play the PBE that gives the payoffs (γ_1, γ_2) .
 If $a_1^0 \neq a_1^*$ (and $a_2^0 = \hat{a}_2$), then play the subgame perfect equilibrium for $G(\delta)$ to achieve the payoffs $(x, y) = G_\epsilon^*$.

By the corollary to Result 1 we can specify an equilibrium for the game $G(\delta)$ that gives any payoffs $(x, y) = G_\epsilon^*$. Below we will place some further restrictions on (x, y) , but at the moment we will note that the pair (x, y) is always restricted to be on the line $g_2 = \alpha + \beta g_1$ shown in Figure 1.

The players' strategies are optimal after period zero: The strategy for player 1 requires her to randomize in period zero; therefore, at a PBE, player 2's priors conditional on $a_1^0 = a_1^*$ will be revised upwards. The choice of q above ensures that his revised priors attach probability μ to the automaton type. Thus conditional on $a_1^0 = a_1^*$, the game $G(\mu, \delta)$ is played from period one onwards. Since (γ_1, γ_2) are by assumption payoffs at a PBE of $G(\mu, \delta)$, the strategies described above certainly

constitute an equilibrium given any history with $a_1^0 = a_1^*$. Moreover, after any history with $a_1^0 \neq a_1^*$, the continuation payoffs correspond to equilibria in $G(\delta)$.¹⁰

The players' strategies are optimal in period zero: Player 1's strategy for period zero is optimal provided she is indifferent between the actions a_1^* and \hat{a}_1 , and provided all other actions in period zero give her a smaller payoff. She is indifferent between a_1^* and \hat{a}_1 if $(1-\delta)\hat{g}_1 + \delta x = (1-\delta)g_1(a_1^*, \hat{a}_2) + \delta\gamma_1$. An action $a_1 \neq a_1^*, \hat{a}_1$ gives her a payoff $(1-\delta)g_1(a_1, \hat{a}_2) + \delta x$, and, by the definition of \hat{a}_1 , $(1-\delta)\hat{g}_1 + \delta x = (1-\delta)g_1(a_1, \hat{a}_2) + \delta x$ for all $a_1 \in A_1$. Indifference between a_1^* and \hat{a}_1 implies

$$(A1) \quad x = \gamma_1 - (1-\delta)\delta^{-1}(\hat{g}_1 - g_1(a_1^*, \hat{a}_2)).$$

But $\hat{g}_1 = g_1(a_1^*, \hat{a}_2)$ and $\gamma_1 = g_1^*$, so (3) and (A1) imply together $x \in [\bar{g}_1 + \varepsilon, g_1^*]$. Thus it is possible to find $(x, y) \in G_\varepsilon^*$ so that $y = \alpha + \beta x$, and hence the above strategy for player 1 is optimal.

Player 2's strategy in period zero is optimal provided his expected payoff from playing \hat{a}_2 exceeds that from any other action. He attaches probability $\pi\mu + (1-\pi\mu)q = \pi$ to player 1 playing action a_1^* and probability $1-\pi$ to her playing \hat{a}_1 , and therefore if he plays an action $a_2^0 \neq \hat{a}_2$ his payoff is bounded above by $(1-\delta)g_2^* + \delta\{\pi\gamma_2 + (1-\pi)(\bar{g}_2 + \beta\varepsilon)\}$, whereas his payoff from the action \hat{a}_2 is $(1-\delta)\{\pi g_2(a_1^*, \hat{a}_2) + (1-\pi)\hat{g}_2\} + \delta\{\pi\gamma_2 + (1-\pi)y\}$. Thus his period zero strategy is optimal provided

$$(A2) \quad \delta(y - \bar{g}_2 - \beta\varepsilon) - (1-\delta)(g_2^* - \hat{g}_2) = \pi\{\delta(y - \bar{g}_2 - \beta\varepsilon) - (1-\delta)(g_2^* - \hat{g}_2) + (1-\delta)(g_2^* - g_2(a_1^*, \hat{a}_2))\}.$$

Next, we shall show that the RHS of (A2) is strictly positive, and the LHS is weakly positive and smaller than the term in braces on the RHS. It is sufficient that $\delta(y - \bar{g}_2 - \beta\varepsilon) - (1-\delta)(g_2^* - \hat{g}_2) = 0$ for this to be true. Since $y = \alpha + \beta x$ and the value of x is determined by (A1), some substitution and rearranging of this latter condition gives

$$\alpha + \beta\gamma_1 = \bar{g}_2 + \beta\varepsilon + (1-\delta)\delta^{-1}\{\beta(\hat{g}_1 - g_1(a_1^*, \hat{a}_2)) + (g_2^* - \hat{g}_2)\}.$$

Make the following substitutions: $\bar{g}_2 = \alpha + \beta\bar{x}$ (for some $\bar{x} = \bar{g}_1$), $g_2^* = \alpha + \beta g_1^*$, and $\hat{g}_2 = \alpha + \beta\hat{x}$ (for some $\hat{x} < g_1^*$). Since $\beta > 0$, by common interests, the above expression now becomes

$$(A3) \quad \gamma_1 = \bar{x} + \varepsilon + (1-\delta)\delta^{-1}\{(\hat{g}_1 - g_1(a_1^*, \hat{a}_2)) + (g_1^* - \hat{x})\}.$$

Condition (3) is sufficient for (A3), since $2M > g_1^* - \hat{x}$, and therefore the RHS of (A2) is strictly positive ($g_2^* > g_2(a_1^*, \hat{a}_2)$) and the LHS is weakly positive and smaller than the term in braces on the RHS.

Using this finding, (A2) is thus equivalent to

¹⁰ If $a_1^0 \neq a_1^*, \hat{a}_1$, then Bayes' rule does not tie down beliefs, and in this case we assume that probability one is put on the normal type; given that the automaton can be thought of as a type for whom it is a dominant strategy to play a_1^* , this is a natural assumption.

$$(A4) \quad \frac{\pi}{1-\pi} \lesssim \frac{\delta(y-\bar{g}_2-\beta\epsilon) - (1-\delta)(g_2^*-\hat{g}_2)}{(1-\delta)(g_2^*-g_2(a_1^*, \hat{a}_2))}.$$

Again, substitute $y=\alpha+\beta x$ and for x from (A1), and this gives the condition (4). We have established that the strategies outlined above constitute an equilibrium for the game $G(\pi, \mu, \delta)$. *Q.E.D.*

Proof of Lemma 1: If the pairs (μ^n, γ^n) $n=0, 1, \dots, N(\delta)+1$ are joined with line segments and the final pair $(\mu^{N(\delta)+1}, \gamma^{N(\delta)+1})$ is joined to $(0, \gamma^{N(\delta)+1})$ with a line segment, this gives a piecewise-linear function $\gamma_\delta(\mu)$ plotted in (μ, γ) -space (shown in Figure 2). The slope of the line segments is given by the ratio $(\gamma^n - \gamma^{n+1})/(\mu^n - \mu^{n+1}) = (\gamma_\delta(\mu^n) - \gamma_\delta(\mu^{n+1})) / (\mu^n - \mu^{n+1})$. From (5), (6) and (7) the ratio can be re-written as follows:

$$(A5) \quad \frac{\gamma_\delta(\mu^n) - \gamma_\delta(\mu^{n+1})}{\mu^n - \mu^{n+1}} = \frac{\gamma_\delta(\mu^n) - ((1-\delta)g_1(a_1^*, \hat{a}_2) + \delta\gamma_\delta(\mu^n))}{\mu^n - \pi^n \mu^n}$$

$$= \left(\frac{\gamma_\delta(\mu^n) - g_1(a_1^*, \hat{a}_2)}{\mu^n} \right) \frac{\delta(\alpha + \beta\gamma_\delta(\mu^n) - \bar{g}_2 - \beta\epsilon) - (1-\delta)(g_2(a_1^*, \hat{a}_2) - \hat{g}_2 + \beta(\hat{g}_1 - g_1(a_1^*, \hat{a}_2)))}{(g_2^* - g_2(a_1^*, \hat{a}_2))}$$

Now restrict the function $\gamma_\delta(\mu)$ to the interval $[\eta, 1]$. On this interval the RHS of (A5) has a finite upper bound which gives us

$$\frac{\gamma_\delta(\mu^n) - \gamma_\delta(\mu^{n+1})}{\mu^n - \mu^{n+1}} \lesssim \frac{(g_1^* - g_1(a_1^*, \hat{a}_2))(g_2^* - \bar{g}_2 - \beta\epsilon)}{\eta(g_2^* - g_2(a_1^*, \hat{a}_2))}.$$

Thus the piecewise linear functions $\gamma_\delta(\mu)$ satisfy a Lipschitz condition and by the Ascoli Theorem converge uniformly to a continuous limit as $\delta \searrow 1$. If we let $\gamma^*(\mu)$ denote this limit, then this will also be the limit of the step functions described by Proposition 2.

Eq. (8) also implies that the function $\gamma^*(\mu)$ is continuously differentiable, provided $\gamma^*(\mu) > \bar{g}_1 + \epsilon$. The RHS of (A5) converges to a continuous finite limit as $\delta \searrow 1$, whilst the LHS of (A5) converges to $d\gamma^*/d\mu$. Thus letting $\delta \searrow 1$ we have the differential equation

$$\frac{d\gamma^*}{d\mu} = \left(\frac{\gamma^*(\mu) - g_1(a_1^*, \hat{a}_2)}{\mu} \right) \frac{\alpha + \beta\gamma^*(\mu) - \bar{g}_2 - \beta\epsilon}{(g_2^* - g_2(a_1^*, \hat{a}_2))}.$$

Given ϵ is chosen so that $\alpha + \beta g_1(a_1^*, \hat{a}_2) - \bar{g}_2 > \beta\epsilon$, this can be solved as follows:

$$\frac{d\gamma^*}{(\gamma^*(\mu) - g_1(a_1^*, \hat{a}_2))(\alpha + \beta\gamma^*(\mu) - \bar{g}_2 - \beta\epsilon)} = \frac{d\mu}{\mu(g_2^* - g_2(a_1^*, \hat{a}_2))},$$

which can be rewritten as

$$\frac{d\gamma^*}{(\gamma^*(\mu) - g_1(a_1^*, \hat{a}_2))} - \frac{\beta d\gamma^*}{(\alpha + \beta\gamma^*(\mu) - \bar{g}_2 - \beta\epsilon)} = \frac{d\mu(\alpha + \beta g_1(a_1^*, \hat{a}_2) - \bar{g}_2 - \beta\epsilon)}{\mu(g_2^* - g_2(a_1^*, \hat{a}_2))}.$$

Integrating,

$$\log\left(\frac{\gamma^*(\mu) - g_1(a_1^*, \hat{a}_2)}{\alpha + \beta\gamma^*(\mu) - \bar{g}_2 - \beta\epsilon}\right) = \left(\frac{\alpha + \beta g_1(a_1^*, \hat{a}_2) - \bar{g}_2 - \beta\epsilon}{g_2^* - g_2(a_1^*, \hat{a}_2)}\right) \log(\mu) + K.$$

By construction $\gamma^*(1) = g_1^*$ and this allows us to determine the arbitrary constant K . Given the assumption $\alpha + \beta g_1(a_1^*, \hat{a}_2) - \bar{g}_2 - \beta\epsilon > 0$ it is now possible to solve the above for $\gamma^*(\mu)$, which gives

$$(A6) \quad \gamma^*(\mu) = g_1(a_1^*, \hat{a}_2) + \left(\frac{A\mu^\kappa}{1 - A\beta\mu^\kappa}\right) \kappa(g_2^* - g_2(a_1^*, \hat{a}_2)),$$

where

$$\kappa = \left(\frac{\alpha + \beta g_1(a_1^*, \hat{a}_2) - \bar{g}_2 - \beta\epsilon}{g_2^* - g_2(a_1^*, \hat{a}_2)}\right) > 0 \text{ and } A = \left(\frac{g_1^* - g_1(a_1^*, \hat{a}_2)}{g_2^* - \bar{g}_2 - \beta\epsilon}\right) > 0.$$

Again the assumption $\alpha + \beta g_1(a_1^*, \hat{a}_2) - \bar{g}_2 - \beta\epsilon > 0$ is necessary and sufficient for $\kappa > 0$ and since $(1 - \beta A)(g_2^* - \bar{g}_2 - \beta\epsilon) = \kappa(g_2^* - g_2(a_1^*, \hat{a}_2))$, it is also necessary and sufficient for $1 - \beta A > 0$. In (A6) if $\kappa > 0$ then as $\mu \rightarrow 0$ so $\gamma^*(\mu) \rightarrow g_1(a_1^*, \hat{a}_2) = \bar{g}_1$, whilst if $\kappa < 0$ then $\gamma^*(\mu) \rightarrow \bar{x} + \epsilon$, where $\alpha + \beta\bar{x} = \bar{g}_2$ and $\bar{x} = \bar{g}_1$.

Thus the constraint (8) is binding on the sequence (μ^n, γ^n) . From (8) as $\delta \rightarrow 1$, so the difference $|\gamma^{N(\delta)} - \bar{g}_1 + (1 - \delta)\delta^{-1}(\hat{g}_1 - g_1(a_1^*, \hat{a}_2)) + \epsilon| \rightarrow 0$ and hence the limit of the points (μ^n, γ^n) is the graph of the function described in the Lemma restricted to the domain $[\eta, 1]$.

The limiting function $\gamma^*(\mu)$ does not depend on η , so as η approaches zero the argument still applies, but for $\eta = 0$ it is possible that there is a limiting discontinuity. *Q.E.D.*

FOOTNOTES

We would like to express our gratitude to Klaus Schmidt, an associate editor and an anonymous referee for useful comments. Our thanks are also due to seminar participants at Bonn, Edinburgh, Erasmus, Exeter and Tilburg for comments. All remaining errors are our responsibility.

1. Even the assumption of *two*-sided uncertainty of the type we assume does not force cooperation in Nash equilibrium. Aumann & Sorin (1989) construct a mixed-strategy counter example (to their main result, which assumes pure strategies) in which cooperation is not approximated as the players become patient.
2. Formally, it is easy to establish that for a fixed probability of the commitment type, and for a given ε , there is a threshold discount factor above which all pure-strategy perfect Bayesian equilibrium payoffs are within ε of the dominant payoff pair. It should be noted, however, that this depends on the assumption that the perturbation *only* involves the above described automaton.
3. Bergin (1989) shows that sequential equilibria have a Markov property; unfortunately this result does not directly have bearing upon the set of payoffs which can be realised in equilibrium. The same can be said for the results of Kalai and Lehrer (1993) and Jordan (1995) who have studied the long-run properties of equilibrium play in contexts more general than the current one.
4. $\text{co}(X)$ denotes the convex hull of the set X .
5. Assumption (ii) can be relaxed at the cost of some additional complications and an appropriate reformulation of Proposition 3. Assumption (iii) is also not essential; the case where the feasible set is one-dimensional was treated in an earlier version, although some of the constructions needed differ. We conjecture that assumption (i) is likewise inessential, although we have not proved this.
6. The automaton can also be thought of as a type with a standard payoff matrix in which the payoffs in the row corresponding to a_1^* are all equal and strictly greater than all other payoffs. At a PBE this type will play a_1^* after every history, including those off the equilibrium path.
7. The reader is referred to Fudenberg and Tirole (1991a, 1991b) for formal definitions of all equilibrium concepts used here. For the purpose of the definitions, the automaton should be interpreted as a payoff-matrix type as described in footnote 6; the strategy of such a type in a PBE must be identical to the automaton strategy. This is in contrast to a Nash equilibrium where the payoff-matrix type need only follow the commitment strategy on the equilibrium path. In the

equilibria we construct, beliefs off-the-equilibrium path put probability zero on the automaton if player 1 has deviated from a_1^* in the past, which is consistent with the idea of an automaton which cannot deviate.

8. This condition is merely to rule out $\kappa=0$, which would change the method of solving the differential equation studied below (but not the conclusion). Note that κ can be negative.

9. An example is constructed in Celentani *et al.* (1996) which is similar to the type of construction we use, and which establishes that payoffs below the Stackelberg payoff can be sustained in a PBE. For their game, however, the Cripps *et al.* (1996) bound is just the minmax payoff.

10. If $a_1^0 \neq a_1^*$, \hat{a}_1 , then Bayes' rule does not tie down beliefs, and in this case we assume that probability one is put on the normal type; given that the automaton can be thought of as a type for whom it is a dominant strategy to play a_1^* , this is a natural assumption.