Limits on Cooperation with Anonymity and Noise

Eric J. Friedman*
Department of Economics, Rutgers University
New Brunswick, NJ 08903.

December 9, 1999

Abstract

We show that social equilibria in the prisoner's dilemma, with anonymous random matching, are not robust to finite probability trembles with large populations. In particular, for reasonable parameter values, defection in every period is the unique social equilibrium; in fact, any other strategy is strictly dominated in the infinitely repeated matching game and thus noncooperation is the only outcome of a rationalizable strategy. We also provide estimates and specific bounds on the parameters under which these results hold.

1 Introduction

Anonymous random matching games provide the most basic models of societal interaction. In particular, the repeated anonymously matched prisoner's dilemma provides a simple model of cooperative behavior. Much work has been done to delineate the conditions under which cooperation will arise (see e.g., Rosenthal 1979, Bendor and Mookherjee 1987, Kandori 1992, Milgrom, North and Weingast 1990, Ellison 1994) or evolve (see e.g., Axelrod 1984, Young 1993, Kandori, Mailath and Rob 1993, Ellison 1993). Social equilibria are also often considered as the basis for social conventions, i.e., only social conventions which are equilibria

*I would like to thank Doug Blair, Rich Mclean and Paul Resnick for helpful conversations and Caltech for its hospitality during part of this work. email: friedman@econ.rutgers.edu, www: http://econ.rutgers.edu/home/friedman

of the repeated matching games can exist in the long run. (We call such equilibria "social equilibria" to distinguish them from stage game equilibria.)

For the prisoner's dilemma the standard analysis is particularly troubling. Although it seems intuitively clear that cooperation won't arise in large populations when players are anonymous, Kandori (1992) showed that for fixed population size and sufficiently low discounting that cooperation can arise as a social equilibrium of "contagion strategies." However, these strategies are quite unappealing aesthetically and appear to be extremely fragile. Kandori (1992) argues that they are unreasonable in the following sense: fix the discount factor, then for a sufficiently large population the contagion strategies will not be an equilibrium. (Similar arguments are in Bendor and Mookherjee 1987, and Milgrom, North and Weingast 1990)

Surprisingly, Ellison (1994) shows that this argument is often not relevant, since for reasonable values of the discount factor and fairly large populations the contagion strategies give rise to cooperative equilibria. He even shows that these equilibria are robust, in a certain sense, against noise (trembles). Thus, cooperation appears to be a reasonable social equilibria of the anonymously repeated game.

In this paper we will argue that the key problem with contagion strategies is the lack of robustness with respect to noise. Rather than considering noise in the limit as tremble probability goes to zero, we consider nonvanishing tremble probabilities.[1] Then for reason-

[1]Note that even in evolutionary analyses, the results can change dramatically when trembles are included. For example, the well known and robust dominance of "tit-for-tat" in evolutionary games (Axelrod 1984) is no longer true when players tremble; instead a new strategy "win-stay lose-shift" becomes important (Nowak and Sigmund, 1993).

able parameter values we show that cooperation can not arise as social equilibrium. In fact, our result is significantly stronger: we show that the strategy "defect in every period" is a strictly dominant strategy *in the repeated game.*

Our result is strong in another sense; we allow for unlimited truthful information transfer (a public record of all previous play). Thus, even when information transfer is easy, cooperation can not arise, showing that the key impediments to cooperation are anonymity and noise, as opposed to information transfer which has typically been the focus of study.

The paper is organized as follows: In Section 2 we present the model and main result. Section 3 provides some numerical estimates, and Section 4 discusses the effects of information transfer. We conclude in Section 5 with a brief survey of generalizations of our analysis.

2 Trembles and Noncooperation

The model that we study is the one used by Kandori (1992) and Ellison (1994) which was based on that of Rosenthal (1979). There are M players, where M is even, who play an infinitely repeated anonymous random matching game: in each period, players are matched at random to play the one-shot prisoner's dilemma. For simplicity we use the following payoff matrix

	C	D
C	$1, 1$	-1.2
D	$2, -1$	$0, 0$

although our qualitative results would be unchanged for any prisoner's dilemma. Each player's total payoff is the discounted sum of their stage game payoffs with discount factor δ.

Following Ellison (1994), we assume that players tremble with probability ϵ, i.e., when a player attempts to choose action C (resp. D) then with probability ϵ they tremble and choose action D (resp. C). We will assume $\epsilon < .5$ throughout the paper to simplify the exposition. We also include a public signal q^t which is i.i.d. uniform on $[0, 1]$.

Lastly, in contrast to these other papers, we assume that the history of play is common knowledge – the actions chosen in every previous stage is publicly observable, e.g., the players know the number of defections in every previous period, but not who chose them.[2] Note that this assumption makes cooperative equilibria more likely and so our nonexistence result in this setting implies nonexistence in games where there is no public record. (In both settings, players know the actions of their matched opponents in previous games.) Let the history available to player i at time t be denoted $h_i^t \in H_i^t$. Formally, a strategy for player s_i is a mapping from histories, h_i^t to mixed strategies over $A_i = \{C, D\}$ and the payoff to that player is given by $U_i(s) = E[\sum_{t=0}^{\infty} \delta^t \pi_i^t \mid s]$ where π_i^t is the payoff to player i in period t.

Our result is based on the following observation: Consider any strategy s_i for player i which chooses a pure strategy C after history h_i^t. Define its mirror strategy to be $\tau(s_i; h^t)$ which is identical to strategy s_i except that on history h_i^t player i plays action D and behaves in the future as if she had played C.

[2]This paper was motivated by a study of cooperation on the Internet (Friedman and Resnick, 1998) in a setting in which players can choose whether or not to be anonymous. In that setting, it is quite common for play histories to be public knowledge since the computer often keeps a complete record of the players previous actions.

Lemma 1 *Choose any $\epsilon > 0$ and $\delta \in (0,1)$. For any opposing strategy vector s_{-i}, and strategy s_i which chooses C after history h_i^t*

$$U_i(\tau(s_i;h^t), s_{-i}) - U_i(s_i, s_{-i}) \geq Pr(h_i^t)[1 - \frac{\delta}{1-\delta} f(M,\epsilon)]$$

where $|f(M,\epsilon)| \leq c(1/\sqrt{M\epsilon} + \epsilon)$ for some $c > 0$, independent of h_i^t and s_{-i}. (Note that $Pr(h_i^t)$, which is the probability of history h_i^t being attained, is the same for (s_i, s_{-i}) and $(\tau(s_i;h^t), s_{-i})$ and that $Pr(h_i^t)$ is strictly greater than 0 for $\epsilon > 0$.)

Proof: This statement follows from the observation that when $M\epsilon$ is large, the distribution of the number of trembles in any period conditional on player i deviating and choosing D is approximately the same as its unconditional distribution, due to the central limit theorem. This implies that the expected future payoff conditional on a deviation is close to its unconditional value. The precise calculation is done explicitly in Friedman and Resnick (1998, Lemma 1); this idea was used in a similar context by Sabourian (1990). $\diamond$

While this lemma suffices to prove our main theorem, it leads to a bound on c which is quite loose, and in particular depends on δ inaccurately. Thus if we are interested in knowing whether cooperation is possible for a specific set of parameters the results will often not be useful. In the next section, we present a more precise method of analysis which will allow us to compute fairly tight bounds on the parameters under which cooperation is possible in social equilibria; however, those results only apply to equilibria and not more general models of rationality or learning.

Using this lemma, it easy to see that for $M\epsilon$ sufficiently large, and ϵ sufficiently small any strategy s_i which plays C after some history h_i^t is strictly dominated by its mirror strategy

$\tau(s_i; h^t)$. However, if we allow mixed strategies, then the result is not as straightforward. For example, consider some strategy, s_i which after some history h_i^t chooses mixed action $(\rho C + (1-\rho)D)$ with $\rho > 0$, then for some choice of M, ϵ this strategy is dominated by the strategy which chooses D with probability 1; however the choice of M, ϵ depends on h_i^t and thus we can not show that all strategies are dominated for fixed M, ϵ and ρ, and using this method of proof, we can only show that such strategies are approximately dominated. However, a more careful analysis shows that this strategy is strictly dominated.

Theorem 1 *There exist constants $\bar{\epsilon}, \beta > 0$ such that for $\epsilon \leq \bar{\epsilon}$ and $M\epsilon \geq \beta$, the strategy "always choose D" is strictly dominant in the repeated game.*

Proof: This is true because any strategy that is not "always choose D" is strictly dominated by a strategy which is made up of a piece of "always choose D". The key insight for proving strict domination for fixed M, ϵ is to notice that the strategy s_i can be reinterpreted as the mixed strategy $\rho s_i^C + (1-\rho)s_i^D$ where s_i^C (resp. s_i^D) is the exactly the same as s_i except that after history h_i^t it chooses C (resp. D) with probability 1. Now consider the strategy $\rho\tau(s_i^C; h^t) + (1-\rho)s_i^D$. Since, for choices of parameters defined by the previous Lemma, $\tau(s_i^C; h^t)$ strictly dominates s_i^C, the convex combination $\rho\tau(s_i^C; h^t) + (1-\rho)s_i^D$ strictly dominates $\rho s_i^C + (1-\rho)s_i^D$. Since we have shown that every strategy except "always choose D" is strictly dominated by some strategy, completeness of the ordering induced by strict domination implies that "always choose D" is strictly dominant. $\Box$

This immediately implies the lack of cooperative social equilibria:

Corollary 1 *For any $\delta < 1$ there exist constants $\bar{\epsilon}, \beta > 0$ such that for $\epsilon \leq \bar{\epsilon}$ and $M\epsilon \geq$*

β, the strategy, "always choose D" for every player, is the unique social equilibrium, i.e., sequential equilibrium of the repeated game.

However, the theorem is significantly stronger as it does not rely on any equilibrium assumptions. It implies no cooperation under a wide variety of assumptions ranging from pure rationality, to adaptive learning to evolutionary dynamics.

Corollary 2 *For any $\delta < 1$ there exist constants $\bar{\epsilon}, \beta > 0$ such that for $\epsilon \geq \bar{\epsilon}$ and $M\epsilon \geq \beta$, the strategy, "always choose D" is the unique rationalizable strategy of the repeated game.*

3 Numerical Estimates

In the previous section we showed that for $M\epsilon$ sufficiently large cooperation can not arise in social equilibria. [3] In this section we address the question: "how large?." This is important since the relevance of our analysis depends on whether the conditions under which it holds are "reasonable".

3.1 A lower bound

First, we wish to compute a lower bound. To do this we consider an extremely simple set of strategies, which are similar to those studied by in games of imperfect public information (Green and Porter 1984) and have been used by Bendor and Mookherjee (1987) to analyze cooperation, among others.

This strategy is a very simple punishment strategy, which begins in the normal phase in which players always cooperate. However, if any player deviates we begin a punishment

[3] Note that ϵ must be sufficiently small, but this constraint is quite simple to deal with, e.g., assuming $\epsilon \leq .1$ is sufficient.

phase in which the players all choose defect. The punishment phase lasts until the first period in which $q^t > \hat{q}$, when a normal phase begins immediately.

The key aspect of this strategy, is that the punishment phase is used to deter players from deviating during a normal phase. This incentive can be increased by choosing larger values of $\hat{q}$, but this in turn increases the severity of the punishment and lowers average payoffs.

If players follow this strategy then the expected (normalized) payoffs to each player are

$$V = \frac{1 - \hat{q}\delta}{1 - \hat{q}\delta P_0} + \epsilon \frac{\hat{q}\delta(3 - P_0) - 2}{1 - \hat{q}\delta P_0},$$

where $P_0 = (1 - \epsilon)^M$ is the probability that no player accidentally deviates in a specific period[4]. Similarly, the expected payoffs conditional on starting in a punishment phase are

$$W = \frac{(1 - \delta)(1 - \hat{q})}{1 - \hat{q}\delta} V.$$

The immediate gain from deviating during a cooperative phase is $1 - \epsilon$, while the future loss is $(V - W)/(1 - \delta)$ *only if no other player deviated that period.* Thus the expected loss is $P_1(V - W)/(1 - \delta)$ where $P_1 = (1 - \epsilon)^{M-1}$ which is the probability that no other player deviated. Thus for this to be an equilibrium we must have

$$P_1(V - W)/(1 - \delta) \geq 1 - 2\epsilon.$$

The maximal payoff arises for the smallest value of $\hat{q}$ for which the strategy is an equilibrium. This gives

$$\hat{q}_{\min} = \frac{(1 - \epsilon)^{2-M}}{\delta(2 - 5\epsilon + \epsilon^2)}$$

[4]The details of this calculation are available from the author. Several similar computations are described in detail in Friedman (1999).

Since, $\hat{q} \leq 1$ this implies that the strategy is an equilibrium for

$$M \leq M_{\max} = -\ln(\frac{(2 - 5\,\epsilon + \epsilon^2)\,\delta}{1 - 2\,\epsilon + \epsilon^2})\,(\ln(1-\epsilon))^{-1}\,.$$

For example when $\epsilon = .01$ and $\delta = .95$ this implies that this is an equilibrium for $M < 63.35$.

Thus, we have proven the following:

Theorem 2 *If* $\hat{q}_{\min}(M, \epsilon, \delta) \leq 1$, *then there exists a social equilibria in which there is at least one intended cooperation.*

Note that even for $M \leq M_{\max}$, when the strategy is an equilibrium, there is a loss in payoffs due to the punishments. Plugging $\hat{q}_{\min}$ into the formula for V yields $V_{\max}$. For $\delta = .95, \epsilon = .01$ we see that for $M = 6$ the normalized expected payoff for each player is .92 and clearly most play is cooperative, but for $M = 40$ the payoff is .5 and so even though the strategies form an equilibrium, approximately half of the periods are punishment periods. Additionally for $M = 63$ the expected payoff drops to .11 while for any larger M the strategies no longer form an equilibrium.

When ϵ is small, we can approximate $M_{\max}$ by its Laurent expansion in ϵ; this yields the formula $M_{\max} \approx \ln(2\delta)/\epsilon$. Thus, to lowest order in ϵ the equilibrium condition is

$$M\epsilon \leq \ln(2\delta)$$

and we see that the key parameter in determining the possibility of cooperation is $M\epsilon$, the average number of total trembles per period. As we will see in the next section, number is important for the set of equilibria in general. Note that this bound is relatively insensitive to the discount factor, when it is close to 1, whereas in Lemma 1 and Theorem 1 the bound is proportional to $\delta/(1-\delta)$ which is extremely sensitve to changes in δ when δ is close to 1. Thus, we now consider a sharpening of Lemma 1 to remove this inconsistency.

3.2 An upper bound

In this section we will construct a an upper bound on the population size for which cooperation is an equilibrium. We will use an idea which is very closely related to regeneration (Abreu, Pearce and Stachetti, 1990) to construct a linear program from which will determine these bounds, both analytically and numerically. However, we will need to make an assumption which is similar in spirit, but weaker, than that traditionally made in this literature, a variation of public perfect equilibria.[5]

Consider the history h_i^t. For each period s with $s < t$, the information can be summarized by four numbers. The first two are public: N_c, the number of cooperations in period s, and N_{cc} the number of games in which both players cooperated. The remaining are private: e_s and e_m which are binary variables representing whether the player, or her matched opponent cooperated. We will argue below that the only one of these which should be necessary to consider to bound the set of equilibrium is N_c as the others should not affect the set of equilibrium since they are not payoff relevant for a related game which has the same set of equilibria as the one we are considering. This motivates the following definition:

Definition 1 *An equilibrium is transformed public perfect (TPP) if all strategies only depend on the history of N_c and not on N_{cc}, e_s and e_m.*

[5] See, e.g. Fudenberg and Tirole (1991).

As we shall see below, these assumptions are apparently much milder than the usual assumption of public perfect equilibria because of the simple structure of our game. The analysis for non TPP equilibria is possible but much more complex (see, e.g., Amarante (1998) for details). We conjecture that the bounds we compute for TPP equilibria actually do hold for any arbitrary equilibria, but have been unable to prove this intuitive result.[6]

3.2.1 The transformed game

Consider the modified stage game with the same set of actions. Given an action vector a^t, where $a_j^t = 1$ if player j chooses C and 0 if she chooses D. Let $x^t = \sum_j a_j^t$, the number of players who chose action C. Let y^t be a random M-vector of 0's and 1's generated i.i.d. where the probability of a 1 in any component is ϵ. Now let z^t be the public signal at time t which is given by $z^t = [\sum_j g(a_j^t, y_j^t)]/M$ where $g(a_j^t, 0) = a_j^t$ and $g(a_j^t, 1) = 1 - a_j^t$.

Define payoffs as follows:

$$\tilde{\pi}_i(a_j^t, z^t) = (1-\epsilon)\frac{M-1}{M}\pi_i(b_\epsilon(a_j^t), z^t) - \frac{1}{M}\pi_i(b_\epsilon(a_j^t), b_\epsilon(a_j^t))$$

where $b_\epsilon(a_j^t)$ is $1-\epsilon$ if $a_j^t = 1$ and ϵ otherwise, and $\pi_i(\alpha, \beta)$ is the payoff to player i in the prisoner's dilemma if player i plays C with probability α and her opponent plays C with probability β.

Now, note that we have constructed this game so that it has the same expected payoffs as our original game. If the players in this modified game only observe the payoff relevant quantities, their own action a_j^t and the the public signal z^t, Then the set of sequential

[6]The complexity arises because the signals N_{cc}, e_s, e_m are slightly correlated when conditioned on N_c. If the correlations were 0 then it would be straightforward to show that the bounds for TPP equilibria corresponded to the general bounds.

equilibria expected payoffs of this game are identical to those in the TPP equilibria of our original game.

The fully general equilibria of our original game corresond to the following modifications of the information available in this game. Let $f : M \to M$ be a matching, $\sigma(j) \neq j$ and $sigma(sigma(j)) = j$ then assume that player j observes both y_j^t and $y_{\sigma(j)}^t$ and the number of pairs $(y_j^t, y_{sigma(j)}^t)$ for which both elements are 1. Conditional on z^t, these signals are not payoff relevant. All they provide is an extremely weak signal on which to coordinate; however, we already have provided a rich signal on which to coordinate, q^t. Thus, as discussed earlier, we conjecture that these signals do not alter our characterizations of the equilibrium payoffs.

3.2.2 Computation of the bound

Given these assumptions we now compute a bound on the amount of cooperation which is possible in any symmetric TPP equilibrium. Let v be the $(1-\delta)$ times the maximal expected average payoff over all symmetric TPP equilibria of this game. Clearly such an equilibrium has all players choosing C in the first period. Also let w be $(1-\delta)$ times the expected average payoff for the same equilibrium beginning in the second period. By maximality of v, $w \leq v$, otherwise we could just start in period 2 and increase v.

Fixing M, ϵ, let $p_i = C_M^i \epsilon^i (1-\epsilon)^{M-i}$, the probability that there are i mistakes in the first period, for $1 \leq i \leq M$, conditional on no player intentionally deviating and $p_i' = C_{M-1}^{i-1} \epsilon^{i-1} (1-\epsilon)^{M-i}$, which is the probability conditional on a single player deviating. Define w_i to be the value of w conditional on i players choosing D in the first period, and note that

$w = \sum_{i=0}^{M} p_i w_i$. Define $w' = \sum_{i=0}^{M} p'_i w_i$.

To deter players from intentionally deviating in the first period it must be true that $(1-\delta) \leq \delta(w - w')$. Thus, we can solve the following linear program to find an upper bound on v.

$$(P) \quad \max v$$

$$s.t.$$

$$(1-\delta) \leq \delta \sum_{i=0}^{M} (p_i - p'_i) w_i$$

$$\forall j : \quad 0 \leq w_j \leq v$$

$$v = (1-\delta)(1-\epsilon) + \delta \sum_{i=0}^{M} p_i w_i$$

Theorem 3 *If P is infeasible then the only symmetric TPP equilibrium consists of players intending to choose D in every period; while if P is feasible then its maximum is an upper bound to the value of the maximal symmetric TPP equilibrium.*

Thus, using this result, we can (numerically) compute an upper bound for the amount of cooperation explicitly as a function of parameters. For example, when $M = 10$ some numerical results are summarized in Table 1.

Note that the dependence on δ (and other parameters) is approximately the same as for the trigger strategies.[7]

[7] When the largest value of ϵ for which the linear program is feasible is plotted against $\log(\delta)$ the resulting graph is a straight line, which is the same dependence as found in the trigger strategy equilibrium.

δ	ϵ	$v_{\max}$
.95	.01	.94
.95	.1	.46
.95	.12	.30
.95	.13	smallest infeasible
.99	.1	.60
.99	.20	smallest infeasible
.999	.28	smallest infeasible
.9999	.33	smallest infeasible

Table 1: Upper bounds on cooperation.

However, the number of variables in this linear program is approximately the same as the population size, M, which makes it difficult to solve when M is large. We now present an analytic bound to determine when this linear program is feasible and cooperation possible in an equilibrium.

Theorem 4 *There can not be any intended cooperations in any symmetric TPP equilibrium if*

$$1/\delta \geq (2P_{\lfloor 2M\epsilon \rfloor} - P'_{\lfloor 2M\epsilon \rfloor}),$$

where $P_i = \sum_{j=0}^{i} p_j$ *and* $P'_i = \sum_{j=1}^{i} p'_j$.

Using this formula we can directly compute estimates on the critical population size for any δ and ϵ without having to solve a linear program with possibly many variables. Some examples are given in Table 2, where we compute the smallest value of δ for which cooperation could arise when M=10,000:

Thus, the dependence on δ is not particularly significant.

$M\epsilon$	$\delta_{\max}$
0.5	.82
0.8	.86
1	.91
2	.96
5	.9956
10	.9997

Table 2: Discount factors for which cooperation can not arise.

4 Information Transfer

Most previous work on this subject has focussed on information transfer as the key impediment to cooperation (e.g., Bendor and Mookherjee 1987, and Milgrom, North and Weingast 1990, Kandori 1992), while in this paper we concentrated on noise (trembles) as the key consideration. This is for two reasons: the first is that with the development of computer networks and information systems information is becoming readily available in many settings. For example, on the Internet, explicit reputation devices are becoming quite common. These range from newsgroup archives, which keep a history of all postings, to explicit reputation databases on auction sites (ebay.com) and gaming sites (fics.org). (See, e.g., Friedman and Resnick 1998 and Kollock 1998 for recent discussions.)

The second reason, is that information transfer arguments are often not relevant as they rely upon limiting cases with unrealistic population sizes, as shown by Ellison (1992). This is because information transfer is actually quite rapid.[8] This implies that while the contagion strategies cannot form an equilibrium for fixed δ as $M \to \infty$, the value of M required is

[8]The amount of time needed for a contagion to spread is on the order of $\log(M)$, which is not particularly large even for populations in the millions.

M	$\delta^{\dagger}_{\max}$	$\epsilon^{\ddagger}_{\min}$ ($\delta = .95$)	$\epsilon^{\ddagger}_{\min}$ ($\delta = .99$)
10	.79	.14	.26
100	.89	.017	.037
1000	.93	.0017	.0041
10^{12}	.99	1.7×10^{-12}	4.1×10^{-12}

Table 3: Barriers to cooperation. (†From Ellison (1994) Table 1. ‡Upper bound from Theorem 4.)

unreasonably large. However, the effect of trembles only relies on the quantity $M\epsilon$ being large, where large here means $O(1)$ and thus even small amounts of noise can have important effects. These are summarized in Table 3 below, where we compare the parameter values for which cooperation can arise.

Most striking is that even with a trillion players, when there are no trembles, cooperation can occur in an equilibrium when $\delta > .99$, which could arise from quarterly interactions when the discount rate is 4%. However, such cooperation can not occur, even with extensive public information, when $\epsilon > 4.1 \times 10^{-12}$, an extremely small number. Thus for this very large population, information transfer imposes a weak restriction on cooperation while trembles provide an apparently insurmountable barrier to cooperation.[9]

5 Extensions

Note that the analysis in this paper can be used to show that for general stage games, which may have many players and be asymmetric, the only sequential equilibria of random matching games are the Nash equilibria of the stage games. An example is the 2 player market game,

[9]Ellison (1994) discusses the idea that finite trembles or small numbers of irrational players might obstruct cooperation, but does not provide any numerical estimates.

where one player is the seller and the other is the buyer and players are divided into two classes, buyers and sellers. A natural generalization of this arises in informal labor markets where each hirer hires a number of laborers. In this case a natural model is one in which there are a set of indistinguishable laborers and another set of indistinguishable hirers. The games played could even vary from period to period. Nonetheless, a suitably modified version of Lemma 1 applies.

A version of Lemma 1 even holds for stage games with an infinite (countable or uncountable) number of actions, subject to mild regularity assumptions. However, in this case the lemma is only approximately true for most players. Nonetheless it still leads to the result that approximate Nash equilibria of the stage game will be played in every period.[10]

It is also clear that these results hold in finitely repeated random matching games, and thus provide some justification for the use of random matching methods used in experiments.

Lastly, we note that while related results formally hold when the stage games are in extensive form, the analysis in such games is much more subtle, and the intuitions from this analysis only apply in special cases, but not in general. This is discussed in detail in Friedman (1999).

A Appendix: Proof of Theorem 4

First we consider an equivalent version of the LP given in the text, without the explicit use of the variable v:

[10]This result closely parallels the main lemma in Fudenberg, Levine and Pesendorfer (1998), in a different setting.

$$(P) \quad \max(1-\delta)(1-\epsilon)+\delta\sum_{i=0}^{M} p_i w_i$$

$$s.t.$$

$$(1-\delta) \leq \delta\sum_{i=0}^{M}(p_i - p_i')w_i$$

$$\forall j: \quad 0 \leq w_j \leq (1-\delta)(1-\epsilon)+\delta\sum_{i=0}^{M} p_i w_i$$

The feasible region for this LP is determined by the inequalities:

$$1 \leq \frac{\delta}{1-\delta}\sum_{i=0}^{M}(p_i - p_i')w_i$$

$$0 \leq w_i \leq (1-\delta)(1-\epsilon)+\delta\sum_{i=0}^{M} p_i w_i$$

Using duality theory of linear programming (Luenberger, 1984), these constraints will be infeasible if the following linear program is unbounded:

$$(D) \quad \max \frac{1-\delta}{\delta}\alpha - \frac{1-\delta}{\delta}\beta$$

$$s.t.$$

$$\forall i: \quad \beta_i/\delta \geq (p_i - p_i')\alpha + p_i\sum_{i=0}^{M}\beta_i$$

$$\forall i: \quad 0 \leq \beta_i$$

The following linear program has the same optimality properties as the one above and is simpler to analyze:

$$(D) \quad \max \alpha - (1-\epsilon)\beta$$

$$s.t.$$

$$\beta \geq \sum_{i=0}^{M} \beta_i$$

$$\forall i: \quad \beta_i/\delta \geq (p_i - p'_i)\alpha + p_i\beta$$

$$\forall i: \quad 0 \leq \beta_i$$

Consider a solution of the form $\beta = \lambda$ and $\alpha = \lambda(1+\gamma)$ with $\gamma > 0$. This will have a payoff of $\lambda\gamma$. Thus, (D) will be unbounded (and thus (P) infeasible) if the following constraints can be satisfied for unbounded $\lambda > 0$:

$$\lambda \geq \sum_{i=0}^{M} \beta_i$$

$$\forall i: \quad \beta_i \geq \delta(2p_i - p'_i + \gamma(p_i - p'_i))\lambda$$

$$\forall i: \quad 0 \leq \beta_i$$

The second constraint is equivalent to

$$\forall i: \quad \beta_i \geq \delta p_i(2 + \gamma - \frac{i(1+\gamma)}{M\epsilon})\lambda.$$

For $i \geq (2+\gamma)M\epsilon/(1+\gamma)$ we set $\beta_i = 0$ while for $i < (2+\gamma)M\epsilon/(1+\gamma)$ we let the constraint bind, so $\beta_i = \delta p_i(2+\gamma - \frac{i(1+\gamma)}{M\epsilon})\lambda$. Thus, for these choices we see that (D) will be unbounded if

$$1 > \delta \sum_{i=0}^{\lfloor 2M\epsilon \rfloor} (2p_i - p'_i),$$

where γ is chosen to be sufficiently small. Letting $P_i = \sum_{j=0}^{i} p_j$ and $P'_i = \sum_{j=1}^{i} p'_j$ completes the proof.

References

[1] D. Abreu, D. Pearce, and E. Stachetti. Toward a theory of discounted repeated games with imperfect monitoring. *Journal of Economic Theory*, 58:1041–1064, 1990.

[2] M. Amarante. A recursive structure for games with private monitoring. mimeo, 1998.

[3] R. Axelrod. *The evolution of cooperation*. Basic Books, New York, 1984.

[4] J. Bendor and D. Mookherjee. Insitutional structure and the logic of ongoing collective action. *American Political Science Review*, 81:129–14?, 1987.

[5] G. Ellison. Learning, local interaction, and coordination. *Econometrica*, 61(5):1047–71, 1993.

[6] G. Ellison. Cooperation in the prisoner's dilemma with anonymous random matching. *Review of Economic Studies*, 61:567–588, 1994.

[7] E. J. Friedman. Robust social norms in bargains and markets. mimeo, 1999.

[8] E. J. Friedman and P. Resnick. The social cost of cheap pseudonyms. mimeo, 1998.

[9] D. Fudenberg, D. Levine, and W. Pesendorfer. When are nonanonymous players negligible? *Journal of Economic Theory*, 79(1):46–71, 1998.

[10] D. Fudenberg and J. Tirole. *Game Theory.* MIT Press, Cambridge, Massachusetts, 1991.

[11] E. Green and R. Porter. Noncooperative collusion under imperfect price information. *Econometrica*, 52:87–100, 1984.

[12] M. Kandori. Social norms and community enforcement. *Review of Economic Studies*, 59:63–80, 1992.

[13] M. Kandori, G. Mailath, and R. Rob. Learning, mutation, and long run equilibria in games. *Econometrica*, 61:29–56, 1993.

[14] P. Kollock. The production of trust in online markets. To appear in: Advances in Group Processes (Vol. 16), eds E.J. Lawler, M. Macy, S. Thyne, and H. A. Walker. Greenwich, CT: JAI Press. 1999, 1998.

[15] David G. Luenberger. *Linear and Nonlinear Programming.* Addison-Wesley Publishing Company, Reading, Massachusetts, second edition, 1984.

[16] P. Milgrom, D. North, and B. Weingast. The role of institutions in the revival of trade: the law merchant, private judges, and the champaign fairs. *Economics and Politics*, 2:1–23, 1990.

[17] M. Nowak and K. Sigmund. A strategy of win-stay, lose-shift that outperforms tit-for-tat in the prisoner's dilemma game. *Nature*, 364(64):56, 1993.

[18] R. Rosenthal. Sequences of games with varying opponents. *Econometrica*, 47(6):1353–1366, 1979.

[19] H. P. Young. The evolution of conventions. *Econometrica*, 61(1):57–84, 1993.