REPEATED GAMES WITH DIFFERENTIAL TIME PREFERENCES: APPENDIX FOR REFEREES

(Proofs of Proposition 2 and Theorem 2)

by

Ehud Lehrer and Ady Pauzner

October, 1997

This appendix contains proofs that were eliminated from the body of the paper for brevity's sake. These proofs will be provided to interested readers upon request.

We start by providing a few formal definitions that are needed for the proofs. We denote by θ^k the stage-*k* public signal, which is a continuous random variable, and denote by $\tilde{\theta}^k$ its realization. A history h_k of length *k* is a list of the realized actions and the realizations of the public signals at previous stages, $\left(\tilde{\theta}^0; a_I^0, a_P^0\right)$, $\left(\tilde{\theta}^1; a_I^1, a_P^1\right)$,..., $\left(\tilde{\theta}^{k-1}; a_I^{k-1}, a_P^{k-1}\right)$. A strategy prescribes to the player which mixed action to take at any stage *k*, for any possible history h_k of length *k* and any realization of θ^k.

To show that the set of equilibrium payoffs (Nash or SP) is closed, we first show that one can replace the continuous coordination devices by finitely valued ones.

Claim 1:

Given any SPE σ, one can construct a SPE σ' in which (1) the two players receive the same expected payoff as in σ and (2) each public signal θ^0, θ^1,... is a random variable that takes one of only three values (with probabilities that may depend on history).

Proof:

We define σ' inductively, as follows. Let (x, y) be the expected continuation payoffs from σ. The pair (x, y) is in the convex hull of all possible continuation payoffs corresponding to possible realizations of θ^0. By the Caratheodory theorem *(reference),* (x, y) is in the convex hull of at most 3 specific continuation payoffs. That is, there are three possible realizations of θ^0: $\tilde{\theta}_1^0$, $\tilde{\theta}_2^0$ and $\tilde{\theta}_3^0$, such that (x, y) is a convex combination of the continuation payoffs of σ after $\tilde{\theta}_1^0$, $\tilde{\theta}_2^0$ and $\tilde{\theta}_3^0$. At the first stage, we replace θ^0 with a correlation device that generates the signals $\tilde{\theta}_1^0$, $\tilde{\theta}_2^0$ and $\tilde{\theta}_3^0$ with the appropriate probabilities. At the first stage we define $\sigma'(\tilde{\theta}_i^0) = \sigma(\tilde{\theta}_i^0)$.

Assume that we have already defined σ' up to stage k, such that after any history shorter than k, the correlation device generates only three signals. Hence, there are only finitely many histories of length k. After every such history h_k, we replace (using again the Caratheodory theorem) the correlating device θ^k with a (history dependent) correlating device θ^{h_k} that generates only three signals, $\tilde{\theta}_1^{h_k}$, $\tilde{\theta}_2^{h_k}$ and $\tilde{\theta}_3^{h_k}$, (with probabilities that may depend on h_k). We define $\sigma'(h_k, \tilde{\theta}_i^{h_k}) = \sigma(h_k, \tilde{\theta}_i^k)$. (Notice that the set of histories over which σ is defined is a superset of the set of histories over which σ' is defined).

We continue this process inductively and obtain the pair of strategies σ'. Note that the continuation payoffs of σ and σ' after any history over which σ' is defined are the same. In particular, the payoffs in the whole repeated game from σ and σ' are the same. Moreover, since there is no profitable one-stage deviation from σ, there is no profitable one-stage deviation from σ'. This implies that σ' is a SPE. $\blacksquare_{\text{(Claim 1)}}$

Claim 2:

Given any Nash equilibrium σ, one can construct a Nash equilibrium σ' in which (1) the two players receive the same expected payoff as in σ and (2) each public signal θ^0, θ^1,... is a random variable that takes one of only three values (with probabilities that may depend on history).

Proof:

We define σ' inductively as in Claim 1, but only for histories h_k that are assigned a positive probability after the k-th step of the inductive process. Over other histories the players minmax each other. More precisely, after a history h_k of length k that is reached with a positive probability under σ', we replace the correlating device θ^k with a (history dependent) correlating device θ^{h_k} that generates only three signals. Note that since the strategies induced by σ after (h_k, θ^k) is almost surely a Nash equilibrium, the signals $\tilde{\theta}_1^{h_k}$, $\tilde{\theta}_2^{h_k}$ and $\tilde{\theta}_3^{h_k}$ can be chosen such that the strategies induced by σ after $(h_k, \tilde{\theta}_i^{h_k})$ are Nash equilibrium. After a history h_k whose probability is zero, σ' instructs the players to minmax their opponents. As in Claim 1, the continuation payoffs after any positive probability history is the same under σ and σ'. In particular, the payoffs in the whole repeated game from σ and σ' are the same.

To see that σ' is a Nash equilibrium, assume that player i deviated from the strategy prescribed in σ' to the pure strategy s_i. This deviation cannot be profitable. To see why, let k be the first stage at which this deviation leads the players off the equilibrium path. From that stage on the players receive, at most, their IR payoffs. If the deviation was profitable, player *i* must have had a profitable deviation from σ: playing according to s_i until stage *k*, and securing her IR payoff thereafter. If, on the other hand, there is no stage at which s_i leads us off the equilibrium path, s_i cannot improve payoffs. To see why, note that if player *i* could gain by deviating to s_i, she would already gain by deviating to s_i for the first *K* stages and then returning to her equilibrium strategy. But every history of length *K* that is assigned positive probability under s_i is also assigned positive probability under *i*'s equilibrium strategy. Hence, no strategy can be more profitable than the equilibrium strategy. $\blacksquare_{\text{(Claim 2)}}$

Claim 3:

For any Nash (SP) equilibrium σ that employs finitely valued and possibly history-dependent randomizing devices, there exists a Nash (SP) equilibrium τ that yields the same payoffs and employs continuous and history-independent randomization devices.

Proof:

Let the randomization devices used by τ, θ^k (k=0,1,...), be all uniformly distributed over [0,1] and i.i.d. We define τ inductively as follows. At the first stage, we divide the unit interval into three intervals $I(\tilde{\theta}_1^0)$, $I(\tilde{\theta}_2^0)$ and $I(\tilde{\theta}_3^0)$, the length of each is equal to the probability assigned to $\tilde{\theta}_1^0$, $\tilde{\theta}_2^0$ and $\tilde{\theta}_3^0$. We let $\tau(\theta^0) = \sigma'(J(\theta^0))$, where $J(\theta^0) = \tilde{\theta}_i^0$ iff $\theta^0 \in I(\tilde{\theta}_i^0)$.

At the second stage, after the history $h_1 = (\theta^0; a_I^0, a_P^0)$, we divide the unit interval according to probabilities assigned by θ^{h_1} to $\tilde{\theta}_1^{h_1}$, $\tilde{\theta}_2^{h_1}$ and $\tilde{\theta}_3^{h_1}$ and define $\tau(h_1, \theta^1) = \tau((\theta^0; a_I^0, a_P^0), \theta^1) = \sigma'((J(\theta^0); a_I^0, a_P^0), J^{h_1}(\theta^1))$, where J^{h_1} identifies θ^1 with $\tilde{\theta}_i^{h_1}$ according to whether θ^1 is in the interval corresponding to $\tilde{\theta}_i^{h_1}$. (Notice that the random variable θ^{h_1} is not history dependent. However, its interpretation by the players does depend on h_1). We continue this process inductively on the finite set of histories and thereby obtain τ. Clearly τ is an equilibrium (Nash or SP, as σ is). $\blacksquare_{\text{(Claim 3)}}$

Lemma A:

The sets of equilibrium payoffs (Nash or SP) are closed.

Proof:

Let $\left(X_n, Y_n\right)_{n=1}^{\infty}$ be a sequence of (Nash or SP) equilibrium payoffs that converges to $\left(X, Y\right)$. Let σ_n be the (Nash or SP) equilibrium that supports the payoff pair $\left(X_n, Y_n\right)$.

By Claims 1 and 2, we can assume that each equilibrium σ_n relies on public signals with a finite number of states. We will construct a (Nash or SP) equilibrium σ that supports the payoff pair (X,Y).

We extract a converging subsequence $(\sigma_{n_k})_{k=1}^{\infty}$. (I.e., given any history of actions and signals, the actions prescribed and the probabilities of the correlating devices converge.) We do so inductively on the set of histories: we extract a subsequence that converges on the first history, from this subsequence we extract a subsequence that converges on the second history, and so on. Finally, we pick the sequence $(\sigma_{n_k})_{k=1}^{\infty}$ using a diagonal technique. This can be done since the set of possible histories is countable: there is a finite number of possible actions and of possible realizations of the public signals. We let the strategy σ be the limit of the sequence $(\sigma_{n_k})_{k=1}^{\infty}$. Since the repeated-game payoffs are continuous, σ is an equilibrium and, moreover, the expected payoffs to the players from following σ is the limit (X,Y).

Note that σ employs history-dependent (and finitely valued) randomization devices. By Claim 3, there exists an equilibrium, τ, with the same payoffs, that employs history-independent and continuous randomization devices. This proves that the limit (X,Y) of the sequence of equilibrium payoffs $(X_n, Y_n)_{n=1}^{\infty}$ is supported by an equilibrium.

■

Lemma B:

The sets of equilibrium payoffs (Nash or SP) are convex.

Proof:

Simply notice that the players can use the first public signal, θ^0, to also jointly mix between equilibria of the repeated game. This means that the sets of equilibrium payoffs are convex. ■

Remark: As mentioned in the paper, the use of a correlating device is not essential to obtaining our results. Without a correlating device the equilibrium payoffs set is "almost" convex: The distance between the equilibrium set and its convex hull converges to zero as Δ goes to 0.

Lemma 2:

For any $\varepsilon > 0$, there exists $\underline{\Delta} > 0$ such that any payoff path, along which all continuation payoffs are in IR^{ε}, can be extended to a subgame perfect equilibrium when $\Delta < \underline{\Delta}$.

Proof of Lemma 2:

Denote player *i*'s minmax action against player *j* by M_i. That is, $ir_j = \max_{a_j} g_j(M_i, a_j)$. Let $L_i = g_i(M_i, M_j)$ be the payoff to player *i* when both players are minmaxing each other and let $H_i = \max_{a_i,a_j} g_i(a_i,a_j) - \min_{a_i,a_j} g_i(a_i,a_j)$ be the maximal difference between all of player *i*'s stage payoffs.

Note that at each stage *k*, the players can attain any payoffs pair in *V* by playing *pure* actions as a function of the stage-*k* signal θ^k. This implies that there is a path of stage-actions that generates the given payoffs path, such that a deviation by any player (from her prescribed pure action) can always be detected by her opponent. To extend this path to a SP equilibrium, we need to specify what players do after a deviation.

Assume that player i has deviated at stage *k*. The players enter a punishment phase, during which both players are minmaxing each other and receive stage payoffs $L_I \le ir_I$ and $L_P \le ir_P$. The length of the punishment phase is random: after each stage there is a probability ρ that the players will terminate the punishment phase and return to stage $k+1$ of the original path (More precisely, the probability of terminating the punishment phase depends on the actions taken by the players at that stage, but its expectations is ρ - see below.) By choosing ρ sufficiently small we can make the (random) number of stages of punishment large enough, so that player *i*'s one stage gain from the deviation (which is less than H_i) is washed away by the loss due to the punishment phase (which is at least ε per stage). By choosing $\underline{\Delta}$ sufficiently small, we can make the random duration of this phase (whose expected value is $\underline{\Delta}/\rho$) arbitrarily short, so that the expected continuation payoff at any stage of the punishment phase is IR.

If mixed actions were observable, possible deviations from the punishment phase could be easily deterred (when Δ is sufficiently small). This can be done, e.g., by reducing to 0 the probability ρ of ending the punishment phase at the stage that follows the deviation. However, we assumed that only *realized* actions are observed. Therefore, a player's deviation from her minmax action to any pure action in the support of the minmax action cannot be detected. To make such a deviation nonprofitable, we make each player indifferent between *all* her pure actions. We do this by letting the probability of terminating punishment phase, ρ, depend on the actions taken by the players.

More precisely, let $W_i^{\Delta}(k)$ be player i's expected continuation payoff during the punishment phase that follows a deviation from the equilibrium path a stage k. (Notice that this payoff does not depend on how many stages of punishment has past, but does depend on when the deviation has occurred.) $W_i^{\Delta}(k)$ is the solution to

$$W_i^{\Delta}(k) = \left(1-\delta_i^{\Delta}\right)\cdot L_i + \delta_i^{\Delta}\cdot\left(\rho\cdot U_i^{\Delta}(k+1) + (1-\rho)\cdot W_i^{\Delta}(k)\right)$$

where $U_i^{\Delta}(k+1)$ is i's continuation payoff of returning to the equilibrium path at stage k+1. Let $D_i^{\Delta}(k) = U_i^{\Delta}(k) - W_i^{\Delta}(k) > 0$ be the i's gain from terminating the punishment phase. If the realized actions at some stage during the punishment phase are a_I and a_P, the players return to the original path with probability $\rho(a_I, a_P)$ (rather than ρ), where

$$\rho(a_I, a_P) = \rho + \frac{(1-\delta_I^{\Delta})\cdot(L_I - g_I(a_I, M_P))}{\delta_I^{\Delta}\cdot D_I^{\Delta}(k)} + \frac{(1-\delta_P^{\Delta})\cdot(L_P - g_P(M_I, a_P))}{\delta_P^{\Delta}\cdot D_P^{\Delta}(k)}.$$

(Note that $\rho(a_I, a_P)$ is between 0 and 1 if Δ is sufficiently small.) The two terms added to ρ exactly offset the difference between the payoff each player actually received at that stage and her expected payoff had she used her minmax action. To see why, lets consider the impatient player. The one-stage gain from playing a_I is $g_I(a_I, M_P) - L_I$. Its effect on I's continuation payoff is $(1-\delta_I^{\Delta})\cdot(g_I(a_I, M_P) - L_I)$. On the other hand, a_I affects the probability $\rho(a_I, a_P)$ through the second summand. (Note that a_I has no affect on the third summand.) The effect of this incremental probability (the second summand) on the continuation payoff is $\left[\frac{(1-\delta_I^{\Delta})\cdot(L_I - g_I(a_I, M_P))}{\delta_I^{\Delta}\cdot D_I^{\Delta}(k)}\right]\cdot\left[\delta_I^{\Delta}\cdot D_I^{\Delta}(k)\right]$. (The gain from terminating the punishment phase, $D_I^{\Delta}(k)$, is multiplied by δ_I^{Δ} since it occurs with a delay of one stage). This exactly offsets the gain from playing a_I rather than M_I (Note that since the expectation of $\rho(a_I, a_P)$ when players play their minmax actions M_i is ρ, our calculation of $W_i^{\Delta}(k)$ and $D_i^{\Delta}(k)$ are correct.). An analogous argument shows that also the patient player has no incentive to deviate from the punishment phase. ■

Proposition 2:

Let B be a convex polygon of feasible stage payoffs and denote the vertices on the Pareto frontier of B as (x_0, y_0) to (x_l, y_l), with $x_0 > ... > x_l$ (and thus $y_1 < ... < y_l$). The Pareto frontier of $F^0(B)$ is the graph of the function $U_P(U_I)$:

$$U_P = y_m + \left(x_m - U_I\right)^{\frac{1}{r}} \left(S_m\right)^{\frac{r-1}{r}} \text{ whenever } \mu_m \geq U_I \geq \mu_{m+1}\ ,\ m = 0 \ldots l-1,$$

where

(i) $\quad r = \dfrac{\log(\delta_I)}{\log(\delta_P)} \quad (>1)$

(ii) $\quad S_m = \sum_{i=m}^{l-1} (x_i - x_{i+1}) (\dfrac{y_i - y_{i+1}}{x_{i+1} - x_i})^{\frac{r}{r-1}}$

(iii) $\quad \mu_0 = x_0$, and $\mu_{m+1} = x_m - S_m (\dfrac{y_m - y_{m+1}}{x_{m+1} - x_m})^{\frac{r}{1-r}}$, $m = 0 \ldots l-1$

Proof:

To find the extreme point of $F^0(B)$ in the direction $(\alpha, 1)$, $\alpha \geq 0$, we solve:

$$Max_{(x(t),y(t))} \int_0^\infty \alpha(-\log\delta_I)\delta_I^t x(t) + (-\log\delta_P)\delta_P^t y(t) \quad s.t.\ (x(t), y(t)) \in B \ \forall t \in [0, \infty)$$

Clearly, an α-optimal path uses only Pareto-optimal vertices of *B*. The α-optimal path starts at some vertex (x_m, y_m), and goes through vertices with increasing indices until, at the tail, it reaches the vertex (x_l, y_l). The optimal path is:

$$(x(t), y(t)) = \begin{cases} (x_m, y_m) & 0 \leq t < T_m \\ (x_{m+1}, y_{m+1}) & T_m \leq t < T_{m+1} \\ \ldots \\ (x_l, y_l) & T_{l-1} \leq t < \infty \end{cases}$$

Where T_i is the solution to

$$\alpha(-\log\delta_I)\delta_I^{T_i} x_i + (-\log\delta_P)\delta_P^{T_i} y_i = \alpha(-\log\delta_I)\delta_I^{T_i} x_{i+1} + (-\log\delta_P)\delta_P^{T_i} y_{i+1}$$

if the equation has a positive solution, and 0 otherwise. Thus, $m = m(\alpha)$ is the first integer such that $\alpha r \dfrac{x_m - x_{m+1}}{y_{m+1} - y_m} > 1$, and for $m \leq i \leq l-1, T_i$ is the solution to

$$\left(\frac{\delta_P}{\delta_I}\right)^{T_i} = \alpha r \frac{x_i - x_{i+1}}{y_{i+1} - y_i}. \text{ I.e., } T_i = \frac{\log\left(\alpha r \dfrac{x_i - x_{i+1}}{y_{i+1} - y_i}\right)}{\log\delta_P - \log\delta_I} \text{ for } m \leq i \leq l-1, \text{ and } T_i = 0 \text{ for } i < m.$$

We now calculate the players' payoffs from this path:

$$U_I = \int_0^{T_m} (-\log\delta_I)\delta_I^t x_m dt + \int_{T_m}^{T_{m+1}} (-\log\delta_I)\delta_I^t x_{m+1} dt + \ldots + \int_{T_{l-1}}^{\infty} (-\log\delta_I)\delta_I^t x_l dt =$$

$$-x_m \delta_I^t \Big|_0^{T_m} - x_{m+1}\delta_I^t \Big|_{T_m}^{T_{m+1}} - \ldots - x_l \delta_I^t \Big|_{T_{l-1}}^{\infty} =$$

$$-x_m(\delta_I^{T_m} - 1) - x_{m+1}(\delta_I^{T_{m+1}} - \delta_I^{T_m}) - \ldots - x_l(\delta_I^{T_{l-1}} - 0) = x_m + \sum_{i=m}^{l-1} \delta_I^{T_i} (x_{i+1} - x_i)$$

Similarly, $U_P = y_m + \sum_{i=m}^{l-1} \delta_P^{T_i} (y_{i+1} - y_i)$.

Notice that for $m \le i \le l-1$

$$\delta_I^{T_i} = e^{\log\delta_I \frac{\log(\alpha r/-s_i)}{\log\delta_P - \log\delta_I}} = e^{\frac{r}{1-r}\log\frac{\alpha r}{-s_i}} = \left(\alpha r \frac{x_i - x_{i+1}}{y_{i+1} - y_i}\right)^{\frac{r}{1-r}} \text{, and}$$

$$\delta_P^{T_i} = e^{\log\delta_P \frac{\log(\alpha r/-s_i)}{\log\delta_P - \log\delta_I}} = \left(\alpha r \frac{x_i - x_{i+1}}{y_{i+1} - y_i}\right)^{\frac{1}{1-r}}.$$

Thus,

$$U_I = x_m + \sum_{i=m}^{l-1}\left(\alpha r \frac{x_i - x_{i+1}}{y_{i+1} - y_i}\right)^{\frac{r}{1-r}} (x_{i+1} - x_i) = x_m - (\alpha r)^{\frac{r}{1-r}} \sum_{i=m}^{l-1} \left(\frac{y_{i+1} - y_i}{x_i - x_{i+1}}\right)^{\frac{r}{r-1}} (x_i - x_{i+1}) \text{, and}$$

$$U_P = y_m + \sum_{i=m}^{l-1}\left(\alpha r \frac{x_i - x_{i+1}}{y_{i+1} - y_i}\right)^{\frac{1}{1-r}} (y_{i+1} - y_i) = y_m + (\alpha r)^{\frac{1}{1-r}} \sum_{i=m}^{l-1} \left(\frac{y_{i+1} - y_i}{x_i - x_{i+1}}\right)^{\frac{1}{r-1}} (y_{i+1} - y_i) =$$

$$y_m + (\alpha r)(\alpha r)^{\frac{r}{1-r}} \sum_{i=m}^{l-1} \left(\frac{y_{i+1} - y_i}{x_i - x_{i+1}}\right)^{\frac{r}{r-1}} \frac{x_i - x_{i+1}}{y_{i+1} - y_i} (y_{i+1} - y_i) =$$

$$y_m + (\alpha r)(\alpha r)^{\frac{r}{1-r}} \sum_{i=m}^{l-1} \left(\frac{y_{i+1} - y_i}{x_i - x_{i+1}}\right)^{\frac{r}{r-1}} (x_i - x_{i+1})$$

This yields a parametric representation of the Pareto frontier,

(*) $\quad \big(U_I(\alpha), U_P(\alpha)\big) = \big(x_{m(\alpha)}, y_{m(\alpha)}\big) + \big(-1, \alpha r\big) C(\alpha) \quad ,0 < \alpha < \infty$

where $C(\alpha) = (\alpha r)^{\frac{r}{1-r}} \sum_{i=m}^{i-1} \left(\frac{y_{i+1} - y_i}{x_i - x_{i+1}}\right)^{\frac{r}{r-1}} (x_i - x_{i+1})$, and

$$m(\alpha) = \begin{cases} 0 & \alpha > \frac{y_1 - y_0}{x_0 - x_1}\Big/ r \\ 1 & \frac{y_1 - y_0}{x_0 - x_1}\Big/ r \ge \alpha > \frac{y_2 - y_1}{x_1 - x_2}\Big/ r \\ \ldots & \\ l & \frac{y_l - y_{l-1}}{x_{l-1} - x_l}\Big/ r \ge \alpha > 0 \end{cases}$$

We proceed towards deriving the explicit frontier formula $U_P(U_I)$. From (*), we have $C(\alpha) = x_{m(\alpha)} - U_I(\alpha)$. Denoting $S_m \equiv \sum_{i=m}^{l-1} \left(\frac{y_{i+1} - y_i}{x_i - x_{i+1}}\right)^{\frac{r}{r-1}} (x_i - x_{i+1})$, we get,

(**) $\quad \alpha = \frac{1}{r}\left(\frac{x_{m(\alpha)} - U_I(\alpha)}{\Sigma_{m(\alpha)}}\right)^{\frac{1-r}{r}}.$

Thus, the ranges: $-s_{m-1}/r \ge \alpha > -s_m/r$ for which $m(\alpha) = m$, translate into

$$m(U_I) = 0 \text{ when } x_0 \geq U_I > x_0 - (-s_0)^{\frac{r}{1-r}} \Sigma_0 \text{, and}$$

$$m(U_I) = m \text{ when } x_{m-1} - (-s_{m-1})^{\frac{r}{1-r}} \Sigma_{m-1} \geq U_I > x_m - (-s_m)^{\frac{r}{1-r}} \Sigma_m,\ m = 1..l-1.$$

From (*), we also have

$$U_P(\alpha) = y_{m(\alpha)} + (\alpha r)C(\alpha) = y_{m(\alpha)} + (\alpha r)(x_{m(\alpha)} - U_I(\alpha))$$

plugging αr from (**) yields the formula:

$$U_P(U_I) = y_{m(U_I)} + \left(x_{m(U_I)} - U_I\right)^{\frac{1}{r}} \left(\Sigma_{m(U_I)}\right)^{\frac{r-1}{r}},$$

where

$$m(U_I) = \begin{cases} 0 & x_0 \geq U_I > x_0 - (-s_0)^{\frac{r}{1-r}} \Sigma_0 \\ 1 & x_0 - (-s_0)^{\frac{r}{1-r}} \Sigma_0 \geq U_I > x_1 - (-s_1)^{\frac{r}{1-r}} \Sigma_1 \\ \ldots & \\ l-1 & x_{l-2} - (-s_{l-2})^{\frac{r}{1-r}} \Sigma_{l-2} \geq U_I > x_{l-1} - (-s_{l-1})^{\frac{r}{1-r}} \Sigma_{l-1} = x_l \end{cases}$$

■

Theorem 2

$interior(W) \subseteq SPE \subseteq E \subseteq W$

where $W = convex\,hull\left(F^0(V \cap IR), IR_P \cap F^0_{NE}(V \cap IR_I), IR \cap F^0_{NW}(V)\right)$

Figure 1 illustrates how to apply Theorem 2 to find the shape of the limit equilibrium sets in different games. The bold lines delineates *V*, and the grey areas are the (limit) equilibrium sets:

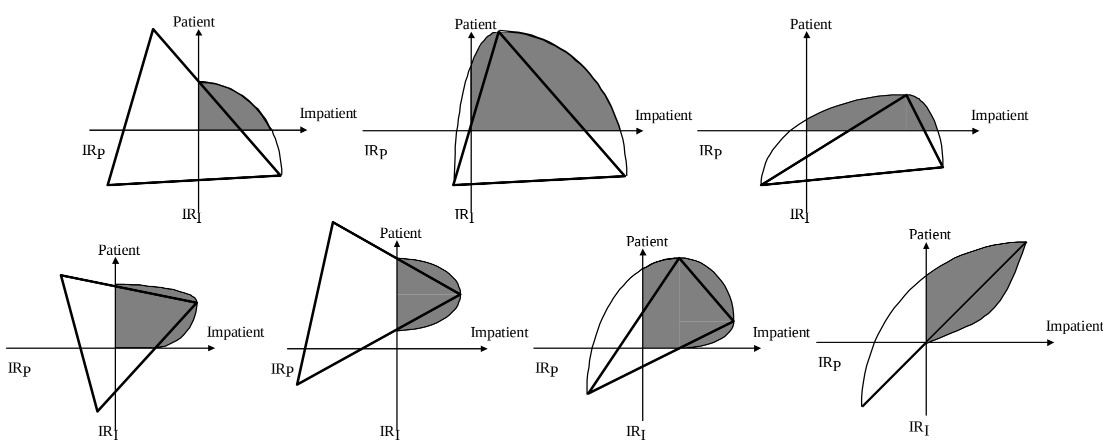

Figure 1

Proof:

The proof uses the technique of Theorem 2. Notice that Lemma 1 already applies to directions α in all quadrants. However, we need some extension of Lemma 3 that covers the southeast quadrant:

Lemma 4:

Suppose that $\alpha_I > 0$, $\alpha_P < 0$, and that the path $\{(X(k),Y(k))\}_{k=0}^{\infty} \in V^{\infty}$ maximizes $\alpha_I U_I^{\Delta}(0) + \alpha_P U_P^{\Delta}(0)$ s.t. $\forall k$, $U_I^{\Delta}(k) \geq ir_I, U_P^{\Delta}(k) \geq ir_P$. If $(U_I^{\Delta}(0), U_P^{\Delta}(0))$ is not strictly Pareto optimal (subject to the same constraints), then $\{(X(k),Y(k))\}_{k=0}^{\infty}$ can be chosen so that $\forall k$, $Y(k) \geq ir_P$.

Proof of Lemma 4:

Observe first that along the optimal path, all stage payoffs must be either on the north-eastern (Pareto) frontier of *V*, or on its southeastern frontier, since otherwise the impatient player's payoff can be increased keeping the patient player's payoff fixed, without violating any constraints.

Assume now that there are some stages k where $Y(k) < ir_P$. If for all such k, $(X(k),Y(k))$ is on the Pareto frontier of *V*, the path would be Pareto optimal, contradicting the assumption. Thus, assume that there is k with $(X(k),Y(k))$ on the southeastern frontier of *V* and $Y(k) < ir_P$. As in Lemma 3, we find some k_0 such that $Y(k_0) < ir_P$ and $Y(k_0+1) > Y(k_0)$.

There may be two cases. If $X(k_0+1) > X(k_0)$, we reach a contradiction using the same modification as in Lemma 3. If $X(k_0+1) \leq X(k_0)$, then the point $(X(k_0+1), Y(k_0+1))$ must be on the Pareto frontier of *V*. This case splits into two:

Case (a): $(X(k_0),Y(k_0))$ and $(X(k_0+1),Y(k_0+1))$ are on a vertical facet of *V*. In this case, $(U_I^{\Delta}(k_0+2), U_P^{\Delta}(k_0+2))$ must be also on the same facet, since otherwise it must be that $U_I^{\Delta}(k_0+2) < X(k_0+1)$, which implies that $U_P^{\Delta}(k_0+2) > Y(k_0+1)$. In such a case, an exchange of payoffs between stages k_0 and k_0+1 can improve on the optimal solution without violating the IR constraint. Thus, $(U_I^{\Delta}(k_0+2), U_P^{\Delta}(k_0+2))$ is indeed on the same facet. Therefore, all the payoffs from k_0 on can be replaced by a constant path which consists of one point on this vertical facet. Specifically, all the patient player's payoffs from k_0 on are above ir_P. Applying the same method again, if necessary, we can find a solution such that all her payoffs are above ir_P, as needed for the lemma.

Case (b): There exists some point $(x,y) \in V$, located to the right of the segment connecting $(X(k_0),Y(k_0))$ and $(X(k_0+1),Y(k_0+1))$. I.e., $Y(k_0+1) > y > Y(k_0)$ and

$x \geq X(k_0+1), X(k_0)$ with at least one inequality strict. Consider the following modification of the path:

$$(\hat{X}(k), \hat{Y}(k)) = (X(k), Y(k)) \text{ for } k \neq k_0, k_0+1$$
$$(\hat{X}(k_0+1), \hat{Y}(k_0+1)) = (1-\varepsilon)(X(k_0+1), Y(k_0+1)) + \varepsilon(x,y)$$
$$(\hat{X}(k_0), \hat{Y}(k_0)) = (1-\varepsilon')(X(k_0), Y(k_0)) + \varepsilon'(x,y)$$

where, as in Lemma 3, $\varepsilon > 0$ is small enough to prevent violation of the IR constraints at stage (k_0+1). $\varepsilon' = \varepsilon\delta_P^{\Delta}\dfrac{Y(k_0+1)-y}{y-Y(k_0)}$, so as to keep $U_P^{\Delta}(0)$ unchanged. Trivially, $U_I^{\Delta}(0)$ is increased, in contradiction to the α-optimality of the original path. ■(Lemma)

Lemma 5:

Suppose that $\alpha_I > 0$, $\alpha_P < 0$. Every α-optimal point in E^{Δ} which is not strictly Pareto-optimal, is α-dominated by a point in $F_{SE}^{\Delta}(V \cap IR)$.

Proof of Lemma 5:

By Lemma 4, an α-optimal point in E^{Δ} is α-dominated by a point supported by a path satisfying $\forall k,\ Y(k) \geq ir_P$. By Lemma 3, this path also satisfies $\forall k,\ Y(k) \geq ir_I$. ■(Lemma)

We return to the **proof of Theorem 2**. Denote,

$\phi(\varepsilon,\Delta) = convex\,hull\left(F^{\Delta}(V \cap IR^{\varepsilon}), IR_P^{\varepsilon} \cap F_{NE}^{\Delta}(V \cap IR_I), IR^{\varepsilon} \cap F_{NW}^{\Delta}(V)\right)$.

By definition, $W = \phi(0,0)$. By Proposition 1, $W^0 \subseteq \bigcup_{\varepsilon>0,\Delta>0}\phi(\varepsilon,\Delta)$.

To prove the theorem, we need to show that $\bigcup_{\varepsilon>0,\Delta>0}\phi(\varepsilon,\Delta) \subseteq SPE \subseteq E \subseteq \phi(0,0)$.

Part 1: $\bigcup_{\varepsilon>0,\Delta>0}\phi(\varepsilon,\Delta) \subseteq SPE$

Every payoff in $F^{\Delta}(V \cap IR^{\varepsilon})$ is supported by a path where every tail payoff is ε-strongly individually rational for both players. By Lemma 1, this is also the case for payoffs in $IR_P^{\varepsilon} \cap F_{NE}^{\Delta}(V \cap IR_I)$ and in $IR^{\varepsilon} \cap F_{NW}^{\Delta}(V)$. Using the technique of Theorem 1, such paths can be extended into subgame perfect equilibria for Δ small enough. The inclusion of the full convex hull holds because SPE is convex, noticing that F^{Δ_1} is included in F^{Δ_2} when Δ_1/Δ_2 is integer.

Part 2: $E \subseteq \phi(0,0)$

We need to show that for any $f \in E^{\Delta}$, $\Delta > 0$ and any direction $\alpha = (\alpha_I, \alpha_P)$ there is a point $x \in \phi(0,0)$ which α-dominates f (i.e., $\alpha \cdot x \geq \alpha \cdot f$).

Case 1 ($\alpha >> 0$): By Theorem 1, every point in E^{Δ} is Pareto dominated by some point in $IR_P \cap F_{NE}^{\Delta}(V \cap IR_I)$. By Proposition 1, this set is contained in $\phi(0,0)$.

Case 2 ($\alpha_I < 0, \alpha_P > 0$): There are two possibilities. If the Pareto frontier of *V* intersects the boundary of IR_I, then the intersection point must α-dominate any point $f \in E^{\Delta}$ (otherwise, there is a point $f \in E^{\Delta}$ that Pareto dominates the intersection point, in contradiction with Theorem 1). If it does not, then $IR \cap F^{\Delta}_{NW}(V)$ is not empty. In this case, every point in E^{Δ} is α-dominated by the α-optimal point in $F^{\Delta}_{NW}(V)$ if that point is in IR, or otherwise by the intersection of $F^{\Delta}_{NW}(V)$ with the boundary of IR. Again, by proposition 1, $F^{\Delta}_{NW}(V)$ is contained in $\phi(0,0)$.

Case 3 ($\alpha_I > 0, \alpha_P < 0$): Similarly to case 2, if the Pareto frontier of *V* intersects the boundary of IR_P, the intersection point of $F^{\Delta}_{NE}(V \cap IR_I)$ and IR_P must α-dominate any $f \in E^{\Delta}$ (again, otherwise there is a point $f \in E^{\Delta}$ that Pareto dominates the intersection point, in contradiction with Theorem 1). If, on the other hand, *V* does not intersect the boundary of IR_P, consider an α-optimal point in E^{Δ}. If this point is Pareto-optimal, it is clearly α-dominated by a point in $F^{\Delta}_{SE}(V \cap IR)$ - the lowest point on the (weak) Pareto-frontier which is still IR. If it is not Pareto-optimal, then by Lemma 5, it is also α-dominated by a point in $F^{\Delta}_{SE}(V \cap IR)$. By Proposition 1 the proof is complete.

Case 4 ($\alpha << 0$): Notice first that *V* must have at least one point in the southwest of (ir_I, ir_P), obtained when both players are minmaxing each other. If the point (ir_I, ir_P) is in *V*, we are done. Otherwise, if *V* lies above (ir_I, ir_P), any $f \in E^{\Delta}$ is α-dominated by the intersection point between the southeast frontier of *V* and the boundary of IR_I. That point is in $F^0(V \cap IR)$. Similarly, if *V* lies below (ir_I, ir_P), the intersection point between the northwest frontier of *V* and the boundary of IR_P does the job.

Case 5 ($\alpha_I = 0$ or $\alpha_P = 0$): The inclusion in this case follows from the four previous cases since the set E^{Δ} is convex. ■