Bargaining, Reputation and Equilibrium Selection in Repeated Games *

Dilip Abreu[†] and David Pearce[‡]

Current Version: September 4, 2002.

Abstract

By arriving at self-enforcing agreements, agents in an ongoing strategic situation create surplus that benefits them both. Little is known about how that surplus will be divided. This paper concerns the role of reputation formation in such environments. It studies a model in which players entertain the slight possibility that their opponents may be one of a variety of behavioral types. When players can move frequently, a continuity requirement together with a weakening of subgame consistency yield a unique solution to the surplus division problem. The solution coincides with the "Nash bargaining with threats" outcome of the stage game.

*We wish to thank Yuliy Sannikov and Timothy Van Zandt for helpful comments on an earlier draft. We are grateful to the National Science Foundation (grant #003693). This paper grew out of ideas developed when the authors were visiting the Russell Sage Foundation.

[†]Princeton University

[‡]Yale University

1 Introduction

Repeated games offer perhaps the simplest model for studying how agents can create surplus by devising self-enforcing agreements. The "folk theorem" (see especially Aumann and Shapley (1976), Rubinstein (1982) and Fudenberg and Maskin (1986)) asserts that there is typically a vast multiplicity of outcome paths that are consistent with the internal logic of a repeated game. This powerful result leaves us with almost no predictive power. With few exceptions, the literature ignores a crucial element of implicit agreements: at the same time as players are trying to create surplus, they are presumably fighting over how the surplus will be divided. Perhaps the nature of this battle has strong predictive implications.

This paper models the struggle over the spoils of potential cooperation in repeated games. It draws on several recent papers that perturb traditional intertemporal bargaining models by introducing slight uncertainty about the motivations of the players (after the style of the seminal paper of Kreps, Milgrom, Roberts and Wilson (1982)). In Abreu and Gul (2000), there are non-optimizing "types" present in the model who make an initial demand (for example, requesting three quarters of the surplus), and stick to it forever. In equilibrium, rational players choose a type to imitate, and there ensues a war of attrition. The war ends when one of the players gives in to the other's demand. The solution is unique, and independent of the fine details of timing of offers.[1] Kambe (1999) modified their model so that all initial demands are chosen optimally, but with the knowledge that later, a player may become irrationally attached to her offer and refuse to alter it. This too is a tractable model, and its pure strategy solution is given an attractive Nash bargaining (Nash (1950)) interpretation by Kambe (1999).

Could the presence of behavioral players resolve the bargaining problem embedded in a repeated game? A natural conjecture is that rational players will find it irresistible to imitate relatively greedy behavioral types, and the first player to give herself away as being rational is the loser (for reasons akin to the Coase conjecture (Coase, 1972)). It is hard to establish such a result. To see what can go wrong, think of a model with one behavioral type on either side, one that takes the same action every period. Assume that if the second player plays a myopic best response to the behavioral type of the first, the first player does better than the second. This model may have many equilibria. For example, rational players may act behavioral forever, if the equilibrium says that if player i reveals herself to be rational and j has not yet done so, j will subsequently reveal herself to be rational also, and thereafter they play an equilibrium that i likes even less than the permanent war. The moral of the story is that perturbations that lead

[1] Abreu and Pearce (1999) "endogenize" the choices of non-optimizing types by having them imitate play that is traditionally observed in games of the kind they are playing; these types are subject to some behavioral bias (so they are proportionately less likely to concede, in any given contingency, than is an average player in the population).

to simple, unique solutions in a classical bargaining model may be overwhelmed by the power of "bootstrapping" (self-fulfilling prophecies) in a repeated game. See Schmidt (1993) for an enlightening investigation of this problem.

Faced with the logical possibility that rational players may not achieve Nash equilibrium (recall the curmudgeonly papers by Bernheim (1984) and Pearce (1984)), game theorists are fond of pointing out that if there is a way of playing a certain game that seems "focal" to everyone, or which has evolved as a tradition, it must be a Nash equilibrium. Here we exploit this perspective by supposing that there is a focal way of splitting the spoils from cooperation in a two-person supergame. What can be said of this arrangement if it is robust to behavioral perturbations of the kind we have been discussing?

More specifically, when players can move often, we look for a payoff pair $u^0 \in \mathbb{R}^2$ such that, if in a supergame perturbed to allow for the slight possibility that each player is behavioral, players expect the continuation value u^0 to apply whenever mutual rationality has newly become common knowledge[2], then something close to u^0 is an equilibrium payoff of the perturbed game.

To put it another way, the full-information supergame appears as a subgame in many parts of the behaviorally-perturbed game. How players think those subgames would be played affects their play, and hence their expected discounted payoffs, in the perturbed game. If players expect payoff u^0 in the full-information game, and if one is willing to impose a form of continuity as one perturbs that game slightly, something close to u^0 should also be an equilibrium possibility in the perturbed game.

We show that for essentially arbitrary two-person games and a rich set of behavioral types, there is exactly one such value u^0, and it coincides with the "Nash bargaining with threats" solution of the stage game. (Recall that Nash (1953) extended his cooperative bargaining theory (Nash, 1950) in a noncooperative setting, endogenizing the disagreement point. In the first of two stages, the players simultaneously choose strategies (threats) that serve as a disagreement point for a Nash demand game to be played in the second period. This two-stage procedure is called "Nash bargaining with threats.")

1.1 Relation to the Literature

The behavioral types used by Abreu and Gul (2000) are generalized to the machines in this paper to allow for complex strategic behavior. Kambe (1999) showed in a modification of the Abreu and Gul model that as perturbation probabilities approach zero, the division of surplus in the slightly perturbed intertemporal bargaining game

[2] Requiring that the same division of surplus is expected regardless of how the full-information subgame is reached is in the spirit of subgame consistency (Selten, 1973). Of course, a full-blooded application of subgame consistency, as a single-valued refinement, would render cooperation in the supergame impossible.

asymptotes to the Nash bargaining solution. Our result is closely related to his, with some major differences due to our repeated game setting. First, as the earlier discussion explained, the bootstrapping features of equilibria are harder to address with behavioral perturbations in repeated games than in bargaining settings where offers are enforceable. Secondly, our result pertains to Nash bargaining *with threats*: a player is rewarded for the ability to hurt her opponent at little cost to herself.

Nash qualified the applicability of his solutions as follows: "...we just assume there is an adequate mechanism for forcing the players to stick to their threats and demands once made; and one to enforce the bargain, once agreed. Thus we need a sort of umpire, who will enforce contracts or commitments." (1953, p. 130). We were intrigued to see his solution emerge in an infinite horizon game without enforceable contracts. It is true that in both Nash's first stage and in pre-agreement play in our game a player wishes to hurt her opponent while making herself as comfortable as possible. But it is striking that balancing these considerations in a quasi-static world leads to the same formula as it does in a dynamic war of attrition with incomplete information.

The bounded recall of the behavioral types introduced here is familiar from the work of Aumann and Sorin (1989) on pure coordination games. Our focus, however, is not on whether cooperation will be achieved, but on the distribution of the available surplus from cooperation.

There is a large literature on reputation formation, almost all of it stimulated originally by three papers: Kreps and Wilson (1982), Kreps, Milgrom, Roberts and Wilson (1982), and Milgrom and Roberts (1982). Work in this vein up until 1992 is authoritatively surveyed by Fudenberg (1992). Many papers since then have addressed the difficulties noted by Schmidt (1993) in extending the celebrated "Stackelberg result" for one long-run player by Fudenberg and Levine (1989)[3] to settings in which long-run reputational players face long-lived but less patient opponents. Noteworthy here are Aoyagi (1996), Celentani, Fudenberg, Levine and Pesendorfer (1996), and Cripps, Schmidt and Thomas (1996). Whereas this literature addresses the issue of how a player recognizes which behavioral type she faces (something we finesse through the device of announcements), it does not treat players who discount payoffs at the same rate.

Our interest is in two-sided reputation formation with equally patient players. Watson (1994) has some success in extending Aumann and Sorin (1989) to the more challenging setting of the prisoners' dilemma; his most interesting results depend on a "no conflict across types" assumption. Watson (1996) considers two-sided reputation formation without equilibrium and without discounting. Cramton (1992) exploits the framework of Admati and Perry (1987) to study two-sided incomplete information in a

[3]Fudenberg and Levine consider an infinite horizon game in which an extremely patient long-run player faces a sequence of one-period opponents. They show in great generality that even slight uncertainty about the long-run player's type allows her to do as well as she would if she could commit to playing any stage game strategy forever.

fully optimizing model of bargaining.

2 The Perturbed Repeated Game

A finite simultaneous game $G = (S_i, u_i)_{i=1}^2$ is played an infinite number of times by two players. Each player i is either "normal" (an optimizer) or with initial probability z_i, "behavioral." A normal player seeks to maximize the present discounted value of her stream of payoffs, using the interest rate $r > 0$ (common to the players). Each period is of length $\Delta > 0$. Player i's action set in G is denoted S_i. Payoff functions are denoted $u_i : S_1 \times S_2 \to \mathbb{R}$. The latter measure *flow* payoffs. Thus when players use actions $(s_1, s_2) \in S_1 \times S_2$ in a given period of length Δ, player i's payoffs in that period are $u_i(s_1, s_2) \int_0^\Delta e^{-rs} ds$. We denote by M_i the set of mixed actions associated with S_i.

At the start of play behavioral players simultaneously adopt and *announce* a repeated game strategy γ_i from some given *finite* set Γ_i. We interpret this as an announcement of a bargaining *position*. Each $\gamma_i \in \Gamma_i$ is a *machine* defined by a *finite* set of states Q_i, an initial state $q_i^0 \in Q_i$, an output function $\xi_i : Q_i \to M_i$, and a transition function $\psi_i : Q_i \times S_j \to Q_i$. Denote by $\pi_i(\gamma_i)$ the conditional probability that a behavioral player i adopts the position/machine γ_i. These conditional probabilities are exogenous to the model.

A normal player i also announces a machine in Γ_i as play begins, but of course she need not subsequently conform to her announcement. More generally, we could allow her to announce something outside Γ_i or to keep quiet altogether. Under our assumptions, however, in equilibrium she never benefits from exercising these additional options. A normal player can condition her choice of mixed action in the t^{th} stage game on the full history of the play in the preceding periods, including both players' earlier mixed actions (assumed observable) and initial announcements.

For a machine γ_i let $Q_i(\gamma_i)$ be its associated set of states and denote by $\gamma_i(q_i)$ the machine γ_i initialized at the state $q_i \in Q_i(\gamma_i)$. We are interested in varying the length of the stage game and eventually letting this time period Δ tend to zero. Starting from any pair of states q_1 and q_2, a pair of positions γ_1 and γ_2 determine a Markov process. The limit of the *average* discounted payoffs as the period length approaches zero is denoted

$$U(\gamma_1(q_1), \gamma_2(q_2)) = \lim_{\Delta \downarrow 0} U(\gamma_1(q_1), \gamma_2(q_2); \Delta) \ .$$

Let $BR_j(\gamma_i(q_i); \Delta)$ denote the set of best responses by j to $\gamma_i(q_i)$ in the *infinitely repeated* game with period length Δ. Let F be the convex hull of feasible payoffs of the stage game G. We will say that $u \in F$ is strongly efficient if there does not exist $u' \in F$ with $u' \geq u$ and $u'_i > u_i$, for some $i = 1, 2$. We say that $u \in F$ is ε-strongly efficient if there does not exist $u' \in F$ with $u' \geq u$ and $u'_i \geq u_i + \varepsilon$, for some $i = 1, 2$.

Assumption 1 *For all $\gamma_i \in \Gamma_i$, $q_i \in Q_i(\gamma_i)$ and $\varepsilon > 0$, there exists $\overline{\Delta} > 0$ such that for all $\Delta \leq \overline{\Delta}$, the payoff vector $U(\gamma_i(q_i), \beta; \Delta)$ is ε-strongly efficient for all $\beta \in BR_j(\gamma_i(q_i); \Delta)$.*

Thus no position $\gamma_i \in \Gamma_i$ is destructive or perverse: if j plays a best response to γ_i, the resulting expected discounted payoff pair is strongly Pareto efficient (in the limit as $\Delta \downarrow 0$). Thus, one might think of player i *demanding* a certain amount of surplus by announcing a position but not requiring gratuitously that surplus be squandered. Let $\overline{U}(\gamma_i(q_i))$ denote the limit of the average discounted payoff vector as Δ tends to zero and player $j \neq i$ plays a best response to γ_i.

Assumption 2 *Each $\gamma_i \in \Gamma_i$ is forgiving: that is, $\overline{U}_j(\gamma_i(q_i))$ is independent of $q_i \in Q_i(\gamma_i)$.*

In other words, in the limit as the period length Δ goes to 0, player j's average discounted payoff from playing a best response to γ_i is the same independently of the state γ_i is initialized at. The above assumptions imply that we can suppress the argument q_i, and simply write $\overline{U}(\gamma_i)$. Furthermore $\overline{U}(\gamma_i)$ is strongly efficient.

Since the behavioral types of each player are entirely mechanical, player i can reveal himself not to be behavioral by doing something no behavioral type does, after a given history. For example, she could announce one position and play something contradictory in the first period. Fixing the sets Γ_1 and Γ_2 and the conditional probabilities of different behavioral strategies, let $G^\infty(z_1, z_2; \Delta)$ denote the infinitely repeated game with period length Δ and initial probabilities of z_1, z_2, respectively that 1 and 2 are behavioral types. For any $u^0 \in \mathbb{R}^2$, modify $G^\infty(z_1, z_2; \Delta)$ as follows to get an infinite-horizon game $G^\infty(z_1, z_2; u^0, \Delta)$: at any information set at which it is newly the case that both players are revealed *not* to be behavioral types, end the game and supplement the payoffs already received by the amount $\frac{u^0}{r}$ (that is, pay them the present discounted value of receiving the payoff flow u^0 forever). We will be interested only in payoffs u^0 that are strongly efficient in the stage game.

Assumption 3 *The continuation payoffs u° used in forming the modified game $G^\infty(z_1, z_2; u^\circ, \Delta)$ are strongly efficient.*

For $(\gamma_1, \gamma_2) \in (\Gamma_1, \Gamma_2)$ let

$$\begin{aligned}(d_1, d_2) &= U(\gamma_1(q_1^0), \gamma_2(q_2^0)) \\ (u_1^1, u_2^1) &= \overline{U}(\gamma_1) \\ (u_1^2, u_2^2) &= \overline{U}(\gamma_2)\end{aligned}$$

where q_i^0 is the (original) initial state of γ_i. In addition define $\hat{u}^i = (\hat{u}_1^i, \hat{u}_2^i)$ as follows:

$$\hat{u}^i = \begin{cases} u^0 & \text{if} \quad u_i^0 \geq u_i^i \\ u^i & \text{if} \quad u_i^0 < u_i^i \end{cases}$$

Lemma 1 below establishes that $\hat{u}^i$ is the flow payoff in equilibrium following concession by player j to a normal player i (who rationally decides whether or not to reveal that she is normal).

The result is related to, though much simpler than, a discussion of one-sided reputation formation in bargaining in Myerson (1991, pp 399-404), and Proposition 4 of Abreu and Gul (2000).

Lemma 1 *For any $\varepsilon > 0$, there exists $\overline{\Delta} > 0$ such that for any $0 < \Delta \leq \overline{\Delta}$, any perfect Bayesian equilibrium σ of $G^\infty(z_1, z_2; u^0; \Delta)$, and after any history for which player j, but not i, is known to be normal, a normal player i's equilibrium payoff lies within ε of $\widehat{u}_i^i$ while a normal player j's equilibrium payoffs lie within ε of $y_i u_j^i + (1 - y_i) \widehat{u}_j^i$ where y_i is the posterior probability that player i is behavioral after the history in question.*

Proof.

Part (i): Suppose $u_i^0 \geq u_i^i$, and consequently that $\widehat{u}^i = u^0$.

Since u^0 and u^i are, by assumption, strongly efficient, it follows that $u_j^0 \leq u_j^i$. For small enough Δ, j can obtain at least $y_i u_j^i + (1 - y_i) u_j^0 - \varepsilon/2$, by playing a best response to γ^i, so long as γ^i continues to be played. If and when player i abandons γ^i (which occurs with probability at most $(1 - y_i)$), player j receives $u_j^0 < u_j^i$. On the other hand if player i is behavioral she will conform with the strategy γ^i forever. In this case j's payoff (for small enough Δ) is at most $u_j^i + \varepsilon/2$. If player i is not behavioral her equilibrium payoff is at least u_i^0 (since i obtains u_i^0 by revealing rationality right away) and j's payoff consequently at most u_j^0 (since u^0 is strongly efficient). Thus j's equilibrium payoff is bounded above by $y_i u_j^i + (1 - y_i)\hat{u}_j^i + \varepsilon$ for small enough δ.

Part (ii): Now suppose $u_i^0 < u_i^i$, consequently $\widehat{u}^i = u^i$.

Definition 1 *Suppose player 1 has not revealed rationality until stage $t_1 - 1$ and that player 2 has. If at stage t_1 player 2 plays an action which is not in the support of behavior consistent with playing an intertemporal best response to γ_1 (initialized at whatever state is reached after the $t_1 - 1$ stage history in question) we will say that player 2 "challenges" player 1 at t_1.*

Suppose player 2 challenges player 1 at t_1 when γ_1 is in state $q_1 \in Q_i(\gamma_1)$. Then, conditional upon player 1 conforming with γ_1 thereafter, player 2's *total* discounted payoff at t_1 is lower than her total discounted payoff from playing a best response to $\gamma_1(q_1)$, by at least $2M(q_1)\Delta$, for some $M(q_1) > 0$. Let $M = \min_{q_1 \in Q_1(\gamma_1)} M(q_1) > 0$.

There exists $\underline{\beta} < 1$ such that if player 1 conforms with γ^1 with probability $\beta \geq \underline{\beta}$ then the loss to player 2 of challenging player 1 at t_1 is at least $M\Delta$ (in terms of total discounted payoffs) conditional upon player 1 continuing to conform with γ_1 thereafter (in the event that player 1 indeed conformed with γ_1 at t_1).

Let $\Delta > 0$ be small enough such that after any $t-1$ stage history for which player 1 has conformed with γ^1 he will continue to conform with γ^1 in stage t, if from stage $t+1$ onwards player 2 will play a best response to γ^1, so long as player 1 conforms at t and continues to conform with γ^1. Such a $\Delta > 0$ exists since $u_1^1 > u_1^0$.

After any t stage history for which player 1 has conformed with γ^1 we will say that player 1 '*wins*' if player 2 plays a best response to γ^1 from stage $(t+1)$ on so long as player 1 continues to conform with γ^1.

Let $\underline{y}_1$ be the infimum over posterior probabilities such that if after any $(t-1)$ stage history player 1 is believed to be behavioral with probability $y_1 \geq \underline{y}_1$ then player 1 wins in any perfect Bayesian equilibrium. Clearly $\underline{y}_1 < 1$. To complete the proof we argue that $\underline{y}_1 = 0$.

Suppose $\underline{y}_1 > 0$. Let $\widetilde{y}_1 \in \left(0, \underline{y}_1\right)$ satisfy $\frac{\widetilde{y}_1}{\underline{\beta}} > \underline{y}_1$. Thus, if at stage $(t-1)$ player 1 is believed to be behavioral with probability $y_1 \geq \widetilde{y}_1$ and is expected to conform with γ^1 in stage t with probability $\beta \leq \underline{\beta}$ then conditional upon conformity in stage t, the posterior probability that player 1 is behavioral exceeds the threshold $\underline{y}_1$. It follows that, in any perfect Bayesian equilibrium, if $y_1 \geq \widetilde{y}_1$ where y_1 is the probability that player 1 is behavioral in round $(t-1)$ then player 1 must be expected to conform with γ^1 with probability $\beta \geq \underline{\beta}$ in stage t.

Claim 1 *Suppose that $y_1 \geq \widetilde{y}_1$ and player 2 challenges at time t_1. Then it must be the case that there exists $t_2 > t_1$ such that, conditional upon player 1 continuing to follow γ_1 until time t_2, challenging player 1 (again) at stage t_2 is in the support of player 2's equilibrium strategy.*

Proof. Suppose that the claim is false. Then normal player 1's equilibrium strategy from time t_1 onwards must necessarily entail conforming with γ_1 forever after, since $u_1^1 > u_1^0$ and Δ is small enough. In round t_1, as noted above prior to the claim, player 1 must conform in equilibrium with γ_1, with probability at least $\underline{\beta}$. But then player 2 can increase his total payoff by at least $M\Delta$ by playing a best response from time t_1. ■

An implication of the above claim (and the argument preceding the claim) is that conditional upon player 1 conforming with γ_1 the posterior probability that player 1 is behavioral cannot exceed the threshold $\underline{y}_1$ (since beyond this threshold there are no further challenges by player 2 in equilibrium).

Claim 2 *Suppose* $y_1 \geq \widetilde{y}_1$. *Then player 2 does not challenge, so long as player 1 continues to conform with* γ_1.

Proof. By the preceding step, if player 2 challenges at t_1 then for any $T > 0$ there exist $t_1 < t_2 < t_3 < ...$ with $t_n - t_{n-1} \geq T$, $n = 2, 3, \ldots$, such that, conditional upon player 1 conforming with γ_1 until time $t_{n-1} - 1$, it is in the support of player 2's equilibrium strategy to challenge at t_{n-1} and challenge at t_n conditional upon player 1 conforming with γ_1 until t_n.

Since the total loss from not playing a best response is at least $M\Delta > 0$, then for T large enough, player 2 must assign probability $(1 - \xi)$, for some $(1 - \xi) > 0$, that player 1 will reveal rationality between t_{n-1} and t_n. Hence if the posterior probability that player 1 is behavioral is y_1^{n-1} at t_{n-1} it is at least $\frac{y_1^{n-1}}{\xi}$ if player 1 conforms with γ_1 until t_n.

But for n large enough, this leads to a contradiction since we must have $y_1^n \leq \underline{y}_1$ for all $n = 1, 2, \ldots$. ■

Starting with the supposition that $\underline{y}_1$ is strictly greater than zero, we are led to Claim 2, which contradicts the definition of $\underline{y}_1$ as an infimum. It follows that $\underline{y}_1 = 0$ and player 1 will conform with γ^1 forever whether or not player 1 is behavioral, and player 2 will play a best response to γ^1 always. ■

Lemma 1 reveals that for small Δ, the game $G^\infty(z_1, z_2; u^0, \Delta)$ is played as a war of attrition. In case (ii), for example, once player j reveals herself rational, i maintains her original stance forever, and hence j should play a best response to i as soon as j reveals her rationality. It is natural to equate revealing rationality to "conceding" to one's opponent and we will freely use this terminology in what follows. To summarize, before rationality is revealed on either side, each player is holding out, hoping the other will concede; this is in effect a simple war of attrition.

Case (i) is only slightly more complicated: j knows what i will do, if rational, following j's revealing herself rational, and what i will do, if behavioral, in the same event. (That is, reveal rationality and stick with the strategy γ^i, respectively.) Because i's posterior concerning j's rationality evolves over time, j's expected payoff from revealing her rationality evolves also. Hence, although they play a war of attrition before anyone concedes, it is not a time-invariant war with (conditionally) stationary strategies.

Calculating the solutions of these wars of attrition is much easier in continuous time. To simplify the analysis, we move now to a war of attrition game in continuous time. Corresponding to the discrete time game $G^\infty(z_1, z_2; u^0, \Delta)$ consider the continuous time game $G_*^\infty(z_1, z_2; u^0)$ in which player i adopts and announces a position, and chooses a time (possibly infinity) at which to "concede" to the other player.

If the initial positions adopted are (γ_1, γ_2), then until one of the players concedes they receive flow payoffs $U(\gamma_1, \gamma_2)$. Now suppose player i is the first to concede at some time t. Then, if player $j \neq i$ reveals rationality, the players receive flow payoffs u^0 thereafter; if not, the players receive flow payoffs $\overline{U}(\gamma_j)$.

As in the repeated game, z_i denotes the initial probability that player i is behavioral, and $\pi_i(\gamma_i)$ is the conditional probability that a behavioral player i adopts position γ_i initially.

Let $\pi_i(\gamma_i)$ and $\mu_i(\gamma_i)$ be the respective probabilities with which behavioral and rational players initially adopt position γ_i. We denote by $F_i(t|\gamma_1, \gamma_2)$ the probability that player i (*unconditional* on whether i is behavioral or normal) will concede by time t, conditional on γ_1 and γ_2 and on $j \neq i$ not conceding before t.

Before defining equilibrium in the overall game we first consider behavior in the continuation game following the choice of positions γ_1 and γ_2 at $t = 0$.

Once γ_1 and γ_2 have been chosen, the posterior probability that i is behavioral is

$$\eta_i(\gamma_i) = \frac{z_i \pi_i(\gamma_i)}{z_i \pi_i(\gamma_i) + (1 - z_i)\mu_i(\gamma_i)}.$$

Since behavioral types never concede,

$$F_i(t \mid \gamma_1, \gamma_2) \leq 1 - \eta_i(\gamma_i) \text{ for all } t \geq 0.$$

Let $y_j(t)$ be the posterior probability that j is behavioral, if neither player has conceded until time t. By Bayes rule

$$y_j(t) = \frac{\eta_j(\gamma_j)}{1 - F_j(t \mid \gamma_1, \gamma_2)} \ .$$

Recall the definitions of u^i (and so on) provided just prior to Lemma 1. Let $\widetilde{u}_i^j(t)$ denote the expected flow payoff of a normal player i who concedes at t. Then

$$\widetilde{u}_i^j(t) = y_j(t) u_i^j + (1 - y_j(t))\hat{u}_i^j \ .$$

The next lemma is similar to Proposition 1 of Abreu and Gul (2000). Here, however, the player conceded to may, if normal, reveal rationality, and this leads to a non-stationary war of attrition.

Note that since u°, u^1 and u^2 are efficient (see assumptions 1 and 3),

$$u_i^\circ \geq u_i^i, i = 1, 2 \Leftrightarrow \hat{u}_i = u^\circ, i = 1, 2 \Rightarrow u_2^1 \geq u_2^2 \text{ and } u_1^2 \geq u_1^1.$$

Lemma 2 *Consider the continuation game following the choice of a pair of positions* (γ_1, γ_2) *such that*

$$u_1^2 > d_1, u_2^1 > d_2 \quad \text{and} \quad \hat{u}^1 \neq \hat{u}^2.$$

This game has a unique perfect Bayesian equilibrium. In that equilibrium, at most one player concedes with positive probability at time zero. Thereafter, both players concede continuously at possibly time-dependent hazard rates $\lambda_j(t) = \frac{r(\widetilde{u}_i^j(t)-d_i)}{\hat{u}_i^i-\hat{u}_i^j}$ $i \neq j,\ i,j = 1,2$ until some common time T^* at which the posterior probabilities that each player i is behavioral reach 1.

Proof. The condition $\hat{u}^1 \neq \hat{u}^2$ is automatically satisfied when $u_2^2 > u_2^1$, or equivalently when $u_1^1 > u_1^2$ (since positions are assumed to be efficient). Suppose that $\mu_i(\gamma_i) > 0, i = 1,2$. Let $T_i \equiv \inf\{t|F_i(t|\gamma_1,\gamma_2) = 1-\eta_i(\gamma_i)\}$. That is, T_i is the time at which the posterior probability that player i is behavioral reaches 1. Then $T_i > 0$ for $i = 1,2$. If not, a profitable deviation is for player i to wait a moment, instantly convince player j that she is behavioral for sure, and consequently be conceded to immediately by a normal player j. Furthermore, $T_1 = T_2$. Suppose, for instance, that $T_2 > T_1$. Then after T_1 only the behavioral type of player 1 is still left in the game. Since behavioral types never concede, a normal player 2 should concede immediately. Hence $T_1 = T_2 = T^*$.

Let $\psi_i(t)$ denote the expected payoff to normal player i from conceding at time t. Then

$$\psi_i(t) = F_j(0)\frac{\hat{u}_i^i}{r} + \textstyle\int_0^t (D_i(s) + e^{-rs}\frac{\hat{u}_i^i}{r})dF_j(s)$$
$$+(1-F_j(t))D_i(t) + e^{-rt}(1-F_j(t))\frac{\widetilde{u}_i^j(t)}{r}$$

where $D_i(s) = d_i\int_0^s e^{-rv}dv$, and $j \neq i$.

Let $A_i = \{t|\psi_i(t) = \sup\psi_i(s)\}$. By standard 'war of attrition' arguments $(0,T^*] \subset A_i$. Hence ψ_i is constant, and thus differentiable, on this range, as must be F_j. Setting $\psi_i'(t) = 0,\ t \in (0,T^*)$ yields a characterization of F_j. That is,

$$\lambda_j(t) \equiv \frac{f_j(t)}{1-F_j(t)} = \frac{r(\widetilde{u}_i^j(t)-d_i)}{(\hat{u}_i^i - \widetilde{u}_i^j(t)) + y_j(t)(u_i^j - \hat{u}_i^j)}$$

where we have used $\dot{y}_j(t) = \lambda_j(t)y_j(t)$, which in turn is obtained by differentiating $y_j(t) = \frac{\eta_j(\gamma_j)}{1-F_j(t)}$. Substituting $\widetilde{u}_i^j(t) = y_j(t)u_i^j + (1-y_j(t))\hat{u}_i^j$ in the denominator yields

$$\lambda_j(t) = \frac{r(\widetilde{u}_i^j(t)-d_i)}{\hat{u}_i^i - \hat{u}_i^j}.$$

Hence, $F_i(t) = 1 - c_i e^{-\int_0^t \lambda_i(s)ds}$ for some constant of integration c_i.

When $c_i < 1$, $F_i(0) > 0$ and player i concedes with positive probability at $t = 0$. In this case, in equilibrium, player $j \neq i$ cannot also concede with positive probability at $t = 0$, and it must be the case that $c_j = 1$. Recall that T_i satisfies

$$F_i(T_i) \equiv 1 - c_i e^{-\int_0^{T^*}\lambda_i(s)ds} = 1 - y_i(\gamma_i)\ .$$

The twin conditions $(1-c_1)(1-c_2)=0$ and $T_1=T_2$ uniquely determine c_1 and c_2. Note that if $\eta_j(\gamma_j)u_i^j+(1-\eta_j(\gamma_j))\hat{u}_i^j \le d_i$ then it must be the case that $\tilde{u}_i^j(0)=y_j(0)u_i^j+(1-y_j(0))\hat{u}_i^j > d_i$ and consequently that $c_j<1, F_j(0)>0$ and $y_j(0)>\eta_j(\gamma_j)$.

Finally suppose that $\mu_i(\gamma_i)>0$ and $\mu_j(\gamma_j)=0, i\neq j$. Then $T^*=T_j=0, c_j=1, c_i<0$, and $F_i(0)=1-\eta_i(\gamma_i)$. If $\mu_i(\gamma_i)=0, i=1,2$, then $T_1=T_2=T^*=0$, and $c_1=c_2=1$. ∎

As noted earlier, a complicating feature of the above analysis is that the hazard rates of concession are time-dependent. It turns out to be helpful to consider a *simplified war of attrition* (henceforth, referred to as such, or as the *simplified game*) in which concession by player i results in the payoff-vector u^j rather than the possibly time-dependent payoff vector $\widetilde{u}^j(t)$. Assume now that $u_2^2>u_2^1$ and $u_1^1>u_1^2$. The analysis of Lemma 2 subsumes the simple case with the constant hazard rates

$$\overline{\lambda}_1=\frac{r(u_2^1-d_2)}{u_2^2-u_2^1}$$

and

$$\overline{\lambda}_2=\frac{r(u_1^2-d_1)}{u_1^1-u_1^2}$$

replacing $\lambda_1(t)$ and $\lambda_2(t)$ of the *actual* war of attrition. In either war, if player i concedes with positive probability at the start of play we will say that player i is "weak". In the simplified game

$$\overline{F}_i(t)=1-c_ie^{-\overline{\lambda}_it} \ .$$

If player 2 is weak,

$$\overline{F}_1(\overline{T}^*)=1-e^{-\overline{\lambda}_1\overline{T}^*}=1-\eta_1(\gamma_1) \ .$$

Hence $\overline{T}^*=\frac{-\log\eta_1(\gamma_1)}{\overline{\lambda}_1}$, and $\overline{c}_2$ is determined by the requirement $1-\overline{c}_2e^{-\overline{\lambda}_2\overline{T}^*}=1-\eta_2(\gamma_2)$. Hence $\overline{c}_2=\frac{\eta_2(\gamma_2)}{(\eta_1(\gamma_1))^{(\overline{\lambda}_2/\overline{\lambda}_1)}}$ and $\overline{F}_2(0)=1-\overline{c}_2$.

Furthermore, it is easy to show that player 2 *is* in fact weak if and only if $\frac{\eta_2(\gamma_2)}{(\eta_1(\gamma_1))^{(\overline{\lambda}_2/\overline{\lambda}_1)}}<1$.

Throughout we will identify corresponding endogenous variables in the simplified game with an upper bar. That is, as $\overline{\lambda}_i$, $\overline{c}_i$, $\overline{F}_i$, and so forth.

Suppose $u_2^0>u_2^2$. Then $\hat{u}^2=u^0$; that is, a normal player 2 who is conceded to by player 1 will reveal rationality. Furthermore, since u^1, u^2 and u^0 are efficient, $u_1^2>\widetilde{u}_1^2(t)>\hat{u}_1^2$ and $u^1=\widetilde{u}^1(t)=\hat{u}^1$ for all $t\in(0,T^*)$.

Then

$$\overline{\lambda}_1=\frac{r(u_2^1-d_2)}{\hat{u}_2^2-\widehat{u}_2^1}=\lambda_1$$

and

$$\overline{\lambda}_2 = \frac{r\left(u_1^2 - d_1\right)}{\widehat{u}_1^1 - \widehat{u}_1^2} > \lambda_2(t) = \frac{r(\widetilde{u}_1^2(t) - d_1))}{\widehat{u}_1^1 - \widehat{u}_1^2} \ .$$

Lemma 3 *Consider the continuation game following the choice of a pair of positions* (γ_1, γ_2). *Suppose* $u_1^1 > u_1^2 > d_1, u_2^2 > u_2^1 > d_2$ *and* $u_2^0 > u_2^2$. *Then if player 2 is 'weak' in the simplified game, he is 'weak' in the actual game. Furthermore he concedes (at the start of play) with higher probability in the actual than in the simplified game. That is* $\overline{F}_2(0) > 0$ *implies* $F_2(0) > \overline{F}_2(0)$.

Proof. Recall from the discussion preceding Lemma 3 that

$$1 - F_i(t) = c_i e^{-\int_0^t \lambda_i(s)ds}$$

and $1 - F_i(T^*) = \eta_i(\gamma_i)$. Consequently, $c_i = y_i e^{\int_0^{T^*} \lambda_i(s)ds}$. Also, from the preceding discussion, $T^* \leq \frac{-\log \eta_1(\gamma_1)}{\lambda_1}$, with strict equality if player 2 is weak. Also, $F_2(0) = 1 - c_2$. We establish the lemma by showing $c_2 < \overline{c}_2$; that is,

$$\int_0^{T^*} \lambda_2(s)ds < \int_0^{\overline{T}^*} \overline{\lambda}_2 ds \ .$$

We have noted earlier that

$$\lambda_2(s) < \frac{\overline{\lambda}_2}{\overline{\lambda}_1} \cdot \lambda_1 \quad s \in [0, T^*) \ .$$

Hence

$$\int_0^{T^*} \lambda_2(s)ds < \frac{\overline{\lambda}_2}{\overline{\lambda}_1} \cdot \lambda_1 \int_0^{T^*} ds \leq \frac{\overline{\lambda}_2}{\overline{\lambda}_1} \left(-\log \eta_1\left(\gamma_1\right)\right) = \int_0^{\overline{T}^*} \overline{\lambda}_2 ds \ .$$

■

Lemma 4 *Consider the continuation game following the choice of positions* γ_1 *and* γ_2 *and suppose that*

$$\hat{u}^1 = \hat{u}^2 \text{ (equivalently } u_i^\circ \geq u_i^i, i = 1, 2\text{)}.$$

Then the unique Perfect Bayesian Equilibrium outcome is u° *if both players are normal and* u^i *if* i *is behavioral and player* $j \neq i$ *is normal.*

Proof. By conceding player i receives a payoff of u_i^j if player j is behavioral and at least $u_i^o \leq u_i^j$ if player j is normal. If player j is behavioral, player i obtains a payoff of at most u_i^j. Conversely for player j . On the other hand a normal player j cannot receive a payoff of less than u_j^o in equilibrium when player i is normal since

a normal player j cannot compensate for this loss by obtaining more than u_j^i from a behavioral player i. It follows that in equilibrium a normal player i obtains u_i^j if player j is behavioral and u_j^o if player j is normal. The result now follows. ■

A strategy for player i in the game $G_*^\infty(z_{1,}z_2; u^\circ)$ is a pair of probability distributions $\pi_i(\cdot)$ and $\mu_i(\cdot)$ over positions (for the behavioral and rational types of player i, respectively); and for every pair of positions $(\gamma_1, \gamma_2) \in \Gamma_1 \times \Gamma_2$, a "distribution function" $F_i(\cdot \mid \gamma_1, \gamma_2)$ over concession times; where $F_i(t \mid \gamma_1, \gamma_2)$ is the probability with which player i (who might be normal or behavioral) concedes by time t, conditional on $j \neq i$ not conceding prior to t.

Fixing a strategy profile, let $\psi_i(t \mid \gamma_i, \gamma_j)$ denote the payoff to a player i who concedes at t, and $\psi_i(\infty \mid \gamma_i, \gamma_j)$ denote the payoff to i from never conceding. Let

$$V_i(\gamma_i) \equiv \sum_{\gamma_j} (z_j \pi_j(\gamma_j) + (1 - z_j)\mu_j(\gamma_j)) \sup_s \psi_i(t \mid \gamma_i, \gamma_j)$$

denote player i's expected payoff from choosing a position γ_i and adopting an optimal concession strategy thereafter. Let

$$A_i(\gamma_i, \gamma_j) \equiv \{t \in \mathbb{R}_+ \cup \{\infty\} \mid \psi_i(t \mid \gamma_i, \gamma_j) = \sup_s \psi_i(s \mid \gamma_i, \gamma_j)\}$$

A strategy profile is an equilibrium of the overall game if for all $\gamma_i \in \Gamma_i, i = 1, 2$

$$\mu_i(\gamma_i) > 0 \Rightarrow \gamma_i \in \arg\max_{\gamma_i' \in \Gamma_i} V_i(\gamma_i),$$

$A_i(\gamma_i, \gamma_j)$ is nonempty and the "support" of the function $F_i(\cdot \mid \gamma_i, \gamma_j)$ is contained in $A_i(\gamma_i, \gamma_j)$. (Since $F_i(t \mid \gamma_i, \gamma_j) \leq 1 - \eta_i(\gamma_i)$ it is formally not a distribution function over concession times.)

3 The Main Characterization

We continue with the analysis of the continuous time war of attrition introduced in Section 2 and focus here on the behavior of the model as the probability of behavioral types goes to zero.

As explained in the Introduction, we wish to impose a continuity requirement on solutions of perturbed games. Suppose that players expect that, whenever mutual rationality becomes common knowledge, their continuation payoffs will be given by some $u^0 \in \mathbb{R}^2$. If the full-information game is perturbed only very slightly (the perturbation probabilities are close to zero), it seems attractive to require that the payoffs expected in the perturbed game are close to u^0. But this is a little too strong. If, for example, $z_1 > 0$ and $z_2 = 0$ then no matter how small z_1 is, Coase-conjecture-like arguments establish that player 1 will do much better than player 2 (if player 1 has any types in Γ_1 that are attractive to imitate). The converse is true if $z_1 = 0$ and $z_2 > 0$, so it is impossible to demand continuity at $z = (0, 0)$. We will confirm, however, that a weaker form of continuity, in which only test sequences with $\frac{z_1}{z_2}$ and $\frac{z_2}{z_1}$ bounded away from zero (”proper sequences”) are considered, *can* be satisfied.

Because the above exercise leads to a ”Nash threats” result, we now provide the relevant definitions. Recall the (standard) Nash (1950) bargaining solution for a convex non-empty bargaining set $F \subseteq \mathbb{R}^2$, relative to a disagreement point $d \in \mathbb{R}^2$, where it is assumed that there exists $u \in F$ such that $u \gg d$. Then the Nash bargaining solution, denoted $u^N(d)$, is the unique solution to the maximization problem

$$\max_{u \in F}(u_1 - d_1)(u_2 - d_2).$$

In the event that there does not exist $u \in F$ such that $u \gg d$, $u^N(d)$ is defined to be the strongly efficient point $u \in F$ which satisfies $u \geq d$.

In Nash (1953) the above solution is derived as the unique limit of solutions to the *non-cooperative* Nash demand game when F is perturbed slightly and the perturbations go to zero. Nash's paper also *endogenizes* the choice of threats, and consequently disagreement point, and it is this second contribution which is particularly relevant here. Starting with a game G, the bargaining set F is taken to be the convex hull of feasible payoffs of G. The threat point d is determined as the non-cooperative (Nash) equilibrium of the following two ‘stage’ game:

Stage 1 The two players independently choose (possibly mixed) threats m_i, $i = 1, 2$. The expected payoff from (m_1, m_2) is the disagreement payoff denoted $d(m_1, m_2)$.

Stage 2 The player's final payoffs are given by the Nash bargaining solution relative to the disagreement point determined in Stage 1.

Thus players choose threats to maximize their stage 2 payoffs given the threats chosen by their opponent. Note that the set of player i's *pure* strategies in the threat game are her set of *mixed* strategies in the game G. Since the Nash bargaining solution yields a strongly efficient feasible payoff as a function of the threat point, the Nash threat game is strictly competitive in the space of pure strategies (of the threat game). Nash shows that the threat game has an equilibrium in pure strategies (i.e., players do not mix over mixed strategy threats), and consequently that all equilibria of the threat game are equivalent and interchangeable. In particular the threat game[4] has a unique equilibrium

[4] We thank Robert Aumann for noting an oversight in our earlier discussion of uniqueness and emphasizing the importance of the existence of pure strategy equilibrium in the Nash-threat game, which a rereading of Nash's paper also served to clarify.

payoff (u_1^*, u_2^*) where $u^* = u^N(d(m_1^*, m_2^*))$ where m_i^* is an equilibrium threat for player i. Our solution essentially yields (u_1^*, u_2^*) as the only equilibrium payoff which survives in *the limit as the probability of behavioral types goes to zero.*

In order to establish the result we need to assume that positions corresponding to the Nash threat strategy and Nash threat equilibrium outcome exist. This may be interpreted as a 'rich set of types' assumption. It is also convenient to modify the definition of a machine, and the specification of the transition function provided there, to allow the latter to condition, in addition, on the set $\{0, 1\}$. We interpret 0 as 'silence' and 1 as the announcement 'I concede.' Behavioral players never announce 1. This harmless modification facilitates the exposition.

We also restrict attention to stage games for which

$$u^N(d(m_1^*, m_2^*)) \equiv u^* \gg d^* \equiv d(m_1^*, m_2^*) \ ,$$

where m^* is an equilibrium profile of the Nash threat game.

Assumption 4 ***Nash-Threat Position***

There exists a position $\gamma_i^* \in \Gamma_i$ such that

$$\overline{U}(\gamma_i^*) = u^*$$

and for all $\gamma_j \in \Gamma_j$,

$$u_i^N(d(\gamma_i^*, \gamma_j) \geq u_i^*$$

For example, the following machine γ_i satisfies all the assumptions (modulo footnote 5) for K large enough. The initial state plays m_i^* and the machine transits out of the initial state only if player j announces '1'. In the latter event the machine transits to a 'sub' machine with a sequence of 'cooperative' states (yielding u^*)[5] and K punishment states. In the event of any 'defection' by player j in the cooperative phase, the K punishment states must be run through before return to the first cooperative state. In the punishment states, the mixed strategy minmaxing player j is played.

Definition 2 *The sequence (z_1^n, z_2^n) $n = 1, 2, ...$ is called a proper sequence if $z_i^n \downarrow 0$ for $i = 1, 2$ and there exist finite, strictly positive scalars $\underline{a}$, $\overline{a}$ such that $\underline{a} < \frac{z_2^n}{z_1^n} < \overline{a}$ for $n = 1, 2 \ldots.$*

[5] If G is a game in which u^* is not a rational convex combination of extreme points of the feasible payoff set, a machine, as defined, will not generate a payoff exactly equal to u^*. This 'technical' difficulty may be dealt with by permitting public randomization, or by rephrasing the preceding assumption and results in terms of an 'almost' Nash threat position.

For a game $G_*^\infty(z;u)$ and a strategy profile σ of $G_*^\infty(z;u)$ let $U(\sigma) \in \mathbb{R}^2$ denote normal players' expected payoffs under σ. Note that the sets Γ_i of positions and the probabilities with which they are adopted by behavioral players, are held fixed throughout.

Proposition 1 says that the "Nash bargaining with threats" payoff u^* satisfies our continuity requirement. Proposition 2 states that u^* is the *only* payoff that survives the continuity test.

Proposition 1 *Consider a proper sequence z^n and corresponding sequences of games $G^\infty(z^n;u^*)$ and perfect Bayesian equilibria $\sigma^n, n = 1, 2...$. Then for any $\varepsilon > 0$ there exists $\overline{n}$ such that $|U(\sigma^n) - u^*| < \varepsilon$ for all $n \geq \overline{n}$.*

Proposition 2 *Fix $u^0 \in \mathbb{R}^2$ and consider a proper sequence z^n and corresponding sequences of games $G^\infty(z^n;u^0)$ and perfect Bayesian equilibria $\sigma^n, n = 1, 2 \ldots$. Suppose $u_i^0 < u_i^*$. Then for any $\varepsilon > 0$ there exists $\overline{n}$ such that for all $n \geq \overline{n}$, $U_i(\sigma^n) \geq u_i^* - \varepsilon$.*

How should the preceding propositions be interpreted? Fudenberg and Maskin (1986) and Fudenberg, Kreps and Levine (1988) offer ample warning of the sensitivity of equilibria of reputationally-perturbed games to the particular perturbation introduced. Sometimes, as in Fudenberg and Levine (1989), there is one 'type' that dominates equilibrium play regardless of what other types are in the support of the initial distribution. With long-lived players on both sides, Schmidt (1993) shows this need no longer be true. In this paper, we need to assume that behavioral types all have certain features (such as sticking to the initial positions they announce up front). But within this class, the "Nash bargaining with threats" type dominates play *regardless of what other types happen to be present in the initial distribution.*

Proof of Propositions

Consider the continuation game after γ_1, γ_2 are chosen and let $u^i \equiv \overline{U}(\gamma_i)$ and $d \equiv U(\gamma_1, \gamma_2)$. Suppose player 1 adopts position γ_1^* and that player 2 adopts some position γ_2. We argue that player 1's payoff in the continuation game following the choice of (γ_1^*, γ_2) is (at least) approximately u_1^* in any perfect Bayesian equilibrium of the overall game, when the prior probabilities that the players are behavioral are sufficiently small.

Case 1 $u_1^1 > u_1^2$ and $u_2^2 > u_2^1$.

Step 1 Since playing a best response to γ_2 yields player 1 u_1^2, it must be the case that $d_1 \leq u_1^2$. Similarly, $d_2 \leq u_2^1$.We will assume until step 5 that $u_j^j > u_j^i > d_j$, $i \neq j$, $i, j = 1, 2$. Finally in step 5 we will allow for the possibility that $u_j^i = d_j$. Furthermore we will first analyze the simplified war of attrition and exploit the discussion following the proof of Lemma 3 to draw conclusions about the actual game. Steps 1 and 2 are concerned with the *simplified* game. Suppose $\overline{\lambda}_1 > \overline{\lambda}_2$

(recall the definitions following the proof of Lemma 2) and suppose that γ_2 is chosen in equilibrium by a normal player 2 with probability $\mu_2^n(\gamma_2) > 0$. Suppose that $\mu_2^n(\gamma_2) \to 2\overline{\mu} > 0$ (this convergence could occur along a subsequence.) Then for n large enough,

$$\eta_1^n \equiv \eta_1^n(\gamma_1) \geq \frac{z_1^n \pi_1(\gamma_1)}{(1 - z_1^n) \cdot 1 + z_1^n \pi_1(\gamma_1)}$$

$$\eta_2^n \equiv \eta_2^n(\gamma_2) \leq \frac{z_2^n \pi_2(\gamma_2)}{(1 - z_2^n)\overline{\mu} + z_2^n \pi_2(\gamma_2)}$$

Hence

$$\log[\eta_2^n/(\eta_1^n)^{\overline{\lambda}_2/\overline{\lambda}_1}] \leq \log \frac{z_2^n}{z^n} \frac{\pi_2(\gamma)}{\pi_1(\gamma)} (z_1^n \pi_1(\gamma_1))^{1-k} + C$$

where C is a finite constant independent of z_1, z_2, and $k \equiv \frac{\overline{\lambda}_2}{\overline{\lambda}_1} < 1$.

$$\log[\eta_2^n/(\eta_1^n)^{\overline{\lambda}_2/\overline{\lambda}_1}] \leq \log \overline{a} \frac{\pi_2(\gamma_2)}{\pi_1(\gamma_1)} + (1-k)\log(z_1^n \pi_1(\tau_1)) + C$$

Clearly, the r.h.s. $\to -\infty$ as $z_1^n \to 0$. Consequently $\frac{\eta_2^n}{(\eta_1^n)^{\overline{\lambda}_2/\overline{\lambda}_1}} \to 0$.

We know from the discussion following lemma 2 that player 2 is 'weak' if and only if $\overline{\omega}_2^n = 1 - [\eta_2^n/(\eta_1^n)^{\overline{\lambda}_2/\overline{\lambda}_1}] > 0$; in this case $\overline{\omega}_2^n$ is the probability with which player 2 concedes to player 1 at time zero. Thus γ_1 wins against γ_2 with probability 1 in the limit so long as γ_2 is used with non-negligible probability by a behavioral player 2, and $\overline{\lambda}_1 > \overline{\lambda}_2$. This key fact was first noted by Kambe (1999). A similar observation appears independently in Compte-Jehiel (2002).

The next step shows that in the simplified game, a normal player 1 can (in the limit as the z_i's $\to 0$) guarantee herself a payoff close to u_1^* by choosing a Nash position γ_1^*.

Step 2 Recall that

$$\overline{\lambda}_1 = \frac{u_2^1 - d_2}{(u_2^2 - u_2^1)/r}$$

$$\overline{\lambda}_2 = \frac{u_1^2 - d_1}{(u_1^1 - u_1^2)/r}$$

Hence

$$\overline{\lambda}_1 > \overline{\lambda}_2 \Leftrightarrow \frac{u_2^1 - d_2}{u_1^2 - d_1} > \frac{u_2^2 - u_2^1}{u_1^1 - u_1^2}$$

Diagrammatically

(Figure 1 here).

$\overline{\lambda}_1 > \overline{\lambda}_2 \Leftrightarrow \theta > s$, where θ is the slope of the line joining (d_1, d_2) and (u_1^2, u_2^1) and s is the negative of the slope of the line joining u^2 and u^1.

Recall the equilibrium of the Nash threat game. It yields the disagreement point d^* and the final outcome u^*. A defining property of u^* is that $\theta^* = s^*$ where θ^* is the slope of the line joining d^* and u^* , and s^* is the absolute value of the slope of some supporting hyperplane to the set of feasible payoffs at u^*.

Consider the Nash position γ_1^* and any position γ_2 such that $u_1^2 \equiv \overline{U}(\gamma_2) < u_1^* - \eta$. Let $d \equiv d(\gamma_1^*, \gamma_2)$ and fix γ_1^*, γ_2 in the discussion below. Diagrammatically we have

(Figure 2 here).

As noted earlier, it must be the case that $d_1 \leq u_1^2$ and $d_2 \leq u_2^1 = u_2^*$. For the moment assume that $u_1^2 > d_1$ and $u_2^2 > d_2$. Let θ be the slope of the line joining d to u^*. Then θ must be at least as large as the (absolute value of the) slope, say s, of *some* supporting hyperplane to F (the set of feasible payoffs) at u^*. If not, $u^N(d)$ would lie strictly to the left of u^*. This implies that player 2 can improve his payoff in the Nash-threat game by playing $m_2 \neq m_2^*$ such that $u(m_1^*, m_2) = d$, a contradiction. Let θ' and s' be as indicated in Figure 2. Then clearly $\theta' \geq \theta$ and $s' \leq s$. It follows that $\theta' > s'$. Hence γ_1^*, γ_2 as specified above yield $\overline{\lambda}_1 > \overline{\lambda}_2$.

By step 1 above $\overline{\lambda}_1 > \overline{\lambda}_2$ implies that player 1 'wins' in the limit with probability 1 ($\overline{\omega}_2 \rightarrow 1$) against γ_2, so long as $\mu_2(\gamma_2) > \underline{\mu} > 0$. Hence by adopting the Nash position player 1 obtains a payoff arbitrarily close to u_1^* with probability $1 - 2\overline{\mu} \cdot (\#\Gamma_2)$. Since $\underline{\mu}$ in step 1 is any strictly positive number, it follows that in the simplified game there exists $\overline{n}$ such that $|U(\overline{\sigma}^n) - u^*| < \varepsilon$ for all $n \geq \overline{n}$.

Step 3 Proof of Proposition 1

Given that player 1 adopts the position γ_1^*, that player 2 adopts a position γ_2 such that $u_1^2 < u_1^*$, and given that $u^0 = u^*$, the actual game reduces to the simplified game. A symmetric argument applies to player 2. Hence Step 2 in fact establishes Proposition 1. ■

Step 4 Proof of Proposition 2 Suppose that $u_1^0 < u_1^*$. If $u_2^2 \geq u_2^0$ then the actual game again reduces to the simplified game and we are done. Now suppose that $u_2^0 > u_2^2$ and that $u_1^0 > d_1$. Then Lemma 3 establishes that $\omega_2 > \overline{\omega}_2$; the conclusions obtained in the simplified game are reinforced in the actual game. ■

Step 5 Finally suppose that player 1 adopts γ_1^*, and that $u_2^1 = d_2$ or $u_1^2 = d_1$. The requirement that $u_2^N(d) \leq u_2^*$ (see Step 2)implies that the only possibility is $u_1^2 = d_1$

while $u_2^1 > d_2$. In this case, unique equilibrium behavior entails normal player 2 conceding to player 1 with probability one immediately; in equilibrium, normal player 1 will never concede to 2 since $u_1^2 = d_1$, so long as there is any possibility that player 2 is normal. ■

Case 2 $u_2^2 \leq u_2^1 = u_2^*$

As in the preceding case we will argue that if $u_1^\circ \leq u_1^*$ the position γ_1^* guarantees a normal player 1 a payoff close to u_1^* as $n \to \infty$.

Step 1 If $u^\circ = u^*$, then Lemma 4 applies, and it follows directly that player 1's payoff converges to u^* as $n \to \infty$.

Now suppose $u_1^\circ < u^*$. If $d_1 \geq u_1^*$ then normal player 1's equilibrium payoff is at least d_1, hence at least u_1^*.

We are left with the case $d_1 < u_1^*$.

Step 2 Truncated Simplified Game

This step is analogous to Lemma 3 and the discussion preceding it.

Choose $\bar{\bar{y}}_2 \in (0,1)$ such that $d_1 < \bar{\bar{y}}_2 u_1^2 + (1-\bar{\bar{y}}_2)u_1^\circ < u_1^1$.

Let $\bar{\bar{u}}^2 \equiv \bar{\bar{y}}_2 u^2 + (1-\bar{\bar{y}}_2)u^\circ$.

Consider a 'truncated' simplified game with constant hazard rates

$$\begin{aligned} \bar{\bar{\lambda}}_1 &= \frac{r(u_2^1 - d_2)}{\bar{\bar{u}}_2^2 - u_2^1} \\ \bar{\bar{\lambda}}_2 &= \frac{r(\bar{\bar{u}}_1^2 - d_1)}{u_1^1 - \bar{\bar{u}}_1^2} \end{aligned}$$

in which player 2 stops conceding when $\bar{\bar{y}}_2(t) = \dfrac{\eta_2(\gamma_2)}{1-\bar{\bar{F}}_2(t)} = \bar{\bar{y}}_2$, i.e. when the posterior probability that player 2 is behavioral reaches $\bar{\bar{y}}_2$. It follows that

$$1 - \bar{\bar{F}}_2(T^*) = \bar{\bar{c}}_2 e^{-\bar{\bar{\lambda}}_2 T^*} = \frac{\mu_2(\gamma_2)}{\bar{\bar{y}}_2},$$

and that player 2 is weak in the truncated simplified game if and only if $\dfrac{\eta_2(\gamma_2)/\bar{\bar{y}}_2}{(\eta_1(\gamma_1))^{\bar{\bar{\lambda}}_2/\bar{\bar{\lambda}}_1}} < 1$. Furthermore, $\bar{\bar{c}}_2 = \min\left\{\dfrac{\eta_2(\gamma_2)/\bar{\bar{y}}_2}{(\eta_1(\gamma_1))^{\bar{\bar{\lambda}}_2/\bar{\bar{\lambda}}_1}}, 1\right\}$.

Finally by an argument analogous to that of Lemma 3, if $1 - \bar{\bar{c}}_2 = \bar{\bar{F}}_2(0) > 0$ in the truncated simplified game, then $F_2(0) > \bar{\bar{F}}_2(0)$ in the actual game.

Step 3 Proceeding as in the proof of Case 1 it may be checked directly that if

$$\mu_2^n(\gamma_2) \to 2\overline{\mu} > 0$$

(i) $\bar{\bar{c}}_2^{\,n} \downarrow 0$ as $n \to \infty$ if $\bar{\bar{\lambda}}_1 > \bar{\bar{\lambda}}_2$

(ii) $\bar{\bar{\lambda}}_1 > \bar{\bar{\lambda}}_2$

(iii) player 1 'wins' with probability approaching 1 as $n \to \infty$. ∎

Proposition 3 *For any game* $G_*^{\infty}(z_1, z_2; u^{\circ}), z_i > 0, i = 1, 2;$ *an equilibrium exists.*

Proof. The argument is analogous to the proof of Proposition 2 of Abreu and Gul (2000), and is only sketched. Let $\mathcal{M}_i$ denote the $|\Gamma_i| - 1$ dimensional simplex. For any $(\mu_1, \mu_2) \in \mathcal{M}_1 \times \mathcal{M}_2$, Lemmas 3 and 4 uniquely pin down equilibrium payoffs and behavior in the continuation games following the choice of positions $(\gamma_1, \gamma_2) \in \Gamma_1 \times \Gamma_2$. It follows directly from these characterizations that the payoffs $V_i(\gamma_i \mid \mu)$ (where we now make explicit the dependence on $\mu = (\mu_1, \mu_2)$) are continuous in μ. Consider the mapping

$$\xi : \mathcal{M}_1 \times \mathcal{M}_2 \to \mathcal{M}_1 \times \mathcal{M}_2$$

where

$$\xi(\mu) = \{\tilde{\mu} | \, \tilde{\mu}_i(\gamma_i) > 0 \Rightarrow \gamma_i \in \arg\max_{\gamma_i'} V_i(\gamma_i' \mid \mu), i - 1, 2\}$$

To complete the proof we need to show that this mapping has a fixed point. Clearly ξ maps to non-empty, convex sets. Furthermore since V_i is continuous in μ, the mapping ξ is upper hemi-continuous. By Kakutani's fixed point theorem, it follows that the required fixed point exists. ∎

4 Conclusion

The excessive power of self-fulfilling prophecies in repeated games stands in the way of obtaining useful predictions about strategic play. Since many of the self-fulfilling prophecies seem arbitrary and fanciful, it is tempting to use some device to limit the "spurious bootstrapping". A modest application of subgame consistency accomplishes that in this paper, and when combined with robustness to behavioral perturbations, yields a unique division of surplus in the supergame.

The central result of the paper could be viewed as a demonstration that division of surplus according to the "Nash threats" solution can arise without any need for external enforcement of threats or agreements (see the quote from Nash (1953) in Section 1.1 above). We are more excited about the potential for application in oligopolistic settings and other infinite-horizon economic problems. It will also be important to understand how robust the Nash threats result is, across different classes of reputational perturbation.

5 References

Abreu, D. and F. Gul (2000), "Bargaining and Reputation," *Econometrica*, 68: 85-117.

Abreu, D. and D. Pearce (1999), "A Behavioral Model of Bargaining with Endogenous Types," mimeo.

Admati, A. and M. Perry (1987), "Strategic Delay in Bargaining," *Review of Economic Studies*, 54: 345-364.

Aoyagi, M. (1996), "Reputation and Dynamic Stackelberg Leadership in Infinitely Repeated Games,"; *Journal of Economic Theory,* 71:3.

Aumann, R. and L. Shapley (1976), "Long Term Competition: A Game Theoretic Analysis," mimeo, Hebrew University.

Aumann, R. and S. Sorin (1989), "Cooperation and Bounded Recall," *Games and Economic Behavior*, 1: 5-39.

Bernheim, D. (1984), "Rationalizable Strategic Behavior," *Econometrica*, 52: 1007-1028.

Celentani, M., D. Fudenberg, D. Levine, and W. Pesendorfer (1996), "Maintaining a Reputation Against a Long-Lived Opponent," *Econometrica* 64: 691-704.

Coase, R. (1972), "Durability and Monopoly," *Journal of Law and Economics*, 15: 143-149.

Compte, O. and P. Jehiel (2002), "On the Role of Outside Options in Bargaining with Obstinate Parties," *Econometrica*, 70(4):1477-1518.

Cramton, P. (1992), "Strategic Delay in Bargaining with Two-Sided Uncertainty," *Review of Economic Studies*, 59(1): 205-225.

Cripps, M., K. Schmidt and J.P. Thomas (1996), "Reputation in Perturbed Repeated Games," *Journal of Economic Theory*, 69: 387-410.

Fudenberg, D. (1992) "Repeated Games: Cooperation and Commitment," in *Advances in Economic Theory* (Sixth World Congress), Vol. 1, ed. J.J. Laffont, Cambridge University Press.

Fudenberg, D., D. Kreps, and D. Levine (1988), "On the Robustness of Equilibrium Refinements," *Journal of Economic Theory*, 44: 354-380.

Fudenberg, D. and D. Levine (1989), "Reputation and Equilibrium Selection in Games with a Patient Player," *Econometrica*, 57: 759-778.

Fudenberg, D. and E. Maskin (1986), "The Folk Theorem in Repeated Games with Discounting or with Incomplete Information," *Econometrica*, 54: 533-556.

Kambe, S. (1999), "Bargaining with Imperfect Commitment," *Games and Economic Behavior*, 28: 217-237.

Kreps, D., P. Milgrom, J. Roberts and R. Wilson (1982), "Rational Cooperation in the Finitely Repeated Prisoners' Dilemma," *Journal of Economic Theory*, 27: 245-252, 486-502.

Kreps, D. and R. Wilson (1982), "Reputation and Imperfect Information," *Journal of Economic Theory*, 27: 253-279.

Milgrom, P. and J. Roberts (1982), "Predation, Reputation and Entry Deterrence," *Journal of Economic Theory*, 27: 280-312.

Nash, J. (1950), "The Bargaining Problem," *Econometrica*, 18: 155-162.

Nash, J. (1953), "Two-Person Cooperative Games," *Econometrica*, 21-128-140.

Pearce, D. (1984), "Rationalizable Strategic Behavior and the Problem of Perfection, 52: 1029-1050.

Rubinstein, A. (1979), "Equilibrium in Supergames with the Overtaking Criterion," *Journal of Economic Theory*, 31: 227-250.

Rubinstein, A. (1982), "Perfect Equilibrium in a Bargaining Model," *Econometrica*, 54: 97-109.

Schmidt, K. (1993), "Reputation and Equilibrium Characterization in Repeated Games of Conflicting Interests," *Econometrica*, 61: 325-351.

Selten, R. (1973), "A Simple Model of Imperfect Competition, where 4 are few and 6 are many," *International Journal of Game Theory* 2, 141-201.

Stahl, I. (1972) *Bargaining Theory*, Stockholm: Economics Research Institute.

Watson, J. (1994), "Cooperation in the Infinitely Repeated Prisoners' Dilemma with Perturbations," *Games and Economic Behavior*, 7: 260-285.

Watson, J. (1996), "Reputation in Repeated Games with No Discounting" *Games and Economic Behavior*, 15: 82-109.

Figure 1:

Figure 2: