

**On Uniquely Implementing Cooperation in the Prisoner's Dilemma:†
Corrigendum**

Bhaskar Chakravorti*

John P. Conley**

and

Bart Taub***

† **Published as** Bhaskar Chakravorti, John Conley and Bart Taub, "On Uniquely Implementing Cooperation in the Prisoner's Dilemma" *Economic Theory*, Vol. 80, 1996, pp. 347-66. The authors wish to thank Dilip Abreu, Robert Aumann, Michael Baye, James Friedman, Richard McLean, Herve Moulin, Ariel Rubinstein, Rajiv Vohra and Simon Wilkie for their comments. Also, participants of seminars and conferences at Arizona, Bellcore, Brown, Illinois, Northwestern, Princeton, Rochester, Vanderbilt and West Virginia have asked stimulating questions. We also thank the referee for many detailed and useful suggestions. The third author's research is currently supported by NSF Grant SES-9213145.

* The Monitor Corporation

** Department of Economics, UIUC, j.p.conley@vanderbilt.edu

*** Department of Economics, UIUC

Abstract

We consider the problem of a principle who wishes to induce two agents playing a one shot prisoner's dilemma to behave cooperatively. We assume that the principle cannot observe the actions of the agents, and is not able to change the strategy sets or payoff functions in the underlying game. The only power the principle has is to randomly delay the arrival of payoffs. Specifically, agents choose their one shot strategies, and then the principle randomly determines whether these are "cheap talk", or if payoffs should be distributed. If the round is cheap talk, then each agent observes the strategy choice of the other and play moves to a new round. This continues until payoffs are distributed. We establish conditions under which the probability of cheap talk can be chosen at the beginning of the induced game in such a way that full cooperation is the only equilibrium outcome. The sufficiency condition is met by a wide class of economic interpretations of the prisoners' dilemma, including those involving strategic complementarities among players.

Keywords: Probabilistic cheap talk, implementation, preference revelation, prisoner's dilemma, folk theorem .

JEL: C72, D43, D62.

1. Introduction

We investigate a class of games which lie on the boundary between repeated games and games with cheap talk. Imagine, for example, that your department is considering the hiring of a new assistant professor. The chairman comes into your office and asks you to serve on the hiring committee. You must assent to or refuse this request without knowing who else has, or will, agree to serve on the committee. Thus, commitments to serve are simultaneous and secret. It is not a certainty, however, that the dean will approve funding for the new position this year. There is only a probability δ that you will actually end up spending December reading the CV's of the candidates. If the dean does not approve funding then the commitment to serve is ex post cheap talk, and it becomes common knowledge which colleagues had agreed to serve. The game is repeated in subsequent years until a new colleague is hired. Cooperative effort results in a good hire that confers external benefits on all the department's faculty. But with skiing in the Alps as the alternative, there is a temptation to free ride. In short, this situation amounts to a prisoner's dilemma in which there is random delay of payoffs. Note that, ex ante, agents' moves have positive expected value. They are ex post cheap talk only in the event that the dean does not approve the position. For this reason we call this a game with probabilistic cheap talk (*PCT*).

In this paper we are primarily interested in problems in which there is a third party, a principal (like the dean), who wishes to induce the agents to cooperate with each other. The *PCT* mechanism we develop complements the standard implementation literature. The implementation tradition considers a designer, mediator or principal who is uninformed about the preferences and other private information known to the agents, but who has complete control over the institutions and procedures embodying the mechanism. In particular, the designer is free to choose the strategy sets and payoff functions of the agents. Such freedom has given rise to a large number of mechanisms, many of which have unnatural strategy spaces and extremely complex payoff rules. In contrast, our approach presumes that the designer has no control over existing economic institutions except for the ability to delay payoffs. In addition, we assume

that it is costly for the principal to monitor the actions of agents directly. Thus, he can neither offer direct rewards to them in return for playing cooperatively, nor choose a probability of ending that is contingent on observed cooperation. On the other hand, we allow our designer to possess some minimal information about the structure of the game. For example, the designer might know the aggregate benefit (but not the individual benefits) of a public project, but may be forced to use the basic voluntary contribution mechanism. We show that within these constraints, if the designer can randomly delay the arrival of payoffs, then he can implement unique outcomes.

A key assumption that we maintain is that actions can be costlessly withdrawn and revised if no payoffs result. For example, if a group wishes to construct a public project, individuals might deposit their contributions into accounts which are held in escrow (possibly by each individual's lawyer) until the construction of the public good is approved by a designated principal. If the principal does not approve of the project, then the deposits are returned and everybody observes the amounts that were deposited. In case of approval, the funds held in escrow are released to be used towards construction of the good. If the funds are sufficient, then any individuals who failed to contribute end up free riding; if the funds are insufficient, then the project is not built and the game ends.¹ Two key aspects of this example are: (i) Depositing funds in the escrow account is costless in the case that the current round is not the terminal one (moreover, the accumulation of interest compensates for real-time discounting by the agents); and (ii) in any given round, the deposits are made simultaneously, and, therefore, the commitment of funds by an agent cannot be made contingent upon the amount committed by any other agent. The commitment, of course, lasts for only one round at a time.

We focus here on a model of a generic symmetric prisoner's dilemma (*PD*) game as a stylized representation of these games. We assume that the principal and agents

¹ We use this escrow game to motivate the idea of costless withdrawal. Strictly speaking, however, the game differs from a prisoner's dilemma in that it is not a dominant strategy to deviate from contributing. We investigate a similar public funding game that does have the structure of a prisoner's dilemma in Chakravorti, Conley and Taub [5].

are patient and do not discount the future. This could be because the process occurs over a sufficiently short period of time or because until approval, economic resources are being held in escrow, so that they accumulate interest over time and neutralize any discounting. Uncertainty regarding the stopping time then transforms the standard *PD* game into one with the essential mathematical structure of a *PD* supergame with discounting. The probability of payoff-relevance provides the counterpart for the discount factor in the supergame model, except that in our context it can be chosen by the principal as opposed to being exogenously determined. The folk theorem for supergames then extends to our situation; it ensures that full cooperation can be supported as an equilibrium of the induced game for some probability of payoff-relevance. As is well known, this result comes with a price: any other individually rational outcome can also be supported if the probability is appropriately low. On the other hand cooperation is impossible if the probability is too high. Thus, the challenge confronting the principal is as follows. Is there an intermediate probability for which the full cooperation outcome is the only plausible outcome?

We answer the question in the affirmative by considering the set of subgame-perfect equilibrium payoffs of the induced *PCT* game satisfying two additional requirements: (i) Pareto dominance in the set of equilibrium payoffs; and (ii) stationarity in the equilibrium path. Requirement (i) is widely accepted as being reasonable.² In our case, it is unavoidable since, as the correspondence with supergame equilibria requires, the non-cooperative outcome remains an equilibrium no matter what the approval probability is. Requirement (ii), we argue, also has heuristic appeal in the context of games with *PCT*. It states that provided no defections from a proposed equilibrium have occurred, each player chooses the same mixed action in every period. In the game induced by our mechanism, if we compare a period t with some other period t' , not only does the future structure of the game appear identical, but the past is also exactly the same, since ex post the history is cheap talk. We appeal to such an invariance property

² There are, of course, situations in which this requirement is less appealing. When there are three or more players, coalition-proof equilibria do not necessarily Pareto dominate the set of equilibria.

to motivate the requirement of a limited notion of stationarity.³ In standard repeated games, by contrast, payoffs are realized in every period. These games are sometimes abstractions of games in which wealth accumulates and consumption occurs in the future. The game in period t differs from that in t' in the sense that the players' wealth has changed. In such games, the notion of stationarity may not be as compelling.

Our main result is as follows. We find that an upper bound on the marginal cost of cooperating when the opponent switches from cooperation to a non-cooperative strategy provides a sufficient condition for unique implementation of full cooperation. This condition is easily satisfied in most natural interpretations of the prisoner's dilemma. These include ones in which players' actions have a strategic complementarity property. The property translates as follows: a player's marginal cost of cooperating should decrease as the other player switches to playing cooperatively.

In the Appendix, we show that the assumption of limited stationarity can be relaxed for the text-book knife-edge case of the *PD* discussed in Section 6. Provided the players choose equilibrium strategies with an arbitrarily large but finite number of Markov states (limited stationarity corresponding to a single such state), interpreted as strategies with a bound on the degree of complexity, the unique implementation property is retained and is identical regardless of the size of the complexity bound. A corollary that is implicit but which we do not prove is that the surfeit of equilibria engendered by the folk theorem is entirely the consequence of admitting nonstationary strategies, strategies that are inadmissible under any notion of bounded rationality.

We proceed as follows. Section 2 lays out the model. Section 3 shows how *PCT* achieves efficiency in the *PD*. Sections 4-6 derive the unique implementation results. The final section concludes. The Appendix contains an extension of the results, by relaxing the assumption of limited stationarity.

³ Such an argument for limited stationarity has been made in other contexts, where similar problems have been studied. Rubinstein and Wolinsky [10] make a case for semi-stationarity in a model of pairwise bargaining. Green and Laffont [7] introduce a notion of posterior implementable equilibria in a model of cheap talk; the role of stationarity in this context is analyzed further in Chakravorti and Conley [5]. Also see Fershtman and Kalai [6] for a defense of stationarity as a realistic assumption in the context of repeated delegation.

2. The Model

We consider the canonical symmetric prisoners' dilemma given by the bimatrix game in Figure 1. Any prisoners' dilemma game with strictly dominated strategies can be transformed into a game of the form given below by subtracting the (b, R) payoffs from all cells and normalizing the marginal cost of choosing the dominated strategy to unity, given that the opponent chooses the dominant strategy.

The players are “Small” (the row player) and “Big” (the column player) whose deterministic choices are $\{above(a), below(b)\}$ and $\{LEFT(L), RIGHT(R)\}$ respectively. The payoffs are in terms of von-Neumann Morgenstern utilities and the class of PD games is characterized by $\{(A, D) \in \mathfrak{R}^2 : 0 < A < D; 2A > D - 1\}$. Also, p and P denote, respectively, the probabilities with which Small chooses a and Big chooses L .

An allocation is a pair of payoffs (x, X) residing in \mathcal{X} , the set of all allocations achievable via mixed strategies. The set of all strictly individually rational allocations in \mathcal{X} is denoted $IR(\mathcal{X}) = \{(x, X) \in \mathcal{X} : (x, X) \gg (0, 0)\}$. The Pareto-efficient frontier of \mathcal{X} is denoted $EFF(\mathcal{X})^4$

It is well known that the unique Nash equilibrium of the PD is $(p, P) = (0, 0)$. Hence, conventionally, $\{a, L\}$ are denoted the “(fully) cooperative” strategies and $\{b, R\}$ are the “non-cooperative” strategies. If the allocation were on $EFF(\mathcal{X})$, then, in general, there is “partial cooperation”, in the sense that at least one player fully cooperates.

We assume that the players' actions result in outputs that accrue to a principal. The payoffs in the PD game are proportional to the profits the principal would make if he approves the output and put it on the market. The actions of the players and the eventual profitability of the resulting output are unobservable to the principal. The objective of the latter is to create an environment where profits are maximized, i.e. the fully cooperative strategies are played. We shall explore a mechanism for implementing this outcome.

⁴ Since the game is a PD , $\mathcal{X} = \text{con}\{(0, 0), (A, A), (-1, D), (D, -1)\}$, i.e. any element of the convex hull of the stage game payoff vectors is attainable through mixed strategies. Another useful consequence of the PD structure is a one-to-one correspondence between the set of mixed strategy vectors and \mathcal{X} .

The mechanism proposed herein is based on the principal's control over the payoff-relevance of actions, which depends on whether or not he has approved the output. We refer to this mechanism as a probabilistic cheap talk (*PCT*) mechanism, and it is played as follows:

The players choose (p, P) , and the principal chooses the probability with which he collects the resulting output. Thus, the principal possesses a randomization device with two realizations: "Cheap Talk" (CT) and "Deadline" (DL). If DL is realized, the output is approved and the *PD* payoffs are generated for the chosen strategies. If CT is realized, then the payoff distribution is put on hold; the chosen strategies are treated as ex post cheap talk; and the process is repeated. The players "observe" all ex post cheap talk. Mixed strategies are assumed observable, i.e. players can distinguish between (p, P) and (p', P') provided they have been played in earlier periods and $(p, P) \neq (p', P')$.⁵

Denote the number of rounds of ex post cheap talk by $t \in \{0, 1, 2, \dots\}$. The history of talk at $t \in \{0, 1, 2, \dots\}$ is denoted h_t , with $h_0 = \emptyset$. Denote by \mathcal{H} the set of all possible histories over all $t \in \{0, 1, 2, \dots\}$. \mathcal{H} is the space of all possible histories at t . A strategy for Small is a profile $\pi = (\pi_t : \mathcal{H} \rightarrow [0, 1])_{t=0}^{\infty}$ and a corresponding strategy for Big is $\Pi = (\Pi_t : \mathcal{H} \rightarrow [0, 1])_{t=0}^{\infty}$. A pair $[\pi_t(h_t), \Pi_t(h_t)]$ specifies the probabilities with which the players choose the cooperative actions in the round t , given the history h_t .

If for some $t' < t$ the players have chosen (p, P) in round t' , we shall say that the resulting history h_t contains (p, P) played at t' . A stationary history is a history $h \in \mathcal{H}$ that satisfies:

$\exists(p, P)$ such that $\forall t \in \{1, 2, \dots\}, \forall t' < t$, it holds that

h_t contains (p, P) played at t' .

Let \mathcal{H}^α denote the sub-class of stationary histories. Correspondingly, a strategy pair,

⁵ The assumption that players can observe mixed strategies has been made to make the exposition easier. Our aim in this paper is to maintain the focus on probabilistic cheap talk; hence, we borrow the most primitive, and simplest, structure from the repeated games literature. The extension to problems where mixed strategies are unobservable can be done by applying the results in the repeated games literature on folk theorems (see Fudenberg and Maskin) that do not use this assumption.

$[p, P]$, satisfies limited stationarity if $[\pi_t(h_t), \Pi_t(h_t)] = [\pi'_t(h'_t), \Pi'_t(h'_t)]$ for all $t' < t$, for all $t \in \{1, 2, \dots\}$, for all $h \in \mathcal{H}^\alpha$.

A *probabilistic cheap talk (PCT)* rule is given by $\delta \in (0, 1]$, where δ is the probability that DL is realized in any particular round conditional on that round being reached. The game induced by a *PCT* rule is denoted $\Gamma(PD; \delta)$. Let $e(\pi, \Pi; \delta, h)$ and $E(\pi, \Pi; \delta, h)$ denote the expected payoffs in \mathcal{X} to Small and Big evaluated at round t , conditional on round t being reached, given the strategies (π, Π) , the *PCT* rule δ and history h . A pair (π, Π) constitutes a subgame-perfect equilibrium in the game $\Gamma(PD; \delta)$ if at every $t \in \{0, 1, 2, \dots\}$, for all (π', Π') , for all $h \in \mathcal{H}$

$$e_t(\pi, \Pi; \delta, h) > e_t(\pi', \Pi; \delta, h)$$

$$E_t(\pi, \Pi; \delta, h) > E_t(\pi, \Pi'; \delta, h).$$

Define

$SPE[\Gamma(PD; \delta)] \equiv \{(x, X) \in \mathcal{X} \exists (\pi, \Pi), h \in \mathcal{H} \text{ s.t. } \forall t \in \{0, 1, 2, \dots\}, \text{ it holds that}$

$[e_t(\pi, \Pi; \delta, h), E_t(\pi, \Pi; \delta, h)] = (x, X), (\pi, \Pi)$ is a subgame-perfect equilibrium of

$$\Gamma(PD; \delta) \text{ and } h = ((\pi_t(h_t), \Pi_t(h_t)))_{t=0}^\infty\}.$$

In words, $SPE[\Gamma(PD; \delta)]$ denotes the set of allocations realized in subgame-perfect equilibria of $\Gamma(PD; \delta)$ regardless of the history of ex post cheap talk.

Remark: The concept given above is actually a refinement of the set of subgame-perfect allocations of $\Gamma(PD; \delta)$ evaluated at $t = 0$. This is ordinarily defined as: $[SPE[\Gamma(PD; \delta)] = \{(x, X) \in \mathcal{X}: \text{there exists } (\pi, \Pi), h \in \mathcal{H} \text{ such that } [e_0(\pi, \Pi; \delta, h), E_0(\pi, \Pi; \delta, h)] = (x, X) \text{ and } (\pi, \Pi) \text{ is a subgame-perfect equilibrium of } \Gamma(PD; \delta)\}$. Given that the players are interested in implementing an allocation, say (x, X) , at any given round, they would like to be assured that the expected payoff from their strategies remains unchanged regardless of the cheap talk that has preceded the current round. Hence, our focus on $SPE[\Gamma(PD; \delta)]$. Note that the strategy pairs that realize payoffs in $SPE[\Gamma(PD; \delta)]$ must satisfy limited stationarity.

3. Efficiency via the PCT Mechanism

In this section, we show that every feasible and individually rational allocation can be attained for some δ as an equilibrium allocation of the mechanism. It illustrates the basic intuition underlying the mechanism, which is a simple application of the Folk Theorem to *PCT* games. This will, in turn, motivate our subsequent objective of resolving the multiplicity problem.

Theorem 1. *For all $(x, X) \in IR(\mathcal{X})$, there exists $\delta \in (0, 1]$ such that $(x, X) \in SPE[\Gamma(PD; \delta)]$.*

Proof/

We shall show that δ can be chosen so that if $(x, X) \in IR(\mathcal{X})$, then $(\hat{x}, \hat{X}) \in SPE[\Gamma(PD; \delta)]$. Choose $(\hat{x}, \hat{X}) \in IR(\mathcal{X})$ and let (\hat{p}, \hat{P}) be the associated stage-game strategies, i.e.

$$\hat{x} = \hat{p}\hat{P}A + \hat{P}(1 - \hat{p})D - (1 - \hat{P})\hat{p}$$

and,

$$\hat{X} = \hat{p}\hat{P}A + \hat{p}(1 - \hat{P})D - (1 - \hat{p})\hat{P}.$$

Choose δ such that

$$\delta < \min \left\{ \frac{\hat{x}}{\hat{P}D}, \frac{\hat{X}}{\hat{p}D} \right\}.$$

We claim that a subgame-perfect equilibrium strategy is $(\hat{\pi}, \hat{\Pi})$ defined by:

$$(\hat{\pi}_t(h_t), \hat{\Pi}_t(h_t)) = \begin{cases} (\hat{p}, \hat{P}) & \text{if } h_t \text{ contains } ((\hat{p}, \hat{P}) \forall t' < t \text{ or if } t = 0 \\ (0, 0) & \text{otherwise} \end{cases}$$

for all $t \in \{0, 1, 2, \dots\}$.

The expected payoff to Small from (\hat{p}, \hat{P}) in the on-the-path subgame beginning at t is

$$\delta\hat{x} + [1 - \delta]\delta\hat{x} + [1 - \delta]^2\delta\hat{x} + \dots = \frac{\delta\hat{x}}{1 - [1 - \delta]} = \hat{x}$$

Also, observe that any deviation from (π, Π) triggers $(0, 0)$ thereafter. Thus, the maximal expected payoff to Small from a unilateral deviation from (\hat{p}, \hat{P}) at any t is $\delta \hat{P}D$. By construction $\delta \hat{P}D < x$. Analogously, a similar argument for Big can be given. Hence $(\hat{x}, \hat{X}) \in SPE[\Gamma(PD; \delta)]$.

■

The mechanism, through the choice of δ , achieves efficient allocations as self-enforcing outcomes. Since δ is history-independent, the choice requires no third-party observation of strategies. However, as is the case with any Folk Theorem-type result, we have an embarrassment of riches – there are too many equilibria to coordinate on. Ideally, the principal would like to single out the fully cooperative outcome as the unique plausible equilibrium. Since δ is the only variable that the principal controls, a natural question for him to ask would be: is there a value of δ for which the set $SPE[\Gamma(PD; \delta)]$ is large enough to include the full cooperation outcome and yet small enough to make all other outcomes implausible? Given that the non-cooperative outcome will always remain an equilibrium for any value of δ , a minimum requirement of plausibility would be Pareto non-dominance within the set of equilibria. Hence, the question reduces to: is there a value of δ such that $SPE[\Gamma(PD; \delta)]$ has a unique Pareto-efficient point which Pareto dominates all other points in $SPE[\Gamma(PD; \delta)]$? Geometrically, this requires that the set $SPE[\Gamma(PD; \delta)]$ be “boxed in” within a square whose principal diagonal is the ray connecting $(0, 0)$ to (A, A) . We refer to this property as unique implementation of full cooperation by a mechanism δ . This occurs when there exists $\delta \in (0, 1]$ such that (i) $SPE(\Gamma(PD; \delta)) \cap EFF(\mathcal{X}) = \{(A, A)\}$ and (ii) $(A, A) > (x, X)$ for all $(x, X) \in SPE[\Gamma(PD; \delta)]$.

Given the resemblance between $1 - \delta$ and the discount factor in a repeated game model, one would hope that the vast literature on the latter subject would give us an answer. The most systematic study of the graph of the equilibrium set as a function of the discount factor is by Stahl [12]. Using the methods of Abreu, Pearce and Stacchetti [2], he finds that the set of correlated equilibria of symmetric PD games is monotonically non-decreasing and upper-hemicontinuous in the discount factor. How-

ever, in general, for no value of the discount factor can it be “boxed in” inside a square. The most striking case in point is the class of *PD* games studied by van Damme [13] (who extends Sorin [10]) using the Nash equilibrium concept. For low values of the discount factor, only the $(0, 0)$ outcome is achieved; but as the factor increases beyond a threshold level every allocation in $\text{IR}(X)$ is an equilibrium. Such a discontinuous blowing up phenomenon would appear to clearly rule out the possibility of boxing in the equilibrium set for any particular value of δ .

However, all is not lost. For our problem, the set of equilibrium allocations is defined by $\text{SPE}[\Gamma(\text{PD}; \delta)]$, each element of which has a corresponding strategy pair which has the limited stationarity property. Hence, the results cited above do not quite address the issue in our context. In the following sections, we shall explore conditions under which δ may be chosen so that the crucial unique implementation property is met despite the negative indications of the repeated games literature.

4. Achieving a Unique Efficient Equilibrium

In this section, we shall show that for any *PD* game the principal can find a value of δ such that $\text{SPE}[\Gamma(\text{PD}; \delta)]$ has a unique Pareto-efficient allocation, namely, the full cooperation outcome. In subsequent sections, we establish conditions under which this equilibrium Pareto-dominates any other equilibrium.

We begin by characterizing the set $\text{SPE}[\Gamma(\text{PD}; \delta)]$ for any value of δ . Define Small’s *acceptable set* by

$$\mathcal{X}_s(\delta) \equiv \{(x, X) \in X \mid \exists(p, P) \in [0, 1] \times [0, 1]\}$$

such that

$$x = pPA + P(1 - p)D - (1 - P)p > \delta PD; X = pPA + p(1 - P)D - (1 - p)P\}.$$

This is the largest set of allocations such that for each element there is a simple trigger strategy for Big such that this allocation would be realized when Small plays a best response. The left-hand side of the inequality in the definition of $\mathcal{X}_s(\delta)$ above is $e_t(\pi, \Pi; \delta, h)$. The right-hand side is Small's maximal payoff in round t from defecting from p , assuming that a grim trigger follows the defection, i.e. both players choose to play non-cooperatively thereafter.⁶

The equality that defines the boundary of $\mathcal{X}_s(\delta)$ in the interior of \mathcal{X} , denoted $\bar{\mathcal{X}}_s(\delta)$, is given by,

$$p(P(A - D + 1) - 1) = PD(\delta - 1)$$

which can be re-arranged to yield a parametric expression in (p, P) :

$$p = \frac{PD(\delta - 1)}{P(A - D + 1) - 1}. \quad (1)$$

The corresponding acceptable set for Big is $\mathcal{X}_B(\delta)$ and the expression for its boundary, denoted $\bar{\mathcal{X}}_B(\delta)$ is,

$$p = \frac{pD(\delta - 1)}{p(A - D + 1) - 1}. \quad (2)$$

Recall that any allocation (x, X) is given by:

$$x = pPA + P(1 - p)D - (1 - P)p \quad (3)$$

$$X = pPA + p(1 - P)D - (1 - p)P \quad (4)$$

⁶ Note that our objective is to capture the largest set of acceptable payoffs for a given value of δ ; hence, assuming that a grim trigger follows a defection is without loss of generality. Given that Small can unilaterally guarantee himself δPD in a given round if (p, P) is played, any other "punishment" strategy would only pick a subset of the acceptable set identified above.

Because there is a one-to-one correspondence between allocations in \mathcal{X} and stationary strategy pairs, $SPE[\Gamma(PD; \delta)] = \mathcal{X}_s(\delta) \cap \mathcal{X}_B(\delta)$. Equations (1)–(4) will be used to examine the effect on $\mathcal{X}_s(\delta) \cap \mathcal{X}_B(\delta)$ as δ varies.

We shall first prove the following results:

Lemma 1. *For all $(x, X) \in IR(\mathcal{X})$, there exists a $\delta \in (0, 1]$ is sufficiently close to 0, such that $(x, X) \in SPE[\Gamma(PD; \delta)] = \mathcal{X}_s(\delta) \cap \mathcal{X}_B(\delta)$.*

Proof/

Refer to proof of Theorem 1.

■

Lemma 2. *For any $\delta \in (0, 1)$, there exists $(\hat{x}, X) \in \mathcal{X}$ such that $\overline{\mathcal{X}}_s(\delta) \cap EFF(\mathcal{X}) = \{(x, X)\}$.*

Proof/

We begin with the observation that $(x, X) \in EFF(\mathcal{X})$ implies that either $p = 1$ or $P = 1$, where (x, X) is related to (p, P) via (3) and (4).

We shall proceed in two steps. First, we show that $\mathcal{X}_s(\delta) \cap EFF(\mathcal{X}) \neq \emptyset$. Second, we shall show that this intersection is a singleton.

Differentiating the expression (1) with respect to P yields:

$$\begin{aligned} \frac{\partial p}{\partial P} &= \frac{D[\delta - 1][P(A - D + 1) - 1] - PD[\delta - 1][A - D + 1]}{[P(A - D + 1) - 1]^2} = \\ &= \frac{-D[\delta - 1]}{[P(A - D + 1) - 1]^2}. \end{aligned}$$

Thus, $\frac{\partial p}{\partial P} \geq 0$. Since the expression (1) is continuous in P , p is increasing continuously in P on the boundary of Small's acceptable set. Let $P = 1$. By (1), one of the following cases must hold:

$$p = \frac{D[\delta - 1]}{(A - D)} \leq 1$$

or

$$p = \frac{D[\delta - 1]}{(A - D)} > 1.$$

In the first case, by (1), observe that $\mathcal{X}_s(\delta) \cap EFF(\mathcal{X}) \neq \emptyset$ and the intersection clearly occurs at

$$(p, P) = \left(\frac{D[\delta - 1]}{A - D}, 1 \right).$$

By (1), if $P = 0$, then $p = 0$ and so in the second case if $P = 1$, then $p > 1$. Thus, by continuity and the fact that $\frac{\partial p}{\partial P} \geq 0$, there must exist $P \in (0, 1)$ such that $p = 1$.

Next, we show that the intersection is unique. We must consider two cases. First, suppose that the intersection occurs at (p, P) where $p = 1$. From (1), we can solve for P :

$$P = \frac{1}{1 + A - D\delta}.$$

Clearly, the expression above yields a unique value for P when $p = 1$. Suppose instead that there is a second intersection at (p', P') . The only other possibility for a second intersection is when $P' = 1$. However, $P' > P$ implies that $p' > p$, since $\frac{\partial p}{\partial P} > 0$. Hence, the second intersection must occur at $p' = 1$. By expression (1),

$$P' = \frac{1}{1 + A - D\delta}.$$

Hence, $(p, P) = (p', P')$.

Next, suppose that there exists an intersection at (p, P) , where $P = 1$. By (1), we can solve for p :

$$p = \frac{D[\delta - 1]}{A - D}.$$

Clearly, this is the unique value of p , when $P = 1$ on the boundary of the acceptable set for Small.

Suppose there is a second intersection at (p', P') . It must be true that $p' = 1$. Also, $p \neq p'$ and $P \neq P'$ must imply that $P' < 1$. However, $P' < P$ implies that $p' < p$, since $\frac{\partial p}{\partial P} \geq 0$. Thus, we must have $p > 1$, a contradiction.

■

Lemma 3. Fix $\delta \in (0, 1]$ and let $\{(x, X)\} = \mathcal{X}_s(\delta) \cap EFF(\mathcal{X})$. Then for all $(x', X') \in EFF(\mathcal{X})$ such that $x' > \hat{x}$, it is the case that $(x', X') \in \mathcal{X}_s(\delta)$.

Proof/

Let (p, P) denote the strategies in the PD game that yield (x, X) via (3) and (4). These strategies must satisfy either $p < 1$ and $P = 1$ or $p = 1$ and $P < 1$. Consider the case where $p < 1$. From (1), we have,

$$\hat{p} = \frac{D[\delta - 1]}{A - D}.$$

Thus, $(p', 1)$ yields a payoff in $\mathcal{X}_s(\delta)$ if

$$p' \leq \frac{D[\delta - 1]}{A - D}.$$

This characterizes the set of payoffs $(x', X') \in EFF(\mathcal{X})$ such that $x' > \hat{x}$.

Next, consider the case where $p = 1$. From (1), we have

$$P = \frac{1}{1 + A - D\delta}.$$

Thus, $(1, P')$ yields a payoff in Small's acceptable set if

$$P \leq \frac{1}{1 + A - D\delta}.$$

In particular, $(p', P') = (1, 1)$ yields a payoff in $\mathcal{X}_s(\delta)$. Hence, setting $P = 1$,

$$p = 1 \leq \frac{D[\delta - 1]}{A - D}.$$

Given $P' = 1$, $(p', 1)$ yields a payoff in $\mathcal{X}_s(\delta)$ for every $p' < 1$. Again, every $(x', X') \in EFF(\mathcal{X})$ such that $x' > x$ must be in $\mathcal{X}_s(\delta)$.

■

Theorem 2. *There exists $\delta \in (0, 1]$ and (x, X) such that $SPE[\Gamma(PD; \delta)] \cap EFF(\mathcal{X}) = \{(x, X)\}$.*

Proof/

By the previous lemmata, the following are true: (i) for any $\delta \in (0, 1]$, it holds that $\bar{\mathcal{X}}_s(\delta) \cap EFF(\mathcal{X})$ is a singleton. (ii) Any point on $EFF(\mathcal{X})$ that Small prefers to this intersection point is an element of $\mathcal{X}_s(\delta)$. (iii) When δ approaches zero, $\bar{\mathcal{X}}_s(\delta)$

approaches the segment $\{(x, X) \in \mathcal{X} \mid x = 0\}$. Correspondingly, we have analogous conclusions regarding Big. Also, $\mathcal{X}_s(\delta) \cap \mathcal{X}_s(\delta) = SPE[\Gamma(PD; \delta)]$.

When δ approaches 1, $\bar{\mathcal{X}}_s(\delta)$ approaches the segment $\{(x, X) \in \mathcal{X} \mid (x, X) \text{ is induced by } (p, P) \text{ with } p = 0\}$ and $\bar{\mathcal{X}}_s(\delta)$ approaches the segment $\{(x, X) \in \mathcal{X} \mid (x, X) \text{ is induced by } (p, P) \text{ with } P = 0\}$. Recall from (1)–(4) that p and P are continuous in δ and x and X are continuous in p and P . By continuity, there exists a $\delta \in (0, 1]$ such that $\bar{\mathcal{X}}_s(\delta) \cap \bar{\mathcal{X}}_B(\delta) \cap EFF(\mathcal{X}) = \mathcal{X}_s(\delta) \cap \mathcal{X}_B(\delta) \cap EFF(\mathcal{X})$, which must be a singleton.

■

Observe that there is a unique “critical” value of δ for which the equality in Theorem 2 holds. Let δ^* denote this value.

5. A Sufficient Condition for Unique Implementation

Define $\theta = (D - A - 1)$ as the measure of the variation in marginal cost of cooperation to a player when the opponent switches from cooperation to the non-cooperative strategy.

The following theorem provides a sufficient condition for unique implementation of full cooperation. The condition is in terms of a bound on the value of q . We shall show, subsequently, that this bound lends itself to an intuitive interpretation.

Theorem 3. *If $\theta < A(D - A + 1)$, then there exists $\delta \in (0, 1]$ such that (i) $SPE(\Gamma(PD; \delta)) \cap EFF(\mathcal{X}) = \{(A, A)\}$ and (ii) $(A, A) > (x, X)$ for all $(x, X) \in SPE[\Gamma(PD; \delta)]$.*

Proof/

We shall characterize $\bar{\mathcal{X}}_s$, showing that it is a monotone increasing function of the payoff x to Small. We show monotonicity by showing that any local maximum, characterized by the derivative, is attained at or beyond the Pareto frontier. First,

observe that, given $\theta = (D - A - 1)$, the sufficiency condition of the theorem can be re-written as:

$$(D - A - 1)(D - A + 1) < A$$

$$\implies (D - A)^2 < 1 + A.$$

Let δ^* denote the *PCT* rule for which Theorem 2 is true. By symmetry, $SPE(\Gamma(PD; \delta)) \cap EFF(\mathcal{X}) = \{(A, A)\}$. Let $S \subseteq \mathcal{X}$ be the square with diagonal $[(0, 0), (A, A)]$. Since $p = 1$ at this equilibrium, and by definition, for all $(x, X) \in \bar{\mathcal{X}}_s(\delta^*)$ and all $P \in [0, 1]$, $x = \delta^* Pd$, we conclude that $\delta^* = \frac{a}{d}$. Since for all $(x, X) \in \bar{\mathcal{X}}_s(\delta^*)$, and all $P \in (0, 1]$, $x = Pa$, we know that $\frac{\partial x}{\partial P} > 0$. Hence, a sufficient condition for $SPE(\Gamma(PD; \delta^*)) \subseteq S$ is that $\frac{\partial S}{\partial P} > 0$, for all $(x, X) \in \bar{\mathcal{X}}_s(\delta^*)$, and all $P \in (0, 1]$.

Substituting the value of δ^* and (1) in the expression (4) for X and, by symmetry:

$$X = \frac{P^2(D - A)(D - A - 1) + P(D - A)D}{P(D - A - 1) + 1} - P \quad (5)$$

Differentiating (5) with respect to P , setting the resulting derivative equal to zero and re-arranging terms, yields the following expression:

$$-P^2(D - A - 1)^2(D - A + 1) - 2P(D - A - 1)(D - A + 1) + D(D - A) - 1 = 0.$$

By showing that the zero-derivative locus lies outside the game frontier and that the derivative is positive inside the frontier, monotonicity of the boundary of the acceptable set is established. This in turn establishes that $SPE(\Gamma(PD; \delta)) \subseteq S$.

For the first-order conditions to be attained exactly on $EFF(\mathcal{X})$, it must be the case that $P = 1$ at that point. In that case, the solution to the first-order condition is $(D - A)^2 = 1 + A$. The general solution for P is given by:

$$P = -\frac{1}{(D - A - 1)}$$

$$\begin{aligned} & \pm \frac{[[2(D-A+1)(D-A-1)]^2 + 4(D-A-1)^2(D-A+1)(D(D-A)-1)]^{\frac{1}{2}}}{-2(D-A-1)^2(D-A+1)} = \\ & = \frac{-(D-A+1) \pm [((D-A)^2 + D-A)(1+D)]^{\frac{1}{2}}}{(D-A)^2 - 1}. \end{aligned} \quad (6)$$

We now establish the parameter ranges in which $P > 1$; in these ranges, the maximum of $\mathcal{X}_s(\delta)$ lies outside the game frontier and the boundary of the set is monotone increasing inside the game frontier. There are two cases.

- (i) $D-A > 1$. In this case, the denominator of (6) is positive, and since the numerator term $-(D-A+1) < 0$, we must take the positive root. For $P > 1$ to hold,

$$\frac{-(D-A+1) + [((D-A)^2 + D-A)(1+D)]^{\frac{1}{2}}}{(D-A)^2 - 1} > 1$$

which reduces to $(D-A)^2 < 1+A$.

- (ii) $D-A < 1$. In this case the denominator of (6) is negative, and since $-(D-A+1) < 0$, it is possible that either the positive or the negative root can yield a positive P . Consider these two sub-cases:

- (a) First consider the negative root. For $P > 1$ to hold,

$$\frac{-(D-A+1) - [((D-A)^2 + D-A)(1+D)]^{\frac{1}{2}}}{(D-A)^2 - 1} < 1$$

which reduces to $-[((D-A)^2 + D-A)(1+D)]^{\frac{1}{2}} < (D-A)(D-A+1)$, which is necessarily true. Therefore, $D-A < 1$ is sufficient for a monotone increasing derivative of the negative root is the correct one.

- (b) Next consider the positive root. For $P > 1$ to hold,

$$\frac{-(D-A+1) + [((D-A)^2 + D-A)(1+D)]^{\frac{1}{2}}}{(D-A)^2 - 1} < 0$$

which reduces to $[((D-A)^2 + D-A)(1+D)]^{\frac{1}{2}} < (D-A)(D-A+1)$, or $(D-A)^2 > 1+A$. But this contradicts $D-A < 1$, and the positive root cannot yield $P > 1$. Therefore, if the positive root is the correct one, the

boundary of $\bar{\mathcal{X}}_s(\delta^*)$ does have a maximum at some $P < 1$. Note also that if the maximum occurs at $P = 0$, we are done. But then from the equation (6),

$$\frac{-(D - A + 1) + [((D - A)^2 + D - A)(1 + D)]^{\frac{1}{2}}}{(D - A)^2 - 1} = P > 0.$$

Since the denominator is negative, the numerator must be negative too. Then after some algebraic manipulation, we conclude that $D - A < 1/D$. We are now ready to show that even in this case, where the boundary of $\mathcal{X}_s(\delta^*)$ has a maximum at some $P \in (0, 1)$, we have $SPE(\Gamma(PD; \delta)) \subseteq S$.

From equation (5),

$$\frac{\partial X}{\partial P} = (D - A) \times \frac{(D - 2P(D - A - 1))(1 + P(D - A - 1)) - P(D - P(D - A - 1))(D - A - 1)}{(1 + P(D - A - 1))^2} - 1.$$

Evaluating at $P = 0$, we find that $\frac{\partial X}{\partial P} = (D - A)D - 1$. Thus, since we know that $D - A < \frac{1}{D}$, we conclude that $\frac{\partial X}{\partial P} < 0$ when $P = 0$. In this case, $\bar{\mathcal{X}}_s(\delta)$ bulges down from the origin before increasing to meet the efficient frontier. Since $\mathcal{X}_s(\delta^*)$ is to the right and below $\bar{\mathcal{X}}_s(\delta^*)$, we have $\mathcal{X}_s(\delta^*) \cap \mathcal{X}_B(\delta^*) = \{(0, 0), (a, A)\}$, for which $SPE(\Gamma(PD; \delta)) \subseteq S$.

The argument for Big is identical, due to symmetry.

■

The games of case (i) are ones in which $\theta > 0$, i.e. the cost of cooperation decreases with the opponent's choice of the non-cooperative strategy; case (ii) corresponds to games in which $\theta < 0$, i.e. the cost increases as the opponent tends towards non-cooperation. The implication of our theorem is that unique implementation of full cooperation is always possible when case (ii) is met. For many situations, this case quite intuitive: given that cooperation is always costly, I lose more from cooperating when the opponent decides not to cooperate compared to the situation when the opponent does cooperate. This is essentially a game with strategic complementarities in the

sense of Milgrom and Roberts [8], and is plausible for the vast majority of economic interpretations of the PD . An example of such an application is provided at the end of this section.

Under case (i), we have a (less easily translatable) condition that guarantees us the implementation result. Beyond the threshold defined by $\theta > A/(D - A + 1)$, the acceptable sets may bulge out of the “box” defined by S ; hence, the Pareto-superiority of (A, A) over all other elements of $SPE(\Gamma(PD; \delta^*))$ cannot be guaranteed.

We close this section with the following example of a free-rider problem with strategic complementarities, which captures the “typical” use of the prisoners’ dilemma-type reasoning in economics.

EXAMPLE: Consider an economy with k consumers of a public good. Each consumer j receives a benefit $V_j \in \mathfrak{R}_{++}$ from the good, which costs $C \in \mathfrak{R}_{++}$ to produce. We make the following assumptions about the economy:

- (A1) $\forall j, C > V_j$ (it is too costly for any single consumer to produce the good alone),
- (A2) $\forall j, \frac{C}{k} < V_j$ (the good yields positive net value when all consumers share in the cost),
- (A3) $\forall j, \frac{C}{k} < C - V_j$ (the size of the population, k , is sufficiently large).

Each consumer has two pure actions: $\{ \textit{Share in the Cost} (S), \textit{Not Share in the Cost} (NS) \}$. The public good is provided if and only if there is a subset of consumers willing to share in the cost. The payoff to i from S , when ℓ ($0 < \ell < k$) other consumers have chosen S is $V_i - \frac{C}{(\ell+1)}$; the payoff to i from choosing (NS) , when ℓ other consumers have chosen S is V_i if $\ell > 0$, and zero if $\ell = 0$.

The game induced by the economy above is a prisoners’ dilemma with strategic complementarities of the form discussed in this section. This simple environment also captures other variants of the problem such as the Tragedy of the Commons, Fish Wars, etc. The strategic complementarity arises from (A3) above. It is well known that the tendency to free-ride in such environments is most severe as the size of the population increases.⁷ Our focus in this paper has been on the classical PD with two players; the

⁷ A simple explanation being the fact that non-cooperative actions by a particular individual are harder to observe or verify by a third-party, when the number of players is relatively large.

results extend to multiple players, as in the example.

6. The Text-book “Knife-Edge” Case

In this section, we consider a class of PD games not covered by the result in the previous section, i.e. ones for which $\theta = 0$. This class is important because it has dominated the literature – virtually any specific example of a PD game in the literature belongs to this class. In such games the cost of choosing any strategy other than the dominant “non-cooperative” strategy is independent of the choice made by the other player. Geometrically, this class is characterized by a parallelogram property of the set of payoffs, i.e.

$$D - A = 1$$

and such PD games shall be referred to as PPD for “parallelogram PD games”.

The following theorem shows that the unique implementation property holds for all PPD games.

Theorem 4. *There exists $\delta \in (0, 1]$ such that (i) $SPE(\Gamma(PPD; \delta)) \cap EFF(\mathcal{X}) = \{(A, A)\}$ and (ii) $(A, A) > (x, X)$ for all $(x, X) \in SPE[\Gamma(PPD; \delta)]$.*

Proof/

The expression (3) may be re-written as

$$x = pP[A - D + 1] - p + PD.$$

Given the parallelogram property, the expression above simplifies to $x = PD - p$.

Writing the corresponding expression for Big and solving for the probabilities yields

$$\begin{pmatrix} p \\ P \end{pmatrix} = \begin{pmatrix} -1 & D \\ D & -1 \end{pmatrix}^{-1} \begin{pmatrix} x \\ X \end{pmatrix}.$$

The solutions are:

$$[p = [D^2 + 1]^{-1}[DX + x] \tag{7}$$

$$[P = [D^2 + 1]^{-1}[Dx + X] \tag{8}$$

Next, choose $\delta \in (0, 1]$. Substituting from (8) into the expression that defines $\mathcal{X}_s(\delta)$, we have:

$$\mathcal{X}_s(\delta) = \{(x, X) \in X : x > \delta D[D^2 + 1]^{-1}[Dx + X]\}.$$

After manipulation of the inequality above we have:

$$\frac{X}{x} \leq \frac{[D^2 - 1]}{\delta D} - D. \tag{9}$$

The corresponding expression for Big is

$$\frac{x}{X} \leq \frac{[D^2 - 1]}{\delta D} - D. \tag{10}$$

When (9) and (10) hold with equality, we have expressions that define $\bar{\mathcal{X}}_s$ and $\bar{\mathcal{X}}_B$ respectively. The appropriate choice of δ , namely δ^* , achieves this equality. From the results of the previous section, δ^* exists and is unique. When both these equations are satisfied, we have a “minimal” non-trivial equilibrium cone given by $\bar{\mathcal{X}}_s(\delta^*) \cap \bar{\mathcal{X}}_B(\delta^*) = \mathcal{X}_s(\delta^*) \cap \mathcal{X}_B(\delta^*)$; geometrically, it is a ray from the origin. The unique intersection with $EFF(\mathcal{X})$ Pareto-dominates all other equilibrium points.

■

7. Conclusion

A game with cheap talk involves payoff-irrelevant moves in initial rounds followed

by a single payoff-relevant round. On the other hand, a repeated game has payoff-relevant moves in every round. We have proposed a mechanism which is on the boundary between these two cases: in every round, moves have payoff-relevance in probability; over the entire horizon, the payoff-relevance of the game is one-shot in expectation. We have shown that such a probabilistic cheap talk mechanism permits a principal to uniquely implement, in a wide class of *PD* problems, his most desired outcome: full cooperation among agents.

The choice of the critical value of δ can be made without observing the agents' actions. Hence, the principal can obtain cooperation without monitoring the agents. Because of the limited stationarity of play, not only are payoffs distributed only once but, in equilibrium, the players choose a cooperative strategy once and for all.

The results obtained here in the context of *PD* games can, presumably, be extended to other situations—e.g. economic applications beyond the simple public good example discussed here. We undertake this task in a companion paper, Chakravorti, Conley and Taub [5].

An objection may be raised to the infinite horizon that is essential to our results. The same criticism has often been leveled against infinitely repeated games. We take the position that the infinite horizon in *PCT* games simply capture the following fact: nobody expects the cheap talk to occur forever; however, at any point in time when an action is taken, no one is certain that payoffs will result and there will be no other opportunities for revision. As Aumann [1] observes, A. Tucker has pointed out that this condition is mathematically equivalent to an infinite sequence of plays. It is essentially the same position adopted by Rubinstein [10] in his response to this criticism of infinitely repeated games.

Finally, it is reasonable to ask the following question: despite the natural focus on limited stationarity given our context, what are the implications of allowing non-stationary play? From the work of Stahl [12] and van Damme [13] cited earlier, we can conjecture that in the absence of some restriction, unique implementation is impossible, in general. In the appendix, we argue that our results are maintained if the assumption

of limited stationarity were to be replaced with a restriction that strategies must have at most a fixed finite number of transitions. This could be interpreted as strategies implementable by a boundedly rational player who is committed to a given bound.

On the other hand, there could be alternative assumptions made of the agency problem in which the principal can enforce stationary play; and our results would hold. We shall sketch two such extensions:

- (i) Suppose that the principal can discern any *revision* of actions over time, even though the actions themselves are unobservable. This would be true in situations where the agents take actions and submit an output to the principal for approval. In subsequent rounds, if any agent requests that the output be returned before it is again considered for approval, then this would signal to the principal that a possible revision of actions is being contemplated. In such situations, the principal can condition δ to be a function of such an observation so that if the output is returned, the probability of DL is set equal to unity, causing the game to revert to one-shot form. This plays the role of a grim trigger, and enforces limited stationary equilibria.
- (ii) A second alternative arises in situations where the principal can control the computational powers of the agents. This is a plausible scenario in a model of an organization in which management hires employees to work together, and controls their work environment, computing equipment, communication channels, etc. This complex set of controls can be modeled by assuming that the principal gives each agent a Moore machine to use to implement strategies in the *PD* game with *PCT*. The machine can be selected to have a bound on the number of states it possesses and the number of transition rules it can process. This is the notion of “size” of a Moore machine offered by Banks and Sundaram [3]. It can be shown that machines can be chosen so that only limited stationary equilibrium strategies of the kind used in this paper can be implemented by the agents. Essentially, such machines have at most two states and three transition rules, and the only equilibria they

can implement are stationary grim trigger strategies.⁸

8. Appendix

In this section, we show that the assumption of limited stationarity of equilibria can be relaxed. For a wide class of *PD* problems, the main properties of the previous sections are preserved even if equilibrium strategies display non-stationarities; however, the number of transitions these non-stationary strategies have must be finite in a sense to be formalized below.

First, define a set of states, N , with cardinality n . A finite-state strategy is a pair (p, P) which for each round of play specifies a list $(p_1, \dots, p_n; P_1, \dots, P_n)$, where (p_i, P_i) denote the probability of cooperation in state i . Each player commits ex ante to a transition probability matrix, \mathcal{T} (whose generic element is written as τ_{ij}), describing transitions between states of the game. \mathcal{T} can be interpreted as a measure of complexity of players' strategies; hence, finiteness of the matrix is an indication that players can implement arbitrarily complex strategies provided they involve a fixed finite number of transitions.

By fixing the transition matrix one can duplicate strategies involving an arbitrarily large but finite number of non-stationarities by an appropriate set of transition probabilities. This is an approximation of the infinitely many variations possible. One might expect that as the approximations improve, the results obtained for the unrestricted case (with infinitely many variations) would be approached asymptotically. We demonstrate here that this is not the case. For parallelogram *PD* games, the unique implementation properties at the critical value of δ are invariant with respect to N .

Our primary set of equations are given by the value recursions arising from the

⁸ A proof of this claim may be obtained from the authors.

postulated strategy and the defection conditions as follows for all $i \in N$:

$$x = \delta(p_i P_i A + P_i(1 - p_i)D - (1 - P_i)p_i) + (1 - d) \sum_{j=1}^n \tau_{ij} x_j \geq \delta P_i D, \quad (11)$$

$$X = \delta(p_i P_i A + p_i(1 - P_i)D - (1 - p_i)P_i) + (1 - d) \sum_{j=1}^n \tau_{ij} x_j \geq \delta p_i D, \quad (12)$$

The value of not deviating from the strategy is the sum of the current payoff weighted by the probability of the game ending, and the expected payoff in the subsequent state conditioned on the current state weighted by the probability that the current round results in cheap talk. The defection is followed by a grim trigger strategy which reduces the continuation payoffs to zero.

The set of subgame-perfect equilibrium allocations of $\Gamma(PD; \delta)$ using finite-state strategies is denoted $SPE^F[\Gamma(PD; \delta)]$ is defined as:

$$\{(x, X) \in \mathcal{X} \mid \exists(\pi, \Pi), h \in \mathcal{H} \text{ such that } (i)[e_0(\pi, \Pi; \delta, h), E_0(\pi, \Pi; \delta, h)] = (x, X);$$

$$(ii)(\pi, \Pi) \text{ is a subgame-perfect equilibrium of } \Gamma(PD; \delta);$$

$$(iii)(p, P) \text{ is a finite-state strategy } \}.$$

In order to characterize the boundary of the equilibrium payoff set, we set the conditions above at equality, and substitute the defection payoffs in the value recursions. We obtain, for all $i, j \in N$:

$$\frac{x_i X_i}{\delta D^2} (A - D - 1) - \frac{X_i}{D} + (1 - \delta) \sum_{j=i}^n \tau_{ij} x_j = 0 \quad (13)$$

$$\frac{x_i X_i}{\delta D^2} (A - D - 1) - \frac{x_i}{D} + (1 - \delta) \sum_{j=i}^n \tau_{ij} X_j = 0 \quad (14)$$

Define $\lambda = 1 - \delta$. The equations above can be written as:

$$\begin{pmatrix} \lambda D\tau_{11} & \cdots & \lambda D\tau_{1n} & -1 & \cdots & 0 \\ \vdots & & \vdots & \vdots & \ddots & \vdots \\ \lambda D\tau_{n1} & \cdots & \lambda D\tau_{nn} & 0 & \cdots & -1 \\ -1 & \cdots & 0 & \lambda D\tau_{11} & \cdots & \lambda D\tau_{1n} \\ \vdots & \ddots & \vdots & \vdots & & \vdots \\ 0 & \cdots & -1 & \lambda D\tau_{n1} & \cdots & \lambda D\tau_{nn} \end{pmatrix} \begin{pmatrix} x_1 \\ \vdots \\ x_n \\ X_1 \\ \vdots \\ X_n \end{pmatrix} = \begin{pmatrix} x_1 X_1 \frac{1}{\delta D} \theta \\ \vdots \\ x_n X_n \frac{1}{\delta D} \theta \\ x_1 X_1 \frac{1}{\delta D} \theta \\ \vdots \\ x_n X_n \frac{1}{\delta D} \theta \end{pmatrix}. \quad (15)$$

Define

$$v \equiv \begin{pmatrix} x_1 X_1 \\ \vdots \\ x_n X_n \end{pmatrix} \text{ and } k \equiv \frac{1}{\delta D} \theta.$$

From (15), we have:

$$[(\lambda DT) - I]x = (1 + \lambda DT)kv = [(\lambda DT)^2 - I]X, \quad (16)$$

which simplifies to

$$[(\lambda DT) - I]x = kv = [(\lambda DT)^2 - I]X. \quad (17)$$

Thus, the conditions (11) and (12) hold with equality only on the 45° line. Writing (17) such that $x_i = X_i$ we have:

$$[(\lambda DT) - I] \begin{pmatrix} x_1 \\ \vdots \\ x_n \end{pmatrix} = k \begin{pmatrix} x_1^2 \\ \vdots \\ x_n^2 \end{pmatrix}. \quad (18)$$

There are two obvious solutions, $x_i = 0$ and $x_i = A$, but there may be other solutions as well.⁹ The essential observation is that there is a unique solution on the

⁹ It can be shown that the solution must use the eigenvalues of \mathcal{T} , but x must not be an eigenvector. This is because we must choose $\frac{1}{\lambda}D$ equal to an eigenvalue; in addition, if we let x be an eigenvector corresponding to the eigenvalue, then the left hand side of (17) becomes zero, which in turn requires $x_i = 0$ for all i . Since it cannot be an eigenvector, any solution of x must be perpendicular to the space defined by the eigenvector corresponding to the eigenvalue in place.

Pareto frontier such that (11) and (12) hold with equality. Thus, the uniqueness of Pareto efficient equilibria at δ^* obtained under the limited stationarity assumption extends to the finite-state strategy equilibrium.

Next, we turn to the question of Pareto dominance of the efficient equilibrium obtained under δ^* . To keep the argument simple, we consider only the case of parallelogram PD games.

In parallelogram PD games, $\theta = 0$. Hence, (15) reduces to:

$$\begin{pmatrix} \lambda D\mathcal{T} & -I \\ -I & \lambda D\mathcal{T} \end{pmatrix} \begin{pmatrix} x \\ X \end{pmatrix} = \begin{pmatrix} 0 \\ \vdots \\ 0 \end{pmatrix}.$$

Our results for the parallelogram PD game are summarized as follows:

Theorem 5. *There exists $\delta \in (0, 1]$ such that (i) $SPE^F[\Gamma(PPD; \delta)] \cap EFF(\mathcal{X}) = \{(A, A)\}$ and (ii) $(A, A) > (x, X)$, for all $(x, X) \in SPE^F[\Gamma(PPD; \delta)]$.*

Proof/

The equilibrium payoffs at the critical value of δ presume that all the defection constraints are binding. For PPD games with $\lambda^* = 1 - \delta^*$, the payoffs can be characterized as:

$$\lambda^* D\mathcal{T}x = X \text{ and } \lambda^* D\mathcal{T}X = x.$$

Combining, we obtain

$$(\mathcal{T}^2 - (\lambda^* D)^{-2} I)x = 0.$$

This is an eigenvalue problem. We can rewrite it as:

$$(\mathcal{T} + (\lambda^* D)^{-1} I)(\mathcal{T} - (\lambda^* D)^{-1} I)x = 0$$

so that it is solved by η , where η is an eigenvalue of \mathcal{T} and by the associated eigenvector x of \mathcal{T} . The solution for the critical λ is then $\lambda^* = (\eta D)^{-1}x = 0$. Given that η and x are an eigenvalue-eigenvector pair, we have:

$$X = \lambda^* D(\lambda^* D)^{-1}Ix \implies \forall i \in N, \frac{X_i}{x_i} = 1.$$

Hence, we have the desired result. In addition, observe that since \mathcal{T} is a stochastic matrix, it has at least one unit eigenvalue; any others must necessarily be fractional. The critical λ (and hence, δ) can be chosen to be associated with the unit eigenvalue, so that $\delta^* = 1 - \lambda^* = 1 - \frac{1}{D}$.

■

References

- [1] Aumann, R., Acceptable points in general cooperative n-person games, in R.D. Luce and A. W. Tucker (eds.) in Contributions to the Theory of Games IV, Annals of Mathematical Studies, No. 40, Princeton University Press, pp 287-324, (1959).
- [2] Abreu, D., D. Pearce and E. Stacchetti: Towards a theory of discounted repeated games with imperfect monitoring, *Econometrica*, 58, 1041-1064, (1990).
- [3] Banks, J. and R. Sundaram, Repeated games, finite automata, and complexity, *Games Economic Behavior*, (1989).
- [4] Chakravorti, B. and J. Conley, Cheap play with no regret, *International Journal of Game Theory*, forthcoming, 1994.
- [5] Chakravorti, B., J. Conley and B. Taub, Probabilistic cheap talk, mimeo, Bellcore and University of Illinois, 1993.
- [6] Fershtman, C. and E. Kalai, Unobserved delegation, mimeo Northwestern University, 1993.
- [7] Green, J. and J.-J. Laffont, Posterior implementability in a two-person decision problem, *Econometrica*, 55, 69-94, (1987).
- [8] Milgrom, P. and J. Roberts, Rationalizability, learning and equilibrium in games with strategic complementarities, *Econometrica*, 58, 1255-1278 (1990).
- [9] Rubinstein, A., Comments on the interpretation of repeated game theory, in J. J. Laffont (ed.) *Advances in Economic Theory*, Sixth World Congress, Cambridge University Press, Cambridge, UK, pp 175-181, (1992).
- [10] Rubinstein, A. and A. Wolinsky, Equilibrium in a market with sequential bargaining, *Econometrica*, 53, 1133-1150, (1985).
- [11] Sorin, S., On repeated games with complete information, *Mathematics of Operations Research*, 11, 147-161, (1986).
- [12] Stahl, D., The graph of prisoners' dilemma supergame payoffs as a function of the discount factor, " *Games and Economic Behavior*, 3, (1991).
- [13] van Damme, E., *Stability and Perfection of Nash Equilibria*, Berlin New York, Springer-Verlag, (1991).