

Economic Applications of Probabilistic Cheap Talk†

Bhaskar Chakravorti*

John P. Conley**

and

Bart Taub***

† **Published as:** Bhaskar Chakravorti, John P. Conley and Bart Taub, "Economic Applications of Probabilistic Cheap Talk" in Topics in *Mathematical Economics and Game Theory: Essays in Honor of Robert Aumann* M. Wooders, editor, Fields Institute Communication, Toronto, pp. 153-170, 1999. We are grateful to the participants of the "Incentives, Implementation, and Applications to Industrial Organization" workshop at SUNY, Stony Brook.

* The Monitor Corporation

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

*** Department of Economics, UIUC

Abstract

In this paper, we apply the concept of “Probabilistic Cheap Talk” (PCT), introduced in Chakravorti, Conley and Taub, (1993) to a variety of simple economic problems in public economics and industrial organization. We focus on problems which have a single Nash (often dominant strategy) equilibrium which is Pareto dominated by other outcomes. We show that adding a PCT structure to the pre-existing game transforms these problems into multi-stage games with a single, Pareto dominant, efficient, subgame perfect equilibrium. Nevertheless, the game remains essentially one-shot. Payoffs are distributed once and only once, and in equilibrium, strategies are chosen once and for all. Our focus is on resolving inefficient outcomes when agents’ strategy spaces and outcome rules are given by pre-existing institutions, and the choices of agents are unverifiable by a third-party. By way of example, we show how PCT can be successfully applied problems involving voluntary provision of public goods, market failure due to externalities, and failures of collusion among oligopolists.

Keywords: Probabilistic cheap talk, implementation, preference revelation, public goods, externalities, Bertrand oligopoly.

JEL: C72, D43, D62, H41, L12, L13.

1. Introduction

In this paper, we apply the concept of “probabilistic cheap talk” (PCT) to a variety of simple economic problems in public economics and industrial organization. This concept was introduced in Chakravorti, Conley and Taub (1993), and explored in the context of abstract prisoners’ dilemma games in Chakravorti, Conley and Taub (1996).

By way of illustration, consider case of joint ventures among regulated firms. Managers of such firms engage in collaborative ventures in the form of investments in mutually beneficial projects. For example, computer or telecommunication firms might jointly developed new technological standards and banks might jointly petition to be allowed to offer new types financial services. Due to the long-lasting nature of these projects, their horizons are often limited to the venture in question. Firms receive payoffs only when the regulator grants permission to bring the output to market. Managers know that approval will be given eventually, but each time they go before the regulatory body, there is a chance that it will be refused. The actions of each firm may be revised prior to the granting of regulatory permission. Once approval is granted, the product is put on the market and the game ends. Managers often have an incentive to free-ride on the investments of others, or follow strategies which would give them an inequitable advantage were the approval to be given. Nevertheless, we still see successful joint ventures despite the one-shot nature of the payoffs. We argue below that it is precisely the uncertain timing of the arrival of payoffs which makes punishment for non-cooperative behavior possible, and provides an explanation for the observed cooperative behavior.

More abstractly, imagine any one shot game. Now suppose that we transform the game as follows. Agents simultaneously and secretly commit to strategies. With some probability, the game ends and payoffs are a distributed according the strategies agent choose. If the game does not end, then the round of play is ex post cheap talk and each agents’ (now payoff irrelevant) strategy choices are publicly revealed. We then move to a new round of play and the process is repeated until eventually the game ends.

The probability with which the game ends in any given round might be determined by some random element in nature (messengers being comprised or communications networks failing) or chosen by a mediator. Note that the game ends and payoffs arrive with certainty if we aggregate these probabilities over the entire time horizon; however, payoffs are received once and only once. This implies the important fact that the standpoint of economic resources used, therefore, the game is one-shot. This is in sharp contrast to other types of repeated games.

We consider several economic problems which have some fundamental features in common. In each application, there is a single Nash equilibrium (which in some cases is also a dominant strategy equilibrium) which is Pareto dominated by other feasible allocations. We show that through the use of a PCT mechanism, the equilibrium set can be expanded to include Pareto efficient outcomes. Better still, we show that the probability of cheap talk is the variable which can be manipulated to generate an equilibrium set with a single Pareto dominant, efficient, individually rational, subgame perfect allocation. In most cases, the equilibrium set can be reduced to consist of only the intuitively appealing Pareto efficient allocation, and the one-shot Nash equilibrium allocation. Moreover, this is accomplished without the necessity of a third-party who can observe the strategy choices of the agents. Thus, the PCT mechanisms allow efficient outcomes to be achieved in environments in which binding contracts cannot be written.

It is instructive to compare our approach to the one usually taken in the related literature on the theory of implementation. Implementing efficient allocations in both public and private goods has a long history.¹ The central premise of implementation theory is that agents' preferences (or other crucial payoff-relevant parameters) are unobservable to the mechanism designer. The designer is a social planner who invites the agents to play a game whose equilibrium outcomes are contained in the set of "desirable" outcomes. The notion of what constitutes desirability and, of course, the equilib-

¹ For a recent example in the public goods context, see Jackson and Moulin (1992).

rium outcomes would depend on the actual preferences of the agents. The appealing aspect of the designer’s game is that this containment relationship holds regardless of what the agents’ true preferences really are. The game generally involves the agents sending messages to the designer who, in turn, allocates resources as a function of the messages. This function and the message space are chosen by the designer.

Our motivation is quite different from that of the implementation literature. We recognize that certain institutions, and therefore the underlying game as well (including strategies and outcome functions), are already in place. The problem is that the existing structure makes it impossible for the agents to achieve an efficient outcome. The PCT mechanism takes the underlying game as historically given and provides the agents a means of communicating with each other prior to playing the game. The game induced by this procedure can be operated by a designer who cannot observe the strategies played by the agents.

Our mechanism has three important features that are lacking in a standard implementation mechanism: (i) The PCT mechanism preserves the pre-existing institution or game and uses it as a building block to induce a new game. In contrast, implementation mechanisms are typically complex and do not generally correspond to any existing economic institution. Also, it is clearly more difficult to find a designer with the power to determine not only the strategy space from which agents must make choices, but also the outcome function. It is much more plausible to imagine that the designer merely modifies certain rules of play in a pre-existing game. (ii) Many actions taken by economic agents are unverifiable to third-parties. The implementation literature presumes that a “message space” can be invented by the designer who can also observe the actual messages chosen. Our method simply requires the agents to choose from the set of actions available in the pre-existing institution; and these actions need not necessarily be observable to an outsider. Thus, the mediator, in our context, can simply be agreed upon by the agents as some form of “noise”. For example, the agents could enter their strategies into computer terminals hooked up to a network that goes down with a given frequency, thereby halting the transmission of the entries to the payoffs

with some probability. The downtime frequency could be agreed upon by the agents. (iii) Finally, many implementation schemes have non-unique equilibria without a clear indication of how the agents would resolve the coordination problem over the Pareto unranked outcomes that would result. Thus, we are left with a generalized Battle of the Sexes problem. Our mechanism also has multiple equilibria, but there is a very simple resolution to the coordination problem.

On the other hand, our mechanisms have certain weaknesses vis-a-vis the implementation mechanisms: (i) In general, our focus is on problems where strategy choice, rather than private information on preferences, is unobservable to third-parties. (ii) In the case of implementation mechanisms, *every* equilibrium is efficient. Our mechanisms do not have such a strong property. We show that, beginning with a game in which there is only one inefficient equilibrium, the game can be modified to expand the equilibrium set in an extremely attractive manner: the expanded set now includes (in addition to the status quo outcome) a unique Pareto dominant, efficient outcome. There are no problems relating to coordination over Pareto unranked outcomes, and it is generally expected that the agents will find this dominant efficient outcome focal. Of course, strict Pareto dominance is not a guarantee that agents will choose such an outcome (for example, if they use coalition-proofness a la Bernheim, Peleg and Whinston (1987) as a solution concept). However, most economic agents will agree that any mechanism that transforms a problem with a single inefficient outcome to one that also has a unique Pareto dominant, efficient outcome transforms the opportunity set in a most dramatic and highly desirable manner.

To summarize, our approach yields results that are weaker in some aspects and stronger in some others as compared to those in standard implementation theory. Also, the class of problems that the two approaches address are quite different.

Before proceeding, we note that we do not explicitly introduce a discount factor. This may be interpreted in different ways depending on the application. Often, it is appropriate to think of the delay in the realization of payoffs as occurring either in non-real time or within a period of time that is so small that there is no discounting. If, on

the other hand, we were to apply our models to environments with strong discounting of the future, recall that while the payoffs are not realized, the underlying economic resources may be assumed to grow (for example, through the accumulation of interest earnings in an escrow account). Thus, it may be assumed that the rate of growth of the resources exactly neutralizes the discounting.

The remainder of the paper is organized as follows. In section 2, we describe a general PCT mechanism. In section 3, we describe a simple public goods problem. We show how the PCT mechanism can be used to generate equal sharing of cost by all agents as the unique Pareto dominant, efficient, subgame perfect equilibrium. In section 4, we look at an externality problem. We show how efficient provision of a positive externality, or abatement of a negative externality, can be generated through the use of a PCT mechanism. In section 5, we look at Bertrand oligopoly. We show that all firms setting price at the monopoly level is the unique Pareto dominant, efficient, subgame perfect equilibrium of a PCT extension of a one-shot game. Section 6 concludes

2. Probabilistic Cheap Talk Mechanisms

In this section we define an abstract PCT mechanism. First, consider a one-shot game $G \equiv \langle N, M, v \rangle$. Let N be the set of agents. Let M^i be the set of *moves* available to agent $i \in N$. Let $v^i : M^1 \times \dots \times M^n \equiv M \rightarrow \mathfrak{R}^n$, be the payoff function for $i \in N$. We m to denote $(m^i)_{i \in N}$ and m^{-i} to denote $(m^j)_{j \in N \setminus \{i\}}$.

Agents simultaneously choose moves. The designer is a third party (a “mediator”, a computer with random down-time, or just plain noise) with the ability to delay the payoffs based on the outcomes of a randomization device with two realizations: “Cheap Talk”(CT) and “DeadLine” (DL). If *DL* is realized, the payoffs are generated for the moves the agents chose. If *CT* is realized, then the payoff distribution is put on hold, the chosen moves are treated as cheap talk, and the process is repeated. In all such repetitions, each agent has two options: i) An agent may opt out of the communications

process by not taking any further actions after the first choice. This is interpreted as a willingness to replay the same move chosen in the previous round of (now payoff irrelevant) play. ii) An agent may choose a new move to replace the previous one. If all agents “opt out” in the first round of play, we refer to this as *event* (α) , otherwise, we have *event* (β) . Note that payoffs are received only once. If event (α) occurs, actions are taken only once and, de facto, there is no communication. Agents “observe” all ex post cheap play but not the current moves of the other agents. The designer can only distinguish between event (α) and event (β) . The ex post cheap talk is unobservable to the designer.

Denote the number of rounds of ex post cheap talk by $t \in \{1, 2, \dots\}$. The history of talk at t is denoted by h_t . Let \mathcal{H} be the set of all possible histories over all $t \in \{1, 2, \dots\}$. Let \mathcal{H}_t be the space of all possible histories at time t . We shall set $h_1 = \emptyset$. A *strategy profile* for $i \in N$ is a mapping $s^i = \{s_t^i : \mathcal{H}_t \rightarrow M^i\}_{t=1}^\infty$. Let S^i be the class of all possible strategy profiles.

If for some $t' < t$ the agents have chosen $m \in M$ in round t' , we shall say that the resulting history h_t *contains* m *at* t' . We shall write this as $m \in_{t'} h_t$. A history $h \in \mathcal{H}$ is said to be *stationary* if

$$\exists m \in M \text{ s.t. } \forall t \in \{1, 2, \dots\}, \text{ and } \forall t' \leq t, m \in_{t'} h_t.$$

If $h \in \mathcal{H}$ is stationary, then the restriction of h to the first t rounds, h_t , is also said to be stationary. Note that event α generates a stationary history. Let \mathcal{H}^α denote the sub-class of stationary histories generated by event α . Let $\mathcal{H} \setminus \mathcal{H}^\alpha = \mathcal{H}^\beta$.

We now define two subclasses of strategies. Let $S_M \in S$ be the class of *stationary trigger strategy profiles* in which agents play a stationary strategy, and respond to any deviation from stationary by going to a “punishment” move the next round:

$$S_M \equiv \{s_{\tilde{m}} \in S\}$$

where for all $h \in \mathcal{H}$, and all $t \in \{1, 2, \dots\}$

$$s_{\tilde{m},t}(h_t) \equiv \begin{cases} s_{\tilde{m},t-1}(h_{t-1}), & \text{if i) } \exists m \in M \text{ s.t. } m \in_{t-2} h^t, \text{ and } m \in_{t-1} h^t \\ & \text{or ii) } t = 2 \\ \tilde{m}, & \text{otherwise.} \end{cases}$$

Let $S_M^M \subset S_M$ be the class of *enforcing trigger strategy profiles* in which agents play a *particular* stationary strategy and respond to any deviation from this strategy by going to a punishment move next round. The difference between these two is that in stationary trigger strategies, an agent may or may not require that other agents make a particular move in the first round. However, he always goes to the punishment move if any agent makes a different move in any subsequent round. In an enforcing trigger strategy, punishment is induced not only by nonstationary play, but also by deviation from a specific move by the other agents, even in the first round. Formally:

$$S_M^M \equiv \{s_{\bar{m}}^{\bar{m}} \in S_M\}$$

where for all $h \in \mathcal{H}$, and all $t \in \{1, 2, \dots\}$,

$$s_{\bar{m},t}^{\bar{m}}(h_t) = \begin{cases} \bar{m}, & \text{if i) } \bar{m} \in_{t-1} h_t; \\ & \text{or ii) } t = 1 \\ \tilde{m}, & \text{otherwise.} \end{cases}$$

Our convention is to have the superscripted move (if any) to be the proposed stationary one, while the subscripted move (if any) is the punishment move. Although our attention will focus on enforcing stationary trigger strategies, the former class is needed for technical reasons in the proofs of the lemmata that follow.

For any $i \in N$, and any $m \in M$ let $m^{i,opt}$ denote the set of *optimal one-shot defections from m for agent i* .

$$m^{i,opt} \equiv \{\hat{m}^i \in M^i \mid \forall \bar{m}^i \in M^i, v^i(m^1, \dots, \hat{m}^i, \dots, m^n) \geq v^i(m^1, \dots, \bar{m}^i, \dots, m^n)\}.$$

In the applications given below, this will always be single valued. In the general case, it is sufficient to take any element of this set to prove the lemmata below. Note that if \tilde{m} is a Nash equilibrium, then $\tilde{m}^{i,opt} = \tilde{m}^i$.

We are now able to give a formal definition the mechanism.² A PCT mechanism is a profile $\delta = \{\delta_t : h_t \rightarrow [0, 1]\}_{t=1}^\infty$ such that for all $t \in \{1, 2, \dots\}$,

$$\delta_t(h_t) = \begin{cases} \delta & \text{if } h_{t-1} \in \mathcal{H}^\alpha \\ 1 & \text{if } h_{t-1} \notin \mathcal{H}^\alpha. \end{cases}$$

² We use δ to mean both the profile of maps from histories into probabilities, and the actual probability with which the game ends if the history is stationary. This is a slight abuse of notation, but it should not cause any confusion in context. We immediately restrict attention to the class of profiles that enforce stationary play.

Given a history h_t , $\delta_t(h_t)$ is the probability that DL is realized at round t . Thus, if the game continues to round t , and the history involves agents opting out up until then (that is, up to $t - 1$), then the game ends with probability δ . However, the game ends with certainty, and the payoffs are distributed, on the first round after any agent is seen revising his strategy. Note that at each t , δ_t is measurable with respect to the information partition $\{\mathcal{H}^\alpha, \mathcal{H}^\beta\}$. Let $\Gamma(G, \delta)$ denote the multi-stage game induced by the PCT mechanism, δ , given the underlying one-shot game, G .

For any agent $i \in N$, participating in a strategy profile $s \in S$ yields the following expected payoff at time t , given history h :

$$\begin{aligned} & \delta_t(h_t)v^i(s_t(h_t)) \\ & + (1 - \delta_t(h_t))\delta_{t+1}(h_{t+1})v^i(s_{t+1}(h_{t+1})) \\ & + (1 - \delta_t(h_t))(1 - \delta_{t+1}(h_{t+1}))\delta_{t+2}(h_{t+2})v^i(s_{t+2}(h_{t+2})) + \dots \end{aligned}$$

Therefore, define the *expected value of strategy s to agent i at time t* , $E_t^i : S \times \mathcal{H}_t \rightarrow \mathfrak{R}_+$ as

$$E_t^i(s, h_t) \equiv \left\{ \delta_t(h_t)v^i(s_t(h_t)) + \sum_{k=t+1}^{\infty} \left[\delta_k(h_k)v^i(s_k(h_k)) \prod_{r=t}^{k-1} (1 - \delta_r(h_r)) \right] \right\}$$

where the histories after t are generated by equilibrium play of s given h_t . A strategy profile $s \in S$ is *subgame perfect equilibrium (SPE) of the game $\Gamma(G, \delta)$* if

$$\forall t \in \{1, 2, \dots\}, \forall i \in N, \forall \bar{s}^i \in S^i, \text{ and } \forall h \in \mathcal{H},$$

$$E_t^i(s, h_t) \geq E_t^i(s^1, \dots, \bar{s}^i, \dots, s^n, h_t).$$

The following lemmata will be useful in proving the results in subsequent sections. The first of these simply says that if a strategy profile is an SPE of a game, then it must be a stationary trigger strategy. This means that the only equilibrium histories are stationary histories.

Lemma 1. *Let $\tilde{m} \in M$ be the unique Nash equilibrium of a one-shot game G . If a strategy, s , is an SPE of $\Gamma(G, \delta)$, then there is $s_{\tilde{m}} \in S_M$ such that $s \equiv s_{\tilde{m}}$.*

Proof/

First, in any SPE strategy, it must be the case that \tilde{m} is played on the round after any nonstationary move occurs. This is because by construction of the PCT mechanism, the game ends with certainty in the next round, and by hypothesis, \tilde{m} is the only Nash equilibrium in such a subgame.

Second, in any SPE, all agents must make the same move each round if the history has been stationary. Suppose instead that it was not optimal to play a stationary strategy. Then suppose that the history h_{t+2} is stationary up to t , with agents playing $\bar{m} \in M$ each round. However at round $t + 1$ the equilibrium move is $\hat{m} \neq \bar{m}$. The expected payoff to agent j from abiding by the “equilibrium” strategy in this subgame is:

$$\delta v^j(\hat{m}) + (1 - \delta)v^j(\tilde{m}).$$

This is because if the PCT game continues until round t , there is a probability of δ that the game ends in round $t + 1$. If the game does not end in round $t + 1$, the mediator ends the game with certainty in $t + 2$ due to the nonstationary play. Thus the probability that the game ends at $t + 2$ is $1 - \delta$. Assume that $\hat{m} \neq \tilde{m}$.

Then the expected payoff to agent j from deviating optimally in round $t + 1$ is

$$\delta v^j(\hat{m}^1, \dots, \hat{m}^{j,opt}, \dots, \hat{m}^n) + (1 - \delta)v^j(\tilde{m}).$$

But since \tilde{m} is the only one-shot Nash equilibrium, for at least one $j \in N$

$$\delta v^j(\hat{m}^1, \dots, \hat{m}^{j,opt}, \dots, \hat{m}^n) + (1 - \delta)v^j(\tilde{m}) > \delta v^j(\hat{m}) + (1 - \delta)v^j(\tilde{m}).$$

and so the strategy could not have been an SPE.

The argument is similar if it happens that $\hat{m} = \tilde{m}$. In this case, agent j can improve his expected payoff by deviating in round t . Since he then receives the Nash equilibrium payoffs in $t + 1$ anyway, there is no incentive not to defect in the previous round t , and so s could not have been an SPE.

■

The next lemma shows that for any enforcing trigger strategy, $s_{\tilde{m}}^{\tilde{m}}$, it is a best response for agents to invoke the punishment move, \tilde{m} , if there is any deviation from the move \tilde{m} . Note that this lemma does not say that \tilde{m}^i is a best response to \tilde{m}^{-i} .

Lemma 2. *Let $\tilde{m} \in M$ be the unique Nash equilibrium of a one-shot game G , and consider any enforcing trigger strategy $s_{\tilde{m}}^{\tilde{m}} \in S\tilde{m}_{\tilde{m}}$. For any $t \in \{1, 2, \dots\}$, suppose also that $h_{t+1} \in \mathcal{H}_{t+1}$ is such that for all $t' < t$ $\tilde{m} \in_{t'} h_{t+1}$, but $\tilde{m} \neq \hat{m} \in_t h_{t+1}$. Then for all $i \in N$, it is a best response in this subgame to abide by the trigger strategy and play \tilde{m}^i in round $t + 1$ and all future rounds.*

Proof/

For $t = 1$, under the hypothesis, $h_2 = \{\hat{m}\}$ where for at least one agent j , $\hat{m}^j \neq \tilde{m}^j$. This history is not consistent with the equilibrium play of the trigger strategy, but is trivially stationary. However, under the trigger strategy, all agents other than j respond to this history by playing \tilde{m}^{-j} in the second round. If the deadline happens not to hit, and the game does not end in the second round, then it certainly ends in the third due to the non-stationary play. Then clearly it is a best response for all agents $i \in N$ to abide by this trigger strategy and play \tilde{m}^i in the second and third rounds.

Finally for any $t > 2$, suppose h_t satisfies the hypothesis. Then the history is stationary at \tilde{m} up until $t - 2$, but at $t - 1$ at least one agent j , makes the move $\hat{m}^j \neq \tilde{m}^j$. Again, if the deadline happens not to hit in round $t - 1$, it certainly ends in round t due to the non-stationary play. Then clearly it is a best response all agents $i \in N$ to abide by this trigger strategy and play \tilde{m}^i in this last round.

■

The next lemma shows that all equilibrium histories can be generated by enforcing trigger strategies, and so attention can be restricted to this class.

Lemma 3. *Let $\tilde{m} \in M$ be the unique Nash equilibrium of a one-shot game, G . Let $s_{\tilde{m}} \in S_M$ be an SPE of $\Gamma(G, \delta)$. Then $h \in \mathcal{H}^\alpha$, with $\hat{m} \in_1 h$, is a possible equilibrium history generated by this strategy if and only if $s_{\tilde{m}}^{\hat{m}} \in S_M^M$ is also an SPE of $\Gamma(G, \delta)$.*

Proof/

It is immediate that $h \in \mathcal{H}^\alpha$ with $\hat{m} \in_1 h$ is a possible equilibrium history generated by some strategy $s_{\hat{m}} \in S_M$ if $s_{\hat{m}}^M \in S_M^M$ is also an SPE since $s_{\hat{m}}^M \in S_M$.

To see the reverse implication, suppose that $h \in \mathcal{H}^\alpha$, with $\hat{m} \in_1 h$, is a possible equilibrium history generated by some strategy $s_{\hat{m}} \in S_M$ but that $s_{\hat{m}}^M \in S_M^M$ was not an SPE. Then for some $t \in \{1, 2, \dots\}$, and some $\bar{h}^t \in \mathcal{H}$, there is an agent $i \in N$ and a strategy s_t^i such that

$$E_t^i(s_{\hat{m},t}^{\hat{m},1}, \dots, s_t^i, \dots, s_{\hat{m},t}^{\hat{m},n}, \bar{h}_t) > E_t^i(s_{\hat{m}}^{\hat{m}}, \bar{h}_t).$$

Suppose first that for all $t' < t$, $\hat{m} \in_{t'} \bar{h}^{t'}$. But since the future play of both $s_{\hat{m}}^{\hat{m}}$ and $s_{\hat{m}}$ are the same given this history, it follows that:

$$\begin{aligned} E_t^i(s_{\hat{m},t}^{\hat{m},1}, \dots, s_t^i, \dots, s_{\hat{m},t}^{\hat{m},n}, \bar{h}_t) &= E_t^i(s_{\hat{m},t}^1, \dots, s_t^i, \dots, s_{\hat{m},t}^n, h_t) \\ &> E_t^i(s_{\hat{m}}^{\hat{m}}, h_t) = E_t^i(s_{\hat{m}}^{\hat{m}}, \bar{h}_t), \end{aligned}$$

which contradicts the hypothesis that $s_{\hat{m}}$ is an SPE.

Suppose instead that for some $t' < t$, $\bar{m} \in_{t'} h^{t'}$, and $\bar{m} \neq \hat{m}$. Without loss of generality, assume that round t' is the first time any move other than \hat{m} is seen. Then by Lemma 2, it is a best response for agent i to abide by the trigger strategy and invoke the punishment move in all rounds after t' .

Thus, $s_{\hat{m}}$ is an SPE, then $s_{\hat{m}}^{\hat{m}}$ also has all the agents playing a best response in every subgame, and is therefore an SPE as well.

■

We close this section with the observation that the construction of δ is such that in equilibrium there is not cheap talk. The agents choose strategies in the first round and then opt out of communication.

3. A Hand Raising Mechanism for the Provision of Discrete Public Goods

Consider the following one-shot game $G^{hr} \equiv \langle N, M, v \rangle$. Each agent $i \in N$, chooses between two moves: $M^i = \{HO, FR\}$ (Help Out, or Free Ride). Let $\#HO : M \rightarrow \{0, \dots, n\}$ be a function that gives the number of agents that agree to help out for in any given profile of moves. Let B^i be the benefit that agent i stands to receive if the public project is undertaken, and C be the cost of the project. Then for all $i \in N$, and all $m \in M$,

$$v^i(m) = \begin{cases} B^i - \frac{C}{\#HO(m)} & \text{if } m^i = HO \\ B^i & \text{if } m^i = FR \text{ and } \#HO(m) \neq 0 \\ 0 & \text{if } \#HO(m) = 0. \end{cases}$$

We assume that for all $i \in N, C > B^i > 0$, and $C < \sum_{i=1}^n B^i$. This means that no single agent would be willing to build the project on his own, but the sum of the benefits to all agents exceeds the cost.

This is a very simple model of a mechanism to produce and pay for a discrete level of a public good. To fix the idea, suppose that a group of neighbors is considering the construction of a community playground for the local children. By raising his hand and agreeing to help out with the playground, an agent promises to show up at the proposed site and share in the effort needed to complete the until the task is completed.

The difficulty, of course, is in overcoming the free rider problem. In the one-shot game, in which the moves are simultaneous, it is a dominant strategy to free ride. Also notice that building a playground is fundamentally a one-shot problem, or at best, a finitely repeated one. It does not make sense to think about building an infinity of playgrounds. We therefore cannot expect agents to arrive at an efficient outcome by using a repeated game argument.

Consider the incentive problem facing any particular agent i when all the other agents play their part of an arbitrary enforcing trigger strategy $s_{\tilde{m}}^M \in S_M^M$. Since $\{FR, \dots, FR\}$ is the only Nash equilibrium of G^{hr} , by Lemma 1, we can assume that the punishment move is $\tilde{m} \equiv (FR, \dots, FR)$. If $\bar{m}^i = FR$, then i is being asked to play his one-shot dominant strategy. Then it is trivially a best response in every subgame

for him to play his part in the trigger strategy equilibrium. On the other hand, if $\bar{m}^i = HO$, he finds that playing his part of this strategy is a best response if and only if the following condition is met:

$$B^i - \frac{C}{\#HO(\bar{m})} \geq \delta B^i.$$

The left hand side of this expression is i 's payoff from playing along with the trigger strategy. By so doing the agent i guarantees himself this payoff whenever the deadline happens to fall. The right hand side is his payoff from the optimal deviation, free riding from the first round of the subgame. If the deadline hits in round one (which is a probability δ event), then he gets the full benefit of the public good, B^i , without paying any of the cost of the project. If the game does not end, then all other agents free ride in the following round. This nonstationary play induces the mediator to end the game with certainty in the third round. Thus, the only possible benefit from defecting from the trigger strategy is free riding in the first round in the event that the deadline happens to hit. Then clearly the agent will abide by the trigger strategy if and only if:

$$\delta \leq 1 - \frac{C}{B^i \#HO(\bar{m})}.$$

First, suppose all agents are identical and so $\forall i \in N, B^i = B$. Let

$$\delta^* = 1 - \frac{C}{Bn}$$

and $\bar{m} \equiv (HO, \dots, HO)$.

Lemma 4. *Suppose all agents in the game G^{hr} are identical. Then $s_{\bar{m}}^{\bar{m}}$ and $\tilde{s}_{\bar{m}}^{\bar{m}}$ are the only SPE of $\Gamma(G^{hr}, \delta^*)$.*

Proof/

Clearly, $s_{\bar{m}}^{\bar{m}}$ is an SPE for any δ since \bar{m} is the only Nash equilibrium of the one-shot game. To see that $\tilde{s}_{\bar{m}}^{\bar{m}}$ is also an SPE, consider any $i \in N$ and any $h \in \mathcal{H}$. This history could have evolved in one of two ways.

Suppose first that the game has not ended at any given t , and suppose that all agents have been playing HO each round. Then by construction of δ^* it is a best response for i to play HO in round t .

Suppose on the other hand that at some $t \in \{1, 2, \dots\}$, some agent j plays FR . Then by Lemma 2, it is a best response for all agents to abide by the trigger strategy and play FR in all remaining rounds.

Finally, suppose there was another trigger strategy $s_{\bar{m}}^{\hat{m}}$ that is an SPE. But then $\#HO(\bar{m}) > \#HO(\hat{m})$ and so

$$\delta^* B = B - \frac{C}{\#HO(\bar{m})} > B - \frac{C}{\#HO(\hat{m})}.$$

Thus it would be optimal for every agent to defect from any other trigger strategy.

■

Theorem 1. *Suppose all agents in the game G^{hr} are identical. Then the only SPE payoffs of $\Gamma(G^{hr}, \delta^*)$ are: $\{(B - \frac{C}{n}, \dots, B - \frac{C}{n}), (0, \dots, 0)\}$.*

Proof/

By Lemma 4, $s_{\bar{m}}^{\bar{m}}$ and $s_{\tilde{m}}^{\tilde{m}}$ are the only SPE. Then by Lemma 3,

$$\bar{h} \equiv \{h \in \mathcal{H} \mid \forall t \in \{1, 2, \dots\}, \bar{m} \in_t h\},$$

and

$$\tilde{h} \equiv \{h \in \mathcal{H} \mid \forall t \in \{1, 2, \dots\}, \tilde{m} \in_t h\},$$

are the only SPE histories. Thus

$$\{(B - \frac{C}{n}, \dots, B - \frac{C}{n}), (0, \dots, 0)\}$$

are the only SPE payoffs.

Thus, if all agents derive the same benefit from the project, we can choose δ high enough so that everyone helping out, and everyone trying to free ride (and the project not being undertaken) are the only an SPE outcomes. Then since everybody helping

out strongly Pareto dominates everybody free riding, it is a focal equilibrium which on which all agents share an interest in coordinating.

Now consider the case of nonidentical agents. We know that an agent is better off helping out than free riding in a trigger strategy enforcing \bar{m} if:

$$B^i - \frac{C}{\#HO(\bar{m})} \geq \delta B^i.$$

Assume that the agents are ordered so that agents with a lower index place more value on the project. For simplicity, assume that all agents receive different benefits. Thus $B^1 > B^2 > \dots > B^n$.

For any $j \in \{1, \dots, n\}$, let $m^{(j)}$ denote the move in which the first j agents help out and the rest free ride:

$$m^{(j)} \equiv (m^{(j),1}, \dots, m^{(j),j}, m^{(j),j+1}, \dots, m^{(j),n},) = (HO, \dots HO, FR, \dots FR).$$

For any $j \in \{1, \dots, n\}$, define:

$$\delta^{(j)} = 1 - \frac{C}{B^{(j)}\#HO(m^{(j)})}.$$

Finally, define j^* to be:

$$j^* \equiv \{j \in \{1, \dots, n\} \mid \forall i \in \{1, \dots, n\}, \delta^{(j)} \geq \delta^{(i)}\}.$$

Note that j^* may not be always be a singleton, but is generically unique. We will assume in the following that j^* is unique.

Lemma 5. *Suppose that j^* is unique. Then the strategies $s_{\bar{m}}^{m^{(j^*)}}$ and $s_{\bar{m}}^{\tilde{m}}$ are the only SPE of $\Gamma(G^{hr}, \delta^{(j^*)})$.*

Proof/

Clearly, $s_{\bar{m}}^{\tilde{m}}$ is an SPE for any δ since \tilde{m} is the only Nash equilibrium of the one-shot game.

For all agents $i > j^*$, playing strategy $s_{\bar{m}}^{m^{(j^*)},i}$ is obviously a best response in any subgame since they play FR in each round. So consider any agent $j \leq j^*$ and any $h \in \mathcal{H}$. This history could have evolved in one of two ways.

Suppose the game has not ended at any given t , and suppose all other agents have been playing the move given by $m^{(j^*)}$ each round. Then by construction of $\delta^{(j^*)}$ it is a best response for j to play HO .

Suppose on the other hand that at some $t \in \{1, 2, \dots\}$, some agent $j \leq j^*$ plays FR . Then by Lemma 2 it is a best response for all agents to abide by the trigger strategy and play FR in all remaining rounds.

Next suppose there is another trigger strategy $s_{\hat{m}}^{m^{(i)}}$ that was an SPE. To see that this could not be observe that for any $i \neq j^*$

$$\delta^{(j^*)} B^i > \delta^{(i)} B^i = B^i - \frac{C}{\#HO(m^{(i)})}.$$

Thus it would be optimal for agent i to defect from this strategy.

It only remains to consider trigger strategies in which a coalition of agents that is not contiguous in the benefit ordering agrees to help out. That is, we must show that a trigger strategy in which the some of the agents who free ride value the public good more than other agents who help out is not an SPE. Let $s_{\hat{m}}^{\hat{m}}$ be any such noncontiguous trigger strategy. Suppose that $\#HO(s_{\hat{m}}^{\hat{m}}) = i$. But we know from the above that $s_{\hat{m}}^{m^{(i)}}$ is not an SPE because it is optimal for i to defect. Let $i' \in N$ be the agent who values the public good the least, but is supposed to help out in the trigger strategy $s_{\hat{m}}^{\hat{m}}$. Then by construction $B^i > B^{i'}$. Therefore,

$$\delta^{(j^*)} B^i > \delta^{(i)} B^i > \delta^{(i)} B^{i'} = B^{i'} - \frac{C}{\#HO(\hat{m})}.$$

Thus it would be optimal for agent i' to defect from this strategy.

■

Intuitively, this is saying is that if the benefits are unequal, and drop off rapidly, then for the critical δ (which induces the shortest game in expectation), only the highest benefiting agents will raise their hands in an SPE. This description of the efficient SPE makes it possible to look at the real world for verification of the model. For example, if parents receive relatively high benefits from the playground, and non-parents receive only a little, it is likely that only the parents will share in the cost in equilibrium.

From a social efficiency, or egalitarian standpoint, it may be important to know how close the benefits of the agents must be to make unanimous hand raising the only Pareto efficient SPE. Below we give a sufficient bound on the rate of decline of the benefit profile for this to be so.

Lemma 6. *Suppose that for all $i \in N$, $B^{\frac{i+1}{B^i}} > \frac{i}{i+1}$, then $s_{\tilde{m}}^{m^{(n)}}$ and $s_{\tilde{m}}^{\tilde{m}}$ are the only SPE of $\Gamma(G^{hr}, \delta^{(n)})$.*

Proof/

In this case $j^* = n$. To see this note that for all $i \in N$:

$$\delta^{(i)} = 1 - \frac{C}{B^i \#HO(m^{(i)})} < 1 - \frac{C}{B^{i+1}(\#HO(m^{(i)}) + 1)} = \delta^{(i+1)}.$$

Then apply Lemma 5.

■

4. Self Enforcing Optimal Urban Planning

Consider the following one-shot game: $G^{up} \equiv \langle N, M, v \rangle$. Here each agent $i \in N$, chooses between two moves: $M^i = \{CC, MB\}$ (Comply with the Code, or Maximize private Benefit). Let $\#CC : M \rightarrow \{0, \dots, n\}$ be a function that gives the number of agents who agree to comply with the code in any given set of moves. Let B^i be the external benefit that agent i receives when any agent undertakes the socially beneficial action (complying with the code), and C^i be the private cost of undertaking this action himself. Then for all $i \in N$, and all $m \in M$,

$$v^i(m) = \begin{cases} B^i \#CC(m) - C^i & \text{if } m^i = CC \\ B^i \#CC(m) & \text{if } m^i = MB. \end{cases}$$

We assume that for all $i \in N$, $C^i > B^i$. This means that the only Nash equilibrium of the one-shot game is for each agent to maximize the private benefit of his property. In fact, this is a dominant strategy equilibrium.

This is a model of an urban planning game. If an agent complies with the code, he generates benefits for himself and his neighbors. However, the private cost of compliance is higher than the private benefit. To fix the idea, we can imagine a group of developers filing plans with the building commission. The random deadline aspect might be generated by uncertainty over when the building commission meets. The game described above is substantially different from the first two. The problem before was to divide the cost of a discrete level of public good. This game is directed toward assuring the provision of an efficient level of a positive externality, or the abatement of a negative externality. The attraction of thinking about this specifically as an urban planning problem is that putting up a development is fundamentally a one-shot proposition characterized by a high probability of delay in finalizing decisions due to bureaucratic procedures.

Let us first consider the case of identical agents. Let

$$\delta^* = \frac{Bn - C}{B(n - 1)}$$

$\tilde{m} \equiv (\tilde{m}^1, \dots, \tilde{m}^n) = (MB, \dots, MB)$, and $\bar{m} \equiv (\bar{m}^1, \dots, \bar{m}^n) = (CC, \dots, CC)$. Lemma 7 says that if all agents are identical, then at the critical δ , everybody undertaking the beneficial action is an efficient SPE which dominates the only other SPE in which no one undertakes the action.

Lemma 7. *Suppose all agents in the game G^{up} are identical. Then $s_{\tilde{m}}^{\tilde{m}}$ and $s_{\bar{m}}^{\bar{m}}$ are the only SPE of $\Gamma(G^{up}, \delta^*)$.*

Proof/

Clearly, $s_{\tilde{m}}^{\tilde{m}}$ is an SPE for any δ since it \tilde{m} is the only Nash equilibrium of the one-shot game.

To see that $s_{\bar{m}}^{\bar{m}}$ is also an SPE consider any $i \in N$ any $h \in \mathcal{H}$. This history could have evolved in one of two ways.

Suppose first that the game has not ended at any given t , that the other agents have been playing \bar{m}^{-i} each round. If i makes the optimal defection from \bar{m} and tries to free ride, then his expected payoff is $\delta^*(n - 1)B$. On the other hand, not defecting

yields a payoff of $Bn - C$ with certainty. But by construction:

$$\delta^* B(n-1) = \frac{Bn - C}{B(n-1)} B(n-1) = Bn - C.$$

Thus, it is a best response for i to play CC since no additional benefit is gained by defecting.

Suppose on the other hand that at some $t \in \{1, 2, \dots\}$, some agent j plays $\hat{m}^j \neq \bar{m}^j$. Then by Lemma 2 it is a best response for all agents to abide by the trigger strategy and play FR in all remaining rounds.

Finally, suppose there was another trigger strategy $s_{\hat{m}}^{\hat{m}}$ that was an SPE. But then $\#CC(\bar{m}) > \#CC(\hat{m})$ and $\delta^* < 1$ so

$$\begin{aligned} \delta^* B(\#CC(\hat{m}) - 1) &= \delta^* B(n-1) - \delta^* B(n - \#CC(\hat{m})) \\ &= Bn - C - \delta^* B(n - \#CC(\hat{m})) > \#CC(\hat{m})B - C. \end{aligned}$$

Thus it would be optimal for every agent to defect from any other trigger strategy.

■

Thus we have:

Theorem 2. *Suppose that all the agents are identical. Then the only SPE payoffs of $\Gamma(G^{up}, \delta^*)$, are: $\{(Bn - C, \dots, Bn - C), (0, \dots, 0)\}$.*

Proof/

By Lemma 7, $s_{\bar{m}}^{\bar{m}}$ and $s_{\tilde{m}}^{\tilde{m}}$ are the only SPE trigger strategies. Then by Lemma 3,

$$\bar{h} \equiv \{h \in \mathcal{H} \mid \forall t \in \{1, 2, \dots\}, \bar{m} \in_t \bar{h}\},$$

and

$$\tilde{h} \equiv \{h \in \mathcal{H} \mid \forall t \in \{1, 2, \dots\}, \tilde{m} \in_t \tilde{h}\},$$

are the only SPE histories. Thus,

$$\{(Bn - C, \dots, Bn - C), (0, \dots, 0)\}$$

are the only SPE payoffs.

■

This result also holds if agents are almost identical, or if at least the ratios of costs to benefits increases sufficiently slowly. It is possible to prove a result similar to Lemma 5 for this game. We will not do so here because the sufficiency condition is less easy to interpret. But it turns out that if we order the agents so that those with a lower index have a lower ratio of costs to benefits then if for all $i \in \{1, \dots, n-1\}$,

$$1 \leq \left(\frac{C^i}{B^i} - \frac{C^{i+1}}{B^{i+1}} \right) + \frac{C^{i+1}}{B^{i+1}}$$

then for

$$\delta = \frac{B^n n - C^n}{B^n (n-1)}$$

the only SPE trigger strategies have all agents abiding by the code or all agents maximizing private benefits. If the cost-benefit ratio increases faster than this, there may also be other SPE trigger strategies.

5. Bertrand Oligopoly

Next we consider the case of constant marginal cost Bertrand oligopolists facing a known demand curve. For simplicity, we will assume that demand is linear, but this is easily generalized. Let the demand be given by

$$Q = \alpha - \beta p,$$

Where $\alpha \geq 0$, and $\beta > 0$. Let C^i be the per unit cost to firm i of making the good.

Now, consider the one-shot game $G^{bo} \equiv \langle N, M, v \rangle$ where for all $i \in N$, $M^i \equiv \mathfrak{R}_+$. Define $LPF : \mathfrak{R}_+^n \rightarrow \{I \mid I \subseteq N\}$:

$$LPF(m) \equiv \{I \subseteq N \mid i \in I \text{ if and only if } \forall j \in N, m^i \leq m^j\}.$$

This correspondence gives the subset of Lowest Priced Firms. The for all $i \in N$ and all $m \in M$ the payoff function is:

$$v^i(m) = \begin{cases} \frac{(\alpha - \beta m^i)(m^i - C^i)}{|LPF(m)|} & \text{if } i \in LPF(m) \\ 0 & \text{otherwise.} \end{cases}$$

Thus, all of the lowest priced firms share the demand equally. Other firms sell nothing. It is well known that the only Nash equilibrium in the case of identical firms is for each firm to price at cost. We restrict attention to this case. Lemma 8 establishes that in equilibrium, all firms will name the same price and this price will be somewhere between cost, and the monopoly price. Thus all firms naming the monopoly price is an efficient SPE that Pareto dominates all other SPE's. Let $\tilde{m} = (C, \dots, C)$.

Lemma 8. *Suppose all firms are identical. Then for $\delta = \frac{1}{n}$ $s_{\tilde{m}}^{\tilde{m}}$ is an SPE of $\Gamma(G^{bo}, \delta)$, if and only if for all $i, j \in N$ $\bar{m}^i = \bar{m}^j$, and $\bar{m} \in [C, \frac{\alpha + \beta C}{2\beta}]$.*

Proof/

First suppose that $\bar{m} = \tilde{m} = (C, \dots, C)$. Then clearly, $s_{\tilde{m}}^{\tilde{m}}$ is an SPE for any δ since it \tilde{m} is the only Nash equilibrium of the one-shot game.

Now suppose that $\bar{m} \in (C, \frac{\alpha + \beta C}{2\beta}]$. To see that $s_{\tilde{m}}^{\tilde{m}}$ is also an SPE consider any $i \in N$ and any $h \in \mathcal{H}$. This history could have evolved in one of two ways.

Suppose first that the game has not ended at any given t , and that the other agents have been playing \bar{m}^{-i} each round. By hypothesis, the equilibrium price is below the monopoly price, $\frac{\alpha + \beta C}{2\beta}$. Thus, optimal defection for any $i \in N$ is to lower the price by ϵ and capture all the demand. His expected payoff in this case is

$$\delta(\alpha - \beta(\bar{m} - \epsilon))(\bar{m} - \epsilon - C).$$

This is because all agents revert to playing C if the game happens not to end in the round that i defects. On the other hand, by not defecting agent i gets a payoff of

$$\frac{(\alpha - \beta\bar{m})(\bar{m} - C)}{n}$$

with certainty. But by construction:

$$\delta(\alpha - \beta(\bar{m} - \epsilon))(\bar{m} - \epsilon - C) < \delta(\alpha - \beta\bar{m})(\bar{m} - C) = \frac{(\alpha - \beta\bar{m})(\bar{m} - C)}{n}.$$

Thus, it is a best response for i to play \bar{m} since no additional benefit is gained by defecting.

Suppose on the other hand that at some $t \in \{1, 2, \dots\}$, some agent j plays $\hat{m}^j \neq \bar{m}^j$. Then by Lemma 2 it is a best response for all agents to abide by the trigger strategy and play $m^i = C$ in all remaining rounds.

Finally, suppose there is another trigger strategy, $s_{\hat{m}}^{\hat{m}}$, that is an SPE but does not satisfy the hypothesis of the Lemma. Note first that if $s_{\hat{m}}^{\hat{m}}$ is an SPE, then for all $i, j \in N$, $\hat{m}^i = \hat{m}^j$. This is because any agent who names a price which is not the lowest receives no profit. Thus for all $\delta > 0$, the optimal defection would necessarily have a positive expected value. Therefore, no agent could name a price above the lowest price in equilibrium. Second, clearly $\hat{m} \geq C$. Otherwise defection would give agents zero profits instead of the negative profits they get from offering to sell below cost. Finally, suppose that $\hat{m} > \frac{\alpha + \beta C}{2\beta}$. In this case, it is optimal to defect to the monopoly price. This gives an expected profit of

$$\delta \left(\alpha - \frac{\alpha + \beta C}{2} \right) \left(\frac{\alpha + \beta C}{2\beta} - C \right).$$

But for all $\hat{m}^i > \frac{\alpha + \beta C}{2\beta}$,

$$\delta \left(\alpha - \frac{\alpha + \beta C}{2} \right) \left(\frac{\alpha + \beta C}{2\beta} - C \right) > \frac{(\alpha - \beta \bar{m}^i)(\bar{m}^i - C)}{n},$$

since it is easy to check that $m = \frac{\alpha + \beta C}{2\beta}$ is profit maximizing.

■

Thus we have:

Theorem 3. *Suppose that all the agents are identical. Then the only SPE payoffs of $\Gamma(G^{bo}, \frac{1}{n})$ are $\left\{ \frac{(\alpha - \beta m)(m - C)}{n}, \dots, \frac{(\alpha - \beta m)(m - C)}{n} \mid m \in [C, \frac{\alpha + \beta C}{2\beta}] \right\}$.*

Proof/

By Lemma 8, $s_{\bar{m}}^{\bar{m}}$ for $\bar{m} \in [C, \frac{\alpha + \beta C}{2\beta}]$ are the only SPE trigger strategies. Then by Lemma 3,

$$\bar{h} \equiv \{h \in \mathcal{H} \mid \forall t \in \{1, 2, \dots\}, \bar{m} \in_t \bar{h}\},$$

for $\bar{m} \in [C, \frac{\alpha + \beta C}{2\beta}]$ are the only SPE histories. Thus

$$\left\{ \frac{(\alpha - \beta m)(m - C)}{n}, \dots, \frac{(\alpha - \beta m)(m - C)}{n} \mid m \in [C, \frac{\alpha + \beta C}{2\beta}] \right\}$$

are the only SPE payoffs.

■

Thus, all the mediator needs to know to choose the appropriate δ is the number of firms in the market. Equal sharing of the monopoly profit is the unique Pareto dominant SPE equilibrium payoff.

It may be possible to sharpen this result by requiring that the defecting firm lower his price by at least a fixed ϵ (instead of the arbitrarily small ϵ here) in order to capture the market. This seems to reduce the equilibrium set to three elements: the monopoly price, the competitive price, and ϵ above the competitive price. The problem is that these last two are also equilibria of the one-shot game. We therefore could not use the lemmata proved in the early sections since the hypothesis that there be only one equilibrium in the one-shot game is not met. Consequently, we do not pursue this further in the current paper.

6. Conclusions

In this paper, we show the usefulness of probabilistic cheap talk as a simple mechanism for resolving a variety of market and coordination failures. A possible criticism of our approach is its requirement that there be an infinite horizon with no bound on the stopping time (we emphasize that although the horizon is infinite, the stopping time is finite with certainty). The same criticism would apply to infinite horizon repeated games. In this regard, we are influenced by Aumann (1959), and, more recently, Rubinstein (1992). We offer the following passage of Aumann, quoted in Rubinstein as commentary:

in the notion of a supergame that will be used in this paper, each superplay consists of an infinite number of plays of the original game G . On the face of it, this would seem to be unrealistic, but actually it is more realistic than the notion in which each superplay consists of a fixed finite (large) number of plays of G Of course when looked at in the large, nobody really expects an infinite number of plays to take place; on the other hand, after each play we do expect that there will be more. A. Tucker has pointed out that this condition is mathematically equivalent to an infinite sequence of plays, so that is what our notion of supergames will consist of.

Rubinstein continues: “By using infinite horizon games we do not assume that the real world is infinite. Taking the view that models are not supposed to be isomorphic with reality, I see the infinitely repeated game model as a tool for analyzing situations where players examine a long term situation without assigning a specific status to the end of the world.”

Though the previous observations are made in the context of *repeated* games, Aumann’s reference to the uncertainty about the possibility of future play translates here into uncertainty about the binding nature of current strategic choices. The “end of the world” in Rubinstein’s comment might easily refer to the actual determination of payoffs. The key element in our model which justifies the infinite horizon is the fact that at any given round of strategy choice agents are aware that they operate in an uncertain world, and due to the exogenous circumstances that determine the arrival of payoffs, they may still have the opportunity to revise their actions.

A second issue involves the partial stationarity restriction on the set of equilibrium strategies. Despite our assertion in the introduction that stationarity is compelling in games with PCT, one may wonder what the implications of non-stationary equilibria would be. We conjecture that the unique implementation results achieved herein would no longer hold. This is based on analysis of the behavior of the set of equilibrium payoffs as the discount factor varies in Stahl (1991), and van Damme (1992) for the infinitely repeated prisoners’ dilemma. In general, there is no value of the discount factor for which the set of Pareto efficient equilibria is a singleton. Hence, we would expect a similar non-uniqueness to extend to PCT games as well. See Chakravorti, B., J. Conley, and B. Taub (1997) for more discussion.

References

- Aumann, R.** (1959): "Acceptable Points in General Cooperative N-Person Games," in *Contributions to the Theory of Games IV*, by R. D. Luce and A. W. Tucher (eds.) Princeton University Press, Princeton NJ.
- Bernheim, D., B. Peleg, and M. Whinston** (1987): "Coalition-Proof Equilibrium," *Journal of Economic Theory*, pp. 1-29.
- Chakravorti, B., J. Conley, and B. Taub** (1993): "On Resolving the One-shot Prisoners' Dilemma through Probabilistic Cheap Talk," Mimeo, University of Illinois and Bellcore.
- Chakravorti, B., J. Conley, and B. Taub** (1996): "On Uniquely Implementing Cooperation in the One-Shot Prisoners' Dilemma," *Economic Theory*, v. 8, pp. 347-66.
- Chakravorti, B., J. Conley, and B. Taub** (1997): "On Uniquely Implementing Cooperation in the Prisoner's Dilemma: Corrigendum," *Economic Theory*, v. 9 pp. 377-8.
- Jackson, M., and H. Moulin** (1992): "Implementing a Public Project and Distributing Its Cost," *Journal of Economic Theory*, pp. 125-140.
- Rubinstein, A.** (1992): "Comments on the Interpretation of Repeated Games," in *Advances in Economic Theory Vol. 1*, by J. J. Laffont (ed) Cambridge University Press, Cambridge.
- Stahl, D.** (1991): "The Graph of the Prisoner's Dilemma Supergame Payoffs as a Function of the Discount Factor," *Games and Economic Behavior*, v. 3 pp. 368-84.
- van Damme, E.** (1991): *Stability and Perfection of Nash Equilibria*. Berlin New York, Springer-Verlag.