

Trading favors: optimal exchange and
forgiveness.

Christine Hauser
University of Rochester

Hugo A. Hopenhayn
Department of Economics, UCLA

Very preliminary

June, 2005

1 Introduction

How is cooperation without immediate reciprocity sustained in a long term relationship? Consider the following example: Two firms are engaged in a joint venture. At random times one of them finds a discovery; if disclosed, the discovering firm's payoffs will be lower but total payoffs higher. In the absence of immediate reciprocity, will the expectation of future cooperation induce disclosure? How much cooperation can be supported? In the case of perfect monitoring -where the arrival of the discovery is jointly observed- full cooperation can be supported with trigger strategies when discounting is not too strong. But in the case of imperfect monitoring -when the arrival of the discovery is privately observed- the first best cannot be supported and the above questions have no definite answer. This paper attempts to fill this gap.

The setup is as follows.¹ Two players interact indefinitely in continuous time. At random arrival times, one player has the possibility of providing a benefit b to the other player at a cost $c < b$.² Following Mobius (2001) we call this a *favor*. This opportunity is privately observed, so a player will be willing to do this favor only if this gives an entitlement to future favors from the other player. Formally, this model is a repeated game with incomplete monitoring with random time intervals (given by the arrival of favors.) We characterize and solve for the Pareto frontier of public perfect equilibria (PPE).(references).

Mobius (2001) considers a simple class of PPE, which we will call the *chips mechanism*.³ Both players start with K chips each. Whenever a player receives a favor, she gives the other player a chip. If a player runs out of chips (so the other one has $2K$), she receives no favors until she gets a chip by granting a favor to the other player. Incentive compatibility (giving favors must be voluntary given the private information) puts a limit on the number of chips. This is obviously a very nice and simple mechanism. However, it has two special features that suggest there is room for improvement. In the first place, the rate of exchange is always one (current) for one (future) favor,

¹Our setup is identical to Mobius (2001), with very minor modifications.

²The motivating example can be easily accommodated by letting b denote the value to firm 2 when firm 1 shares information and c the decrease in firm's 1's profits when it discloses the information rather than keeping it secret.

³This type of mechanism is used in Skrzypacz and Hopenhayn (2004) in repeated auction with incomplete information.

so it is independent of the distribution of chips. Due to discounting, a player that is entitled to many future favors will value a marginal favor less. This suggests that the rate of exchange (or relative price) of favors should depend on current entitlements. Secondly, entitlements do not change unless a favor is granted (e.g. chips do not *jump* from one player to the other.) This is a special feature that rules out the possibility of interest or depreciation of entitlements. As we show in the paper, relaxing these two features allows for higher payoffs.

Our analysis proceeds in several steps. As usual in the literature, the recursive approach introduced by Abreu, Pearce and Stacchetti (APS) is used. We first establish that the set of Pareto optimal PPE is a self-generating set. This is not true in general and relies on some special features of our formulation that we discuss. Moreover, it also guarantees that the equilibrium is renegotiation proof. As a consequence of our result, the recursive formulation reduces to a one dimensional dynamic programming problem which is solved by a simple algorithm. Optimal PPE have two key features: 1) the relative price of favors decreases with a player's entitlement and 2) the entitlements change over time even in periods with no trade. As a consequence of the first result, starting from an initial symmetric point (the analogue of equal number of chips) if a player receives a number of consecutive favors, he must pay back considerably more to return to the initial point. We solve the model numerically for a large set of parameter values and find that the gains relative to the chips mechanism can be quite large (in some cases over 30% higher). Interestingly, in all our numerical simulations the disadvantaged player's utility increases over time during periods of no trade, so in the optimal equilibria there is forgiveness.

Our model is a continuous time-repeated game with imperfect monitoring. This is a class of games that had not been previously analyzed. An exception is Sannikov (2004), which in independent recent work studies games within this class where the stochastic component follows independent diffusion processes and provides a differential equation that characterizes the boundary of the set of PPE. Our model does not fit exactly in that class since our stochastic process is a jump process (Poisson arrivals), yet we also derive a differential formulation to characterize the boundary of the set of PPE. In addition we establish that the set of Pareto optimal PPE is self-generating, which as far as we know is a new result in the literature on repeated games of imperfect monitoring.

The paper is organized as follows. Section 2 describes the model. Section

3 describes in more detail the chips mechanism. Section 4 develops the recursive formulation. Section 5 describes the solution algorithm and provides numerical results.

2 The model

We analyze an infinite horizon, two-agent partnership. Time is continuous and agents discount future utility at rate r . There are two symmetric and independent Poisson processes -one for each agent- with arrival rates α representing the opportunity of producing a favor. We assume that favors are perfectly divisible, so partners can provide fractional favors.⁴ Agents' utilities and costs are linear in the amount of favors exchanged. The cost per unit of a favor is c and the corresponding benefit to the other player $b > c$. Letting x_i represent a favor granted by player i and x_j a favor received, the utility for player i is given by:

$$U_i(x_i, x_j) = -cx_i + bx_j.$$

Since arrivals are Poisson and independent, only one player is able to grant a favor at a point in time, so $x_i(t) > 0$ implies $x_j(t) = 0$. We assume that arrivals are privately observed by each player, so the ability of providing a favor is private information. Since the cost of providing a favor is less than the benefit generated, it is socially optimal for agents to grant favors. Indeed, in the absence of informational constraints, a public perfect equilibrium would exist that achieves this optimum through a simple Nash reversion strategy: An agent grants favors whenever she can, as long as her partner has done so in the past, and stops granting favors whenever her partner has defected. This equilibrium can be supported for

$$c < \left(\frac{\alpha}{r + \alpha} \right) b$$

The problem arises when the ability to do a favor is private information, so an agent observes whether her partner has provided a favor or not, but is unable to detect a deviation where her partner has passed the opportunity to do a favor. The question then becomes: How to ensure the maximum

⁴Alternatively, given our assumption of linear utilities, we can assume there is public randomization for the provision of favors.

cooperation and exchange of favors between agents given these informational constraints?

3 A Simple Debt Accounting Mechanism

In his paper, Mobius considers equilibria of a simple class. The equilibrium proposed is Markov perfect, where the state variable is the difference k between the number of favors granted by agent 1, and those received from agent 2. For $-K \leq k \leq K$, the agents obey the following strategies:

- Agent 1 grants favors if $k < K$, and stops granting favors if $k = K$
- Agent 2 grants favors if $k > -K$, and stops granting favors if $k = -K$

It is obvious that the choice of K is crucial in determining the expected payoffs of players. Since it is efficient to have favors granted whenever possible, an efficiency loss occurs when agents reach the boundaries K and $-K$, where only the indebted agent is granting favors. So the larger K is, the lower is the incidence of this situation, and the larger are the expected payoffs of the agents. To understand how K is determined, note that because of discounting the marginal value of the right to an extra favor diminishes with the current entitlement of favors the player has; $2K$ is the largest number such that this marginal value exceeds the cost c .

The scheme proposed by Mobius' is very simple. Moreover it is asymptotically efficient (as $\alpha/r \rightarrow \infty$). To see this, note that in the long run the distribution over the states $k = \{-K, -K + 1, \dots, 0, 1, \dots, K - 1, K\}$ is uniform, so the probability that a favor is not granted is $1/K$. It is easy to verify that $K \rightarrow \infty$ as $\alpha/r \rightarrow \infty$, so this probability goes to zero.

There are two special features of this scheme that suggest there is room for improvement. The first special feature is that the rate of exchange of current for future favors is the same (equal to one) regardless of entitlements. Relaxing this constraint could reduce the region of inefficiency for two reasons. First, consider the case where the state is K so agent one has the maximum entitlement of favors. As we argued before, at this point agent's one marginal value to an entitlement of an extra favor is lower than c . But there is still room for incentives if agent two were to promise more than one favor in exchange. Secondly, in Mobius' scheme the incentive constraints only bind at the extremes, but are slack in between, where the marginal value

to future favors exceeds c . A lower rate of exchange could allow to expand the number of possible favors.

The second special feature is that agents' continuation values do not change unless one or the other agent grants a favor. This is restrictive, and rules out the possibilities of appreciation (charging interest) or depreciation (forgiveness) of entitlements and punishment in case "not enough" cooperation is observed.

In the following section, strategies are not limited to a particular scheme. Instead, we characterize the optimal Perfect Public equilibria

4 Characterizing the optimal Perfect Public Equilibrium

The game described above falls in the class of repeated games with imperfect monitoring. As usual in this literature, we restrict our analysis to Public Perfect Equilibrium (PPE), where strategies are functions of the public history only and equilibrium is perfect Bayesian.

A public history up to time t , denoted by h^t , consists of agents' past favors including size and date. A strategy $x_{it} : h^t \rightarrow [0, 1]$ for player i specifies for every history and time period the size of favor the agent grants if the opportunity to do so arises. A public perfect equilibrium is a pair of strategies $\{x_{1t}, x_{2t}\}_{t \geq 0}$ that constitute a perfect Bayesian equilibrium. To analyze this game, we consider a recursive representation following the formulation of Abreu, Pearce and Stacchetti (APS). Let

$$\mathbf{V}^* = \{(v_1, v_2) \mid \exists \text{ PPE that achieves these values}\},$$

the set of PPE values. Given the linearity of payoffs and convexity of strategy sets, it follows that the set V^* is convex. The set V^* is a subset of the set of feasible and individually rational payoffs, which is given in figure 1.

The set of feasible payoffs is given by the outer x - y axis. Values have been multiplied by r/α so they are expressed as flow-equivalents. The extreme points are given by the vectors $(-c, b)$, $(b, -c)$, $(b - c, b - c)$, $(0, 0)$. The first two correspond to the situation where only one player is giving favors and doing so at every possible time; the third vector corresponds to full cooperation by both players and the last one to no cooperation. The feasible set of payoffs is obtained by convex combinations of these points and are




Figure 1: Feasible and individually rational payoffs (flow equivalent)

derived by combining the endpoints over time. Individually rational payoffs are obtained by restricting this set to the positive orthant. The set V^* is a proper subset of this set.

4.1 Factorization

To characterize the set V^* we follow the general idea of APS, which decompose (factorize) equilibrium values into strategies for the current periods and continuation values for each possible end of period public signal. The difficulty in our case is that there is no *current period* in our model.

Let \mathbf{W} denote an initial set of vectors of continuation values for the players (v, w) . Adapting the approach in Abreu, Pearce and Stacchetti (APS), equilibrium payoffs can be factorized in the following way. Let t denote the (random) time at which the next favor occurs. Consider any $T > 0$. Factorization is given by functions $x_1(t), x_2(t), v_1(t), v_2(t), w_1(t), w_2(t)$ and values v_T, w_T with the following interpretation: If the first favor occurs at time t and is given by player i , then $x_i(t) \in [0, 1]$ specifies the size of the favor, $v_i(t)$ the continuation value for player one and $w_i(t)$ the continuation value for player two, where for each t the vector $(v_i(t), w_i(t)) \in \mathbf{W}$. If no favor occurs until time T , the respective continuation values are v_T, w_T . The

strategies and continuation values give the following utility to player one:

$$v = \int_0^T e^{-rt} \left\{ \frac{x_2(t)b + v_2(t) - x_1(t)c + v_1(t)}{2} \right\} p(t) dt + e^{-(r+2\alpha)t} v_T$$

where $p(t)$ denotes the density of the first arrival occurring at time t . This is the density of an exponential distribution with coefficient 2α (the total arrival rate). Letting $\beta = e^{-(r+2\alpha)}$ and $z_1(t) = x_2(t)b - x_1(t)c + v_1(t) + v_2(t)$, the above equation simplifies to:

$$v = \alpha \int_0^T \beta^t z_1(t) dt + \beta^T v_T. \quad (1)$$

Similarly, letting $z_2(t) = x_1(t)b - x_2(t)c + w_1(t) + w_2(t)$ one obtains the value w for player two:

$$w = \alpha \int_0^T \beta^t z_2(t) dt + \beta^T w_T. \quad (2)$$

The incentive compatibility condition requires that for all t ,

$$v_1(t) - x_1(t)c \geq \alpha \int_t^T \beta^{s-t} z_1(s) ds + \beta^{T-t} v_T \quad (3)$$

$$w_2(t) - x_2(t)c \geq \alpha \int_t^T \beta^{s-t} z_2(s) ds + \beta^{T-t} w_T \quad (4)$$

The left hand side gives the net utility of giving a favor at time t and the right hand side the continuation utility if the agent passes this opportunity.

Starting with a set $\mathbf{W} \subset \mathfrak{R}_+^2$ of values, this factorization gives a new set of values $B_T(\mathbf{W})$ given by all pairs (v, w) such that there exist functions $x_1(t), x_2(t), v_1(t), v_2(t), w_1(t), w_2(t)$, where $(v_i(t), w_i(t)) \in \mathbf{W}$ and (1), (2), (3), (4) are satisfied. Following APS, a set of values \mathbf{W} is *self generating* if $\mathbf{W} \subset B_T(\mathbf{W})$. If $(v, w) \in \mathbf{W}$, then there exists a PPE that gives the players initial payoffs (v, w) . The set of PPE V^* is the largest set \mathbf{V} such that $\mathbf{V} = B_T(\mathbf{V})$. The corresponding Pareto frontier of values can be characterized by the following program:

$$\begin{aligned} W(V) = \max & \alpha \int_0^T \beta^t \{x_1(t)b - x_2(t)c + w_1(t) + w_2(t)\} dt + \beta^T w_T \\ & \text{subject to (1), (3), (4), } (v_i(t), w_i(t)) \in \mathbf{V} \text{ and } (v_T, w_T) \in \mathbf{V} \end{aligned}$$

In general, payoffs in the Pareto frontier may require *inefficient equilibria* (i.e. equilibria with dominated payoffs) after some histories. The following proposition shows that this is not needed in our repeated game.

Proposition 1 *The Pareto set of values $\{(v, w) \in \mathbf{V}^*$ such that $w = W(v)\}$ is self-generating. The domain is given by an interval $[0, v_h]$ where $v_h \leq \bar{v}$. The Pareto frontier is concave and its slope lies in the interval $[-b/c, -c/b]$.*

Proof. see appendix 1. ■

Self-generation implies that any point in the Pareto frontier can be obtained by relying on continuation values that are also in the frontier. This implies that PPE supporting the Pareto frontier are renegotiation proof.

As in APS, an algorithm of successive approximations can be defined by iterating on the operator B_T , starting from a set containing V^* (such as the set of feasible and individually rational payoffs defined above.) This procedure converges monotonically (by set inclusion) to V^* . The algorithm can be simplified in our case, restricting to iterations on a value function defined by the frontier of values. It can be shown that starting from the frontier of the set of feasible and individually rational payoffs, convergence to the frontier of V^* is monotonic. There is a difficulty with this algorithm, since for each iteration the optimal strategies are the solution to an optimal control problem rather than a simple optimization problem, as in APS. Appendix 2 provides an alternative simplified algorithm, where incentive constraints are relaxed, that gives monotone convergence (from above) to the Pareto frontier relying on an elementary optimization problem.

The following section develops an alternative procedure to characterize the Pareto frontier that relies on the differentiable structure of the game.⁵

5 A differential approach

We follow a heuristic approach. For small T the equation

$$W(v(0)) = \alpha \int_0^T \beta^t \{x_1(t)b - x_2(t)c + W(v_1(t)) + W(v_2(t))\} dt + \beta^T W(v(T))$$

⁵A related procedure was developed by Sannikov (2004) for a continuous time game with stochastic diffusion processes.

can be approximated by:

$$W(v(0)) - W(v(T)) = \alpha T \{x_1 b - x_2 c + W(v_1) + W(v_2)\} + (\beta^T - 1) W(v_T)$$

Dividing by T , taking limits as $T \rightarrow 0$ and letting $v(0) = v$,

$$W'(v)\dot{v} = -\alpha \{x_1 b - x_2 c + W(v_1) + W(v_2)\} + (r + 2\alpha) W(v)$$

Following a similar procedure in equation (1) gives:

$$\dot{v} = -\alpha \{x_2 b - x_1 c + v_1 + v_2\} + (r + 2\alpha) v$$

Finally, taking limits the incentive constraints read:

$$\begin{aligned} v_1 - x_1 c &\geq v \\ W(v_2) - x_2 c &\geq W(v). \end{aligned}$$

In this continuous time problem, the choice variables are x_1, x_2, v_1, v_2 and \dot{v} , where the first four variables are the analogue of the controls $x_1(t), x_2(t), v_1(t), v_2(t)$ in the previous problem, and \dot{v} is the analogue of choosing $v(T)$. The optimization problem can be rewritten as:

$$rW(v) = \max_{x_1, x_2, v_1, v_2, \dot{v}} \alpha(x_1 b - x_2 c) \quad (5)$$

$$+ \alpha (W(v_1) - W(v) + W(v_2) - W(v)) + W'(v) \dot{v} \quad (6)$$

subject to :

$$rv = \alpha (x_2 b - x_1 c + v_1 - v + v_2 - v) + \dot{v} \quad (7)$$

$$v \leq v_1 - x_1 c \quad (8)$$

$$W(v) \leq W(v_2) - x_2 c. \quad (9)$$

The following properties can be established as a result of the concavity of the function W .

Proposition 2 *The solution to the optimization problem defined by (5)-(9) has the following properties:*

1. *Both incentive constraints bind.*
2. $x_1 = \min((v_h - v)/c); x_2 = \min((W(0) - W(v))/c).$

As in Mobius', favors are done while the values of the players are away from the boundary. Full favors are given unless there is not enough utility in the set to compensate for the cost of the player giving the favor. In that case, the size of the favor given by the player is limited by the distance to this boundary (divided by the unit cost). In contrast to Mobius, the lower bound on player values is actually zero, so the individual rationality constraint binds. Also in contrast to Mobius, the relative price of favors depends on v , i.e. on the entitlements of the players. This rate of exchange is approximately $|W'(v)| \in [c/b, b/c]$. In the simulations reported below, the extremes of this set are attained (at least approximately), giving rise to a range of relative prices of the order of $(b/c)^2$. Both of these features, i.e. the larger domain of values and the variable relative prices of favors, can potentially accommodate a considerably larger number of favors before reaching the frontiers.

The following Proposition gives properties of the optimal \dot{v} .

Proposition 3 *In the optimal policy,*

1. $\dot{v}(0) = \dot{v}(v^*) = \dot{v}(v_h) = 0$, where v^* is the unique point satisfying $v^* = W(v^*)$;
2. $\dot{v}'(0) > 0$, $\dot{v}'(\bar{v}) < 0$.

This Proposition shows that $\dot{v} > 0$ in some region to the left of v^* and $\dot{v} < 0$ in some region to the right of this point. Our simulations reported below suggests that this property indeed holds in the whole region (except the extremes), so that \dot{v} drifts towards the *equal treatment* point v^* : the player with a lower entitlement is gradually rewarded.

In the equilibrium described above, the player's values provide an accounting device of past history and current entitlements. An alternative equivalent accounting is given below. Let T_i represent the expected discounted number of favors that player i will give in the rest of the game. It obviously follows that:

$$\begin{aligned} v &= T_2 b - T_1 c \\ w &= T_1 b - T_2 c \end{aligned}$$

which solving gives:

$$T_1 = \frac{vc + bw}{b^2 - c^2}$$

$$T_2 = \frac{wc + bv}{b^2 - c^2}$$

and

$$T_1 - T_2 = \frac{v - w}{b + c}.$$

This difference has the interpretation of a net balance between the player's assets and liabilities. When $T_1 < T_2$, the player is in a net debt position. As we find, $\dot{v} > 0$ in this region, which has the interpretation of debt forgiveness.

6 Numerical results

The optimal strategies differ in several dimensions from the very simple strategies proposed by Mobius. How important is this? This section provides some numerical computations to examine this question.

There are 4 parameters in the model: r, α, c, b . In comparing the performance of different alternatives, two normalizations can be made where all that matters in these comparisons are the values of c/b and r/α . In the next tables, the following normalizations are used: $b = 1$ and $r = 0.01$.

The following tables give a measure of how far each alternative scheme is from the first best at the symmetric point of the boundary where players get equal utilities. The first column gives the % values for the optimal scheme described above; the second column gives values for an optimal scheme with the added restriction that $\dot{v} = 0$; the third column gives the values for Mobius' scheme.

%difference with optimum				%difference with optimum			
α	0.8			c	0.5		
c	$\dot{V} \neq 0$	$\dot{V} = 0$	Mobius	α	$\dot{V} \neq 0$	$\dot{V} = 0$	Mobius
0.4	2.7	3	6.4	0.2	6	6.5	15.5
0.5	3.1	3.3	7.6	0.3	4.9	5.3	12.9
0.65	3.4	3.5	10.5	0.4	4.3	4.6	11
0.75	3.5	3.6	13.5	0.5	3.8	4.1	9.8
0.95	10.3	10.3	32.3	0.6	3.5	3.7	9.1

The performance of all these schemes decreases with c and increases with the arrival rate α . There can be substantial improvements over Mobius' scheme: e.g. for $\alpha = 0.8$ and $c = 0.95$, the second best is 10% within the first best, while Mobius' scheme is 30% apart. It is also interesting to observe that restricting $\dot{v} = 0$ does not have a substantial impact on performance.

The following two tables give the maximum number of consecutive favors starting at the midpoint that lead to the boundary. The optimal scheme can accommodate a much larger number of favors (between 5 to over 10 times more.) In part this is due to flexible relative prices, which depart significantly from one.

Number of favors		Avg. price left	
α	0.8	of midpoint	
c	$\dot{V} \neq 0$	Mobius	
0.4	167.1	13.2	2.32
0.5	119	11.2	1.87
0.65	70.1	8.1	1.46
0.75	45.7	6.2	1.29
0.95	6.5	2.1	1.39

Number of favors		Avg. price left	
c	0.5	of midpoint	
α	$\dot{V} \neq 0$	Mobius	
0.2	29	5.1	1.87
0.3	44	6.2	1.87
0.5	74	8.5	1.87
0.6	89	9.2	1.87

References

- [1] Abreu, Dilip, David Pearce and Ennio Stacchetti, "Toward a Theory of Discounted Repeated Games with Imperfect Monitoring," *Econometrica* 58(5), September 1990, 1041-1063.
- [2] Mobius, Markus (2001) "Trading favors ", manuscript May 2001.
- [3] Sannikov, Yuliy "Games with Imperfectly Observable Actions in Continuous Time,"

7 Appendix I: Proof of Proposition 1

Take a point (v, w) in the Pareto frontier of V^* . For any T this is factored by strategies and continuation values $\{x_i(t), v_i(t), w_i(t), v_T, w_T\}$. Suppose (v_T, w_T) is not in the Pareto frontier so there exists $\varepsilon > 0$ such that $(v_T + \varepsilon, w_T + \varepsilon) \in V^*$. Define new paths $\tilde{x}_i(t) = \max(0, x_i(t) - d(t))$, where $d(t) = \frac{\varepsilon}{c} e^{\alpha(T-t)} \beta^{T-t}$. We prove these paths together with

$$\{v_i(t), w_i(t), v_T + \varepsilon, w_T + \varepsilon\}$$

are admissible with respect to V . Letting

$$z(s) = x_1(s)b - x_2(s)c + w_1(s) + w_2(s),$$

$$\begin{aligned} & \alpha \int_t^T \beta^{s-t} \{\tilde{x}_2(s)b - \tilde{x}_1(s)c + v_1(s) + v_2(s)\} ds + \beta^{T-t} (v_T + \varepsilon) \\ \leq & \alpha \int_t^T \beta^{s-t} \{z(s) + \varepsilon e^{\alpha(T-s)} \beta^{T-s}\} ds + \beta^{T-t} (v_T + \varepsilon) \\ = & \alpha \int_t^T \beta^{s-t} z(s) ds + \beta^{T-t} (v_T + \varepsilon) + \alpha \varepsilon \beta^{T-t} \left(\frac{-1 + e^{\alpha(T-t)}}{\alpha} \right) \\ = & \int_t^T \beta^{s-t} z(s) ds + \beta^{T-t} v_T + d(t) c \\ \leq & v_1(t) - x_1(t) c + d(t) c = v_1(t) - \tilde{x}_1(t) c \end{aligned}$$

so the incentive constraint for player one is satisfied. Similar argument shows that the same holds for player two. Let (\tilde{v}, \tilde{w}) denote the values associated to the new admissible path. Then

$$\begin{aligned} \tilde{v} & \geq \alpha \int_0^T \beta^t z(s) ds + \beta^t (v_T + \varepsilon) - \alpha \int_0^T \beta^t b d(t) dt \\ & = v + \beta^t \varepsilon - \alpha \int_0^T \beta^t b d(t) dt. \end{aligned}$$

Since $d(t)$ is bounded, the right hand side exceeds v for T small. Similar argument shows that $\tilde{w} > w$, contradicting the hypothesis that (v, w) are in the Pareto frontier of V^* . This proves that (v_T, w_T) must belong to the Pareto frontier of V^* .

We now show that $(v_i(t), w_i(t))$ also belong to the frontier. Suppose towards a contradiction that this was not true. Let

$$b = \min \{v + w \mid (v, w) \text{ are in the Pareto frontier of } V^*\}.$$

From the concavity and symmetry of the boundary of V^* , $b = W(0)$ which by lemma 4 is strictly positive. It follows that $v_T + w_T \geq b$. Let T_i be the set of time periods such that $v_i(t), w_i(t)$ are not in the Pareto frontier of V^* . We construct an alternative admissible path that improves on the given one as follows. Let $\hat{v}_i(t)$ be the value such that the pair $(\hat{v}_i(t), w_i(t))$ is in the frontier. Similarly define $\hat{w}_i(t)$. Without loss of generality (by choice of T),

$$\int_0^T \beta^t (\hat{v}_i(t) - v_i(t)) dt + \int_0^T \beta^t (\hat{w}_i(t) - w_i(t)) dt < \beta^T b.$$

It is thus possible to construct paths $\{\tilde{v}_i(t), \tilde{w}_i(t)\}$ where for each t and i either $(\tilde{v}_i(t), \tilde{w}_i(t))$ equals $(v_i(t), w_i(t))$ or it equals $(\hat{v}_i(t), w_i(t))$ such that

$$\varepsilon_1 = \int_0^T \beta^t [\tilde{v}_1(t) - v_1(t) + \tilde{v}_2(t) - v_2(t)] dt < v_T$$

and

$$\varepsilon_2 = \int_0^T \beta^t [\tilde{w}_1(t) - w_1(t) + \tilde{w}_2(t) - w_2(t)] dt < w_T.$$

Define $\tilde{v}_T = v_T - \varepsilon_1$ and $\tilde{w}_T = w_T - \varepsilon_2$. Note that $(\tilde{v}_T, \tilde{w}_T) \in \text{int}(V^*)$. We now show that the path $\{x_i(t), \tilde{v}_i(t), \tilde{w}_i(t), \tilde{v}_T, \tilde{w}_T\}$ is admissible. Let

$$\delta_1(t) = \tilde{v}_1(t) - v_1(t) + \tilde{v}_2(t) - v_2(t)$$

and

$$\delta_2(t) = \tilde{w}_1(t) - w_1(t) + \tilde{w}_2(t) - w_2(t).$$

$$\begin{aligned} & \alpha \int_t^T \beta^{s-t} \tilde{z}_i(t) dt + \beta^{T-t} \tilde{v}_T \\ &= \alpha \int_t^T \beta^{s-t} z_i(t) dt + \beta^{T-t} v_T + \int_t^T \beta^{s-t} \delta_1(t) dt - \beta^{T-t} \varepsilon_1 \\ &= \alpha \int_t^T \beta^{s-t} z_i(t) dt + \beta^{T-t} v_T + \beta^{-t} \left(\int_t^T \beta^s \delta_1(t) dt - \beta^T \varepsilon_1 \right) \\ &\leq \alpha \int_t^T \beta^{s-t} z_i(t) dt + \beta^{T-t} v_T \\ &\leq v_1(t) - x_1(t) c \\ &\leq \tilde{v}_1(t) - x_1(t) c. \end{aligned}$$

A similar argument can be used to verify incentive compatibility for player two. We have constructed paths where $(\tilde{v}_i(t), \tilde{w}_i(t))$ are in the Pareto frontier for all t . To end the proof, note that since $(\tilde{v}_T, \tilde{w}_T)$ are not in the frontier, by our previous argument the path can be further improved by an alternative one that takes values in the Pareto frontier (without modifying $(\tilde{v}_t(t), \tilde{w}_t(t))$).

Lemma 4 *The boundary of V^* has slope in the set $[-b/c, -c/b]$.*

Proof. We show that if the set V satisfies this property, so will the set $U = B_T(V)$ for all T . Let (v, w) be a point in the boundary of U where $v > 0$. If there is a set of positive Lebesgue measure where $x_2(t) > 0$ then for any $0 < \varepsilon < \int \beta^t x_2(t) dt$ there is a new path $\hat{x}_2(t) \leq x_2(t)$ with

$$\int \beta^t \hat{x}_2(t) dt = \int \beta^t x_2(t) dt - \varepsilon.$$

This gives rise to the points $(v - \varepsilon b, w + \varepsilon c)$ in $B_T(V)$. On the contrary suppose that $x_2(t) = 0$ for all $t \in [0, T]$. Since $v > 0$, either $v_T > 0$ or there exists a set of positive Lebesgue measure in $[0, T]$ where $v_1(t)$ or $v_2(t)$ is strictly positive. Since $x_2(t) = 0$ for all t , the points (v_T, w_T) , $(v_1(t), w_1(t))$ and $(v_2(t), w_2(t))$ must be in the Pareto frontier of V . Suppose $v_T > 0$. Then for any $0 < \varepsilon < v_T$ there exists values (\hat{v}_T, \hat{w}_T) in the frontier of V where $\hat{v}_T = v_T - \varepsilon$ and $\varepsilon \cdot c/b \leq \hat{w}_T - w_T$. These terminal values together with the original path are self-generating (recall that $x_2(t) = 0$ for $t \in [0, T]$) and give rise to a point $(\hat{v}, \hat{w}) \in B_T(V)$ where $(\hat{w} - w) / (\hat{v} - v) \leq -c/b$. Similar argument can be used in case there exists a set of positive Lebesgue measure in $[0, T]$ where $v_1(t)$ or $v_2(t)$ is strictly positive. This proves that the boundary of $B_T(V)$ has slope less or equal to $-c/b$. A symmetric argument shows that the slope is greater or equal to $-b/c$. ■

8 Appendix 2: Algorithm

We consider a relaxed problem. More precisely, we will say that

$$\left(\{x_i(t), v_i(t), w_i(t)\}_{t=0}^T, v_T, w_T \right)$$

is weakly admissible with respect to V if the above conditions hold with the weaker incentive constraints:

$$\begin{aligned} v_1(t) - x_1(t) c &\geq \beta^T v_T \\ w_2(t) - x_2(t) c &\geq \beta^T w_T \end{aligned}$$

Lemma 5 *Suppose that $(\{x_i(t), v_i(t), w_i(t)\}_{t=0}^T, v_T, w_T)$ is weakly admissible with respect to convex V . then the constant paths*

$$(x_1, x_2, v_1, v_2, w_1, w_2, v_T, w_T)$$

defined by:

$$\begin{aligned} x_i &= \frac{1}{\int_0^T \beta^t dt} \int_0^T \beta^t x_i(t) dt \\ v_i &= \frac{1}{\int_0^T \beta^t dt} \int_0^T \beta^t v_i(t) dt \\ w_i &= \frac{1}{\int_0^T \beta^t dt} \int_0^T \beta^t w_i(t) dt \end{aligned}$$

are weakly admissible with respect to V .

Proof. By convexity of V , the continuation values lie in V and it is also obvious that $0 \leq x_i \leq 1$. So we only need to verify (weak) incentive compatibility. The incentive verify immediately. ■

Denote by \hat{B}_T the associated APS operator. Note that by definition, the *averaged* path gives rise to the same values v, w . So without loss of generality in defining \hat{B}_T we can restrict to constant paths. Because the incentive constraints defining \hat{B}_T are weaker than those defining B_T , its largest fixed point contains the set of PPE values.