Long-Run Implementation in Repeated Public Good Games with Incomplete Information*

Ricardo Serrano-Padial†
University of Wisconsin-Madison

September 2005

Abstract

Despite predictions of complete free riding in one shot public good games, repeated interaction allows for any level of public provision, so long as it is feasible and Pareto superior to no provision at all. I investigate the long run effects of weakening the information players have about each others' preferences. I find that when agents are patient enough any provision level can be attained in the long run.

JEL Classification: *C73, D82, H41.*
Keywords: *Public good games, repeated games, incomplete information, belief-invariant Nash equilibrium.*

1 Introduction

The Folk Theorem asserts that cooperation can be sustained in non-zero sum repeated games through punishment threats. This result is based on two informational assumptions. First, payoff functions are observable. Thus, players are aware of what payoffs can be implemented in equilibrium and how to punish a deviation. Second, players have enough information about others' past actions to detect with high probability deviations from the equilibrium path.[1] While imperfect monitoring of

*I am deeply grateful to Vince Crawford, Navin Kartik, Nagore Iriberri, Jason Murray, Adam Sanjurjo, Joel Watson, and especially to Joel Sobel and two anonymous referees, for their valuable advice.

†Email: rserrano@ssc.wisc.edu; Web: www.ssc.wisc.edu/~rserrano; Address: Department of Economics, Social Science Building, 1180 Observatory Drive, Madison, WI 53706-1393.

[1] For the Folk Theorem in games of complete information, see Aumann and Shapley (1976) and Rubinstein (1977, 1979) for games with no discounting and Fudenberg and Maskin (1986) for the discounted case. Fudenberg, Levine, and Maskin (1994) study games with imperfect monitoring.

players' actions has received much attention, little is known about what happens when there is incomplete information about players' payoff functions, which depend on privately observed types. In this context, players may not be able to induce adherence to some strategy profiles satisfying feasibility and individual rationality in the complete information game. Ignorance of a player's payoff function may preclude other players punishing a deviation. Moreover, detecting a deviation may not be possible in strategies that require the revelation of private information. We investigate the latter case in a general class of infinitely repeated public good games with discounting. These games have the property that the existence of private information hinders players' ability to detect a deviation but it does not affect their ability to find suitable punishments. Hence, we are able to isolate the effect of the detection problem on the set of equilibrium payoffs from enforcement issues.

We study long run implementation in a class of dynamic public good games defined as follows. A finite number of players play repeatedly a stage game in which they choose how much to contribute to a public good. Individual contributions are observed at the end of each stage (perfect monitoring). A player's payoff in the stage game is a function of both the vector of contributions and the player's type, which she privately observes before the start of the game and remains fixed throughout.[2] Given any combination of types, free riding (i.e. zero contribution) is a strictly dominant strategy in the stage game. Payoffs of the overall game are represented by the discounted sum of expected stage game payoffs. These so called infinitely repeated games with incomplete information are classical Bayesian games: Nature picks a type profile according to a commonly known prior probability, and players play repeatedly the stage game selected by Nature.

We find that long run implementation is possible as long as there is little but strictly positive discounting: After a *screening* phase in which contributions may not be close to the desired level, any feasible and individually rational provision can be sustained. This asymptotic result holds even if priors over type profiles are private knowledge, provided the strategies proposed are best responses regardless of the beliefs players hold about other players' types. That is, they represent a *belief-invariant Nash equilibrium.*

With the exception of Chakrabarti (2003), who studies a Cournot-type oligopoly game with privately observed costs, the existing literature on repeated games with incomplete information has focused on two-player games. Aumann, Maschler, and Stearns (1968) introduce the analysis of games with no discounting. In their model one player has complete information while the other neither knows her type nor observes her payoffs after each stage. In this case, equilibrium may even fail to exist.[3] As for the discounted case, Cripps and Thomas (2003) study the "known own

[2] In particular, for any given type, individual payoffs depend positively on the overall sum of contributions and negatively on own contributions. This implies that players can punish a player by contributing zero regardless of her type. Thus, detected deviations are always punishable.

[3] See Hart (1985) for a characterization of the equilibrium payoff set. On the other hand,

payoffs" case (i.e. the uninformed player's payoffs are invariant to Nature's choice of types). They find that when players are patient enough and the probability of one of the types goes to one, the Folk Theorem holds.[4]

2 The Model

Consider a game $\Gamma(P, \delta)$ with a finite number of players $I \geq 2$, defined as follows. According to the nondegenerate probability distribution P, Nature picks a type profile $\theta = (\theta_1, \theta_2, ..., \theta_I)$ from the finite support of P, denoted by Θ^I, where Θ is the set of individual types. Any given player i only observes her own type $\theta_i \in \Theta$. After Nature's move, players enter the infinitely repeated game $\Gamma(\theta, \delta)$, in which payoffs of the stage game $G(\theta)$ are determined by the *realized* type profile θ. Players' actions are observed at the end of each stage.

$G(\theta) = \left\{A, \{U(\theta_i)\}_{i=1}^I\right\}$ is a public good game, in which $A = [0, c] \subset \mathbb{R}_+$ is the set of possible individual contributions, and $U(\theta_i) : A^I \longrightarrow \mathbb{R}$ is the payoff function, which is a function of both a player's own contribution (a_i), and the level of provision of the public good $(\sum\limits_{j=1}^{I} a_j)$:

$$U(a|\theta_i) = U(a_i, \sum_{j=1}^{I} a_j|\theta_i).$$

Assumption 1 *For any $\theta_i \in \Theta$, $U(.,.|\theta_i)$ is continuous, strictly decreasing in its first argument and strictly increasing in its second argument.*

Assumption 2 *(Incentive to free ride) For any $\theta_i \in \Theta$, $U(r, g + r|\theta_i) > U(r', g + r'|\theta_i)$ for any $g \geq 0$, and any $r' > r \geq 0$.*

The last assumption implies that, for any profile θ, contributing zero would be a strictly dominant strategy in the stage game $G(\theta)$.

Let $\underline{u}(\theta_i) = \min_{a_{-i} \in A^{I-1}} \max_{a_i \in A} U(a_i, a_{-i}|\theta_i)$ be player i's minmax payoff when her type is θ_i. We have that $\underline{u}(\theta_i) = U(0, 0|\theta_i)$ for any $\theta_i \in \Theta$, by *Assumptions 1* and *2*. Denote opponents' minmaxing strategy by $\underline{a}_{-i}(\theta_i)$. Note that $\underline{a}_{-i}(\theta_i) = (0, ..., 0)$ for all $\theta_i \in \Theta$, so $\Gamma(P, \delta)$ is a game with type-independent minmaxing. This property plays an important role in our construction. It allows players to punish a detected deviation regardless of the information they have about the deviator's type.

Neyman and Sorin (1998) find that the equilibrium set is non-empty when both players remain uninformed.

[4]This results does not hold with no discounting. Shalev (1994) shows that the set of equilibrium payoffs is strictly smaller than the set of feasible and individually rational payoffs.

A payoff vector $v = (v_1, ..., v_I)$ is individually rational (IR) under θ if $v_i \geq \underline{u}(\theta_i)$ $\forall i$. We will also say that a contribution vector is IR, meaning that its associated payoff vector is IR.

A pure strategy for player i in $\Gamma(\theta, \delta)$ is a sequence of maps $s_i = \{s_i^t\}_{t=1}^{\infty}$, $s_i^t : H^{t-1} \to A$, where H^t denotes the set of all possible histories of play $h_i^t = (\theta_i, a^1, a^2, ..., a^t)$.[5] Players discount future payoffs according to discount factors $\{\delta_i\}_{i=1}^{I}$, with $0 < \delta \leq \delta_i < 1$ for any i. $\{\delta_i\}_{i=1}^{I}$ are common knowledge. The payoff for player i when players follow strategy profile $s = (s_1, ..., s_I)$ and the realized type vector is θ is given by

$$v_i(s|\theta) = (1 - \delta_i) \sum_{t=1}^{\infty} \delta_i^{t-1} U(a^t|\theta_i).$$

3 Belief-Invariant Nash Equilibrium

Our characterization of $\Gamma(P, \delta)$'s payoff set applies to both PBE and Nash equilibrium (NE). In addition, we introduce the concept of *belief-invariant* Nash equilibrium (BINE), which we use in the long run implementation result stated in the next section.

A strategy profile s is a NE of $\Gamma(P, \delta)$ if

$$E_P [v_i(s|\theta)|\theta_i] \geq E_P [v_i(s_i', s_{-i}|\theta)|\theta_i] \;\; \forall s_i' \text{ and all } \theta_i \in \Theta, \;\; i = 1, 2, ..., I. \qquad (1)$$

Since P is common knowledge, a NE strategy profile only needs to satisfy (1) for the prior probability distribution P. However, strategies satisfying (1) for any well-defined probability distribution constitute a NE even if players have private and differing beliefs about the likelihood of each $\theta \in \Theta^I$. Thus, NE implemented with these *belief-invariant* strategies are robust to weakening standard Bayesian assumptions about players' knowledge of P.

Definition 1 *A strategy s_i is a belief-invariant best response to s_{-i} for type θ_i player i if $E_P [v_i(s_i, s_{-i}|\theta)|\theta_i] \geq E_P [v_i(s_i', s_{-i}|\theta)|\theta_i]$ for all s_i' and all $P \in \Delta(\Theta^I)$, where $\Delta(\Theta^I)$ is the space of probability distributions over type profiles.*[6]

[5] Since the set of contribution vectors is convex, we do not need to take into account mixed or correlated strategies of the stage game, given that for any such strategy profile we can always find a contribution vector that yields the same expected payoffs.

[6] *Belief-invariant best responses are logically related to the "belief free strategies" defined in Ely, Hörner, and Olszewski (2005). The latter refer to best responses that are independent of the opponent's private histories in two-player games with imperfect private monitoring. In such games, each player observes her own action and a private signal about the other player's action. Notice that in $\Gamma(P, \delta)$ actions are observable. Thus, the only private information is the player's type. Since belief-invariant best responses are independent of the other players' types, they are trivially independent of the "private histories" of the other players and could therefore be considered belief*

Definition 2 *A strategy profile* s *is a belief-invariant Nash equilibrium (BINE) of* $\Gamma(P,\delta)$ *if* $E_{P_i}[v_i(s|\theta)|\theta_i] \geq E_{P_i}[v_i(s_i', s_{-i}|\theta)|\theta_i]$ *for all* s_i'*, all* $P_i \in \Delta(\Theta^I)$ *and all* $\theta_i \in \Theta$*,* $i = 1, 2, ..., I$*.*

In a BINE, players on the equilibrium path are best responding to opponents' strategies given any well-defined probability distribution over opponents' types. In other words, BINE strategies are invariant to individual beliefs about players' types.

4 Long Run Implementation

Let $\mathcal{U}_\theta$ be the convex hull of the set of feasible payoff vectors of $G(\theta)$ that satisfy strict IR for all players. According to the Folk Theorem, any payoff vector in $\mathcal{U}_\theta$ can be supported as the average payoff vector in a subgame perfect equilibrium of $\Gamma(\theta,\delta)$, provided δ is high enough (Fudenberg and Maskin (1986)). However, in $\Gamma(P,\delta)$, the set of feasible and strictly IR payoffs depends on which type profile is selected by Nature. Accordingly, we define a mapping (*payoff target*), which uniquely assigns a payoff vector to any possible type profile, as the object of our implementation results.

Definition 3 *A payoff target for* $\Gamma(P,\delta)$ *is a mapping* $w : \Theta^I \to \mathbb{R}^I$ *such that* $w(\theta) \in \mathcal{U}_\theta$ *for any* $\theta \in \Theta^I$*.*

We show in *Proposition 1* that any payoff target can be supported in the long run by a BINE, provided $\Gamma(P,\delta)$ satisfies two mild assumptions. *Assumption 3* is a weak separability condition, which is implied in games where the standard single crossing property holds. It ensures that no two distinct types equally rank all contribution vectors belonging to any I-dimensional convex subset of A^I. Before stating this assumption, we define *utility equivalent* types in a set X of contribution vectors as types that share the same ranking over all contribution vectors belonging to X.[7] For any $\varepsilon > 0$, let $B(a;\varepsilon)$ be the ε -neighborhood of contribution vector a.

Definition 4 *If for a given set* $X \subseteq A^I$*,* $U(a|\theta_i) = b + dU(a|\theta_i')$ *with* $d > 0$ *for all* $a \in X$*, then types* θ_i'*,*θ_i *are utility equivalent in* X*.*

Assumption 3 *Given any contribution vector* $a \in interior(A^I)$ *and any scalar* $\varepsilon > 0$*, there are no two distinct types in* Θ *that are utility equivalent in* $B(a;\varepsilon)$*.*

---

free. Both equilibrium concepts rule out sequential equilibria in which players best responses depend on the particular beliefs about opponents private histories (either types in this setting or signals in the imperfect monitoring case).

[7] This is a local version of the definition provided by Abreu, Dutta, and Smith (1994).

That is, for any pair of types and a contribution vector a, we can always find two contribution vectors close to a such that the first type strictly prefers one while the second type prefers the other.

The following assumption implies the existence of a positive contribution vector that is strictly Pareto superior to zero provision under any possible type profile. This assumption is satisfied, for instance, if there is a strictly IR lump-sum contribution vector when nature selects all players to have the lowest WTC among all types in Θ.

Assumption 4 *There exists a contribution vector $a^* >> 0$ that is strictly IR for all players under any $\theta \in \Theta^I$.*

Proposition 1 states that, if the game satisfies *Assumptions 1-4* and δ is close to one, any payoff target w can be the continuation payoff in a BINE for any period after $T(\delta)$. That is, the continuation payoff vector at any period $t > T$ is $w(\theta)$, where θ is the realized type profile. That is, any feasible provision scheme can be sustained in the long run provided it is Pareto superior to zero provision.

Proposition 1 *Assume $\Gamma(P, \delta)$ satisfies Assumptions 1-4. Then, for any payoff target w, there is a discount factor $\underline{\delta} < 1$ and a finite period $T(\delta)$ such that, for all $\delta \in (\underline{\delta}, 1)$, there is a BINE of $\Gamma(P, \delta)$ with continuation payoffs at any $t > T$ equal to $w(\theta)$ for any $\theta \in \Theta$.*[8]

This result is based both on the existence of type-independent minmaxing (*Assumptions 1* and *2*), and on the introduction of rewards for truthful reporting at the beginning of the game (*Assumptions 3* and *4*). While the former enables players to punish observable deviations (contributions different from those on the equilibrium path), the latter prevents undetected deviations (i.e. untruthful reporting).

The strategies used to prove *Proposition 1* consist of three different phases of play on the equilibrium path, coupled with off-equilibrium minmax threats. In the first period, each player reveals her type (θ_i) through her initial contribution. This *communication* phase is followed by a *screening* phase in which a rewarding cycle is played repeatedly. Each cycle consists of I periods, one per player. In each period $i = 1, .., I$ of the cycle, players choose a strictly IR contribution vector $a(\theta_i; i)$, which depends solely on player i's report. This vector has the property that player i's payoff will be higher given player i's truthful revelation in the communication phase ($u(a(\theta_i; i)|\theta_i) > u(a(\theta_i'; i)|\theta_i)$ for all $\theta_i' \neq \theta_i$). By *Assumptions 3* and *4*, $a(\theta_i; i)$ exist for all $\theta_i \in \Theta$. Once this phase is over, players play any feasible and strictly IR contribution scheme (*implementation* phase). In the presence of positive discounting, players can deter untruthful reporting by choosing the appropriate length of the screening phase. If the screening phase is played long enough, the discounted value

[8] *The proof is in the Appendix.*

of rewards for telling the truth will outweigh any discounted potential gains from lying to be obtained during the implementation phase. In addition, by strict IR of $\{a(\theta_i; i)\}$ any detected deviation can be deterred with minmax threats for patient enough players (δ close to one).

The belief-invariant feature of these strategies comes from the fact that a player's incentives for truthful reporting are independent of other players' types (i.e. $a(\theta_i; i)$ does not depend on θ_j, for any $j \neq i$). Therefore, no matter what a player believes about Nature's choice of types, it is always a best response to report her true type. Once all players have reported their types there is no uncertainty about the sequence of contributions to be played.

Remark 1 *(Consistency off the equilibrium path) By Assumption 2 punishment threats are credible: not deviating when minmaxing a player is a belief-invariant best response. Therefore, the BINE strategies proposed show a higher consistency than that required in PBE.*[9]

Remark 2 *Proposition 1 does not hold in the undiscounted case as payoffs in the implementation phase are the only relevant factor for players with $\delta_i = 1$. Thus, truthful reporting may not be incentive compatible (IC) for certain payoff targets.*

We conclude by pointing out two important limitations of *Proposition 1*. First, average payoffs may not be close to the payoff target. Second, the length of the screening phase depends on the discount factor. Hence, it is natural to ask whether an approximate Folk Theorem is possible. The answer to this question is not straightforward: It depends on the particular parameters of the game. For instance, Cripps and Thomas (2003) show that when there are two players, one of them with only one possible type while the other has two types, an approximate folk theorem holds if the probability of one of second player's types is sufficiently close to zero. However, a proof of a folk theorem type of result is not available for the general case.

Appendix

Proof of *Proposition 1*. Consider the following strategy profile s:

On the equilibrium path:

A) *Communication period*: In the first period each player i chooses a contribution $m(\theta_i)$ so as to signal her type θ_i.[10] The vector of contributions in the first period is denoted by $m(\theta)$.

[9] *The notion of PBE requires that equilibrium strategies are best responses under beliefs consistent with Bayesian updating. This means that initial beliefs about types have to be derived from P. In this case, beliefs are irrelevant, so any belief system will support on- and off-equilibrium play.*

[10] Since there are infinitely many actions in A and finite types, it is possible to uniquely assign a different contribution level $m(\theta_i)$ to each type θ_i.

B) *Screening phase*: Players play cyclically (N times, where N is a function of δ) the sequence of I contribution vectors $(a(\theta;1), a(\theta;2), ..., a(\theta;I))$ where $\theta = (\theta_1, ..., \theta_I)$ are the reported types in the first period. Each vector $a(\theta;i)$ satisfies the following conditions:[11]

(i) *Individual Rationality (IR)*: $U(a(\theta;i)|\theta_j) \geq \underline{u}(\theta_j)\ \forall j$, with strict inequality when $i = j$.

(ii) *Incentive Compatibility (IC)*: $U(a(\theta;i)|\theta_i) > U(a(\theta_i', \theta_{-i};i)|\theta_i)$ for any $\theta_i' \neq \theta_i$, where $a(\theta_i', \theta_{-i};i)$ would be the contribution vector to be played if the reported types were $(\theta_1, ..., \theta_{i-1}, \theta_i', \theta_{i+1}, ..., \theta_I)$.

(iii) θ_{-i}*-Independence*: $a(\theta;i)$ is independent of the report of any player $j \neq i$. That is, $a(\theta;i) = a(\theta_i;i)$ for all $\theta_i \in \Theta$.

C) *Implementation phase*: If no deviation has occurred, a contribution vector yielding $w(\theta)$ is played.

Off the equilibrium path:

A) If a player i does not send a meaningful signal (i.e. with positive probability on the equilibrium path), players do not contribute from $t = 2$ on.

B) If a player deviates at the screening or the implementation phase, players do not contribute from the period after the deviation on.

By *Assumption* 4, we know that there exists a vector a^* strictly IR for any report θ. Since $U(.,.|.)$ is continuous, there exists an $\varepsilon > 0$ such that any vector $a \in B(\varepsilon, a^*)$ is also strictly IR.

By *Assumption 3*, for any pair of distinct types $\theta_i, \theta_i' \in \Theta$ we can find two contribution vectors $\mathsf{a}(\theta_i, \theta_i'), \mathsf{a}(\theta_i', \theta_i) \in B(\varepsilon, a^*)$ such that $U(\mathsf{a}(\theta_i, \theta_i')|\theta_i) > U(\mathsf{a}(\theta_i', \theta_i)|\theta_i)$ and $U(\mathsf{a}(\theta_i', \theta_i)|\theta_i') > U(\mathsf{a}(\theta_i, \theta_i')|\theta_i')$.[12]

For each type θ_i we select a set $A^*(\theta_i) = \{\mathsf{a}(\theta_i, \theta_i')\}_{\theta_i' \neq \theta_i}$ of $\|\Theta\| - 1$ contribution vectors satisfying these inequalities, one for each element of $\Theta \setminus \{\theta_i\}$. We then order the $\|\Theta\| (\|\Theta\| - 1)$ vectors in $A_i^* = \cup_{\theta_i \in \Theta} A^*(\theta_i)$ in ascending order with respect to type-θ_i utility, breaking ties arbitrarily, and assign them strictly increasing weights adding up to one. Finally, we define $a(\theta;i)$ as the resulting convex combination of the vectors in A_i^*.[13] By defining $a(\theta;i)$ in this way we ensure that it satisfies (i)-(iii). Note that, by (iii), player i's report cannot affect the selection of $a(\theta;j)$ for $j \neq i$.

[11] Hence, the length of the screening phase is $N \times I$.

[12] This property is similar to the *payoff asymmetry* property stated in Abreu, Dutta, and Smith (1994). The main difference is that we deal with a continuous action space, whereas in their model the action space is discrete.

[13] This construction is similar to the one used by Abreu, Dutta, and Smith (1994) in their proof of *Lemma 2*.

Thus, strategy profile s exists. Now we need to check that it is a BINE. After observing θ_i, player i's payoff at $t = 1$ on the equilibrium path is given by $E_P[v_i(s|\theta)|\theta_i]$, with

$$v_i(s|\theta) = (1 - \delta_i)U(m(\theta)|\theta_i) + \delta_i \pi_i(\theta) + \delta_i^T w_i(\theta).$$

In this expression, $T = N \times I + 1$ and $\pi_i(\theta)$ represents the payoff obtained during the screening phase:

$$\begin{aligned} \pi_i(\theta) &= (1 - \delta_i) \sum_{j=1}^{I} \sum_{s=0}^{N} \delta_i^{(sI+j-1)} U(a(\theta; j)|\theta_i) \\ &= (1 - \delta_i) \frac{1 - \delta_i^{IN}}{1 - \delta_i^{I}} \sum_{j=1}^{I} \delta_i^{(j-1)} U(a(\theta; j)|\theta_i). \end{aligned} \quad (2)$$

It suffices to show that, for all i, , all $\theta \in \Theta$, and all strategies s'_i we have that $v_i(s(\theta)|\theta) \geq v_i(s'_i, s_{-i}|\theta)$. In such a case, $E_P[v_i(s|\theta)|\theta_i] \geq E_p[v_i(s'_i, s_{-i}|\theta_i)|\theta_i]$ for all $P \in \Delta(\Theta^I)$, which implies that s is a BINE.

We have two possible scenarios in which a player might deviate:

1) *Undetected deviation*: Player i reports a type θ'_i different from the true one. In this scenario, her payoff satisfies the following inequality for any realized type profile $\theta \in \Theta$:

$$v_i(s_i^1, s_{-i}|\theta) \leq (1 - \delta_i) \max_{a \in A^I} U(a|\theta_i) + \delta_i \pi'_i(\theta) + \delta_i^T \max_{a \in A^I} U(a|\theta_i),$$

where $\pi'_i(\theta)$ corresponds to the same expression as (4), except in that it uses $a(\theta'_i, \theta_{-i}; i)$ instead of $a(\theta; i)$. By *IC* and θ_{-i}-*Independence* of vectors $\{a(\theta; i)\}$, we have that $\pi'_i(\theta) < \pi_i(\theta)$.

2) *Detected deviation*: Player i deviates in any other circumstance. If she deviates in period $t \geq 1$, her continuation payoffs will be:

$$v_{i,t}(s_i^2, s_{-i}|\theta) \leq (1 - \delta_i) \max_{a \in A^I} U(a|\theta_i) + \delta_i \underline{u}(\theta_i).$$

The upper bound on payoffs is higher for undetected deviations as long as $\pi'_i(\theta) \geq \underline{u}(\theta_i)$. Assume this is true without loss of generality, so that we need only prove that reporting a false type is not a best response.

Assume that there is a type $\theta'_i \neq \theta_i$ for player i such that, by not reporting truthfully, she gets

$$v_i(s_i^1, s_{-i}|\theta) = (1 - \delta_i) \max_{a \in A^I} U(a|\theta_i) + \delta_i \max_{a(\theta'_i, \theta_{-i}; i)} \pi'_i(\theta) + \delta_i^T \max_{a \in A^I} U(a|\theta_i). \quad (3)$$

If we show that there is a discount factor $\underline{\delta}_i < 1$ such that, for any $\delta_i > \underline{\delta}_i$, there is a finite $T(\delta_i) = N(\delta_i) \times I + 1$ such that $v_i(s|\theta) \geq v_i(s_i^1, s_{-i}|\theta)$, then it would not be a best response for player i to report a false type.

First, note that $v_i(s|\theta) > (1 - \delta_i) \min_{a \in A^I} U(a|\theta_i) + \delta_i \pi_i(\theta) + \delta_i^T \underline{u}(\theta_i)$.[14] Subtracting

[14] The inequality is strict since $w_i(\theta) > \underline{u}(\theta_i)$.

(5) from this inequality we get

$$v_i(s|\theta) - v_i(s_i^1, s_{-i}|\theta) > (1-\delta_i)\left[\min_{a\in A^I} U(a|\theta_i) - \max_{a\in A^I} U(a|\theta_i)\right] +$$
$$+\delta_i\left[\pi_i(\theta) - \max_{a(\theta_i',\theta_{-i};i)} \pi_i'(\theta)\right] + \delta_i^T\left[\underline{u}(\theta_i) - \max_{a\in A^I} U(a|\theta_i)\right]. \tag{4}$$

Further,

$$\pi_i(\theta) - \max \pi_i'(\theta) \geq (1-\delta_i)\frac{1-\delta_i^{IN}}{1-\delta_i^I}\delta_i^{(I-1)}\left[U(a(\theta;i)|\theta_i) - \max_{\theta_i'\neq\theta_i} U(a(\theta_i', \theta_{-i}; i)|\theta_i)\right].$$

Since $U(a(\theta;i)|\theta_i) > U(a(\theta_i', \theta_{-i}; i)|\theta_i)\ \forall\theta_i' \neq \theta_i$ by IC, the right-hand side of the above inequality is strictly greater than zero for $\delta_i < 1$ and N finite. Thus, the first and the third terms on the right hand side of (6) are strictly negative, whereas the second is strictly positive.

As $\delta_i \uparrow 1$, we have that: $(1-\delta_i) \to 0$; $(1-\delta_i^{IN}) \to 0$ at a slower rate than $(1-\delta_i)$ if $I \times N > 1$; $\frac{(1-\delta_i)}{1-\delta_i^I}\delta_i^I \to 1/I$ and; $\delta_i^T \to 1$.

In addition, for $\delta_i < 1$, as $N \to \infty$ the following is true: $(1-\delta_i^{IN}) \to 1$ and; $\delta_i^T = \delta_i^{NI+1} \to 0$.

Therefore, for a big enough discount factor δ_i (strictly smaller than one) there is a finite $N(\delta_i)$ such that the right hand side of (6) is nonnegative. This implies $v_i(s|\theta) - v_i(s_i^1, s_{-i}|\theta) > 0$.

In order to determine $\underline{\delta}$, pick the lowest discount factor so that

$$\frac{(1-\delta_i)\delta_i^I}{1-\delta_i^I}\left[U(a(\theta;i)|\theta_i) - \max_{\theta_i'\neq\theta_i} U(a^i(\theta_i', \theta_{-i})|\theta_i)\right] >$$
$$> (1-\delta_i)\left[\max_{a\in A^I} U(a|\theta_i) - \min_{a\in A^I} U(a|\theta_i)\right]$$

is satisfied for any type profile θ, $i = 1, ..., I$. Next, for any game $\Gamma(P, \delta)$ with discount factors $(\delta_1, ..., \delta_I)$ such that $\underline{\delta} < \delta_i < 1\ \forall i$, choose $N(\delta_i)$ so that (6) is satisfied $\forall\theta$. Finally, pick the highest $N(\delta_i)$ of all players.[15] ■

References

Abreu, D., P. J. Dutta, and L. Smith (1994): "The Folk Theorem for Repeated Games: A Neu Condition," *Econometrica*, 62, 939–948.

[15]If inequality (6) holds for $N(\delta_i)$, it also holds for any $N > N(\delta_i)$.

AUMANN, R., M. MASCHLER, AND R. STEARNS (1968): "Repeated Games of Incomplete Information: An Approach to the Non-Zero-Sum Case," *Mathematica*, ST-116(III), 287–403.

AUMANN, R., AND L. SHAPLEY (1976): "Long Term Competition: A Game Theoretic Analysis," Mimeo. Hebrew University.

CHAKRABARTI, S. (2003): "Information Revelation and Collusion in Oligopolies with Unknown Costs," Mimeo. Indiana University-Purdue University Indianapolis.

CRIPPS, M., AND J. THOMAS (2003): "Some Asymptotic Results in Discounted Repeated Games of One-Sided Incomplete Information," *Mathematics of Operations Research*, 28, 433–462.

ELY, J. C., J. HÖRNER, AND W. OLSZEWSKI (2005): "Belief-Free Equilibria in Repeated Games," *Econometrica*, 73, 377–415.

FUDENBERG, D., D. LEVINE, AND E. MASKIN (1994): "The Folk Theorem with Imperfect Public Information," *Econometrica*, 62, 997–1039.

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

HART, S. (1985): "Nonzero-Sum Two-Person Repeated Games with Incomplete Information," *Mathematics of Operation Research*, 10, 117–153.

NEYMAN, A., AND S. SORIN (1998): "Equilibria in Repeated Games of Incomplete Information: The General Symmetric Case," *International Journal of Game Theory*, 27, 201–210.

RUBINSTEIN, A. (1977): "Equilibrium in Supergames," Center for Mathematical Economics and Game Theory. Hebrew University of Jerusalem.

——— (1979): "Equilibrium in Supergames with the Overtaking Criterion," *Journal of Economic Theory*, 21, 1–9.

SHALEV, J. (1994): "Nonzero-Sum Two-Person Repeated Games with Incomplete Information and Known-Own Payoffs," *Games and Economic Behavior*, 7, 246–259.