Nash Bargaining Part I - The Continuous Case

Prajit K. Dutta Satoru Takahashi

January 15, 2013

Abstract

This paper considers finite horizon alternating move two player bargaining with a *continuum* of agreements. Two models are studied. The first is a variant of the Rubinstein model in which an agreement, once reached, is never broken. The second model is new to the literature and allows for repeated bargains. In this model, an agreement can last for a few periods, be followed by periods of disagreement and then, potentially, a new and different agreement can emerge. It is shown that, if the Pareto frontier is concave, then in both models the unique limiting equilbrium is the Nash Solution, the agreement that maximizes the product of player payoffs.

1 Introduction

This paper considers finite horizon two player bargaining, an environment in which the players are trying to arrive at a split of the pie, with share α going to player 1 and $1-\alpha$ going to player 2. It examines the case where the potential agreements form a *continuum*, i.e., $\alpha \in [0, 1]$ and every one of those agreements is a Nash equilibrium for single-stage bargaining.

This multiplicity - that "any agreement is possible" - gives rise to the celebrated bilateral monopoly or bargaining problem**,** a problem that the pre-game theory generation of economists (including Edgeworth and Hicks) considered essentially indeterminate. The game-theoretic **bargaining question** then is: *are there bargaining protocols, modeled as non-cooperative games, that yield a unique equilibrium, a single agreement*?

To be sure, the first such proposal that yielded a unique prediction was set not in a non-cooperative framework but rather in a cooperative game; by Nash (1950) in his axiomatic derivation of the well-known "Nash Solution." to the bargaining problem[1] The best known non-cooperative answer is Rubinstein (1982) who showed that if players take turns bargaining (with the player at time t either accepting the offer on the table or proposing an alternative share), then there is a unique subgame perfect equilibrium to the bargaining game. Moreover, when

[1] The Nash Solution picks that share α for which the product of the two players' payoffs is *maximized.*

the players' discount factor approaches 1, then the unique equilibrium converges to the "Nash Solution".

There is by now a large literature studying variations of the Rubinstein game. In this paper, we study two models and the first one, which we call the *Irreversible Agreements* model, is a variation on Rubinstein. That model has the same alternating move structure, that a player at time t either accepts the existing offer or counter-proposes. It also presumes that once an agreement is reached, it cannot be undone and a flow of payoffs corresponding to the agreement is received every period till the end of the game.[2] The one difference between our *Irreversible Agreements* model and Rubinstein's is that our agreement remains in force for the remaining periods of a *finite horizon* whereas, in Rubinstein's case, it remains in force for the infinite future. Another difference is that we allow very general payoffs.

The real departure from the literature though is our second model, one of *Reversible Agreements.* This is a model of repeated bargaining, in much the same sense in which we have models of repeated oligopoly or repeated trade agreements. It captures the idea that when two negotiators sit down to negotiate an agreement they do not always expect it to be the last such negotiation. They believe that the agreement they are about to hammer out will last a certain number of periods, say ζ, and they also know that yet another agreement will need to be negotiated thereafter (and, of course, that one might end up being identical to the current agreement). This feature of repeated bargaining is, we believe, pervasive. When the auto companies negotiate with unions they are aware that the agreements are short-term and a precursor to subsequent negotiations. When two countries negotiate a bilateral trade or defense or industrial treaty, again those are typically short-term treaties.

Pervasive as repeated bargaining is, best that we can tell, it has been ignored by the literature. This paper is a first step towards an analysis.

As a first step, there are two modeling choices that we make. First, we take the length of the current agreement ζ to be one, i.e., each agreement, once reached, lasts one period. This is easily generalized.

Second, the current agreement becomes - by assumption - the status quo option for the subsequent negotiation. Suppose that an agreement has been reached at time t; player 1 responded to the offer on the table, $1-\widehat{\alpha}$ for player 2, by agreeing to take share $\widehat{\alpha}$. In the next period $t+1$, this "move" of player 1, the share $\widehat{\alpha}$ then becomes the status quo proposal to which player 2 has to respond. If he agrees to take $1-\widehat{\alpha}$, the agreement continues for one more period - with associated positive payoffs. Alternatively, if he responds with $1-\widetilde{\alpha}$ there is no longer an agreement in period $t+1$ (and either way, the period $t+1$ agreement proposal $1-\widehat{\alpha}$ or the disagreement proposal, $1-\widetilde{\alpha}$, is the offer on the table at period $t+2$). And so on.

That the agreement in period t forms the status quo for period $t+1$ is one

[2]The standard interpretation of the Rubinstein infinite horizon model is that play ends once agreement is reached with a terminal payoff corresponding to the agreement. That, of course, is equivalent to saying that flow payoffs continue to be accumulated every one of the remaining periods at a rate that adds up to the terminal payoff.

way for the different periods or bargains to be linked through time. If $t+1$ was a "fresh start" then, of course, the Irreversible Agreements model would be simply a sequence of Reversible Agreements models and the results of the latter would apply to the former. There are certainly other modeling assumptions that could link periods.

Now to the results. In what follows let T stand for the game horizon. It should be clear that uniqueness results will only hold for T large. When T is small, we are likely to get multiplicity since we know, for example, that when T is one, we have an infinity of possible equilibria. Indeed, given that the agreements form a continuum, it should not surprise the reader that we will uniquely choose from the continuum only when $T \longrightarrow \infty$. In what follows, an "equilibrium when $T \longrightarrow \infty$" should be understood to be the limit of equilibria for finite T. Furthermore, the continuum of agreements in the stage game will be assumed to define a Pareto frontier f that is downward sloping (since a large share for player 1 is, by definition, a smaller share for player 2).

In the *Irreversible Agreements - or single bargain* - model we prove three results, in each case when $T \longrightarrow \infty$. First, we show that as long as f is continous, there is a unique selection, i.e., although there may be multiple shares that are equilibrium shares when T is finite, there is only one α^* that is the limiting share (of player 1). An example then shows that it is possible that this limit depends on which player is the last to move.[3] Second, if f is continuously differentiable, it must be the case that each of these two possible limits is a local maximum for the product of payoffs function, i.e., we must always get a *local Nash Solution*. Third, if f is additionally, concave, then there is a unique limit and it corresponds to the global Nash Solution.

We then turn to the *Reversible Agreements - or multiple bargains* - model. We first present two examples - first when the Pareto frontier is linear and second when the Pareto frontier is circular. We show that when there are, say, τ periods left, the mover in that period will only accept agreements that give her intermediate payoffs, i.e., she will block payoffs that are two low for her and will also block payoffs that are too high for her. The first - blocking bad agreements - is also a feature of the Irreversible Agreements model. The interesting new feature with repeated bargaining is the second - blocking agreements that are too good!

In particular, there are shares $\alpha_1(\tau) < \alpha_2(\tau)$ - which we compute - such that only agreements that lie in $[\alpha_1(\tau), \alpha_2(\tau)]$ are accepted when there are τ periods left. Moreover, the acceptance region shrinks with the horizon length because $\alpha_1(\tau+1) \geqslant \alpha_1(\tau)$ and $\alpha_2(\tau+1) \leqslant \alpha_2(\tau)$. In other words, the longer is the game horizon the less willing is a player to accept a deal. The reason why advantageous agreements are rejected is that a player understands that the very good agreement might be short-lived. For instance, if player 1 accepts $\widetilde{\alpha} > \alpha_2(\tau+1)$, she fears that the short-term high payoff would then be followed by a break-down in the agreement next period. Player 2, rather than take a

[3] The player that is last to move accepts any share - when she finds herself in the last period. That, consequently, puts her at a dis-advantage from a bargaining point of view.

continuance of low payoffs associated with the share $1-\widetilde{\alpha}$ would take a short-term loss of payoffs - due to a disagreement - and instead look for the higher share implied by $1-\alpha_1(\tau-1)$. So, ongoing intermediate payoffs are better than the "stop-start" of a high initial payoff followed by a long stretch of low payoffs.

The main theorem of the paper - and the general result of the second model - then shows that what is shown by computation in the examples is more broadly true. Provided f is continuously differentiable and concave, only intermediate agreements are equilibria. Furthermore, as $T \longrightarrow \infty$, the unique limit of the intermediate acceptance region is a singleton - and that singleton is again the Nash Solution. In other words, when T is large but finite, there could possibly be multiple equilibrium agreements but they all have to be (arbitrarily) close to the Nash Solution, fitting in ever more tightly around that solution as we take longer horizons.

In Summary, alternating offer bargaining produces the Nash Solution as the unique outcome and does so not only when there is a single bargain but even when there are multiple bargains.

The astute reader will wonder why we have analyzed a finite horizon problem rather than an infinite horizon problem. The reason is that the infinite horizon would have a folk theorem and certainly there would be no chance of a unique equilibrium selection; indeed, the bargaining multiplicity would dimensionally increase since not only would every agreement emerge as an equilibrium but so would dis-agreement outcomes that are individually rational. That there is a folk theorem follows from the fact that any alternating move game is a stochastic game where the state variable at any time is the fixed action of the non-mover. By the result of Dutta (1986) such a stochastic game would have a folk theorem.[4]

The other modeling detail is that, in this paper, we are considering a continuum of agreements. In a companion paper - Dutta and Takahashi (2012) - we have analyzed the case where there are only a finite number of agreements. The advantage of the finite case is that it is a little easier to understand in that setting why the Nash solution emerges in equilibrium. We explain that reasoning now. Consider the following stage game:

$1\backslash 2$	$1-\alpha_1$	$1-\alpha_2$
α_1	h_1, l_2	$0, 0$
α_2	$0, 0$	l_1, h_2

with $h_i > l_i > 0$ for $i = 1, 2$.

We proceed by backward induction. In the last period, it is a best response to accept whatever share is offered. Similarly, as long as there are only a few remaining periods. Things are different when the number of remaining periods

[4] Strictly speaking, the Dutta (1986) result only applies when the state space is finite whereas the state space - the set of possible shares - in this setting is a continuum. However, at the risk of considerably heightened technical analysis, the Dutta result can be extended to this setting.

becomes larger than T_i, where T_i is defined as

$$(T_i - 1)h_i = T_i l_i$$

or $T_i = h_i/(h_i - l_i)$. Now, by blocking a bad agreement, and counter-proposing the better one, a player receives a payoff of 0 in the current period but continues with payoffs of h_i in the future, whereas by accepting a bad agreement, he receives a payoff of l_i in each of the remaining periods.

Suppose for example that $T_1 > T_2$. Then when the horizon increases past T_2 player 2 rejects α_1. This leads to the equilibrium outcome $\alpha_2, 1 - \alpha_2$ which favors player 2. Conversely, if $T_1 < T_2$.

What determines $T_1 \gtrless T_2$? Simple algebra shows that $T_1 > T_2$ if and only if $l_1 h_2 > h_1 l_2$, i.e., comparison between two Nash products.[5]

The Model is presented in Section 2. In Section 3 we present two examples to illustrate possible equilibria in the Irreversible Agreements model while in Section 4 we present the three main results within that model. Sections 5 and 6 parallels that in the Reversible Agreements case; Section 5 presents two examples while Section 6 presents the main theorem. Generalizations and extensions are collected in Section 7.

2 Model

In this section, we present two finite horizon bargaining models - one with *irreversible agreements* and the second with *reversible agreements*. But first we discuss in some detail the bargaining stage game that we will focus on and the timing structure within which we study that game.

2.1 Bargaining Stage Game

Consider a two player bargaining environment in which the players are trying to arrive at a split of the pie, with share α going to player 1 and $1 - \alpha$ going to player 2. Suppose that the potential agreements α form a *continuum*, i.e., $\alpha \in [0, 1]$. Suppose further that $x(\alpha)$ represents the agreement payoff of player 1 whilst $y(\alpha) = f(x(\alpha))$ is the agreement payoff of player 2. (Frequently, we will suppress the argument α and simply write agreement payoffs as $(x, f(x))$). Players' actions are "proposals" on shares. The following assumptions are made about this bargaining stage game:

Definition *The stage game is defined by strategy sets: player 1 can ask for a share $\widehat{\alpha} \in [0,1]$ and player 2 can ask for a share $1 - \widetilde{\alpha} \in [0,1]$. The asks are compatible if $\widehat{\alpha} = \widetilde{\alpha}$ and we say then that there is an agreement. The asks are incompatible if $\widehat{\alpha} \neq \widetilde{\alpha}$ and then we say that we have a disagreement. The following assumptions are made on the associated payoffs:*

[5] There are a number of knotty "integer problems" however that bedevil the finite agreements case. For further details, see Dutta and Takahashi (2012).

1. *disagreement strategy tuples have a constant payoff (normalized to* 0*), i.e, the payoff to* $\widehat{\alpha}, 1-\widetilde{\alpha}$*, where* $\widehat{\alpha} \neq \widetilde{\alpha}$ *is* (0,0);

2. *each agreement strategy tuple is a strict Nash equilibrium, i.e., whenever* $\widehat{\alpha} = \widetilde{\alpha}$, $x(\widehat{\alpha}) > 0$ *and* $f(x(\widehat{\alpha})) > 0$*;*

3. *the Pareto frontier is downward sloping, i.e.,* $x(.)$ *is strictly increasing in* α *while* f *is strictly decreasing.*

An example will help clarify the definition.

Example 1 *With* $x \in [0,1]$ *suppose that* $f(x) = 1 - x$.

This is the linear bargaining problem that has been studied in many papers including (.) Here, $\alpha = x$ and so we can equivalently think of the choice as being either x or α. In fact, given the monotonicity assumptions in 3. above, in general we can equivalently think of the choice as being either of those variables. Within this model section we will continue to talk of the choice as being the share α. In the next sections, when we get to the analysis, we will often talk about the (direct) choice of payoff x (by player 1 and $y = f(x)$ by player 2).

Obviously this bargaining problem - sometimes also referred to as the bilateral monopoly problem - has a continuum of equilibrium outcomes when the stage game is played simultaneously by the two players; each of the strategy tuples $(\widehat{\alpha}, 1-\widehat{\alpha})$ is a Nash equilibrium with associated payoffs $(x(\widehat{\alpha}), f(x(\widehat{\alpha})))$. Hence, the (*Nash*) *bargaining problem* that arises is: which of the agreements will emerge from strategic rational play?

John Nash first proposed, in an axiomatic derivation, the following (Nash bargaining) solution:[6]

Definition *The Nash bargaining solution in a bargaining stage game is the agreement* $x^m, f(x^m)$ *that maximizes the product of the players' payoffs, i.e.,*[7]

$$x^m f(x^m) = \max_x \; x.f(x)$$

2.2 Alternating Offers and Finite Horizon

The multiplicity of bargaining equilibria worsens if bargaining takes place over many periods. If the same stage game is played simultaneously and repeatedly, i.e., an agreement is sought afresh every period, then the well-known folk theorems apply in both the infinite as well as the finite horizon versions. (Fudenberg and Maskin (1986) and Benoit and Krishna (1985).) Not only can agreements emerge then as equilibria but also periods of disagreements (since every individually rational payoff can emerge as subgame perfect equilibrium payoff).

[6] And subsequently as a limit solution to a class of non-cooperative one period simultaneous bargaining games. See Nash (1950, 1953).

[7] Note that in Nash's formulation what is maximized is the product of net payoffs, where each player's net payoff is the gross payoff less a status quo payoff. Here the status quo, or disagreement, payoff is zero.

Rubinstein (1982) was the first paper to cut through this indeterminacy. He restricted attention to a single bargain, i.e., an attempt to find one agreement but did so in a model of alternating offers. He considered a model in which in every period only one of the two players moves. She does so by responding to an offer that has been made in the previous period, an offer that specifies a share for the proposer. A response is a counter-proposal that might or might not lead to an agreement. If the proposal is, say, $\widehat{\alpha}$, the responder could then say "yes" with a counter-proposal of $1-\widehat{\alpha}$ i.e., agree. On the other hand, the mover could say "no" via a counter-proposal of $1-\widetilde{\alpha}$ where $\widetilde{\alpha} \neq \widehat{\alpha}$, i.e., disagree, and that proposal would then become the tabled offer for the subsequent period.

In Rubinstein's model, whenever an agreement is reached, the game ends with each player receiving the agreement payoffs.

The above will be exactly the first of two models that we will study. We will call that the *Irreversible Agreements* model because - once an agreement is reached - it cannot be undone. The one difference between our *Irreversible Agreements* model and Rubinstein's will be that our agreement will remain in force for the remaining periods of a finite horizon whereas, in Rubinstein's case, it remains in force for the infinite future.[8]

The second model will be one in which an agreement is in force for the period in which it is reached but need not be in force in subsequent period(s). In that model, if the counter-proposal is $1-\widehat{\alpha}$ then both players get agreement payoffs in that period. Next period, the mover responds to the existing (non-mover's) previous offer - $1-\widehat{\alpha}$. If she responds with $\widehat{\alpha}$ then the agreement continues for one more period - with associated positive payoffs. Alternatively, she is free to respond with $\widetilde{\alpha}$ thereby collecting one period of zero disagreement payoffs but possibly moving play to a better - for her - agreement in the future.

Since we study finite horizon games that end after period T $(<\infty)$ we will need, to begin with, a move convention about who moves in the last period. One possibility is that it is player 1 who moves last and hence in the penultimate period it is player 2 who moves etc. In other words, player 1 is the mover when there are an odd number of periods left whilst player 2 is the mover when there are an even number of periods left. An alternative convention is to consider the opposite possibility with player 2 being the mover when there is only one period left to go.

It will be seen, that when there are very few periods of bargaining left, the chosen convention will give asymmetric bargaining power to the player who does *not* move last. By looking at long enough bargaining periods, however, we will search for equilibria where this convention plays no role. If it does we will call such an equilibria *local* and if it does not we will call it *global.* (More below.)

Lifetime payoffs will be evaluated according to the (undiscounted) *average payoff;* player $i's$ evaluation of the payoff stream $\pi_i(a_t, b_t)$, $t = 1, 2, \ldots T$ will

[8]The standard interpretation of the Rubinstein infinite horizon model is that play ends with payoffs $(x(\widehat{\alpha}), f(x(\widehat{\alpha})))$. That, of course, is equivalent to saying that flow payoffs $(1-\delta)(x(\widehat{\alpha}), f(x(\widehat{\alpha})))$ continue to be accumulated every one of the remaining infinite periods. (δ is the discount factor in the above.)

be given by[9]

$$\frac{1}{T}\sum_{t=1}^{T}\pi_i(a_t, b_t) \tag{1}$$

One further specification that is required is how the game starts, i.e., what is the "initial state" of the game, the offer that is already on the table at the beginning of period 1. There are two standard ways to do this: one is to be agnostic and allow all possible initial states. A second is to consider a period 0 and allow the "mover" in that period to choose the initial state (for period 1). Our results do not depend on which of these specifications is chosen. For concreteness, we will allow any possible initial state and will look for equilibris that are independent of the initial state.

2.3 Model 1 - Irreversible Agreements

Definition: *A bargaining game (with horizon T and stage game G) is said to have* ***irreversible agreements*** *if, starting with an offer $\widehat{\alpha}$ (respectively, $1-\widehat{\alpha}$), provided the counter-offer is $1-\widehat{\alpha}$ (respectively, $\widehat{\alpha}$), i.e., if an agreement is reached in period t, it then remains in force for the remaining $T-t$ periods, i.e., the strategy tuple played in periods $t, t+1, \ldots T$ is also given by $(\widehat{\alpha}, 1-\widehat{\alpha})$ with associated payoffs $x(\widehat{\alpha}), f(x(\widehat{\alpha}))$ every period.*

Put differently, if an agreement is reached, there is no further opportunity to make offers (and counter-offers). In the sense that there is only one agreement, this model is essentially the finite horizon Rubinstein model.

2.4 Model 2 - Reversible Agreements

The more dynamic model is one of reversible agreements.

Definition: *A bargaining game (with horizon T and stage game G) is said to have reversible agreements if, when an offer $\widehat{\alpha}$ (respectively, $1-\widehat{\alpha}$) is made and followed by the agreement counter-offer $1-\widehat{\alpha}$ (respectively, $\widehat{\alpha}$)* in period t, *it remains in force for only that one period. In period $t+1$, the acceptance offer of the other player, $1-\widehat{\alpha}$ (respectively, $\widehat{\alpha}$), becomes the offer to which the responder in that subsequent period responds. Similarly, if the offer was not accepted in period t, then the counter-offer constitutes the offer to which the responder in period $t+1$ responds.*

Let us consider some examples:

Example 1 *Agreement - Initially and Forever:* Supose the initial state is $\widehat{\alpha}$. Player 2, the mover in period 1 say, responds by playing $1-\widehat{\alpha}$. This constitutes an agreement with associated positive payoffs in period 1. In period 2, the responding player, now player 1, picks $\widehat{\alpha}$ in response to the current offer of $1-\widehat{\alpha}$ and the players earn the same payoff in period 2 as well. And this continues

[9] Here $a_t = \widehat{\alpha}$ for player 1 and $b_t = 1-\widetilde{\alpha}$ for player 2. Furthermore, $\pi_i(a_t, b_t) = x(\widehat{\alpha})$ or $f(x(\widehat{\alpha}))$ respectively, if there is an agreement and equals 0 otherwise. Our analysis also applies to the case of discounting, for any discount factor $\delta < 1$. In that case the appropriate evaluation of lifetime payoffs is $\frac{1-\delta}{1-\delta^T}\sum_{t=1}^{T}\delta^{t-1}\pi_i(a_t)$.

till the end of the game garnering them total payoffs of $T[x(\widehat{\alpha}), f(x(\widehat{\alpha}))]$ or per-period average payoffs of $x(\widehat{\alpha}), f(x(\widehat{\alpha}))$.

Example 2 *Disagreement Transiting to An Agreement:* Again, supose the initial state is $\widehat{\alpha}$. Player 2 responds by playing $1-\widetilde{\alpha}$, $\widehat{\alpha} \neq \widetilde{\alpha}$. This constitutes a disagreement with associated zero payoffs in period 1. However, player 2 might thereby have put on the table an offer whose agreement, if reached, is far more favorable for her. There might follow a set of offers and counter-offers none of them constituting an agreement and so stage payoffs might remain zero. However, the series of rejections might then narrow the set of potential agreements to something that is a better "compromise" for both parties, say α'. Under the twin pressures of finally getting an offer that is more reasonable for the responder - and that the end of the game is nearing - the two players might then reach the agreement $\alpha', 1-\alpha'$ for, say, the last $T-\zeta$ periods.

Note that both models are examples of dynamic or stochastic games with the payoff-relevant state at time t being the fixed action in that period (to which the responder responds) and the number of remaining periods in the game.

2.5 Equilibrium - Global and Local Convergence

All actions are observable. Strategies, in both models, are hence defined in the usual way as history-dependent action choices. A $t-th$ period strategy for a mover in that period is a history-dependent (mixed) action choice and a complete strategy is a specification for every move period (and after every history). A strategy vector - one strategy for every player - defines in the usual way a (possibly probabilistic) action choice and hence an expected payoff for the $t-th$ period for each player.

Since the game has a finite horizon and since this is a game of perfect information, Subgame Perfect Equilibrium (SPE) is determined via backward induction in both models. When there is a single period left, the mover in that period, player 1 say, plays her best reponse; call that $\alpha^*(\beta, 1)$ where - henceforth - $\beta = 1-\alpha$ is the previous period proposal of player 2. Evidently, the two variables - the offer on the table β and the remaining number of periods, 1, is all that is payoff-relevant in selecting a best response. Similarly, at the penultimate period, the mover - player 2 - has a best response $\beta^*(\alpha, 2)$ based on the proposal on the table, α, the number of remaining periods, 2, and the best response function $\alpha^*(\beta, 1)$ that will govern the subsequent period's action. In this fashion, we derive best response functions for player 1, $\alpha^*(\beta, \tau)$, whenever there are an odd number of periods τ remaining and best response functions $\beta^*(\alpha, \tau)$ for player 2 when τ is even. Note that in this perfect information setting optimal choices are, typically, Markovian and unique. The only way that a mixed strategy SPE - or a history-dependent SPE - can arise is from ties in the mover's payoffs.[10]

[10] Note that in both models, the *set* of SPE payoffs is always $history-independent$; the set of payoff vectors that can arise as a SPE in the subgame that starts at date t and a particular history only depends on the current state, i.e., the offer on the table (and the number of remaining periods τ).

This paper will investigate properties the equilibrium choice and value sets must satisfy, when τ is "large". Towards that end, let us define the equilibrium agreements

$$\begin{aligned} A^*(\tau) &= \{\alpha : \beta^*(\alpha, \tau) = 1 - \alpha\},\ \tau \ even \\ A^*(\tau) &= \{\alpha^*(\beta, \tau) : \exists \beta \ s.t. \ \alpha^*(\beta, \tau) = 1 - \beta\},\ \tau \ odd \end{aligned}$$

As should be clear from the definitions, A^* is the set of agreements that can arise from equilibrium play, either because player 1 will agree to take that share - when τ is *odd* or because that is a share that player 2 would be willing to give player 1 - when τ is *even*. Finally, define the limit

$$A^* = \cap_\tau A^*(\tau)$$

whenever it exists. In particular, if the limit exists and is, furthermore, a singleton, then we shall say that the bargaining game has a unique equilibrium agreement for "long" horizons:

Definition 2 *The alternating move, finite horizon Irreversible Agreement game - respectively, Rreversible Agreement game - has a unique equilibrium if A^* is a singleton.*

Let A^* have a singleton and suppose it is α^*. It should be clear in the Irreversible Agreement game that, if the horizon is long enough, then the only initial proposals, that will get accepted (immediately and thereafter stay in place by definition) are proposals in the neighborhood of α^*. The longer the horizon the smaller is that acceptance neighborhood.

It should also be clear that even in the Reversible Agreement game that, if the horizon is long enough, then the only initial proposals, that will get accepted (immediately and thereafter stay in place through equilibrium behavior) are proposals in the neighborhood of α^*. Again, the longer the horizon the smaller is that acceptance neighborhood. In this instance, these equilibria with agreements that are never broken, are, in principal, a subset of more complicated equilibrium behavior that might involve, for instance, agreement-disgreement cycles. We will comment on that in Section 5.

When we first defined the finite horizon game, we had two parameters - initial state and horizon length - and one convention to specify - the last mover. In the definitions and arguments immediately above, we have found concepts where the two parameters cease to be relevant - since the limit acceptance set is independent of them. However, we still have the convention that the last mover be player 1. Clearly, an alternative convention would be that the last mover is player 2. Everything that we have done so far can be re-cast in this parallel world: with one period to go, player 2 has a best response function $\beta^{'}(\alpha, 1)$ and based on that player 1 has a best response function with two periods to go $\alpha^{'}(\beta, 2)$. In this fashion, we derive best response functions for player 1, $\alpha^{'}(\beta, \tau)$, whenever there are an even number of periods τ remaining and best response functions $\beta^{'}(\alpha, \tau)$ for player 2 when τ is odd.

Finally, let us define the equilibrium agreements

$$\begin{aligned} A^{'}(\tau) &= \{\alpha : \beta^{'}(\alpha,\tau) = 1-\alpha\},\ \tau \ odd \\ A^{'}(\tau) &= \{\alpha^{'}(\beta,\tau) : \exists\beta \ s.t.\ \alpha^{'}(\beta,\tau) = 1-\beta\},\ \tau \ even \end{aligned}$$

and the associated limit

$$A^{'} = \cap_{\tau} A^{'}(\tau)$$

whenever it exists. In particular, if both limits exist and are, furthermore, singletons equal to each other, then we shall say that the bargaining game has a unique **global equilibrium** agreement:

Definition 3 *The alternating move, finite horizon Irreversible Agreement game - respectively, Rreversible Agreement game - has a unique global equilibrium if* A^* *and* A' *are singletons and* $A^* = A'$.

3 Irreversible Agreements: Examples of Global and Local Convergence

In this section we provide two examples within Model 1, the model where agreement, once reached, is irreversible. The first is an example where the Pareto frontier is linear and this example will exhibit global convergence; regardless of whether it is player 1 or 2 who is the last mover, there is a unique agreement in the limit. In fact, in this example, the agreement is reached within two periods. Moreover, that unique agreement is nothing but the Nash bargaining solution. The example will also illustrate an algorithm that will re-appear in a more general result which is to be found in the next section, a result that will apply to broader classes of Pareto frontiers - beyond the linear - that have the property that the Nash product $x.f(x)$ is single-peaked.

The second example will be one in which the Pareto frontier is piece-wise linear and convex (i.e., kinked "inwards"). Consequently, the Nash product will be multi-peaked. In this example, there will be local convergence but not global convergence in that there will be two candidate limits, say α^* and α', $\alpha^* < \alpha'$ and which of them is reached will depend on which player is the last mover. If player 1 is the last to move that puts her at a dis-advantage and consequently the limit is α^*. On the other hand, if player 2 is the last to move that places greater bargaining power in the hands of player 1 and leads to an agreement α' that is more advantageous for her. Furthermore, one of the two limts will be seen to be the Nash bargaining solution and the other will be seen to be a "local" Nash bargaining solution, i.e., a local maximizer of the Nash product $x.f(x)$. The example will also illustrate a more general result which is to be found in the next section. The general result will show that bargaining leads always to a local Nash bargaining solution.

3.1 Linear Pareto Frontier

Suppose that

$$f(x) = -mx + c, \ m > 0, c > 0$$

Note that $argmax_x \ xf(x) = \frac{c}{2m}$. We will use this example to investigate global convergence. In terms of the analysis, it will turn out to be easier to imagine that what each player chooses is a payoff - rather than a share of the pie. For example, player 1 accepts all agreements that give her a certain set of payoffs. Since the payoff function $x(\alpha)$ is monotonically increasing in the share α, it clearly does not matter whether we have the player choosing x or α. Similarly, for player 2 we will imagine that he chooses $y = f(x)$ rather than the share $1 - \alpha$ that yields that payoff.

Within this example, we will prove the following:

Proposition 4 *If f is linear, then there is global convergence. Explicit computation shows that the equilibrium strategies are of the form: there are payoffs $\underline{x}(\tau)$ and $\overline{x}(\tau)$, $\underline{x}(\tau) < \overline{x}(\tau)$, that form the boundary of the (payoff) agreement region. If it is player 1's turn to move, with τ periods to go, she will only agree provided the proposal on the table gives her at least $\underline{x}(\tau)$ as payoff. Similarly, if it is player 2's turn to move, with τ periods to go, he will only agree provided the proposal on the table gives him at least $f(\overline{x}(\tau))$. Furthermore, $\underline{x}(\tau) < \frac{c}{2m} < \overline{x}(\tau)$. Also, $\underline{x}(\tau + 2) \geqslant \underline{x}(\tau)$ for all τ odd. Similarly, $\overline{x}(\tau + 2) \leqslant \overline{x}(\tau)$ for all τ even. Finally, $\lim_{t\to\infty} \underline{x}(\tau) = \lim_{t\to\infty} \overline{x}(\tau) = \frac{c}{2m}$.*

Recall that τ is the number of remaining periods and player 1 is the player who moves in the last period.

As noted, $\underline{x}(\tau)$ for $\tau = 1, 3, 5$ ... denotes the (worst) agreement payoff that player 1 is willing to accept at those periods (and let the associated payoff for player 2 be denoted $\overline{y}(\tau)$, i.e., $\overline{y}(\tau) = f(\underline{x}(\tau))$). Let $\underline{y}(\tau)$ for $\tau = 2, 4, 6,$ ... denote the worst agreement that player 2 will similarly accept when it is his turn to move. Let $\overline{x}(\tau) = \frac{\underline{y}(\tau) - c}{-m}$ denote the associated payoff for player 1, i.e., that is the best agreement that player 2 will receive in even periods.

We will now compute $\underline{x}(\tau)$ and $\overline{x}(\tau)$ and show that

$$\lim_{\tau\to\infty} \underline{x}(\tau) = \lim_{\tau\to\infty} \overline{x}(\tau) = \frac{c}{2m}$$

Hence, the state-independent SPE is the Nash solution.

Proof. We will compute the agreement regions. We start with ■

$\tau = 1$ It is easy to see that player 1 will agree to any share when there is a single period left because agreement gives her a non-negarive payoff (and, in fact,a strictly positive payoff if the proposed share is positive). Hence, $\underline{x}(1) = 0$ implying $\overline{y}(1) = c$

$\tau = 2$ With two periods to go, player 2 can guarantee himself $\overline{y}(1) = c$ by turning down a proposal in this period and waiting for the next with a counter-proposal that gives him everything. Hence, the worst payoff that he will accept is given by the equality $c = 2\underline{y}(2)$ implying that $\underline{y}(2) = \frac{c}{2}$ implying that $\overline{x}(2) = \frac{c}{2m}$

$\tau = 3$ With three periods to go, player 1 can guarantee herself $\overline{x}(2) = \frac{c}{2m}$ by turning down a proposal on the table and waiting for her best payoff for two periods. Hence, the worst payoff that she will accept is given by the equality $\frac{c}{m} = 3\underline{x}(3)$ implying that $\underline{x}(3) = \frac{1}{3}\frac{c}{m}$ i.e., that $\overline{y}(3) = \frac{2}{3}c$

$\tau = 4$ Continuing in that fashion, $\underline{y}(4)$ can be derived from $2c = 4\underline{y}(4)$ implying that $\underline{y}(4) = \frac{c}{2}$ implying that $\overline{x}(4) = \frac{c}{2m}$

$\tau = 5$ Again, the worst payoff the mover, player 1, will take is given by $\frac{2c}{m} = 5\underline{x}(5)$ implying $\underline{x}(5) = \frac{2}{5}\frac{c}{m}$ implying $\overline{y}(5) = \frac{3}{5}c$

$\tau = 6$ $3c = 6\underline{y}(6)$ implying that $\underline{y}(6) = \frac{c}{2}$ implying that $\overline{x}(6) = \frac{c}{2m}$

$\tau = 7$ $\frac{3c}{m} = 7\underline{x}(7)$ implying $\underline{x}(7) = \frac{3}{7}\frac{c}{m}$ implying $\overline{y}(7) = \frac{4}{7}c$

$\tau = 8$ $4c = 8\underline{y}(8)$ implying that $\underline{y}(8) = \frac{c}{2}$ implying that $\overline{x}(8) = \frac{c}{2m}$

$\tau = 9$ $\frac{4c}{m} = 9\underline{x}(9)$ implying $\underline{x}(9) = \frac{4}{9}\frac{c}{m}$ implying $\overline{y}(9) = \frac{5}{9}c$

$\tau = 10$ $5c = 10\underline{y}(10)$ implying that $\underline{y}(10) = \frac{c}{2}$ implying that $\overline{x}(10) = \frac{c}{2m}$

It is clear that the odd-period worst agreement payoffs of player 1, $\underline{x}(\tau)$ - after factoring out the $\frac{c}{m}$ - are

$$0, \frac{1}{3}, \frac{2}{5}, \frac{3}{7}, \frac{4}{9}, \ ...\frac{n-1}{2n+1}, ...$$

for $n = 1, 2, ..$ Of course that sequence converges to $\frac{1}{2}$. It is also clear that the best payoffs of player 1, $\overline{x}(\tau) = \frac{c}{2m}$ for all even numbered τ. Hence, the unique limit payoff for player 1 is $\frac{c}{2m}$, the Nash bargaining solution.

Something that the astute reader would have already noted - and that will be useful in the sequel - is the following observation. In the odd periods, when it is player 1's turn to move the lowest payoff that she will accept is given by the following equation

$$\tau\underline{x}(\tau) = (\tau - 1)\overline{x}(\tau - 1)$$

and, similarly, in the even periods, the lowest payoff that player 2 would be willing to take is given by

$$\tau f(\overline{x}(\tau)) = (\tau - 1)f(\underline{x}(\tau - 1))$$

This equations will be extensively used in the general analysis that follows.

Let us now turn to global convergence - that the same limit would hold had we considered the case in which player 2 is the last mover. Easy computation, again using equations above, tell us that the odd period payoffs of player 1, now $\overline{x}(\tau)$, are - after factoring out $\frac{c}{m}$ - given by

$$1, \frac{2}{3}, \frac{3}{5}, \frac{4}{7}, \frac{5}{9}, \ ...\frac{n}{2n-1}, ...$$

for $n = 1, 2, ..$ Of course that sequence converges to $\frac{1}{2}$. It is also clear that the worst payoffs of player 1, $\underline{x}(\tau) = \frac{c}{2m}$ for all even numbered τ. Hence, again, the unique limit payoff for player 1 is $\frac{c}{2m}$, the Nash bargaining solution and we have global convergence.■

3.2 Convex Pareto Frontier

We turn now to a second example. In this case, the Pareto frontier will be piece-wise linear and kinked "inwards" thereby creating a convex frontier. Suppose that

$$\begin{aligned} f(x) &= -2x+10,\ x \leqslant \frac{8}{3} \\ f(x) &= -0.5x+6,\ x \geqslant \frac{8}{3} \end{aligned}$$

Note that the Nash product. $xf(x)$, is double-peaked, achieveing a local maximum in the left line-segment at 2.5 and another at 6 - which is where it also achieves the global maximum, i.e., $argmax_x\ xf(x) = 6$. We will use this example to show - a) there need not be global convergence, i.e., that the agreement that is eventually reached might be different depending on whether it is player 1 or 2 who moves last but that, b) it must be a local maximum of the Nash product function.

Proposition 5 *Consider the convex f example above. Explicit computation shows that player 1's equilibrium worst and best payoffs $\underline{x}(\tau)$ and $\overline{x}(\tau)$, $\underline{x}(\tau) < \overline{x}(\tau)$, have the following property:* $\lim_{t\to\infty} \underline{x}(\tau) = \lim_{t\to\infty} \overline{x}(\tau) = 2.5$ *if player 1 is the last mover but* $\lim_{t\to\infty} \underline{x}(\tau) = \lim_{t\to\infty} \overline{x}(\tau) = 6$ *if, instead, player 2 is the last mover. So there is no global convergence but there is salience to the Nash product function in that a) there is a limiting agreement and b) that limit has to be at least a local maximum of the Nash product function.*

Proof. As in the previous example, we will compute the agreement regions. We start with ■

$\tau = 1$ Clearly, $\underline{x}(1) = 0$ implying $\overline{y}(1) = 10$.

$\tau = 2$ With two periods to go, since player 2 can guarantee himself $\overline{y}(1) = 10$ by turning down a proposal in this period, the worst payoff that he will accept is given by the equality $10 = 2\underline{y}(2)$ implying that $\underline{y}(2) = 5$ implying that $\overline{x}(2) =\ 2.5$.

$\tau = 3$ With three periods to go, since player 1 can guarantee herself $\overline{x}(2) =\ 2.5$ by turning down a proposal on the table, the worst payoff that she will accept is given by the equality $5 = 3\underline{x}(3)$ implying that $\underline{x}(3) = \frac{5}{3}$ i.e., that $\overline{y}(3) = \frac{20}{3}$.

$\tau = 4$ Continuing in that fashion, $\underline{y}(4)$ can be derived from $20 = 4\underline{y}(4)$ implying that $\underline{y}(4) = 5$ implying that $\overline{x}(4) =\ 2.5$.

$\tau = 5$ Again, the worst payoff the mover, player 1, will take is given by $10 = 5\underline{x}(5)$ implying $\underline{x}(5) = 2$ implying $\overline{y}(5) = 6$.

$\tau = 6$ $30 = 6\underline{y}(6)$ implying that $\underline{y}(6) = 5$ implying that $\overline{x}(6) =\ 2.5$.

$\tau = 7$ $15 = 7\underline{x}(7)$ implying $\underline{x}(7) = \frac{15}{7}$ implying $\overline{y}(7) = \frac{40}{7}$.

$\tau = 8$ $40 = 8\underline{y}(8)$ implying that $\underline{y}(8) = 5$ implying that $\overline{x}(8) =\ 2.5$.

$\tau = 9$ $20 = 9\underline{x}(9)$ implying $\underline{x}(9) = \frac{20}{9}$ implying $\overline{y}(9) = \frac{50}{9}$.

$\tau = 10$ $50 = 10\underline{y}(10)$ implying that $\underline{y}(10) = 5$ implying that $\overline{x}(10) = 2.5$.

$\tau = 11 \qquad 25 = 11\underline{x}(11)$ implying $\underline{x}(11) = \frac{25}{11}$

It is clear that the odd-period worst agreement payoffs of player 1 are

$$0, \frac{5}{3}, \frac{10}{5}, \frac{15}{7}, \frac{20}{9}, \ \frac{25}{11} ... \frac{5(n-1)}{2n-1}, ...$$

for $n = 1, 2, ..$ Of course that sequence converges to 2.5. It is also clear that the best payoffs of player 1, $\overline{x}(\tau) =$ 2.5 for all even numbered τ. Hence, the unique limit payoff for player 1 is 2.5, which is a local maximum for the Nash product function but not the Nash bargaining solution.

Let us now turn to the limit that would hold had we considered the case in which player 2 is the last mover. Easy computation, again using equations above, tell us that the odd period payoffs of player 1, now $\overline{x}(\tau)$, are given by

$$\frac{12}{1}, \frac{24}{3}, \frac{36}{5}, \frac{48}{7}, \ ... \frac{12n}{2n-1}, ...$$

for $n = 1, 2, ..$ Of course that sequence converges to 6. It is also eaasily shown that the worst payoffs of player 1, $\underline{x}(\tau) =$ 6 for all even numbered τ. Hence, again, the unique limit payoff for player 1 is 6, and this is the Nash bargaining solution.

To summarize, in this example, there are two local maxima of the Nash product function. The one that is preferred by player 1 is approached if player 2 is the last to move - and vice-versa. However, those are the only possible equilibrium agreements. ■

4 Irreversible Agreements: Results on Global and Local Convergence

In this section we provide results on convergence, i.e., unique selection of a bargaining agreement. There are three results. The first says that as long as f is continous, then the sequences of best and worst agreement, $\overline{x}(\tau)$ and $\underline{x}(\tau)$ must converge regardless of whether player 1 or player 2 is the last to move. Call the limit when player 1 is the last to move x^* and that when player 2 is the last to move x'. Note that in the two examples of the previous section we were able to compute these limits directly. The first result of this section will therefore establish that - even when they are not computable - they must exist.

The second result says that both x^* and x' have to be local maxima of the Nash product function $x.f(x)$ as long as f is known to be continously differentiable. This was illustrated by the second example of the previous section. An immediate corollary of that result is that when the Nash product is single-peaked - for example if f is concave - then the two limits have to coincide and be the Nash bargaining solution. This was illustrated by what we computed when the Pareto frontier, f, is linear in the first example of the previous section.

4.1 Convergence

Theorem 6 *Suppose that f is continuous. Then, there is a limit to the best and worst agreements of player 1, $\overline{x}(\tau)$ and $\underline{x}(\tau)$, and the two limits coincide. This is true regardless of whether player 1 is the last to move or player 2.*
Proof. *As we have seen above, the two sequences are derived via the following equations*

$$\underline{x}(\tau+1) = \frac{\tau}{\tau+1}\overline{x}(\tau),$$
$$f(\overline{x}(\tau+1)) = \frac{\tau}{\tau+1} f(\underline{x}(\tau)).$$

From the first equation above it follows that

$$\frac{\tau+2}{\tau+1}\underline{x}(\tau+2) = f^{-1}\left(\frac{\tau}{\tau+1} f(\underline{x}(\tau))\right)$$

■

Recall that the sequence $\underline{x}(.)$ is defined with intervals of 2,; if player 1 is the last to move then it is of the form $\underline{x}(1), \underline{x}(3), ...$ $\underline{x}(\tau)$, $\underline{x}(\tau+2)...$ for all odd τ while it is of the form $\underline{x}(2)$, $\underline{x}(4), ...\underline{x}(\tau)$, $\underline{x}(\tau+2)...$ for all even τ when player 2 is the last to move. We will now show that in both cases this sequence is convergent. It is clearly a sequence drawn from a compact set and hence has a limit on a subsequence, say τ_n, and call the subsequential limit x_∞. On that subsequence, using the equation above, we must have

$$\lim_{t_n \longrightarrow \infty} \underline{x}(\tau_n + 2) = x_\infty$$

given the continuity of f. Since for every convergent subsequent it follows that the subsequence formed by the elements $\tau_n + 2$ converge to the same limit, it must be the case that the original sequence itself converges to x_∞. Evidently this argument shows that there is a limit regardless of whether we take the odd sequence or the even one, though the two limits need not be the same. Call them, as above, x^* and x' respectively.

Note also that

$$f(\overline{x}(\tau+1)) = \frac{\tau}{\tau+1} f(\underline{x}(\tau))$$

and from that it follows that the sequence $\overline{x}(\tau)$ must itself have a limit and that limit must be the same as that for $\underline{x}(\tau+1)$. ■

4.2 Local Convergence

Theorem 7 *Suppose that f is continuously differentiable and furthermore that $-\infty < f' < 0$. Then, the two limits x^* and x' must both be local maxima of the Nash product function $x.f(x)$.*

Proof. By eliminating $\overline{x}(\cdot)$ in the equation above we have

$$f\left(\frac{\tau+2}{\tau+1}\underline{x}(\tau+2)\right)=\frac{\tau}{\tau+1}f(\underline{x}(\tau)). \tag{2}$$

■

Given the assumptions on f,[11] and considering one of the limits, say x^*, what we know is that in a small enough neighborhood of x^*, if $x<x^*$, then $xf'(x)+f(x)>0$; if $x>x^*$, then $xf'(x)+f(x)<0$.

Notice that $\underline{x}(\tau+2)-\underline{x}(\tau)=O(1/\tau)$. Thus, by linearizing (2), we have

$$f(\underline{x}(\tau))+(\underline{x}(\tau+2)-\underline{x}(\tau))f'(\underline{x}(\tau))+\frac{\underline{x}(\tau+2)f'(\underline{x}(\tau))}{\tau}=f(\underline{x}(\tau))-\frac{f(\underline{x}(\tau))}{\tau}+O\left(\frac{1}{\tau^2}\right),$$

i.e.,

$$\underline{x}(\tau+2)-\underline{x}(\tau)=-\frac{\underline{x}(\tau)f'(\underline{x}(\tau))+f(\underline{x}(\tau))}{\tau f'(\underline{x}(\tau))}+O\left(\frac{1}{\tau^2}\right). \tag{3}$$

Fix any $\varepsilon>0$. Then for sufficiently large τ, $|\underline{x}(\tau+2)-\underline{x}(\tau)|<\varepsilon$. Given that f' is continuously differentiable and bounded above it follows that $\underline{x}(\tau)f'(\underline{x}(\tau))+f(\underline{x}(\tau))$ is arbitrarily small. Running $\varepsilon\downarrow 0$ we can conclude that x^* must be a local maximum of the Nash product function $x.f(x)$. Clearly this argument holds for any limit point of $\underline{x}(\tau)$ and, in particular, for x'.■

4.3 Global Convergence

Theorem 8 *Suppose that f is continuously differentiable, that $-\infty<f'<0$ and furthermore that $x.f(x)$ is single-peaked. Then, the two limits x^* and x' coincide at the global maxima of the Nash product function $x.f(x)$, i.e., equilibrium play uniquely selects the Nash bargaining solution.*

Proof. The proof follows as an immediate corollary of the result in the previous sub-section. When f is single-peaked it has, by assumption, only one (local/global) maxima. Hence, the two limits x^* and x' coincide at the global maximum which is, again from the definition, the Nash bargaining solution.■

Remark 9 *A sufficient condition for single-peakedness of the Nash product function is when f is concave. Hence, we have found the general result for the linear Pareto frontier case whose equilibrium was computed in the previous section.*

■

[11] The following may not be the weakest possible assumptions. Say, continuous differentiability might be relaxed to mere differentiability, etc.

5 Reversible Agreements: Two Examples of Global Convergence

In this section we switch to Model 2 - the Reversible Agreements model - in which an agreement once reached is only good for that period. In the subsequent period, the player who then has the mover can over-turn the agreement by proposing a different split of the pie. He would, of course, suffer the short-term payoff consequence of zero disagreement payoffs but might wish to trade that in for a more attractive payoff in the future. On the face of it, it would appear that equiibrium behavior could be ceaselessly complex. Agreements might appear, last for a bit, and then be replaced by other agreements. Or there might be long periods of time over which there is simply a tug and fro of bargains between the two players with no agreements reached.

We start by providing two examples where the behavior is remarkably ordered. The first is a return to our initial example where the Pareto frontier is linear. Within this example we will exhibit that yet again there is global convergence; regardless of whether it is player 1 or 2 who is the last mover, there is a unique agreement in the limit. In fact, in this example, the agreement sets will turn out to be identical to those in the Irreversible case! Consider the case where player 1 is the last to move. Recall that there are worst payoff agreements for her, $\underline{x}(\tau)$, for $\tau = 1, 3, 5...$; with odd number of periods τ left, player 1 is willing to accept any agreement that gives her a payoff of at least $\underline{x}(\tau)$. Similarly there are best payoff agreements for player 1, $\overline{x}(\tau)$, for $\tau = 2, 4, 6...$; with even number of periods τ left, player 2 is willing to accept any agreement that gives him a payoff of at least $f(\overline{x}(\tau))$, or $\overline{x}(\tau)$, for $\tau = 2, 4, 6...$ is the most that player 1 can hope to achieve. It turns out in the examples that follow that there is **two-sided blocking** - when there are an odd number of periods left to play, player 1 will accept any agreement that gives her a payoff of at least $\underline{x}(\tau)$ but no more than $\overline{x}(\tau - 1)$. The second part is the new phenomenon; even if player 1 is confronted with a particularly attractive proposal - one that gives her more than $\overline{x}(\tau - 1)$ - she will say, no thank you! As will become clear, this anticipatory block is instigated from knowing that such an agreement is very short-term, will be over-turned next period, and when it is over-turned will be replaced by a considerably worse agreement.

The second example will show that the above result is not an artfact of risk-neutrality. In it we will consider a (quarter) circle - and hence concave - Pareto frontier. In this example, there will again be global convergence; again, a player will block not only agreements that are dis-advantageous for her but also agreements that are too good.[12]

In both examples, the set of agreements will converge uniquely and globally

[12] Indeed, this example, like the linear example, generates identical sequences $\underline{x}(\tau)$ and $\overline{x}(\tau)$ in both the Irreversible and Reversible models. The only difference is that in the Irreversible model, player 1 blocks anything below $\underline{x}(\tau)$ while in the Reversible model she also blocks anything above $\overline{x}(\tau - 1)$. Similarly, player 2 only blocks anything below $f(\overline{x}(\tau))$ in the Irreversible model but additionally blocks every agreement that gives him above $f(\underline{x}(\tau - 1))$ in the Reversible case.

to the Nash bargaining solution, i.e., the unique maximizer of the Nash product $x.f(x)$. These example will also illustrate a more general result which is to be found in the next section. The general result will show that Reversible Agreement bargaining leads always to the Nash bargaining solution if f is concave.

5.1 Linear Pareto Frontier

Suppose that

$$f(x) = -mx + c,\ m > 0, c > 0$$

Note again that $argmax_x\ xf(x) = \frac{c}{2m}$. Let us start with the case where player 1 is the last to move.

As before, $\underline{x}(\tau)$ for $\tau = 1, 3, 5$ ... will denote the (worst) agreement payoff that player 1 is willing to accept at those periods (and let the associated payoff for player 2 be denoted $\overline{y}(\tau)$, i.e., $\overline{y}(\tau) = f(\underline{x}(\tau))$). Let $\underline{y}(\tau)$ for $\tau = 2, 4, 6$, ... denote the worst agreement that player 2 will similarly accept when it is his turn to move. Let $\overline{x}(\tau) = \frac{\underline{y}(\tau)-c}{-m}$ denote the associated payoff for player 1. These will be computed in exactly the same fashion as in the Irreversible Agreement case and will have the exact same values. We will then show that player 1 - when it is her turn to move - will accept any agreement *between* $\underline{x}(\tau)$ and $\overline{x}(\tau-1)$ and we will, without loss write $\overline{x}(\tau) = \overline{x}(\tau-1)$. We will also show that player 2 - when it is his turn to move - will accept any agreement between $\underline{y}(\tau)$ and $\overline{y}(\tau-1)$ and we will, without loss write $\overline{y}(\tau) = \overline{y}(\tau-1)$. In other words, when it is player 2's move, player 1's payoffs will again be nested within $\underline{x}(\tau)$ and $\overline{x}(\tau)$.

We already know that

$$\lim_{\tau\to\infty} \underline{x}(\tau) = \lim_{\tau\to\infty} \overline{x}(\tau) = \frac{c}{2m}$$

and hence, even in this case, the limiting SPE will turn out to be the Nash bargaining solution.

$\tau = 1$ With one period to go, player 1 will accept anything, i.e. the lowest payoff she will take $\underline{x}(1) = 0$. At the same time there is no payoff that she will deem too good and refuse, i.e., $\overline{x}(1) = \frac{c}{m}$. Hence, $\underline{x}(1) = 0, \overline{x}(1) = \frac{c}{m}$ with associated payoffs of player 2 $\overline{y}(1) = c$ and $\underline{y}(1) = 0$ respectively.

$\tau = 2$ With two periods to go, player 2 can guarantee himself $\overline{y}(1) = c$ by turning down a proposal in this period and waiting for the next with a counter-proposal that gives him everything. Hence, the worst payoff that he will accept is given by the equality $c = 2\underline{y}(2)$ implying that $\underline{y}(2) = \frac{c}{2}$ implying that $\overline{x}(2) = \frac{c}{2m}$. Additionally, since player 1 does not block anything in the subsequent period, he has no reason to block any agreements that give him too much, i.e., $\overline{y}(1) = c$ and hence $\underline{x}(1) = 0$. Collecting all this we have, $\underline{x}(2) = 0, \overline{x}(2) = \frac{c}{2m}$.

$\tau = 3$ With three periods to go, player 1 can guarantee herself $\overline{x}(2) = \frac{c}{2m}$ by turning down a proposal on the table and waiting for her best payoff for two periods. Hence, the worst payoff that she will accept is given by the equality $\frac{c}{m} = 3\underline{x}(3)$ implying that $\underline{x}(3) = \frac{1}{3}\frac{c}{m}$ i.e., that $\overline{y}(3) = \frac{2}{3}c$. Additionally, if player 1 is offered any agreement in which her share is greater than $\frac{1}{2}\frac{c}{m}$ then she blocks it. This is because accepting such an agreement would lead to a block by player 2 at $\tau = 2$ thereby leading to a payoff only in the current period. That payoff is less than $2.\frac{c}{2m}$ which she can ensure by counter-offering with $\frac{c}{2m}$ which will then get played for the last two periods. Hence, $\underline{x}(3) = \frac{1}{3}\frac{c}{m}$, $\overline{x}(3) = \frac{1}{2}\frac{c}{m}$.

$\tau = 4$ Continuing in that fashion, $\underline{y}(4)$ can be derived from $2c = 4\underline{y}(4)$ implying that $\underline{y}(4) = \frac{c}{2}$ implying that $\overline{x}(4) = \frac{c}{2m}$. Additionally, if player 2 is offered any agreement y in which his share is greater than $\frac{2}{3}c$ (and player 1's is less than $\frac{1}{3}\frac{c}{m}$) then he blocks it. This is because accepting such an agreement would lead to a block by player 1 at $\tau = 3$ thereby leading to a total payoff of $y + 2\frac{c}{2}$. That payoff is less than $3.\frac{2}{3}c$ which he can ensure by counter-offering with $\frac{c}{2m}$ which will then get played for the last three periods. Hence, $\underline{x}(4) = \frac{1}{3}\frac{c}{m}$, $\overline{x}(4) = \frac{1}{2}\frac{c}{m}$.

$\tau = 5$ Again, the worst payoff the mover, player 1, will take is given by $\frac{2c}{m} = 5\underline{x}(5)$ implying $\underline{x}(5) = \frac{2}{5}\frac{c}{m}$ implying $\overline{y}(5) = \frac{3}{5}c$. By similar logic as above, $\overline{x}(5) = \frac{1}{2}\frac{c}{m}$. Hence, $\underline{x}(5) = \frac{2}{5}\frac{c}{m}, \overline{x}(5) = \frac{1}{2}\frac{c}{m}$

$\tau = 6$ $3c = 6\underline{y}(6)$ implying that $\underline{y}(6) = \frac{c}{2}$ implying that $\overline{x}(6) = \frac{c}{2m}$. Again there is no point for player 2 in accepting any agreement y in which he gets more than $\frac{3}{5}c$ - and player 1 gets less than $\frac{2}{5}\frac{c}{m}$ - in the current period. Such an agreement would get blocked in the next period and the total payoff to player 2 - $y + 4.\frac{1}{2}c$ - would be less than if he blocked by counter-offering with $\frac{3}{5}c$ which would be accepted next period and lead to a total payoff of $5.\frac{3}{5}c$. Hence, $\underline{x}(6) = \frac{2}{5}\frac{c}{m}, \overline{x}(6) = \frac{c}{2m}$.

$\tau = 7$ $\frac{3c}{m} = 7\underline{x}(7)$ implying $\underline{x}(7) = \frac{3}{7}\frac{c}{m}$ implying $\overline{y}(7) = \frac{4}{7}c$. By similar logic as above, $\overline{x}(7) = \frac{1}{2}\frac{c}{m}$. Hence, $\underline{x}(7) = \frac{3}{7}\frac{c}{m}, \overline{x}(7) = \frac{1}{2}\frac{c}{m}$

$\tau = 8$ $4c = 8\underline{y}(8)$ implying that $\underline{y}(8) = \frac{c}{2}$ implying that $\overline{x}(8) = \frac{c}{2m}$. If player 2 accepts any agreement in which player 1 currently gets less than $\frac{3}{7}\frac{c}{m}$ then she promptly blocks it in the next period and the total payoff for player 2 in that sequence of play is $y + 6.\frac{c}{2}$. Instead by pre-emptively blocking and offering player 1 $\frac{3}{7}\frac{c}{m}$ player 2 can guarantee than he will have positive payoffs from the next period onwards. The total payoff in that sequence is $7.\frac{4}{7}c = 4c$ which is greater. Hence, $\underline{x}(8) = \frac{3}{7}\frac{c}{m}, \overline{x}(8) = \frac{c}{2m}$.

It is clear that the worst agreement shares of player 1 $\underline{x}(\tau)$ - after factoring out the $\frac{c}{m}$ - are

$$0, 0, \frac{1}{3}, \frac{1}{3}, \frac{2}{5}, \frac{2}{5}, \frac{3}{7}, \frac{3}{7}, \frac{4}{9}, \frac{4}{9} \; ... \frac{n-1}{2n}, \frac{n-1}{2n}, ..$$

for $n = 1, 3, 5, ..$ Of course that sequence converges to $\frac{1}{2}$. It is also clear that the best agreement payoffs $\overline{x}(\tau) = \frac{c}{2m}$ for $\tau > 1$.

Let us now turn to global convergence - that the same limit would hold had we considered the case in which player 2 is the last mover. Easy computation, again using equations above, tell us that the odd period payoffs of player 1, now $\overline{x}(\tau)$, are - after factoring out $\frac{c}{m}$ - given by

$$1, \frac{2}{3}, \frac{2}{3}, \frac{3}{5}, \frac{3}{5}, \frac{4}{7}, \frac{4}{7}, \frac{5}{9}, \frac{5}{9}, \; ... \frac{n+2}{2(n+1)}, \frac{n+2}{2(n+1)} ...$$

for $n = 2, 4, ..$ Of course that sequence converges to $\frac{1}{2}$. It is also clear that the worst payoffs of player 1, $\underline{x}(\tau) = \frac{c}{2m}$ for all even numbered τ. Hence, again, the unique limit payoff for player 1 is $\frac{c}{2m}$, the Nash bargaining solution and we have global convergence.■

5.2 Circular Pareto Frontier

Suppose instead that

$$mx^2 + y^2 = c,\ m > 0, c > 0$$

and hence that

$$y = \sqrt{c - mx^2},\ x = \sqrt{\frac{1}{m}\left(c - x^2\right)}$$

Note that $argmax_x\ xf(x) = \frac{1}{\sqrt{2}}\frac{c}{m}$.

Again we will start with the case where player 1 is the last to move. Computing $\underline{x}(\tau)$ for $\tau = 1, 3, 5$ ... and $\overline{x}(\tau) =$ for $\tau = 2, 4, 6$ in exactly the same fashion as above, we will show that player 1 - when it is her turn to move - will accept any agreement *between* $\underline{x}(\tau)$ and $\overline{x}(\tau)$. Similarly, when it is player 2's move, player 1's payoffs will again be nested within $\underline{x}(\tau)$ and $\overline{x}(\tau)$.

We will now compute $\underline{x}(\tau)$ and $\overline{x}(\tau)$ and show that

$$\lim_{\tau\to\infty} \underline{x}(\tau) = \lim_{\tau\to\infty} \overline{x}(\tau) = \frac{1}{\sqrt{2}}\sqrt{\frac{c}{m}}$$

Hence, in this case as well, the state-independent SPE is the Nash solution.

$\tau = 1$ $\quad \underline{x}(1) = 0$ implying $\overline{y}(1) = \sqrt{c}$. Also, $\overline{x}(1) = \sqrt{\frac{c}{m}}$. Hence, $\underline{x}(1) = 0, \overline{x}(1) = \sqrt{\frac{c}{m}}$.

$\tau = 2$ $\quad \frac{1}{2}\sqrt{c} = \underline{y}(2)$ implying that $\overline{x}(2) = \sqrt{\frac{3}{4}}\sqrt{\frac{c}{m}}$. Also, $\underline{x}(2) = 0$. Hence, $\underline{x}(2) = 0, \overline{x}(2) = \sqrt{\frac{3}{4}}\sqrt{\frac{c}{m}}$.

$\tau = 3$ $\quad \frac{2}{3}\sqrt{\frac{3}{4}}\sqrt{\frac{c}{m}} = \underline{x}(3)$ implying $\underline{x}(3) = \sqrt{\frac{1}{3}}\sqrt{\frac{c}{m}}$ implying $\overline{y}(3) = \sqrt{\frac{2}{3}}\sqrt{c}$. For any $x > \overline{x}(2)$, if player 1 accepts such an agreement, it will be blocked by player 2 in the next period and hence the total payoff to player 1 from such an acceptance will be x. Blocking the agreement and proposing $\overline{x}(2)$ will lead to a total payoff of $2.\overline{x}(2)$ which is greater than x. Hence, $\overline{x}(3) = \overline{x}(2)$. There is clearly no point in blocking an agreement between $\underline{x}(2)$ and $\overline{x}(2)$ since such agreements are not blocked in the future and - from the definition of $\underline{x}(2)$ - the total over the remaining three periods is greater than the highest payoff from blocking, i.e., $2\overline{x}(2)$. Hence, $\underline{x}(3) = \sqrt{\frac{1}{3}}\sqrt{\frac{c}{m}}, \overline{x}(3) = \overline{x}(2) = \sqrt{\frac{3}{4}}\sqrt{\frac{c}{m}}$.

$\tau = 4$ $\quad \frac{3}{4}\sqrt{\frac{2}{3}}\sqrt{c} = \underline{y}(4)$ implying that $\underline{y}(4) = \sqrt{\frac{3}{4}}\sqrt{c}$ implying that $\overline{x}(4) = \sqrt{\frac{5}{8}}\sqrt{\frac{c}{m}}$. For any good agreement for player 2, say y, in which 1's payoff share is less than $\underline{x}(3)$, by accepting it player 2 knows that it will be blocked in the subsequent period (with a counter-offer of $\overline{x}(2)$). Hence the total payoff to accepting that good agreement is $y + 2.\frac{1}{2}\sqrt{c}$. In contrast, by blocking that agreement with a counter-offer of $\underline{x}(3)$ - which will get accepted in the next period - player 2 can guarantee $3.\sqrt{\frac{2}{3}}\sqrt{c}$ and that is clearly higher. There is clearly no point in blocking an agreement between $\underline{x}(3)$ and $\overline{x}(4)$ since such agreements are not blocked in the future and - from the definition of $\overline{x}(4)$ - the total over the remaining four periods is greater than the highest payoff from blocking. Hence, $\underline{x}(4) = \underline{x}(3) = \sqrt{\frac{1}{3}}\sqrt{\frac{c}{m}}$, $\overline{x}(4) = \sqrt{\frac{5}{8}}\sqrt{\frac{c}{m}}$.

$\tau = 5$ $\quad \frac{4}{5}\sqrt{\frac{5}{8}}\sqrt{\frac{c}{m}} = \underline{x}(5)$ implying $\underline{x}(5) = \sqrt{\frac{2}{5}}\sqrt{\frac{c}{m}}$ and $\overline{y}(5) = \sqrt{\frac{3}{5}}\sqrt{c}$. For any $x > \overline{x}(4)$, if player 1 accepts such an agreement it will be blocked by player 2 in the next period and hence the total payoff to player 1 from such an acceptance will be $x + 3\underline{x}(3) = x + 3\sqrt{\frac{1}{3}}\sqrt{\frac{c}{m}}$. Blocking the agreement and proposing $\overline{x}(4)$ instead will lead to a total payoff of $4.\overline{x}(4) = 4\sqrt{\frac{5}{8}}\sqrt{\frac{c}{m}}$ which is greater. Hence, $\overline{x}(5) = \overline{x}(4)$. There is clearly no point in blocking an agreement between $\underline{x}(4)$ and $\overline{x}(4)$ since such agreements are not blocked in the future and - from the definition of $\underline{x}(4)$ - the total over the remaining three periods is greater than the highest payoff from blocking, i.e., $4\overline{x}(4)$. Hence, $\underline{x}(5) = \sqrt{\frac{2}{5}}\sqrt{\frac{c}{m}}$, $\overline{x}(5) = \overline{x}(4) = \sqrt{\frac{5}{8}}\sqrt{\frac{c}{m}}$.

And so on. Below are the computations - pasted in from the irreversible case - of the lower threshold in each period. The upper threshold can be computed by continuing the arguments for $t = 1, ...4$.

$\tau = 6$ $\quad \frac{5}{6}\sqrt{\frac{3}{5}}\sqrt{c} = \underline{y}(6)$ implying that $\underline{y}(6) = \sqrt{\frac{5}{12}}\ \sqrt{c}$ implying that $\overline{x}(6) = \sqrt{\frac{7}{12}}\sqrt{\frac{c}{m}}$

$\tau = 7$ $\quad \frac{6}{7}\sqrt{\frac{7}{12}}\sqrt{\frac{c}{m}} = \underline{x}(7)$ implying $\underline{x}(7) = \sqrt{\frac{3}{7}}\sqrt{\frac{c}{m}}$ implying $\overline{y}(7) = \sqrt{\frac{4}{7}}\sqrt{c}$

$\tau = 8$ $\quad \frac{7}{8}\sqrt{\frac{4}{7}}\sqrt{c} = \underline{y}(8)$ implying that $\underline{y}(8) = \sqrt{\frac{5}{12}}\ \sqrt{c}$ implying that $\overline{x}(8) = \sqrt{\frac{9}{16}}\sqrt{\frac{c}{m}}$

$\tau = 9$ $\quad \frac{8}{9}\sqrt{\frac{9}{16}}\sqrt{\frac{c}{m}} = \underline{x}(9)$ implying $\underline{x}(9) = \sqrt{\frac{4}{9}}\sqrt{\frac{c}{m}}$ implying $\overline{y}(9) = \sqrt{\frac{5}{9}}\sqrt{c}$

$\tau = 10$ $\quad \frac{9}{10}\sqrt{\frac{5}{9}}\sqrt{c} = \underline{y}(10)$ implying that $\underline{y}(10) = \sqrt{\frac{9}{20}}\ \sqrt{c}$ implying that $\overline{x}(10) = \sqrt{\frac{11}{20}}\sqrt{\frac{c}{m}}$

It is clear that the worst agreement shares of player 1, $\underline{x}(\tau)$ - after factoring out $\sqrt{\frac{c}{m}}$ - are

$$0, 0, \sqrt{\frac{1}{3}}, \sqrt{\frac{1}{3}}\sqrt{\frac{2}{5}}, \sqrt{\frac{2}{5}}\sqrt{\frac{3}{7}}, \sqrt{\frac{3}{7}}\sqrt{\frac{4}{9}}, \sqrt{\frac{4}{9}} \ ...\sqrt{\frac{n-1}{2n}}, .\sqrt{\frac{n-1}{2n}}..$$

for $n = 1, 3, ..$ Of course that sequence converges to $\sqrt{\frac{1}{2}}$. It is also clear that the best agreements $\overline{x}(\tau)$, after again factoring out $\sqrt{\frac{c}{m}}$, are

$$1, \sqrt{\frac{3}{4}}, \sqrt{\frac{3}{4}}\sqrt{\frac{5}{8}}, \sqrt{\frac{5}{8}}, \sqrt{\frac{7}{12}}, \sqrt{\frac{7}{12}}, \sqrt{\frac{9}{16}}, \sqrt{\frac{9}{16}}, \ldots \sqrt{\frac{n+1}{2n}}, \sqrt{\frac{n+1}{2n}}..$$

for $n = 2, 4, ..$ Of course that sequence also converges to $\sqrt{\frac{1}{2}}$.

We can then turn to the case in which player 2 is the last to move. It is tedious but not difficult to show that there too the limit is the Nash bargaing solution.■

6 Reversible Agreements: A General Result on Global Convergence

In this section we prove a general result that shows why the examples in the previous section worked. In particular, we show that whenever f is concave - in addition to being C^1 and downward sloping - then we have global convergence even in the Reversible Agreements model:

Theorem 10 *Suppose that f is continuously differentiable, that $-\infty < f' < 0$ and furthermore that it is concave with the boundary conditions, $\alpha = 0 \Longrightarrow x = 0, f(0) > 0$, $\alpha = 1 \Longrightarrow x = \overline{x} > 0, f(\overline{x}) = 0$. Then, best and worst payoff sequences are monotonic - the best payoff $\overline{x}(\tau)$ decreases in τ while the worst payoff $\underline{x}(\tau)$ increases. Furthermore, the two sequences converge to the Nash bargaining solution as $\tau \longrightarrow \infty$. These statements are true regardless of whether player 1 or 2 is the last to move.*

Proof. Start with the case where player 1 is the last to move. We will establish that - as in the Examples in the previous section - for each τ, the mover will only accept deals in the "middle", in which player 1's payoffs are at least $\underline{x}(\tau)$ and no more than $\overline{x}(\tau)$.

Consider $\tau = 1$. Clearly, all agreements are accepted by this player and hence, $\underline{x}(1) = 0$, $\overline{x}(1) = \overline{x}$.

Consider instead $\tau = 2$. Evidently, player 2 determines his lowest acceptable payoff via

$$0 < \frac{1}{2}f(0) = f(\overline{x}(2))$$

implying $\overline{x}(2) < \overline{x}(1) = \overline{x}$. On the other hand, player 2 has no reason to block any agreements where he receives a "high" payoff (because player 1 will accept all agreements in the subsequent period). So, $\underline{x}(2) = 0$.

To summarize, for the first two horizon lengths, we do have $\underline{x}(2) \geqslant \underline{x}(1)$, $\overline{x}(2) \leqslant \overline{x}(1)$. Make the induction hypothesis that the inequalities hold for all $\tau \leqslant \widehat{\tau}$, i.e., $\underline{x}(\tau) < \overline{x}(\tau)$ and for all $\tau \leqslant \widehat{\tau} - 1$ $\underline{x}(\tau + 1) \geqslant \underline{x}(\tau)$, $\overline{x}(\tau + 1) \leqslant \overline{x}(\tau)$.

Now consider $\tau = \widehat{\tau} + 1$. Suppose the mover in that period is player 2.

Case 1 - Consider any proposal in which player 1's current payoff is $x > \overline{x}(\widehat{\tau})$ - or $f(x) < f(\overline{x}(\widehat{\tau}))$. If player 2 accepts it, by definition of $\overline{x}(\widehat{\tau})$, it then gets blocked in the next period by player 1 with a counter-offer of $\overline{x}(\widehat{\tau}-1)$ which will, again by hypothesis, get accepted by player 2 two periods from now. The life-time payoff to acceptance then is

$$f(x)+(\widehat{\tau}-1).f(\overline{x}(\widehat{\tau}-1))$$

An alternative would be for player 2 to reject the agreement and instead counter-offer with an agreement that would yield player 1 $\underline{x}(\widehat{\tau})$ and that would, by hypothesis, be accepted in the next period. The life-time payoff to player 2 for this counter-proposal would be

$$\widehat{\tau}.f(\underline{x}(\widehat{\tau}))$$

Since $\overline{x}(\widehat{\tau}-1)>\underline{x}(\widehat{\tau}-1)\geqslant\underline{x}(\widehat{\tau})$, and hence, $f(\overline{x}(\widehat{\tau}-1))<f(\underline{x}(\widehat{\tau}-1)\leqslant f(\underline{x}(\widehat{\tau}))$ it follows that

$$\widehat{\tau}.f(\underline{x}(\widehat{\tau}))>f(x)+(\widehat{\tau}-1).f(\overline{x}(\widehat{\tau}-1)).$$

Put differently, player 2 rejects all proposals in which $x>\overline{x}(\widehat{\tau})$ - or $f(x)<f(\overline{x}(\widehat{\tau}))$.

Case 2 - Consider any proposal in which player 1's current payoff is $x<\overline{x}(\widehat{\tau})$ - or $f(x)>f(\overline{x}(\widehat{\tau}))$. By the same logic as above, it is clear that the lifetime payoff to accepting the proposal today is

$$f(x)+(\widehat{\tau}-1)f(\overline{x}(\widehat{\tau}-1))$$

■

while the lifetime payoff to counter-offering with a proposal that gives player 1 $\underline{x}(\widehat{\tau}))$ in the next period is

$$\widehat{\tau}.f(\underline{x}(\widehat{\tau}))$$

Note that by definition of $\underline{x}(\widehat{\tau})$

$$\widehat{\tau}.\underline{x}(\widehat{\tau})=(\widehat{\tau}-1).\overline{x}(\widehat{\tau}-1)$$

which can equivalently be written as

$$\underline{x}(\widehat{\tau})=\frac{\widehat{\tau}-1}{\widehat{\tau}}\overline{x}(\widehat{\tau}-1)+\frac{1}{\widehat{\tau}}0$$

Since f is concave we know that

$$\begin{aligned} f(\underline{x}(\widehat{\tau})) &\geqslant \frac{\widehat{\tau}-1}{\widehat{\tau}}f(\overline{x}(\widehat{\tau}-1))+\frac{1}{\widehat{\tau}}f(0) \\ &> \frac{\widehat{\tau}-1}{\widehat{\tau}}f(\overline{x}(\widehat{\tau}-1))+\frac{1}{\widehat{\tau}}f(x) \end{aligned}$$

(the inequality being strict as long as $f(x)<f(0)$). So in this Case too, player 2 is better off from rejecting the proposal.

Put differently, the two cases have shown that player 2 will block all proposals in which player 1 gets $x \leqslant \underline{x}(\widehat{\tau})$ or $x \geqslant \overline{x}(\widehat{\tau})$.

Case 3 - Consider any proposal in which player 1's current payoff is $x \in [\underline{x}(\widehat{\tau}), \overline{x}(\widehat{\tau})]$. By hypothesis, any such agreement lasts till the end of the game and hence yields life-time payoffs of $(\widehat{\tau}+1)f(x)$ to Player 2. The worst such agreements, from Player 2's perspective, are $x = \overline{x}(\widehat{\tau}) - \varepsilon$, for small ε. These need to be compared against the life-time payoff to rejecting the proposal whose best payoff - as seen above - is $\widehat{\tau}.f(\underline{x}(\widehat{\tau}))$. In the event that,

$$(\widehat{\tau}+1)f(\overline{x}(\widehat{\tau})) \geqslant \widehat{\tau}.f(\underline{x}(\widehat{\tau}))$$

accepting $x \in [\underline{x}(\widehat{\tau}), \overline{x}(\widehat{\tau})]$ is better than rejecting that agreement. Hence, the optimal response is to accept in that region and reject everything outside that region, i.e., $\underline{x}(\widehat{\tau}+1) = \underline{x}(\widehat{\tau}) < \overline{x}(\widehat{\tau}) = \overline{x}(\widehat{\tau}+1)$ On the other hand, In the event that,

$$(\widehat{\tau}+1)f(\overline{x}(\widehat{\tau})) < \widehat{\tau}.f(\underline{x}(\widehat{\tau}))$$

define $\overline{x}(\widehat{\tau}+1)$ by the equality

$$(\widehat{\tau}+1)f(\overline{x}(\widehat{\tau}+1)) = \widehat{\tau}.f(\underline{x}(\widehat{\tau}))$$

It follows that $\overline{x}(\widehat{\tau}+1) > \underline{x}(\widehat{\tau}) = \overline{x}(\widehat{\tau}+1)$. The theorem is proved.■

7 Some Extensions and Generalizations

7.1 Choosing When to Propose

7.2 Three Players