

Federal Reserve Bank of Minneapolis
Research Department Staff Report 287

March 2001

Finite Memory and Imperfect Monitoring*

Harold L. Cole

University of California, Los Angeles
and Federal Reserve Bank of Minneapolis

Narayana R. Kocherlakota

University of Minnesota
and Federal Reserve Bank of Minneapolis

ABSTRACT

In this paper, we consider a class of infinitely repeated games with imperfect public monitoring. We look at symmetric perfect public equilibria with memory K : equilibria in which strategies are restricted to depend only on the last K observations of public signals. Define Γ_K to be the set of payoffs of equilibria with memory K . We show that for some parameter settings, $\Gamma_K = \Gamma_\infty$ for sufficiently large K . However, for other parameter settings, we show that not only is $\lim_{K \rightarrow \infty} \Gamma_K \neq \Gamma_\infty$, but that Γ_k is completely degenerate. Moreover, this last result is essentially independent of the discount factor.

*We thank Stephen Morris, Steven Tadelis, and participants of the Yale Theory discussion group for comments and suggestions. The views expressed herein are the those of the authors and not necessarily those of the Federal Reserve Bank of Minneapolis or the Federal Reserve System.

1. Introduction

For the last ten years, the analysis of Abreu, Pearce, and Stacchetti (APS) (1990) has been the foundation for the analysis of repeated games with imperfect public monitoring. They develop a recursive representation for the set of equilibrium payoffs in such games. However, this representation relies crucially on players' being able to use strategies that depend on arbitrarily long histories of past events. For example, they demonstrate in some games, an equilibrium with high initial payoffs for all players involves a particular realization of the public signal's triggering infinite repetition of a stage-game equilibrium.

There are at least two concerns with such equilibria. The first is obvious: is it plausible that players keep track of a long sequence of events? The second arises from recent work by Mailath and Morris (2000). They perturb repeated games with public monitoring by adding a small amount of idiosyncratic noise. They show that strategies that exhibit infinite history dependence are not robust to this type of perturbation. These concerns suggest a natural question: does requiring strategies to exhibit finite history dependence radically change the set of equilibria if we allow the extent of history dependence to be arbitrarily long?

This paper seeks to address this question within a simple repeated game with imperfect public monitoring. We examine the extent to which the set of equilibrium payoffs with infinite-memory strategies is a good approximation to the set of equilibrium payoffs with arbitrarily long finite-memory strategies. (Throughout, we use the terms *memory* and *history dependence* equivalently.) In particular, we look at strongly symmetric perfect public equilibria with memory K : equilibria in which strategies are restricted to depend only on the last K observations of public signals. Define Γ_K to be the set of payoffs of equilibria with memory K . We show that for some specifications of the parameters of the stage game, $\Gamma_K = \Gamma_\infty$ for

sufficiently large K . However, for other specifications of the stage game, we show that not only is $\lim_{K \rightarrow \infty} \Gamma_K \neq \Gamma_\infty$, but that Γ_k is a singleton. Moreover, this last result depends only on the parameters of the stage game, and so is independent of the discount factor.

Our arguments are similar in spirit to those of Bhaskar (1998). He shows that in an overlapping generations economy, with one player in each cohort, there is a unique equilibrium to the Hammond transfer game when players know only a finite number of periods of past play. Our results extend those of Bhaskar to a class of repeated games with imperfect public monitoring, at least for the case of perfect public equilibria.

2. A Class of Games

We describe a class of repeated games with imperfect public monitoring.

A. Stage Game

Consider the following stage game, which is similar to the partnership game considered by Radner, Myerson and Maskin (1986). There are two players. Player 1 and player 2's action sets are both $\{C, D\}$. Player i 's payoffs are given by:

$$y - c, \text{ if } a_i = C$$

$$y, \text{ if } a_i = D$$

The variable y is random, and has support $\{0, 1\}$. The probability distribution of y depends on action choices:

$$\Pr(y_2 = 1 | a_1 = a_2 = C) = p_2$$

$$\Pr(y_2 = 1 | a_i = D, a_j = C) = p_1$$

$$\Pr(y_2 = 1 | a_1 = a_2 = D) = p_0$$

Throughout, we assume that:

$$1 > p_2 > p_1 > p_2 - c > p_0 > p_1 - c$$

These inequalities guarantee that the probability of receiving a good payoff is increasing in the number of players that choose C . They also guarantee that both players playing C is Pareto superior to their playing D , and that both players' playing D is a unique equilibrium of the stage game.

B. Information Structure and Equilibrium

The stage is infinitely repeated; players have a common discount factor δ , $0 < \delta < 1$. We also assume that there is a public randomizing device; specifically, let $\{\theta_t\}_{t=0}^{\infty}$ be a collection of independent random variables, each uniformly distributed on the unit interval. We define $\theta^t = (\theta_0, \dots, \theta_t)$ and $y^t = (y_1, \dots, y_t)$.

We assume that player i 's action choices are unobservable, but the outcome of y is observable to both players. Hence, player i 's history after period t is given by $h_i^t = ((a_{is})_{s=1}^t, y^t, \theta^t)$. The public history after period t is $h^t = (y^t, \theta^t)$. We denote by $y_s(h^t)$ and $\theta_s(h^t)$, $s \leq t$, the realizations of y_s and θ_s in public history h^t . We use the notation (y_s^r, θ_s^r) to represent $(y_t, \theta_t)_{t=s}^r$.

In this world, a strategy for player i is a mapping σ_i from the collection of possible histories for player i into $\{C, D\}$. A *public strategy* σ_i is a strategy which maps any two histories for player i with the same public history into the same action.

Given these notions of strategies, we restrict attention to strongly symmetric public equilibria, in which both players use the same public strategy. Thus, an equilibrium is a public strategy σ such that σ is a player's optimal strategy, given that the other player is using σ .

C. Finite Memory Equilibria

We are interested in exploring equilibria in which the players' strategies are restricted to depend only in a limited way upon histories. A public strategy with memory K is a public strategy such that $\sigma(h^t) = \sigma(h^{t'})$ if $y_{t-s}(h^t) = y_{t-s}(h^{t'})$ and $\theta_{t-s}(h^t) = \theta_{t-s}(h^{t'})$, for all s such that $0 \leq s \leq \min(K, t) - 1$. Thus, the strategy can only depend on (at most) the last K realizations of the public signals. Correspondingly, an equilibrium with memory K is an equilibrium in which the strategy has memory K . (Thus, definitionally, an equilibrium with infinite memory is the same as an equilibrium.)¹

In any equilibrium, all players receive the same expected utility at any stage of the game. We use the notation Γ_K to refer to the set of payoffs delivered by equilibria of memory K . The key propositions that follow are about the question: does $\lim_{K \rightarrow \infty} \Gamma_K = \Gamma_\infty$?

3. Equilibrium Payoffs with Infinite Memory

From APS (1990), we know that Γ_∞ is a closed interval. It is also straightforward to show that the minimax payoff in the stage game is p_0 , which is also an equilibrium payoff in the stage game. Hence, the lower bound of Γ_∞ is given by $v_{\min} \equiv p_0/(1 - \delta)$.

¹Note that in this definition of an equilibrium with memory k , we have not imposed limited recall on the players. Hence, players can contemplate using arbitrary functions of past histories, but choose in equilibrium to use strategies that depend only on the last k lags of the public signal. In contrast, in a game with recall limit k , players can only contemplate using strategies that are measurable with respect to what they have seen in the last k periods. We discuss allowing for bounded recall later in the paper.

What about the upper bound, v_{\max} , of Γ_{∞} ? We know from APS (1990) that if $v_{\max} > v_{\min}$, then the equilibrium that delivers v_{\max} has the form:

$$(1) \quad \begin{aligned} \sigma(h^t) &= D \text{ if for some } s \leq t, y_s(h^t) = 0 \text{ and } \theta_s(h^t) \geq \pi, \\ \sigma(h^t) &= C \text{ otherwise.} \end{aligned}$$

Verbally, we can think of two possible phases in this equilibrium: a “cooperate” phase and a “non-cooperate” phase. Players start in the cooperate phase, and stay there until they observe $y = 0$ and a sufficiently high realization of θ . Then they start playing a permanent “non-cooperate” phase in which they both play D forever. The possibility of switching from the cooperative to the noncooperative phase whenever the outcome is $y = 0$ is in effect a punishment for low output, and this punishment is what induces the players to play C in the cooperation phase even though it is costly.

The continuation payoff in the cooperate phase is v_{\max} , and the continuation payoff in the non-cooperate phase is v_{\min} . Hence, we can see that:

$$(2) \quad v_{\max} = p_2(1 + \delta v_{\max}) + (1 - p_2)[\delta(\pi v_{\max} + (1 - \pi)v_{\min})] - c$$

For the strategy to be an equilibrium, it must be true that in the non-cooperate phase, players prefer to play D rather than deviate to C

$$v_{\min} \geq p_1[1 + \delta v_{\min}] + (1 - p_1)\delta v_{\min} - c$$

but this is satisfied trivially because $(p_1 - c) < p_0$. As well, it must be true that in the cooperate phase, players prefer to play C rather than deviate to D :

$$(3) \quad v_{\max} \geq p_1(1 + \delta v_{\max}) + (1 - p_1)\delta[\pi v_{\max} + (1 - \pi)v_{\min}]$$

Moreover, for v_{\max} to be the maximal element of Γ_{∞} , the latter inequality must be an equality. Otherwise, we can increase π and thereby increase the value of v_{\max} implied by the flow equation (2), without violating the equilibrium requirement (3).

In the strategy supporting v_{\max} , the punishment for realizing a $y = 0$, which we denote by v_{pun} , is given by

$$v_{pun} = \pi v_{\max} + (1 - \pi)v_{\min}.$$

From the above discussion, it follows that, given (p_2, p_1, c) , (v_{\max}, v_{pun}) are the solutions to the two equations:

$$v_{\max} = p_2(1 + \delta v_{\max}) + (1 - p_2)\delta v_{pun} - c$$

$$v_{\max} = p_1(1 + \delta v_{\max}) + (1 - p_1)\delta v_{pun}$$

Since $p_2 - c < p_1$, $v_{pun} < v_{\max}$. Hence, $\Gamma_{\infty} = [v_{\min}, v_{\max}]$ if and only if $v_{pun} \geq v_{\min}$. It is tedious but simple to show that this is equivalent to assuming that $p_2 - p_1 - c + \delta p_1 c - \delta p_0 p_2 + \delta p_1 p_0 > 0$.

When we switch to finite memory equilibria, we will find that the key to generating v_{\max} is being able to credibly threaten to punish low output levels with v_{pun} while respecting the memory constraint on the equilibrium strategies.

4. Equilibrium Payoffs with Finite Memory

With finite memory, the structure of analysis is substantially different. Let \bar{v}_K and \underline{v}_K denote the upper and lower bounds of the set of payoffs when the equilibrium strategies only depend upon at most memory K . Clearly since playing D was an equilibrium of the stage game, playing D forever is an equilibrium with memory K . Hence the lower bound of the

payoff set is unchanged and $\underline{v}_K = p_0/(1 - \delta)$.

However, we can no longer enforce the playing of C during the cooperative phase by threatening to punish low output levels with the possibility of switching to playing D forever. The reason is that switching to playing D forever after some event would involve us keeping track of the fact that the event had occurred for all of the subsequent periods. Instead, with memory K , if some event triggers switching from playing C to playing D , then if that event does not recur within the next K periods, the players switch back to C . Therefore, the restriction to memory K strategies has two effects: First, it ties together the probability of switching from the cooperative phase to the noncooperative phase, and vice versa. Second, it makes the incentive constraints during the noncooperative phase more difficult to satisfy since players might have an incentive to deviate in order to make the switch to the cooperative phase more likely. These two restrictions make it harder to support v_{\max} with a memory K strategy because of the difficulties associated with being able to credibly threaten a payoff of v_{pun} when $y = 0$ while respecting the memory constraint.

To illustrate these points, consider the following strategy:

$$(4) \quad \begin{aligned} \sigma(h^t) &= D \text{ if for some } t - K + 1 \leq s \leq t, y_s(h^t) = 0 \text{ and } \theta_s(h^t) \geq \tilde{\pi}, \\ \sigma(h^t) &= C \text{ otherwise.} \end{aligned}$$

Again, verbally, we can think of two possible phases in this equilibrium: a “cooperate” phase and a “non-cooperate” phase. Players start in the cooperate phase, and stay there until they observe $y = 0$ *and* a sufficiently high realization of θ . Then they start playing a “non-

cooperate” phase in which they both play D which lasts until they have not observed both $y = 0$ and a sufficiently low realization of θ in any of the last K periods.

This strategy is similar to the infinite memory strategy used to support the best equilibrium payoff: (i) high output realizations cause the cooperate phase to be extended for sure, and (ii) low output realizations can cause the players to switch to the noncooperate phase. It differs in that the noncooperative phase is not permanent. This aspect is an inevitable consequence of (i). But a property like (i) is necessary if high payoff levels are to be achieved. Despite property (i), this strategy has the potential to induce cooperative behavior in the same circumstances as the infinite memory strategy since, as K gets large and $\tilde{\pi}$ approaches zero, the continuation payoff from $y = 0$ in the cooperative phase approaches v_{\min} . Hence it would seem to offer the prospect of delivering the appropriate level of v_{pun} by the appropriate choice of $\tilde{\pi}$.

However, with this strategy, players can influence the probability of switching from the noncooperative to the cooperative phase. If the current period’s outcome was $y = 0$ and $\theta \leq \tilde{\pi}$, then the likelihood that they will switch back to being in the cooperative phase $K + 1$ periods from now is $[p_0 + (1 - p_0)(1 - \tilde{\pi})]^K$. However, if one of the players deviated and played C instead of D in the next period, the likelihood of switching back to the cooperative phase in $K + 1$ periods rises to $[p_1 + (1 - p_1)(1 - \tilde{\pi})][p_0 + (1 - p_0)(1 - \tilde{\pi})]^{K-1}$. The possibility of influencing the possibility of switching back to cooperation may induce a deviation at some point during the noncooperative phase and hence undercut the possibility of this strategy credibly threatening v_{\min} . We turn next to showing that this can lead to the set of equilibrium payoffs under finite memory being much more restricted than under infinite memory.

5. Finite Memory: A Non-Convergence Result

Our first result is to show that there exists an open set of parameters such that $\lim_{K \rightarrow \infty} \Gamma_K \neq \Gamma_\infty$. In fact, as the following proposition shows, if p_1 is sufficiently close to p_2 , always playing D is the only equilibrium with memory K , for any finite K .

PROPOSITION 1. *If:*

$$(5) \quad p_1 > 0.5(p_2 + p_0)$$

then $\Gamma_K = \{p_0/(1 - \delta)\}$ for all K .

To understand the logic of this proposition, consider two different histories of length K of the public signal y :

$$(y_{t-1}, \dots, y_{t-K}) = (1, 1, \dots, 1, 1)$$

$$(y'_{t-1}, \dots, y'_{t-K}) = (1, 1, \dots, 1, 0)$$

Trying to support an equilibrium other than always playing D with K memory strategies means that we must be able to generate different outcomes from two histories such as these. For example, in the strategy in (4), it must be a continuation equilibrium for players to play D after the first type of history, and play C after the second type of history. But the difference between these histories vanishes after this period. Hence, the players' continuation payoffs are the *same* function of y_t after both these histories. In order to generate these two different continuation equilibria, we need to be able to choose continuation values v_1 and v_0 so as to make it an equilibrium to choose C or choose D . The essence of the proposition is that if $p_1 > (p_2 + p_0)/2$, this cannot be done.

Proof. The first part is that to show that the only equilibrium with memory 1 is to always choose D for all public histories. The second part is to assume inductively that the only equilibrium with memory $(K - 1)$ is to always choose D . Then, we show that, given an equilibrium with memory K , the equilibrium strategies must be independent of the K th lag of the public signal. Hence, an equilibrium with memory K must be an equilibrium with memory $(K - 1)$, and so, by induction, the only equilibrium with memory K , for any K , is to always choose D .

Part 1: If always playing D is not the only equilibrium with memory 1, then there exists some period t such that $\sigma(h^{t-2}, y_{t-1}, \theta_{t-1}) = C$ and $\sigma(h^{t-2'}, y'_{t-1}, \theta'_{t-1}) = D$. Define $v_{1,t+1}$ to be the expected continuation payoff in period $(t + 1)$ if $y_t = 1$, and $v_{0,t+1}$ to be the expected continuation payoff if $y_t = 0$ (where the expectations are over θ_t).

Then:

$$p_2[1 + \delta v_{1t}] + (1 - p_2)\delta v_{0t} - c \geq p_1[1 + \delta v_{1t}] + (1 - p_1)\delta v_{0t}$$

$$(1 - p_0)\delta v_{0t} + p_0[1 + \delta v_{1t}] \geq p_1[1 + \delta v_{1t}] + (1 - p_1)\delta v_{0t} - c$$

The first inequality guarantees that C is an equilibrium. The second inequality guarantees that D is an equilibrium. Together, these inequalities imply that:

$$(p_2 - c - p_1)/(p_2 - p_1) \geq \delta(v_{0t} - v_{1t}) \geq (p_1 - p_0 - c)/(p_1 - p_0)$$

But this implies that:

$$1 - c/(p_2 - p_1) \geq 1 - c/(p_1 - p_0)$$

or:

$$(p_2 - p_1) \geq (p_1 - p_0)$$

which violates $p_1 > (p_2 + p_0)/2$.

Part 2: Now, we show that in any equilibrium with memory K , the strategies must be independent of the K th lag of the public signals. Suppose not, and:

$$\sigma(y_{t-K}^{t-1}, \theta_{t-K}^{t-1}) = C$$

$$\sigma(y_{t-K}^{t-1'}, \theta_{t-K}^{t-1'}) = D$$

$$(y_{t-K-1}^{t-1}, \theta_{t-K-1}^{t-1}) = (y_{t-K-1}^{t-1'}, \theta_{t-K-1}^{t-1'})$$

where $(y_s^r, \theta_s^r) = (y_i, \theta_i)_{i=s}^r$. Define:

$$v_1 = E_{\theta_t} v_t(y_{t-K}^{t-1}, \theta_{t-K}^{t-1}, y_t = 1, \theta_t)$$

$$v_0 = E_{\theta_t} v_t(y_{t-K}^{t-1'}, \theta_{t-K}^{t-1'}, y_t = 0, \theta_t)$$

It follows that if playing C is weakly preferred to D at history $(y_{t-K}^{t-1}, \theta_{t-K}^{t-1})$, then:

$$p_2(1 + \delta v_1) + (1 - p_2)\delta v_0 - c \geq p_1[1 + \delta v_1] + (1 - p_1)\delta v_0.$$

Similarly, if playing D is weakly preferred to C at history $(y_{t-K}^{t-1'}, \theta_{t-K}^{t-1'})$, then:

$$(1 - p_0)\delta v_0 + p_0\delta v_1 \geq p_1[1 + \delta v_1] + (1 - p_1)\delta v_0 - c.$$

Together, these inequalities imply that $p_2 - p_1 \leq p_1 - p_0$, which is a contradiction.

The proposition then follows inductively. \square

It is worth emphasizing that this proposition holds regardless of the size of δ .

6. Finite Memory: Convergence

We now show that if $p_1 \leq (p_2 + p_0)/2$ and $v_{pun} > v_{\min}$, then there exists K^* such that if $K \geq K^*$, then $\Gamma_K = \Gamma_\infty$.

A. A Convergence Result for Maximal Payoffs

To show this, we first show that for K sufficiently large, we can construct an equilibrium with memory K that has payoff v_{\max} .

To do so, consider the following strategy with memory K . The strategy is of the form:

$$(6) \quad \sigma_K^{\max}(h^t) = D \text{ if } \theta_{t-k^*+1}(h^t) \geq \pi_K, \text{ where}$$

$$k^* = \min \{k \in \{1, \dots, K\} : y_{t-k+1}(h^t) = 0\}$$

$$\sigma_K^{\max}(h^t) = C \text{ otherwise.}$$

Verbally, we can think of two possible phases in this equilibrium: a “cooperate” phase and a “non-cooperate” phase. Players start in the cooperate phase, and stay there until they observe $y = 0$ and a sufficiently high realization of θ . Then they start playing a “non-cooperate” phase in which they both play D which lasts until either they have not observed $y = 0$, or in the most recent period in which $y = 0$ the realization of θ is sufficiently low.

Remark: This strategy is similar to that in (4), though the punishment phase is somewhat less severe for a given level of memory K . This could potentially reduce the set of equilibrium payoffs for a given K relative to that which could be supported by (4). However, since we are interested in the set of equilibrium payoffs as K gets large, and since the payoff in the punishment phase also converges to v_{\min} as $K \rightarrow \infty$ and $\pi_K \rightarrow 0$, this aspect is inessential. We use (6) rather than (4) because of its analytic tractability.

When can we find an equilibrium of this type that delivers payoff v_{\max} ? It must be true that:

$$v_{\max} = p_2(1 + \delta v_{\max}) + (1 - p_2)\delta v'_{pun} - c$$

where v'_{pun} is given by

$$(7) \quad v'_{pun} = \pi_K v_{\max} + (1 - \pi_K) X_K$$

$$(8) \quad X_K = (p_0 + (1 - p_0)\delta v'_{pun}) \sum_{i=1}^{K-1} (p_0\delta)^i + \delta^K p_0^K v_{\max}$$

For this strategy to be a viable equilibrium, we need to verify three things. First, we need to make sure that the players find it weakly optimal to play C in the cooperate phase of the equilibrium. Note, though, from the definition of v_{pun} , that $v_{pun} = v'_{pun}$ and also that:

$$v_{\max} = p_1(1 + \delta v_{\max}) + (1 - p_1)\delta v_{pun}$$

and so in the cooperate phase, players are indifferent between playing C or not.

Second, we need to verify that players are willing to play D in the non-cooperate phase of the equilibrium. Consider a history $h^{t-1} = (y_{t-1}, \dots, y_{t-K})$ in which $y_{t-k^*} = 0$, $\theta_{t-k^*} > \pi$, and $y_{t-k} = 1$ for all $k < k^*$. The equilibrium payoff in this history is determined by the value of k^* , and is equal to X_{K-k^*+1} , where X_j is as defined in (8). Thus, for each k^* if playing D is weakly preferred to C it must be the case that

$$X_{K-k^*+1} \geq p_1(1 + \delta X_{K-k^*}) + (1 - p_1)v_{pun} - c,$$

where $X_0 \equiv v_{\max}$. Note that X_{K-k^*+1} satisfies the recursive equation:

$$X_{K-k^*+1} = p_0(1 + \delta X_{K-k^*}) + (1 - p_0)v_{pun},$$

Subtracting the recursion from the incentive condition yields

$$(p_1 - p_0)(1 + \delta X_{K-k^*} - \delta v_{pun}) \leq c.$$

Thus, to make sure that players want to play D in the non-cooperate phase, we must verify the above inequality for all $k^* \in \{1, \dots, K\}$.

We verify this inequality as follows. We know that:

$$p_2(1 + \delta v_{\max}) + (1 - p_2)\delta v_{pun} - c = p_1(1 + \delta v_{\max}) + (1 - p_1)\delta v_{pun},$$

or, equivalently:

$$\delta(v_{\max} - v_{pun}) = c/(p_2 - p_1) - 1$$

It is trivial to see that X_k is decreasing in k . Hence, it follows that for any k^* :

$$\begin{aligned} & (1 + \delta X_{K-k^*} - \delta v_{pun}) \\ & \leq (1 + \delta X_0 - \delta v_{pun}) \\ & = (1 + \delta v_{\max} - \delta v_{pun}) \text{ (by definition of } X_0) \\ & = c/(p_2 - p_1) \\ & \leq c/(p_1 - p_0) \end{aligned}$$

Thus, because $(p_2 - p_1) \geq (p_1 - p_0)$, players are willing to play D throughout the punishment phase.

Finally, we need to find K so that $0 \leq \pi_K$. Again, X_K is decreasing in K and hence we can conclude that

$$\pi_K = \frac{v_{pun} - X_K}{v_{\max} - X_K}$$

is increasing in K . Furthermore, note that

$$\begin{aligned} \lim_{K \rightarrow \infty} X_K &= [p_0 + (1 - p_0)\delta v_{pun}] / (1 - p_0\delta) \\ &= \frac{(1 - \delta)}{(1 - p_0\delta)} \left(\frac{p_0}{1 - \delta} \right) + \frac{(1 - p_0)\delta}{(1 - p_0\delta)} v_{pun}, \end{aligned}$$

and so $\lim_{K \rightarrow \infty} X_K$ is a convex combination of $p_0/(1 - \delta)$ and v_{pun} . This implies that if $v_{pun} > p_0/(1 - \delta)$, then there exists K^* , such that for all $K \geq K^*$, $X_K < v_{pun}$.

This analysis verifies the following proposition.

PROPOSITION 2. *If $(p_2 + p_0)/2 \geq p_1$, and $v_{pun} > v_{\min}$, then there exists K^* such that for all $K \geq K^*$, the maximal element of Γ_K is v_{\max} .*

The crux of this proposition is that a K -period non-cooperate phase, if $K \geq K^*$, is sufficiently harsh to induce cooperation. Crucially, as long as $(p_2 + p_0)/2 \geq p_1$, players are willing to play non-cooperate.

B. A Convergence Result for the Equilibrium Payoff Set

We have seen that under the conditions of Proposition 2, the maximal element of Γ_K is v_{\max} for sufficiently large K . The minimal element of Γ_K is v_{\min} for any K . But is Γ_K connected or are there holes in Γ_K ? When $K = \infty$, we can use the initial draw of the public randomization device to create any payoff between the endpoints of Γ_∞ . But with finite memory, this permanent randomization between equilibria is no longer possible. In this subsection, we show that when the conditions of proposition 2 are satisfied, $\Gamma_K = \Gamma_\infty$ for K sufficiently large.

Suppose that the conditions of proposition 2 are satisfied, and $K > K^*$. Let $\gamma \in \Gamma_\infty$,

and consider the following specification of strategies. Let $\tau \in \{0, 1, 2, \dots, \infty\}$ be such that

$$\frac{p_0(1 - \delta^\tau)}{1 - \delta} + \delta^\tau v_{\max} \leq \gamma < \frac{p_0(1 - \delta^{\tau+1})}{1 - \delta} + \delta^{\tau+1} v_{\max}.$$

Let π_τ such that

$$\pi_\tau \left[\frac{p_0(1 - \delta^\tau)}{1 - \delta} + \delta^\tau v_{\max} \right] + (1 - \pi_\tau) \left[\frac{p_0(1 - \delta^{\tau+1})}{1 - \delta} + \delta^{\tau+1} v_{\max} \right] = \gamma.$$

Denote the strategy that supports payoff v_{\max} by $\sigma_{K^*}^{\max}$. Then we can define the strategy σ_γ that supports payoff γ as follows.

$$\sigma_\gamma(y^{t-1}, \theta^{t-1}) = D \text{ for all } t < \tau.$$

$$\sigma_\gamma(y^{\tau-1}, \theta^{\tau-1}) = C \text{ if } \theta_{\tau-1} \leq \pi_\tau$$

$$\sigma_\gamma(y^{\tau-1}, \theta^{\tau-1}) = D \text{ if } \theta_{\tau-1} > \pi_\tau$$

For $t \geq \tau$,

$$\sigma_\gamma(y^t, \theta^t) = \sigma_{K^*}^{\max}(y_\tau^t, \theta_\tau^t) \text{ if } \theta_{\tau-1} \leq \pi_\tau$$

$$\sigma_\gamma(y^t, \theta^t) = \sigma_{K^*}^{\max}(y_{\tau+1}^t, \theta_{\tau+1}^t) \text{ if } \theta_{\tau-1} > \pi_\tau$$

The basic idea of this strategy is that the players play D through period $(\tau - 1)$. Then, in period t , they switch to playing σ_{\max} if $\theta_{\tau-1}$ is low. Otherwise, they play D in period t , and switch to playing σ_{\max} in period $(t + 1)$.

By construction, σ_γ delivers payoff γ . We need to verify that σ_γ is indeed an equilibrium with memory K . Note that in the histories in which the strategy specifies that the players choose D , their actions have no effects on future payoffs. Hence, playing D is weakly optimal. Also, we know that $\sigma_{K^*}^{\max}$ is an equilibrium, so that playing according to $\sigma_{K^*}^{\max}$ is weakly optimal whenever the strategy makes this specification.

We still need to verify that σ_γ is a strategy with memory K . Since $\sigma_{K^*}^{\max}$ is a strategy with memory $K^* < K$, it follows that $\sigma_{\max}(y_\tau^t, \theta_\tau^t) = \sigma_{\max}(y_{\tau+1}^t, \theta_{\tau+1}^t)$ for $t \geq (\tau + K^*)$. Thus, the realization of θ_τ does not affect play after period $(\tau + K^*)$, and σ_γ is a strategy with memory $K > K^*$.

7. Discussion

In this section, we discuss the robustness of our results to two different perturbations of the setup. First, we consider what happens if players have bounded recall. Second, we consider whether our results for perfect public equilibrium extend to sequential equilibrium.

A. Bounded Recall

Throughout, we have assumed that players have perfect recall. This means that, while equilibrium strategies are required to be functions of K lags of past history, players can contemplate deviations from equilibrium play that are arbitrary functions of past history. In contrast, if players' recall is limited, then they can only contemplate using strategies that are functions of the last K periods of history.

Let Γ_K^{br} be the set of perfect public equilibrium payoffs when players have recall limit K . Then, we can demonstrate a result analogous to Proposition 1: if $(p_2 + p_0)/2 < p_1$, then $\Gamma_K^{br} = \{p_0/(1 - \delta)\}$ for all K . The proof is identical to that of Proposition 1. Intuitively, in the proof of Proposition 1, we eliminate the possibility of other equilibria by contemplating the possibility of players' deviating to strategies consistent with bounded recall.

We do not have a direct analogy to Proposition 2. However, it is simple to see that $\Gamma_K^{br} \supseteq \Gamma_K$ (because there are fewer possible deviations with recall K). Hence, we know that,

under the assumptions of Proposition 2:

$$\lim_{K \rightarrow \infty} \Gamma_K^{br} \supseteq \Gamma_\infty$$

B. Sequential Equilibrium

Following much of the literature on repeated games with imperfect public monitoring, in this paper we use perfect public equilibrium as the equilibrium concept. In such an equilibrium, players' strategies are a function only of public history. In contrast, in a sequential equilibrium, players' strategies can be arbitrary functions of both public and private history. Define a sequential equilibrium with memory K to be a sequential equilibrium in which a player's equilibrium strategy is a function of K lags of private and public history. Let Γ_K^{se} denote the set of payoffs of strongly symmetric sequential equilibria with memory K . It is straightforward to see that $\Gamma_K^{se} \supseteq \Gamma_K$; we know from APS (1990) that $\Gamma_\infty^{se} = \Gamma_\infty$.

Proposition 1 does not extend to sequential equilibria: rather remarkably, $\Gamma_1^{se} = \Gamma_\infty$ for any parameter settings.² Consider any element of Γ_∞ . From results in APS(1990), we know that it can be supported as an equilibrium payoff by a public strategy of the form:

$$\begin{aligned} \sigma(y^t, \theta^t, (a_i)^t) &= D \text{ if } \theta_0 \geq \pi \text{ OR} \\ &\quad \text{if } \theta_t \geq \pi \text{ and } y_t = 0 \\ &= C \text{ otherwise} \end{aligned}$$

Then, consider the following strategy:

$$\sigma^*(y^t, \theta^t, (a_i)^t) = D \text{ if } a_{it} = D \text{ OR } y_t = 0 \text{ and } \theta_t \geq \pi$$

²We thank Stephen Morris for pointing this out to us.

$$= C \text{ otherwise}$$

Note that σ^* is a strategy with memory 1. Our goal is to show that σ^* is a strongly symmetric sequential equilibrium with the same payoff as σ .

To see this, define the following function:

$$\Phi(y^t, \theta^t; \sigma) = \sigma(y^t, \theta^t, \sigma(y^{t-1}, \theta^{t-1}))$$

to be a player's actions as a function of the public history, given he uses σ . The definition of σ^* guarantees that:

$$\Phi(y^t, \theta^t; \sigma) = \Phi(y^t, \theta^t; \sigma^*)$$

This means that whether player i plays according to σ or σ^* , his actions are the same function of public history. But this means that if player i uses σ^* or σ , player j is indifferent between using σ and σ^* . Since σ is a best response to σ at any history, σ^* is a best response to σ^* at any history, and σ^* is a sequential equilibrium.

The trick here is that a player's past action serves as a summary statistic that encodes whether game play is in a "cooperate" or "non-cooperate" phase. Implicitly, one lag of private actions encodes the relevant portion of the full public history.

References

- [1] Abreu, D., Pearce, D., and Stacchetti, E., 1990, Toward a theory of discounted repeated games with imperfect monitoring, *Econometrica* 58, 1041-1063.
- [2] Bhaskar, V., 1998, Informational constraints and the overlapping generations model: Folk and Anti-Folk theorems, *Review of Economic Studies* 65, 135-149.
- [3] Mailath, G., and Morris, S., 2000, Repeated games with almost-perfect monitoring, University of Pennsylvania working paper.
- [4] Radner, R., Myerson, R., and Maskin, E., 1986, An example of a repeated partnership game with discounting and uniformly inefficient equilibria, *Review of Economic Studies* 53, 59-70.