

90-F-22

SOCIAL NORMS AND RANDOM MATCHING GAMES

by

Masahiro Okuno-Fujiwara
Faculty of Economics
University of Tokyo

and

Andrew Postlewaite
Department of Economics
University of Pennsylvania

November 1990

Discussion Papers are a series of manuscripts in their draft form. They are not intended for circulation or distribution except as indicated by the author. For that reason Discussion Papers may not be reproduced or distributed without the written consent of the author.

SOCIAL NORMS AND RANDOM MATCHING GAMES

by

Masahiro Okuno-Fujiwara
Faculty of Economics
University of Tokyo

and

Andrew Postlewaite
Department of Economics
University of Pennsylvania

October, 1989
(this draft, October, 1990)

***/** We are grateful to Michi Kandori, Shinsuke Kambe, and Akihiko Matsui for their helpful comments and pointing out several mistakes in earlier versions and to Ehud Kalai for helpful conversations. Financial support of the Research-in-aid Program from the Japanese Ministry of Education 63301075 and from the National Science Foundation is gratefully acknowledged.

ABSTRACT

Nash equilibrium has been tremendously useful in understanding economic problems in which strategic behavior is important. The theoretical foundations of the solution concept often include the assumption that the game to be played is common knowledge, an assumption that often seems not realistic, particularly in games involving large numbers of players. In this paper we introduce the concept of norm equilibrium for random matching games. A norm equilibrium essentially a Nash equilibrium that implicitly relies on substantially less information than the traditional common knowledge of the game. We use norm equilibria to provide a type of folk theorem for random matching games and to analyze the effect of increasing the numbers of players in a random matching game.

1. Introduction

As incomplete information and strategic behavior have become important topics in understanding economic phenomena, game theory has come to play a critical role in economic thinking. Many applications of game theory have proven useful in conceptualizing and analyzing problems in conflict situations.

Nevertheless, in many applications of game theory to economics there has been an increasing uneasiness in interpreting the results and making predictions based on the theory. The vast majority of the papers of the work using game theory to analyze economic problems uses Nash (or Bayes Nash) equilibrium as the solution concept. This solution concept is sometimes justified on the grounds that it is the only rational choice if players can solve the infinite regress of the chain argument, "I think that you think that I think that..." provided that all the structure of the game is common knowledge.

This complete information interpretation, though a viable and consistent justification for Nash equilibrium, has weaknesses.¹ First, if the game has multiple equilibria, this interpretation does not specify how players coordinate or focus on a specific equilibrium. This is an especially serious problem in super (or repeated) games, which tend to have a large (generally infinite) set of Nash equilibria.

Second, the assumption that the game (its strategies, payoff functions, information sets, etc. of all the players in the game) is common knowledge often appears very restrictive, particularly for applications involving asymmetric information. Here, each player is assumed to have full knowledge of what possible information each player might have, and each player is assumed to know that all other players have such full knowledge, and so on. In short, the exact structure of how information is not perfect is common knowledge.

Even in the case that there is no uncertainty fundamental to the problem, the assumption of common knowledge often seems unrealistic, especially for games with many players. For example, as in Aumann [1987] we can think of the entire economy as a game and all economic outcomes as a Nash (or correlated) equilibrium. The assumption that we know the set of players, let alone their preferences and/or strategies, seems heroic.²

This paper proposes an alternative theory of human behavior in conflict situations. Our theory will stress the importance of social norms and standards of behavior in individual decision-making. The norms and standards will serve primarily to facilitate coordination among players. Our model will provide an alternative interpretation of Nash equilibrium which, we believe, is less susceptible to some (but certainly

¹ For more general and fundamental criticisms of this interpretation of Nash equilibrium, see, for example, Binmore [1987].

² For a related discussion and a definition of a solution concept motivated by concerns similar to these, see Kaneko [1987].

not all) criticisms leveled against existing interpretations.

We should stress that we are not proposing a solution concept different from Nash equilibrium, but rather we are proposing an alternative to existing interpretations. We do not reject the standard interpretations in all cases; for many problems it may be appropriate, particularly for games that are new and novel to the players and which are isolated from other games in which the players may be involved.

In the next section, we will provide definitions of the basic components of our model. We will present the formal model in section 3 and our concept of norm equilibrium in section 4. Section 5 contains several simple examples of norm equilibria and of our results. Section 6 extends the model to the case of a finite population while section 7 contains a discussion of our results including related literature.

2. Motivation

Our view in this paper is that to understand human behavior in a situation in a society, one should consider the behavior of all the members of the society in all possible situations simultaneously. Traditional game theory can, in principle, solve such a problem. Assuming all the relevant data of the problem are common knowledge among all the members of the society, the Nash equilibrium (or even better, the correlated equilibrium) concept can be applied. But it seems unreasonable to assume that all the details of every encounter are common knowledge in a society consisting perhaps of millions of people.

Our aim is to propose a theory with weaker requirements concerning the knowledge of decision-makers. In our model a decision-maker, faced with a conflict situation, considers society's norm as well as the characteristics of the particular situation. Given the particular conflict situation including the characteristics of the other participants, the decision maker will predict opponents' behavior on the basis of the prevailing social norm, or social standard of behavior. Since the opponent's predicted action is specified with the help of a social standard of behavior, an individual's problem is a simple maximization problem of choosing an optimal action given the predicted behavior.

An individual will not treat a conflict situation in isolation, however. His action/choice in this situation will affect his position in future encounters. We will use the term status to summarize the information about a person which will affect future encounters and we will model the effect of an individual's current actions on his future status via a transition mapping. The problem an individual in a conflict situation faces is now the following simple problem: choose an action that will maximize his life-long payoff consisting of the immediate payoff and the value of the future given the resulting position in which he will find himself given the social standard of behavior and the transition mapping.

We shall call a pair consisting of a social standard of behavior and a transition mapping a social norm. Not every social norm can be maintained, however. As in the traditional treatment of game theory, we assume individuals act in their own self-interest. They calculate whether they are better off following

the social norm or violating it. The decision will depend on, among other things, the distribution of the status levels of the others in the society. We will call a social norm a norm equilibrium if it is self-fulfilling, that is, that (i) each individual in the society finds it in his interest to follow the social standard of behavior, and (ii) the distribution of status levels in the society is stationary.

In this interpretation of individual behavior, we can distinguish between two kinds of information that a decision maker needs. He needs local information about the immediate situation, including the status of his opponent, and knowledge of the social norm and the current status distribution. We stress that the entire society need not be common knowledge; in fact, a decision maker need not even have complete information about many aspects of society. What is needed as common knowledge in our formulation is the prevailing social norm and the underlying (stationary) status distribution. In the succeeding sections, we shall formalize these ideas in an extremely simplified society in which two classes of individuals are matched randomly to play a symmetric stage game in each period.

3. Model

The basic model we will use begins with a random matching model of the sort used by Rosenthal [1979]. A society will consist of two sets of players I_1 and I_2 , sets of the same size. Until section 6, we shall assume I_i ($i=1,2$) is a continuum, $[0,1]$. In each period $t=1,2,\dots$, each player from I_1 is matched randomly with a player in I_2 (and vice versa) to play a stage game Γ . We assume that the probability that a currently matched pair of players will be matched again is zero. In section 6, we shall discuss the generalization of this model to the case of a finite population.

Players of both types seek to maximize the expected discounted sum of stage game payoffs. There is a discount factor $\delta \in (0,1)$ which is common to all players. The stage game is a pair $\Gamma = \{A, \pi\}$ where $A = A_1 \times A_2$ and $\pi: A \rightarrow \mathbb{R}^2$. A_i is the set of actions available to a player of type i in the stage game and $\pi_i(a)$ denotes the stage game payoff to a player of type i when action pair $a \in A$ is chosen. The payoff to a player of type i , π_i , is said to be individually rational if it is at least as large as the level he can guarantee for himself, i.e.,
$$\underline{u}_i = \min_{a_j \in A_j} \max_{a_i \in A_i} \pi_i(a_i, a_j), (i \neq j).$$

Throughout the paper, we shall assume:

A.1 The set of individually rational payoffs is bounded.

We shall sometimes denote a random matching game by $\Gamma^\omega(\delta)$ when its stage game is Γ , its discount factor is δ and I_1 is a continuum.

In each period t , each player of I_1 is assigned an element x of a finite set $X_1 = \{x_1, \dots, x_{k_1}\}$ which we will refer to as his status or status level. A status assignment χ is a pair (χ_1, χ_2) where χ_1

is a tuple $(x_i(h))_{h \in I_1}$, specifying the status level $x_i(h)$ that the player $h \in I_1$ possesses at a particular time. We will assume that when a pair is matched, their status levels will be common knowledge to the pair. In particular, their action choices will typically be functions of the pair of status levels.

A player's status level $x \in X_i$ is updated in each period by a predetermined rule, the transition mapping $\tau_i: X \times A_i = X_1 \times X_j \times A_1 \rightarrow X_i$.³ That is, τ_i specifies the status level of a player of type i in the next period, $\tau_i(x, z, a) \in X_i$, when his current status level is $x \in X_i$, the matched player's current status level is $z \in X_j$ ($i \neq j$) and i 's current action is $a \in A_i$. We write $\tau = (\tau_1, \tau_2)$.

In each period, there is a status distribution $p_i \in \Delta_{K_i-1}$, a probability measure on X_i or an element of the K_i-1 dimensional simplex, specifying the proportion of the population of status $x \in X_i$ by $p_i(x)$ in that period. We denote $p = (p_1, p_2) \in \Delta := \Delta_{K_1-1} \times \Delta_{K_2-1}$.

4. Norm Equilibrium

A pure (Markov) strategy for a player of type i is a mapping $s_i: X \rightarrow A_i$ ⁴ specifying a choice of action $s_i(x, z) \in A_i$ in each stage game when players with status levels $x \in X_i$ and $z \in X_j$ ($i \neq j$) are matched. The set of all pure strategies for a player of type i is denoted by S_i . We will call a pair of strategies (s_1, s_2) prescribed to all players of types 1 and 2 respectively a social standard of behavior and denote it by $\sigma = (\sigma_1, \sigma_2)$. A pair $\beta = (\tau, \sigma)$ will be referred to as a social norm.

It is straightforward to extend the definition of a social standard of behavior to allow for the prescription of a random action for some matchings. To avoid the additional notational complexity, we will generally restrict attention to pure strategies. An exception will be that in both theorems below, we allow the SSB to assign random actions for some (disequilibrium) matchings. We emphasize that in the case that a random action is prescribed, the transition function is a function of the realized action taken.

In each period and in each matching, in addition to the transition mapping and the social standard of behavior, only the status levels of the matched players are common knowledge to the pair. Hence, the history of the plays a player has chosen in the past and the current status distribution may not be known precisely; they become known to each player only to the extent they are reflected in the status levels of the matched players.

³ The transition mapping may depend upon A_j , the opponent's action, as well. For equilibrium behavior, this restriction is made without loss of generality in our framework because our equilibrium concept includes the social standard of behavior, that is, the specification of the choice of action as a function of status levels. However, when considering out of equilibrium behavior, allowing this possibility may affect our analysis as it increases the ways status may depend upon what the opponent did.

⁴ To be precise, a Markov strategy is a mapping $\sigma_i: X \times \Delta \rightarrow S_i$. However, this generalization is not necessary as long as disequilibrium behavior is not explicitly treated. See, also, section 6 below.

Given τ and σ , the characteristic function, $\xi_i: X_i \times X \rightarrow \{0,1\}$, is defined as:

$$\xi_i(y,x,z) = \begin{cases} 1 & \text{if } y = \tau_i(x,z,\sigma_i(x,z)) \\ 0 & \text{otherwise.} \end{cases}$$

If a player of type i chooses σ_i , the transition probability from the status level x to the status level y is defined as a function of the opposite status distribution as:

$$q_{xy}^i(\tau,\sigma_i,p_j) = \sum_{z \in X_j} p_j(z) \xi_i(y,x,z) \quad (i \neq j),$$

Let $Q_i(\tau,\sigma_i,p_j)$ be the $K_i \times K_i$ matrix, an element of which is denoted $q_{xy}^i(\tau,\sigma_i,p_j)$, and let $Q(\tau,\sigma,p) = (Q_1(\tau,\sigma_1,p_2), Q_2(\tau,\sigma_2,p_1))$. Given $\beta = (\tau,\sigma)$, if all the players follow a SSB σ , the transition probability $Q_i(\tau,\sigma,p) = Q_i(\beta,p)$ unambiguously characterizes the future status distribution for each of the two player sets. If the current status distribution is p , then the distribution of status levels k periods from now, $p^{(k)}(\beta,p)$, can be defined inductively as:

$$p_i^{(1)}(\beta,p) = p_i Q_i(\tau,\sigma_i,p_j), \text{ and } p_i^{(k)}(\beta,p) = p_i^{(k-1)} Q_i(\tau,\sigma_i,p_j^{(k-1)}).$$

We will denote by $p^{(k)}(x,s_i;\beta,p)$ the status distribution in the k -th period from now if the SSB is σ , the current status distribution is p , but a player of type i with status level $x \in X_i$ changes his strategy to $s_i \in S_i$. Since each player set is a continuum and each player is of measure zero, $p^{(k)}(x,s_i;\beta,p) = p^{(k)}(\beta,p)$ always holds. If the player sets were finite, however, an individual player's deviation from the social standard of behavior would alter the probability distribution of status levels in future periods. We consider this case in section 6.

We shall say a social norm β is stationary at a status distribution p if $p^{(k)}(\beta,p) = p$ for all $k=1,2,\dots$. Given (β,p) , if a player of type i with status level x chooses a strategy $s_i \in S_i$, his expected payoff in each period is defined by,

$$\Pi_i(x,s_i;\beta,p) = \sum_{z \in X_j} p_j(z) \pi_i(s_i(x,z), \sigma_j(z,x)).$$

When s_i is the SSB itself, i.e., $s_i = \sigma_i$, then his immediate expected payoff is denoted as $\Pi_i(x,\sigma_i;\beta,p)$.

Suppose that the social norm $\beta = (\tau,\sigma)$ is stationary at the status distribution p . Then for each $i=1,2$, there is a well-defined associated present discount payoff (or value) for each status x and for each Markov strategy s_i . These payoffs are defined by simultaneously solving for all $x \in X_i$:

$$v_i^o(x,s_i;\beta,p) = \Pi_i(x,s_i;\beta,p) + \delta \sum_{z \in X_j} p_j(z) v_i^o[\tau_i(x,z,s_i(x,z)), s_i, \beta, p] \quad (j \neq i).$$

If s_i is the SSB itself, i.e., when $s_i = \sigma_i$, then his present discounted payoff is denoted as:

$$\begin{aligned}
v_i^o(x, \sigma_i; \beta, p) &= \Pi_i(x, \sigma_i; \beta, p) + \delta \sum_{z \in X_j} p_j(z) v_i^o[\tau_i(x, z, \sigma_i(x, z)), \sigma_i; \beta, p]. \\
&= \Pi_i(x, \sigma_i; \beta, p) + \delta \sum_{y \in X_i} Q_{xy}^i(\tau, \sigma, p) v_i^o(y, \sigma_i; \beta, p)
\end{aligned}$$

Definition: A triplet $(\beta^*, p^*) = (\tau^*, \sigma^*, p^*)$ is called a norm equilibrium of $\Gamma^o(\delta)$ if

- (a) β^* is stationary at p^* ,
- (b) for all $i=1,2$, for all $x \in X_i$ and for all $s_i \in S_i$,
$$v_i^o(x, \sigma_i^*; \beta^*, p^*) \geq v_i^o(x, s_i; \beta^*, p^*).$$

Two remarks are in order. First, a norm equilibrium is defined as a SSB in which no player can find any other Markov strategy which unilaterally improves his payoff. It is well known, however, that a Markov deterministic strategy satisfying (b) is an optimal strategy among all mixed general (including non-Markovian) strategies (see, e.g., Derman [1970] Theorem 3.1). Second, by the criterion of unimprovability of Markov decision theory (see, e.g., Howard [1960] or Whittle [1985]), condition (b) of norm equilibrium can be expressed in terms of unimprovability. To do so, for a social norm β , define $\alpha_i(x, z, a, \beta)$ to be the net gain in a stage game payoff when a player of type i with status $x \in X_i$ is matched with player with $z \in X_j$ and plays $a \in A_i$ instead of the action prescribed by the SSB, $\sigma_i(x, z)$;

$$\alpha_i(x, z, a, \beta) = \pi_i(a, \sigma_j(z, x)) - \pi_i(\sigma_i(x, z), \sigma_j(z, x)).$$

Then the next lemma follows immediately.

Lemma 1: $(\beta^*, p^*) = (\tau^*, \sigma^*, p^*)$ is a norm equilibrium if and only if

- (a) β^* is stationary at p^* ,
- (b) for $i, j=1,2$ ($i \neq j$), for all $x \in X_i$ and $z \in X_j$, and $a \in A_i$,
$$\alpha_i(x, z, a, \beta^*) \leq \delta [v_i^o(\tau_i(x, z, \sigma_i^*(x, z)), \sigma_i^*; \beta^*, p^*) - v_i^o(\tau_i(x, z, a), \sigma_i^*; \beta^*, p^*)].$$

Lemma 1 simply states that for any pair of matched status levels, the immediate gain from deviating from the SSB must be less than the resulting loss in the future due to a change in status, that is, the continuation value. This resulting loss, however, is evaluated along the equilibrium path, or evaluated by the value function $v_i^o(\cdot, \sigma_i^*; \beta^*, p^*)$.

5. Folk Theorem

To illustrate the concept of norm equilibrium, we will present several examples. Suppose $K_1 = K_2 = 2$, i.e., there are only two status levels for each type and $X_i = \{G, B\}$. In this section, we shall

use the following stage game Γ_1 with $M > 0$.

	C	D	P
C	4,4	0,5	-1,-100
D	5,0	1,1	0,-M
P	-100,-1	-M,0	-100,-100

The game is a simple variant of the prisoner's dilemma game. The actions C and D are the usual cooperate and deviate (or defect) actions. There is an additional action, P, which is essentially a "punishment" action. The maximal symmetric payoff is obtained by the players cooperating, that is, by each player playing C. We are interested in determining when this cooperative behavior can be supported as a norm equilibrium by a social norm (β, p) . Also note that the security level, \underline{u}_i , is zero for both types.

Example 1: Consider the social norm (β, p) defined as follows.

$$\tau_i(x, z, a) = \begin{cases} G & \text{if } (x, z, a) = (G, G, C) \text{ or } (G, B, D), \\ B & \text{otherwise.} \end{cases}$$

$$\sigma_i(x, z) = \begin{cases} C & \text{if } x = z = G, \\ D & \text{otherwise.} \end{cases}$$

$$p_i(G) = 1 \text{ and } p_i(B) = 0.$$

In words, the social norm prescribes that a player should choose C if both he and his opponent are good, i.e., status G, and should defect (choose D) if either is bad (status B). A player's status is revised according to τ . A player with status G remains a G so long as he follows the prescribed social standard of behavior but changes to bad, B, if he deviates from it. The status level B is "absorbing" in the sense that a B remains a B regardless of his action.

For this social norm, the present discounted payoff for a player of type i with status x along the equilibrium path, $v_i^o(x)$, (we suppress other arguments of v_i^o for ease of notation) is;

$$v_i^o(G) = 4/(1-\delta) \text{ and } v_i^o(B) = 1/(1-\delta).$$

In view of lemma 1, the triplet (β, p) is a norm equilibrium if $\delta \geq 1/4$. This is exactly the condition necessary to make the prescribed behavior (i.e., the one-shot Nash reversion) a perfect equilibrium for the fixed player repeated game with this stage game.

We would like to make several comments about this example.

First, the fact that the probability distribution over status levels is degenerate is not important. If for $i=1,2$, $p_i(B) \leq r$ and $p_i(G) \geq 1-r$, $r > 0$ then similar calculations reveal that (β, p) would be a norm equilibrium for $\delta \geq 1/[4-3r]$. Thus, the presence of a small proportion of status B people in the society increases the threshold discount factor which is consistent with this norm being an equilibrium, but continuously. This is to be contrasted with example 3 below.

The second observation is that while the outcome is the same as at a conventional Nash equilibrium, playing the equilibrium strategy may require a vast amount of information if the information about status levels is not available. Suppose it is possible to obtain records of a player's past plays. Suppose further the player who is matched in a certain period has played D sometime in the past, say in period t . This does not necessarily imply that he has deviated from the equilibrium path since this is the prescribed behavior against some opponents. To check whether he has deviated or not, we must check whether or not the player he was matched in period t had played D before. If this player had, we must check the history of the player with which this player was matched in the period, and so on. To play this equilibrium strategy without knowing status levels, essentially knowledge about the history of the entire society would be required.

The last observation is that the social norm in this example is not optimal, in the sense that for some parameter values it will not support cooperation for which other social norms could support cooperation. The following example demonstrates this fact.

Example 2: Consider the following pair of social norm and status distribution $(\beta, p) = (\tau, \sigma, p)$ with the same stage game Γ_1 .

$$\tau_1(x, z, a) = \begin{cases} G & \text{if } (x, z, a) = (G, G, C) \text{ or } (G, B, P), \\ B & \text{otherwise} \end{cases}$$

$$\sigma_1(x, z) = \begin{cases} C & \text{if } (x, z) = (G, G), \\ P & \text{if } (x, z) = (G, B), \\ D & \text{otherwise.} \end{cases}$$

$$p_1(G) = 1 \text{ and } p_1(B) = 0.$$

This social norm differs from that in the previous example in that a good player is to "punish" his opponent if the opponent is bad (status B) by playing action P. As before, a player with status G

retains that status so long as he follows the social standard of behavior and reverts to B if he deviates from it. The status level B is absorbing, as before.

The corresponding discounted payoffs are;

$$v_1^a(G) = 4/(1-\delta) \text{ and } v_1^a(B) = 0.$$

In light of lemma 1, three inequalities must be satisfied for this to be a norm equilibrium, associated matchings of a G with another G, a G with a B, and a B with a G. The inequality associated with the last match, a B meeting a G, is vacuously satisfied since the SSB prescribes that the B play a one shot best response in this case. The other two inequalities are respectively, $1 \leq \delta(4/(1-\delta)-0)$ and $M+1 \leq \delta(4/(1-\delta)-0)$. It can be easily checked that these constraints are satisfied if $\delta \geq \max\{(1+M)/(5+M), 1/5\}$. Hence, there are pairs of M and δ for which this social norm can support cooperation, (C,C), while the social norm of example 1 cannot.

There are three points we want to make about this example. First, note that in the social norm of example 2,

$$\tau_1(G,B,P) = G, \text{ but } \tau_1(G,G,P) = \tau_1(B,B,P) = B.$$

Status transition does not depend only upon current status and current action, but depends also upon the opponent's status level. Unlike the status transition function in example 1, $\sigma_1(B,G)$, $\sigma_1(G,G)$ and $\sigma_1(G,B)$ are all distinct. This aspect distinguishes our model from reputation models (e.g., Kreps [1989], Kreps and Wilson [1982], Kreps, *et. al.* [1982] and Rosenthal [1979]), where the choice of action depends only upon the reputation of one player and the new reputation does not depend upon the opponent's reputation. (See, however, Rosenthal and Landau [1979]).

The second point is that as in the previous example, while the probability distribution over status levels is degenerate, the social norm would still have been part of an equilibrium with a positive proportion of people having status B. As in that example, the minimum δ that was consistent with equilibrium would increase as the proportion of people with status B increased.

The third point we wish to make is that it is easier to support cooperation with random matching than if there were a fixed match. To see this, note that when a player deviates from the SSB, he will become status B; the SSB prescribes that in this case, he is to be punished forever, that is, if a player of status G is matched with a B the G player is to play P. If $M > 1$, this is not a best response to D, hence there is a cost to the G. For a fixed matching, this cost would be borne by the same status G player each period, while in the random matching case, a given G player will not be matched each period with this given status B player, and hence will not incur the cost of punishing every period. As a result, the constraint associated with a G matched with a B will be more easily satisfied in the random matching case than in the fixed matching case. That is, for a given M, there will be a larger set of δ 's that satisfy the constraint in the random matching case.

Example 2 suggests that a version of the Folk theorem may hold with two status levels. For $i = 1, 2$, let $\underline{u}_i = \min_{a_j} \max_{a_i} \pi_i(a_i, a_j)$, i.e., \underline{u}_i is the security level player i can guarantee himself in the stage-game payoff with pure strategies.

Theorem 1: If $\pi_i(a_i^*, a_j^*) > \underline{u}_i$ for $i = 1, 2$, then (a_i^*, a_j^*) is supported as a norm equilibrium outcome with two status levels when δ is sufficiently close to 1.

Proof: Let $X_i = \{G, B\}$. Define a_i^1 and a_j^1 so that $\underline{u}_i = \pi_i(a_i^1, a_j^1)$. Let (q_1^*, q_2^*) be the stage-game one-shot (possibly mixed strategy) Nash equilibrium where $q_i^*(a_i)$ denotes the probability of playing a_i . For any arbitrary pair (a_i^*, a_j^*) satisfying the condition, let

$$\tau_i(x, z, a) = \begin{cases} G & \text{if } (x, z, a) = (G, G, a_i^*) \text{ or } (G, B, a_j^1), \\ B & \text{otherwise.} \end{cases}$$

$$\sigma_i(x, z) = \begin{cases} a_i^* & \text{if } (x, z) = (G, G), \\ a_j^1 & \text{if } (x, z) = (G, B), \\ a_i^1 & \text{if } (x, z) = (B, G), \\ a_i & \text{with probability } q_i^*(a_i) \text{ if } (x, y) = (B, B). \end{cases}$$

$$p_i(G) = 1 \text{ and } p_i(B) = 0 \quad i = 1, 2.$$

Clearly, $v_i^o(G) = \pi_i(a_i^*, a_j^*)/(1-\delta)$ and $v_i^o(B) = \underline{u}_i/(1-\delta)$. It follows from lemma 1 that the triplet (τ, σ, p) is a norm equilibrium. Q.E.D.

By allowing coordination through time, any strictly individually rational payoff vector can be approximated as the average payoff outcome with two status levels when δ is sufficiently close to 1.

If the discount factor, δ , is not sufficiently close to 1, however, allowing more status levels may support a more efficient outcome. The following example with three status levels, $\{G, B, H\}$, but with the same stage-game Γ_1 as in examples 1 and 2 will illustrate this point.

Example 3: Consider the following triplet, (τ, σ, p) .

$$\tau_i(x, z, a) = \begin{cases} G & \text{if } (x, z, a) = (G, G, C) \text{ or } (G, H, C), \\ H & \text{if } (x, z, a) = (G, B, P) \text{ or } (H, \cdot, D), \\ B & \text{otherwise.} \end{cases}$$

$$\sigma_i(x,z) = \begin{cases} C & \text{if } (x,z) = (G,G) \text{ or } (G,H), \\ P & \text{if } (x,z) = (G,B), \\ D & \text{otherwise.} \end{cases}$$

$$p_i(G) = 1, \text{ and } p_i(H) = p_i(B) = 0, \quad i=1,2.$$

One can interpret the additional status level, H, as that reserved for a "hero". A person with status level G becomes H by punishing a B if he meets one. Once a person is an H, the social standard of behavior is that he is to play D regardless of his opponent. The addition of such a status level allows a person who has punished a B to be rewarded. Since the act of punishing a B is costly to the player doing the punishing, there must be an incentive for the player to do so. In example 2 we used the "threat" of turning a G player into status B if he failed to punish. But there are limits to the threat of punishing this way; we can think of the addition of the H as using a carrot in addition to the stick to provide the proper incentives for punishing.⁵

The corresponding discounted payoffs are:

$$v_i^\sigma(G) = 4/(1-\delta), v_i^\sigma(H) = 5/(1-\delta), \text{ and } v_i^\sigma(B) = 0.$$

As in the previous example, there are two relevant constraints. The constraint that a G meeting another G plays C rather than D is the same as before, and will be satisfied when $\delta \geq 1/5$. The constraint stemming from a meeting of a G with a B differs in that the future loss from not following σ is now larger by $1/(1-\delta)$. This is because following σ results in the G being "upgraded" to status H which yields a per-period payoff of 5 instead of 4. This alters the constraint associated with a G meeting a B, yielding the inequality $\delta \geq (1+M)/(6+M)$ rather than $\delta \geq (1+M)/(5+M)$ as in example 2. Thus for $0 < M < 2/3$ and $(1+M)/(5+M) > \delta > (1+M)/(6+M)$, the norm equilibrium of example 3 with three status levels supports cooperation while that of example 2 with two status levels cannot.

The first comment on this example is that there is no reason in general that we should believe that three status levels is the maximal number which can be useful. In general, one should expect to be able to construct games for which arbitrarily large numbers of status levels are necessary in order to support a particular outcome as an equilibrium.

⁵ The practice of giving war veterans preference in hiring might be an example of such incentives.

A second comment is that for this example, unlike the two previous examples, the degenerate probability distribution over the players' status levels is important. Suppose that there is a positive proportion of type 1 that is changed to status B. Then in each period, these players will be matched with that proportion of players of type 2, which will result in those type 2 players being converted to status H if the SSB were followed. Thus, with probability 1, the distribution of the status levels for type 2 players would converge to that with only H's. This clearly would violate the stationarity of p that is part of the definition of a norm equilibrium.

The plausibility of the social norm is somewhat less in that case as well. Part of the definition of a norm equilibrium guarantees that a player of any status would have an incentive to follow the SSB. This incentive is assured only for the given probability distribution of status levels, p , however. If the distribution of status levels is not stationary, these incentives may change. In fact, it is clear that in example 3 this is the case. The constraint that a player of status G meeting a G includes $v(G)$, the status he will have if he follows the SSB. But the value of being a G rather than a B is that when meeting an G, the payoff will be higher. On the other hand, the payoff is higher to a B than to a G in following the SSB when matched with an H. If the distribution of type 2 players asymptotically has no G's, then asymptotically there is no advantage to a type 1 player in being a G rather than a B. Thus, as the proportion of G's decreases, $v_1(G)$ decreases; at some point the constraint associated with the match of a type 1 status G player matched with a type 2 status G will be violated and the type 1 G will not find it in his interests to follow the SSB.

The above paragraph outlines how the SSB would "unravel" if there is a positive proportion of B status people in the society initially. At some point, the difference between the continuation values for status levels G and B will be less than the immediate gain from playing D when a type 1 G is matched with a type 2 G. We have avoided the problems associated with changing values associated with various status levels by insisting that the distribution of status levels p be stationary when all players are following the SSB. We will come back to this point in the section on finite player problems.

The unraveling described above stemmed from the fact that whenever a G was matched with a B, he became an H forever if he followed the SSB. We could have avoided the unraveling if instead of having H be an "absorbing" status, H was granted for a single period only, that is, if the transition function τ was altered to:

$$\tau_1(x,z,a) = \begin{cases} G & \text{if } (x,z,a) = (G,G,C), (G,H,C), \text{ or } (H, \cdot, D), \\ H & \text{if } (x,z,a) = (G,B,P), \\ B & \text{otherwise.} \end{cases}$$

For the case in which the distribution of status levels is $p_1(G) = 1$, and $p_1(H) = p_1(B) = 0$,

$i=1,2$, the discounted payoffs would now be:

$$v_1^{\sigma}(G) = 4/(1-\delta), v_1^{\sigma}(H) = 5 + 4\delta/(1-\delta), \text{ and } v_1^{\sigma}(B) = 0.$$

As in example 2, the constraint associated with a G meeting another G is unchanged and will be satisfied when $\delta \geq 1/5$. The constraint stemming from a meeting of a G with a B differs in that the future loss from not following σ is now larger by δ than in the social norm of example 2. This is because following σ results in his having status H for one period, yielding a payoff of 5 instead of 4. For $2/3 > M > 0$, the binding constraint on the minimum δ for which (τ, σ, p) constitutes a norm equilibrium is that arising from the incentive for a G to play D rather than P when meeting a B , that is, when he is called upon to punish a past defector. The three-status social norm here has relaxed that constraint and hence, for small but positive M , there will be δ 's such that this social norm supports cooperation while that of example 2 with two status levels does not.

Thus, the change to a "one-period" reward as an H still reduces the minimum δ for which cooperation can be supported below that of example 2, although not as much as if the H status were permanent. However, for this case in which the status H is accorded for a single period only, the social norm will not unravel as it does when the status H is permanent. Even if there is a positive proportion of B 's, the proportion of H 's in the next period can be at most the proportion of B 's this period (at most because some of the B 's will be matched with each other and not generate an H next period).

6. Finite Player Sets

In the previous sections we assumed that the player set for each type was a continuum. Since each player is of measure zero, no unilateral deviation from the social norm by a single player will alter the status distribution. If there is a finite number of players, deviations from the social norm by a single player may alter the status distribution even when the equilibrium status distribution is stationary. The social norm and discussion of example 3 point out that a deviation by a single player from the SSB might have very different consequences in a continuum model than in a model with a large but finite number of players. This leads us to a direct investigation of social norms in finite societies.

Until now, we have assumed that a player's choice of action in any matching depended only on the pair of statuses in that matching. In general, a Markov strategy allows the choice to depend on the status distribution as well. We did not explicitly consider this additional aspect because any unilateral deviation could have no effect on the status distribution in the model with a continuum population.⁶ Since in the

⁶ In the case of continuum, one may also allow a Markov strategy to be contingent on the underlying distribution. However, we could alter the SSB's in the previous sections so that, if a positive measure of the population has deviated from equilibrium behavior (and hence, the distribution differs from the equilibrium distribution), the social norm requires all the players to play the one-shot Nash equilibrium. This is clearly a perfect equilibrium.

finite population case individual deviations may alter the status distribution, we must take explicit account of the fact that the choice of action in each period may depend upon the underlying status distribution.⁷

In this section, we shall prove a result analogous to Theorem 1 for finite societies. For our result, we will strengthen somewhat our definition of a norm equilibrium. The definition of a norm equilibrium guarantees that a player of any status level meeting a player of any other status level will find it in his best interest to follow the SSB. This should hold even if the status distribution does not have full support, and hence some matchings are impossible. The reason that we insist that following the SSB be optimal even for matchings that are impossible given the status distribution is similar to the argument for perfect Nash equilibrium. If a player plays differently from the SSB, the status distribution will change; matchings that are impossible given the (proposed) equilibrium status distribution will not necessarily be impossible following a deviation from the SSB. An SSB that prescribes for some matchings actions that are not optimal for a player are the same as the non-credible threats that subgame perfection eliminates.

We note that in our definition of a norm equilibrium, we have not gone as far as subgame perfection does toward this end. In the continuum case, we do not ask that the strategy described by the SSB in a norm equilibrium be optimal for all distributions, only for the stationary distribution that is part of the definition of a norm equilibrium. Since any countable number of deviations cannot alter the status distribution, one might think of the definition of a norm equilibrium as requiring the prescribed strategy to be optimal as long as there are no more than a countable number of deviations from the strategy.⁸

The finite society case is different in that a single deviation may alter the distribution. As example 3 demonstrates, a single deviation in an arbitrarily large finite society asymptotically can have a large effect on the status distribution. For the finite case, we would like to preserve the property described above: a single deviation from the SSB should never lead to the matching of a pair, one of whom would want to deviate from the SSB. We must strengthen the definition of a norm equilibrium to do this. We will strengthen the definition by (1) requiring that the social norm be such that any single deviation not have a large effect on the status distribution, even asymptotically, and (2) requiring that the SSB be optimal not only for a given status distribution, but also for all distributions "close" to it.

Let player set I_i ($i=1,2$) be a finite set $\{1, \dots, n\}$. Let Θ be the set of all permutations of $\{1, \dots, n\}$. In each period $t=0,1,2, \dots$, $\theta_t \in \Theta$ is chosen randomly and, for each $h \in I_1$, $(h, \theta_t(h))$ are

⁷ It is clear that a strategy that depends on the actual status distribution is somewhat inconsistent with the limited information available to players with which we motivated our interest in this model. However, distribution dependent strategies will only be considered by players contemplating deviations from an SSB; any constraint on the strategies allowed players that restricts the information dependence would only strengthen our results.

⁸ Kalai and Neme [1989] have considered the implications of asking for subgame perfection for a restricted set of subgames.

matched. A matching history from t to $t+k$ is a k -tuple, $\theta^{t,t+k}=(\theta_t, \dots, \theta_{t+k-1})$. The probability that the currently matched players will be matched again in a given future period is $1/n$. The set of status distributions in this section is denoted as:

$$\Delta_n = \{p=(p_1, p_2) \in \Delta_{K_1-1} \times \Delta_{K_2-1} \mid p_i(x) = k/n \text{ for some } k=0,1, \dots, n\}.$$

We shall denote the random matching game with stage game Γ and the population size n by $\Gamma^n(\delta)$ when the discount factor is δ .

A distribution dependent Markov-strategy for a player of type i is a mapping $s_i: X \times \Delta_n \rightarrow A_i$ specifying a choice of action $s_i(x, z, p) \in A_i$ in a stage game when players with status levels $x \in X_i$ and $z \in X_j$ are matched and the current (announced) status distribution is p . Namely, s_i depends upon the state of each matched game defined as a triple (x, z, p) . In this section S_i will denote the set of all such strategies. We call a strategy distribution independent if it is independent of the distribution p . An SSB is a pair of distribution independent strategies $\sigma=(\sigma_1, \sigma_2)$, where $\sigma_i: X \rightarrow A_i$. We shall allow, however, that each individual can deviate from a social norm by taking any distribution dependent strategy.

Suppose a social norm $\beta=(\tau, \sigma)$ is given. In each period τ , the current status assignment χ and the realized permutation θ_t determine the status of each player's opponent; these and the choice of actions in the match determine the new status assignment for the subsequent period. Thus, even if all players of each type choose the SSB, (σ_1, σ_2) , the resulting status distribution in (the beginning of) period $t+k$ depends upon the current (period t) status distribution, p_t , and the matching history, $\theta^{t,t+k}=(\theta_t, \dots, \theta_{t+k-1})$. We shall denote this stochastic process by $P_n^{(k)}(\beta, p_t) = (P_{n1}^{(k)}(\beta, p_t), P_{n2}^{(k)}(\beta, p_t))$, which is dictated by the realization of $\theta^{t,t+k}$.

Suppose a player of type i has a status level $x \in X_i$ in t and the underlying status distribution is $p \in \Delta_n$. Even if all players follow the social norm β , this player's future status level is similarly determined by the matching history $\theta^{t,t+k}$. Thus, we denote by $x_{ni}^{(k)}(x, p; \beta)$ the stochastic process of the status of a player of type i if he currently has the status x and the underlying population is p . We shall denote by $Q_{ni}^{(k)}(x', p'; x, p, \beta)$ the probability that the two stochastic processes $P_n^{(k)}(\beta, p)$ and $x_{ni}^{(k)}(x, p; \beta)$ take the values p' and x' .

Suppose a social norm $\beta=(\tau, \sigma)$ is given and the current status distribution is p . Suppose further that a player of type i with a status level $x \in X_i$ chooses a distribution dependent strategy s_i instead of the SSB σ_i . This deviation from the social norm will change the underlying stochastic process from $P_n^{(k)}(\beta, p)$ to $P_n^{(k)}(x, s_i; \beta, p)$ and $Q_{ni}^{(k)}(x', p'; x, p, \beta)$ to $Q_{ni}^{(k)}(x', p'; s_i, x, p, \beta)$ in an obvious manner.

Given a social norm, β , and the current status distribution, p , the expected discounted payoff $v_i^n(x, s_i; \beta, p)$ for a player of type i with a status x if he uses a Markov strategy $s_i \in S_i$ is defined by simultaneously solving for all $(x, p) \in X_i \times \Delta_n$:

$$v_i^n(x, s_i; \beta, p) = \Pi_i(x, s_i; \beta, p) + \sum_k \delta^k \sum_{(x^k, p^k)} Q_{ni}^{(k)}(x^k, p^k; s_i, x, p, \beta) v_i(x^k, s_i; \beta, p^k).$$

Similarly for all $x \in X_1$ and $p \in \Delta_n$, define his payoff if he follows the SSB as:

$$v_1^n(x, \sigma_1; \beta, p) = \Pi_1(x, \sigma_1; \beta, p) + \delta \sum_{(x', p')} Q_{ni}^{(1)}(x', p'; \sigma_1, x, p, \beta) v_1^n(x', \sigma_1, \beta, p').$$

A social norm $\beta = (\tau, \sigma)$ is called stationary at a status distribution p if the underlying stochastic process is deterministic and constant at p . Then for $\Gamma^n(\delta)$, the stage game Γ played in a society with a finite number of individuals, the analogue of our definition of a norm equilibrium for $\Gamma^n(\delta)$, the continuum of players case, would require that for a social norm β^* and status distribution p^* , (1) β^* should be stationary at p^* , and (2) for all i , s_i and x , $v_1^n(x, \sigma_1^*; \beta^*, p^*) \geq v_1^n(x, s_i; \beta^*, p^*)$.

This straightforward extension of the definition of a norm equilibrium is unsatisfactory, however. As the following example will show, an SSB that is part of a norm equilibrium may not even be consistent with Nash equilibrium.

Example 4: Let the stage game be:

	N	C	D
N	1,1	0,0	0,0
C	0,0	4,4	0,0
D	0,0	0,0	0,0

Note that the stage game has three pure strategy one-shot Nash equilibria. Action N is the play to achieve the inferior Nash equilibrium, while C corresponds to the superior Nash equilibrium. We call the action D deviation. Let the social norm be as follows:

$$\tau_1(x, z, a) = \begin{cases} G & \text{if } (x, z, a) = (G, G, N) \\ B & \text{otherwise.} \end{cases}$$

$$\sigma_1(x, z) = \begin{cases} N & \text{if } (x, z) = (G, G) \\ C & \text{if } (x, z) = (B, B) \\ D & \text{otherwise.} \end{cases}$$

$$p_1(G) = 1 \text{ and } p_1(B) = 0.$$

For this social norm and status distribution, $v_1^n(G) = 1/(1-\delta)$ and $v_1^n(B) = 0$.

In words, the social norm can be described as follows. There are two status levels, G (good) and B (bad). If two G's are matched, they are to play the intermediate level Nash equilibrium and, if they do so, each will retain his G status. If either deviates, his status will be altered to B. If two B's are matched, they are to play C, the best of the Nash equilibria. When a B is matched with a G, both are to play D. An important feature of this social norm is that, when the matching is either a B and a G or two B's, regardless of their choice they will be assigned B status in the next period. Thus B is an absorbing status level and whenever a G is matched with a B, the G will be converted to this absorbing status.

Clearly, this social norm satisfies the definition of norm equilibrium for any $\delta < 1$. However, in a society with finite population, it may not be stable. To see this, consider a society consisting of just one individual in each type (i.e., $n=1$). This society would be the same as the usual fixed player repeated game, as the same two players will be matched in each period. In the social norm above, each player will receive $1/(1-\delta)$. If either player unilaterally deviates, he will lose 1 in the period in which he deviates and in addition he will lose 1 in the next period while converting his opponent to status level B. However, from the third period on, his payoff each period will be 4 rather than 1, his payoff per period had he not deviated. Hence if $\delta \geq 1/2$, each player has incentive to deviate from the first norm.

If n is large, the expected time it takes to "convert" sufficiently many players to increase one's per period payoff increases; nevertheless, if δ is sufficiently close to 1, each player will have an incentive to deviate from the social norm, knowing that the distribution of status levels in society will (with very high probability) ultimately change in a way that makes the initial deviation worthwhile (if the discount factor β is sufficiently close to 1).

The above example clearly shows that for a social norm to be plausible for finite societies, we must take into account the changes in the distribution of status levels that result from an individual's deviation from the SSB. There are several ways in which we can strengthen the definition of a norm equilibrium that take these changes into account. The "strongest" concept would ask that following the prescribed behavior of an SSB in a social norm be optimal not only for the assumed stationary distribution p^* , but for all status distributions.

Definition: (β^*, p^*) is a strong norm equilibrium of $\Gamma^n(\delta)$ if;

- (1) β^* is stationary at p^* , and
- (2) for all $i=1,2$, $s_i \in S_i$, $x \in X_i$, and $p \in \Delta_n$,

$$v_i^n(x, \sigma_i; \beta^*, p) \geq v_i^n(x, s_i; \beta^*, p).$$

However, we use a weaker concept than this. We do not use this stronger concept of strong norm equilibrium because it requires that a social norm provide proper incentives to follow the SSB even at

distributions that might never arise for any given actions the players take. We could ask it be optimal for a player of any status to follow the SSB at all status distributions that could result from some combination of actions the players might choose, but even this is stronger than we will ask. Asking that a SSB provide correct incentives for every distribution that could result from some set of actions is essentially asking that the prescribed strategies constitute a subgame perfect equilibrium. This requirement seems overly strong as the next example illustrates.

Example 5: This example will be a modification of example 2. Let the stage game be as in that example with $M=2$:

	C	D	P
C	4,4	0,5	-1,-100
D	5,0	1,1	0,-2
P	-100,-1	-2,0	-100,-100

The transition function and SSB are somewhat different than in that example; here the status B has been replaced with two status levels, B1 and B2.

$$\tau_i(x,z,a) = \begin{cases} G & \text{if } (x,z,a) = (G,G,C), (G,B1,P), (G,B2,P) \text{ or } (B2, \cdot, \cdot) \\ B2 & \text{if } (x,z,a) = (B1, \cdot, \cdot) \\ B1 & \text{otherwise.} \end{cases}$$

$$\sigma_i(x,z) = \begin{cases} C & \text{if } (x,z) = (G,G) \\ P & \text{if } (x,z) = (G,B1) \text{ or } (G,B2) \\ D & \text{otherwise.} \end{cases}$$

$$p_i(G) = 1, \text{ and } p_i(B1) = p_i(B2) = 0, \quad i=1,2.$$

In words, the SSB prescribes that if a G is matched with another G he is to play C; if a G is matched with a B1 or B2, the G is to choose the punishment action P. Independent of his match a B1 or a B2 is to choose D. The transition rule specifies that a G deviating from the SSB is to be punished for two periods, that is, he is to become a B1 for one period, promoted to a B2 for one period and then will become a G, remaining so as long as he follows the SSB.

For this social norm, $v_1(G)=4/(1-\delta)$, $v_1(B1)=\delta^2[4/(1-\delta)]$, and $v_1(B2)=\delta[4/(1-\delta)]$; it is straightforward to calculate that for $\delta \geq 1/2$, this is a norm equilibrium. If we change the status distribution

to allow for small proportions of status levels other than G , it will still be optimal to follow the SSB if δ sufficiently close to 1.

An interesting feature of this social norm is that whatever the status distribution is, if all players follow the SSB, all players will be status G in at most two periods. Thus, in a given match in a particular period t , when a player is deciding whether or not to follow the SSB, the current distribution will have no affect on his decision except in period $t+1$. Consider then the decision facing a G who finds himself matched with a $B1$ or $B2$. Following the SSB, choosing P , in this case is not a best reponse if he considers the current payoff only; it yields him an immediate payoff of -2 rather than the payoff of 1 that he could get by playing D . For it to be optimal for a type 1 status G player to follow the SSB it must be that the discounted future (expected) payoff from playing P exceeds that from playing D by at least 3 . This holds for our given distribution in which all type 2 players are of status G . It is easy to see that it also holds if nearly all type 2 players currently are of status G or $B2$ and δ is sufficiently close to 1 since, as we remarked above, the distribution can only affect a person's decision through its affect on his payoff in the next period and any person of status G or $B2$ will be status G next period if they follow the SSB. Hence the value to a type 1 player of status G depends only on the proportion of type 2 players that are of status $B1$, $P_2(B1)$.

Hence if $\delta \geq 1/2$ and the distribution of statuses is such that there are no type 2 players of status $B1$ it is optimal for a type 1 player in any match to follow the SSB. On the other hand, if some type 2 players are of status $B1$, a type 1 player of status G may not have an incentive to play P against a $B1$ or $B2$. In fact a simple calculation shows that for any δ such that $1/2 \leq \delta < 1$, if the proportion of type 2 players of status $B1$ exceeds $(4\delta^2 + 4\delta - 3)/7\delta$, such a type 1 G player will have an incentive to deviate from the SSB.

Thus, the only reason for such a player not to follow the SSB would be if the proportion of players of the opposite type of status $B1$ was above this threshold. Since a player will remain of status $B1$ for a single period before becoming a $B2$, this can only happen (given our initial distribution) if a non-negligible proportion of players simultaneously deviates from the SSB. As we consider increasingly large (finite) societies, a given proportion of people simultaneously deviating entails a simultaneous deviation by an increasingly large number of players.

Hence, the strategy associated with the above social norm will not be optimal for all status distributions that can result from some given actions the players might take. In other words, subgame perfection doesn't hold. One might argue that the above social norm is nevertheless quite plausible for large societies since the subgames in which subgame perfection fails result only from a large number of simultaneous deviations.

This discussion leads us to the following definitions.

Definition: We say that β^* is asymptotically stationary at p^* if:

- (a) β^* is stationary at p^*
- (b) for any $\epsilon > 0$, $\exists \epsilon'$ with $0 < \epsilon' \leq \epsilon$ and n_0 such that for $i=1,2$, $\forall x \in X_i$, $\forall s_i \in S_i$, for $n > n_0$, if $p \in \Delta_n$ and $d(p, p^*) \leq \epsilon'$, then $\text{Prob} \{P_n^{(k)}(x, s_i; \beta^*, p) = p' \mid d(p', p^*) \leq \epsilon\} = 1$ for all $k=1,2,\dots$, where $d(p, p')$ is the distance between p and p' .

That is, β^* is asymptotically stationary at p^* if it is stationary at p^* and if whenever the status distribution is initially close to p^* , any status distribution that can result from a single player's deviating from the social norm will also be close to p^* .

Definition: We say that the pair (β^*, p^*) is an asymptotically stationary norm equilibrium if:

- (a) β^* is asymptotically stationary at p^*
- (b) $\exists \epsilon > 0$ such that for $i=1,2$, for all $x \in X_i$, and for all $s_i \in S_i$, if $p \in \Delta_n$ satisfies $d(p, p^*) \leq \epsilon$, then $v_i^n(x, \sigma_i^*; \beta^*, p) \geq v_i^n(x, s_i; \beta^*, p)$.

In words, (β^*, p^*) is an asymptotically stationary norm equilibrium if β^* is asymptotically stationary at p^* and if for any status distribution close to p^* , every player must be at least as well off by following the SSB as he would be by deviating from it.

We can now present the analog of Theorem 1 for the case of finite populations.

Theorem 2: If $\pi_i(a_i^*, a_j^*) > \underline{u}_i$ for $i = 1,2$, then (a_i^*, a_j^*) is supported an asymptotically stationary norm equilibrium outcome of $\Gamma^n(\delta)$ with two status levels if δ is sufficiently large.

Proof: Consider again the social norm β^* and distribution of statuses used in the proof of Theorem 1,

$$\tau_i^*(x, z, a) = \begin{cases} G & \text{if } (x, z, a) = (G, G, a_i^*) \text{ or } (G, B, a_i^j), \\ B & \text{otherwise.} \end{cases}$$

$$\sigma_i^*(x, z) = \begin{cases} a_i^* & \text{if } (x, z) = (G, G), \\ a_i^j & \text{if } (x, z) = (G, B), \\ a_i^i & \text{if } (x, z) = (B, G), \\ a_i & \text{with probability } q_i^*(a_i) \text{ if } (x, y) = (B, B) \end{cases}$$

$$p_i^*(G) = 1 \text{ and } p_i^*(B) = 0 \quad i = 1,2.$$

This social norm is stationary for every probability distribution, and hence is asymptotically stationary. It is also clear that $v_1^n(G, \sigma_1^*; \beta^*, p)$ is arbitrarily close to $\pi_1(a_1^*, a_1^*)/(1-\delta)$ and for any $s_1 \in S_1$ $v_1^n(B, s_1; \beta^*, p)$ cannot exceed \underline{u}_1 by more than an arbitrarily small amount if p is sufficiently close to the above distribution p^* . It follows then, that for any $s_1 \in S_1$ $v_1^n(G, \sigma_1^*; \beta^*, p) \geq v_1^n(B, s_1; \beta^*, p)$ if δ is sufficiently close to 1. For a status B player, the distribution of statuses, p , will be unaffected by his choice of strategy, s_1 , if all other players follow β^* . Since σ prescribes a best response for every matching a status B player is in, $v_1^n(B, \sigma_1^*; \beta^*, p) \geq v_1^n(B, s_1; \beta^*, p)$. Thus, (b) is satisfied also and (β^*, p^*) is an asymptotically stationary norm equilibrium. Q.E.D.

Our asymptotically stationary norm equilibrium is something like "local perfect" equilibrium in the sense that we do not require perfection for all disequilibrium paths, but only for a restricted set: those associated with a status distribution which is within ϵ of the equilibrium distribution. Kandori [1988] has shown that it is possible to extend our result (in particular theorem 2) to perfect norm equilibrium. The basic idea is as follows. First allow any finite number of statuses. Then it is relatively straightforward to show the Folk theorem holds because it is well-known that punishments for finite but sufficiently many periods is enough to support cooperation in (fixed player) repeated games.⁹ More than two status levels are required for this result as different statuses are necessary to "count" the number of periods a deviant has been punished. If we allow the transition mapping to be stochastic, however, we can support cooperation with only two status levels. If the probability of regaining good status is small even when a bad player accepts punishment, having bad status results in more severe punishments than otherwise. Then assigning sufficiently small probability of regaining good status will support cooperation.¹⁰

7. Discussion and Conclusion

In this paper, we have introduced a concept of norm equilibrium. Norm equilibrium is a way to reinterpret Nash equilibrium within the context of an economy. Traditional game theory justifies Nash equilibrium as a predicted outcome when "rational" players are involved in a situation with a possible conflict of interest. Norm equilibrium, on the other hand, is a stationary societal equilibrium (in the language of Rosenthal [1979]) in which players choose their best actions, but information is limited and only the social norm (and the status distribution) are common knowledge. In order to illustrate norm equilibrium, we have employed the framework of random matching games in this paper. Several remarks are in order in this respect.

⁹ See, for example, Fudenberg and Maskin [1986].

¹⁰ For more detail, see Kandori [1988].

First, as we have emphasized, the informational needs implicit in a norm equilibrium are less than those associated with Nash equilibrium in general. The informational requirements of a social norm are still significant, however. A player needs to know, among other things, the status level of his matched opponent in order to know what action to choose. Even if not formally modelled, the interest in the model hinges on the plausibility of this information being available.

One possible way that the information might be transmitted is through a third party. Milgrom, North and Weingast [1988] (hereafter MNW) analyze a model whose theoretical structure is quite similar in spirit to our model. MNW are interested in the process of trade in the absence of a governmental legal system that can enforce contracts. In their model, players are randomly matched for a single encounter in which they might trade; without some mechanism to enforce contracts, trade will not be possible. MNW suggest that in response to the inefficiency caused by the lack of a governmental court system to enforce contracts, there developed an institution called the medieval law judge, a prominent role for whom was to serve as recorder and conduit of information concerning an individual's past transactions. Although we have not done so, one could augment our model to incorporate formally an institution of this sort to perform the information transmission role we exogenously assume in our model.

In the absence of a formal institution to carry out this informational role, we can imagine more casual transmission mechanisms. For some types of trials, it is a common practice among opposing lawyers to discuss, prior to trial, various aspects of the arguments to be made. This eliminates the necessity of both of the lawyers preparing rebuttals to all possible points that might be brought up, resulting in a reduction of work for both lawyers without altering the outcome of the case. When asked why there were not lawyers who would mislead the opposing lawyer to gain an advantage in the trial, Ms. Blackmon said that in fact there were such lawyers. In practice, when a lawyer is matched with an unknown opponent, she asks various senior colleagues about the opponent. Some opponents are known as opportunists who cannot be trusted; if she is matched with one, she simply doesn't cooperate.¹¹ This informal mechanism provides an almost perfect example of the way in which the current status of a person can precede him/her in random matchings.

We should emphasize that we don't want to hinge the interest in the model that we have presented on the two examples above. First, we have not modelled the process by which a player learns his opponent's status level, and second, even in the examples it may be farfetched to believe that the status levels are common knowledge. However, we do believe in the basic precepts of the model. The conflict situations a player finds himself in are not independent; how he behaves in a situation today will affect his future interactions. Precisely how the information about one encounter is relayed on to future players may be difficult to model, but we believe that the inter-play linkage is important to understanding some problems

¹¹ We want to thank Leslie Blackmon for this remark.

such as the nature of the Japanese firm and how Congress works. A very interesting research problem would be the careful modelling of the process by which information is transferred. Such work could presumably allow for uncertainty about the status of one's opponent.

We have provided a number of problems to which we think our model applies. For many of these problems, we don't believe that the distribution of the status levels in the society is stationary. We should point out that there are relatively simple extensions of our definitions that would allow for non-stationary. This also would be an interesting research project.

A third feature of our model we wish to discuss is the assumption that the underlying population is given and all players live forever. It is straightforward to introduce overlapping generations with finite (expected) life in the model. For example, we may assume that in each period a fixed proportion, say $1-\delta$, of the entire population (selected at random) will die. An equal number of people will be born in that period, keeping the size of the population constant. Moreover, if the pure discount factor is δ , then the discounting factor which takes account of this probabilistic death is exactly δ and our model can be easily extended to this stochastic overlapping generation model.

With this interpretation in a continuum population world, the result in section 6 may be interpreted as follows. Even if some fraction of population actually deviates from the social norm, as long as the size of deviant population is small and the condition of Theorem 2 is satisfied, the fraction of the population that follows the social norm never falls below a certain level and eventually goes back to 1. In this sense, the social norm is dynamically stable.

The next comment is on the relation of our work to that of Abreu [1988]. In a game played repeatedly by a fixed set of players, Abreu used the notion of optimal punishment schemes to characterize the degree to which players might be able to achieve cooperation in the game. In several of our examples, the social norm bears some similarity to the equilibria Abreu considers. There is one important difference, however. While the optimal punishment schemes considered by Abreu are simple to describe, they might be complicated to carry out. Recall that in our model the action taken in any encounter can depend only on the status levels of the two players in the encounter. If one wanted to encode a punishment following some history that had a player taking a particular action for n periods, one would need at least n status levels to "count" what stage of the punishment the player is in. We aren't suggesting that this sort of complexity makes Abreu's punishment strategies uninteresting; rather, for the random matching games involving large numbers of players in which we are interested, there should be a role for the type of social norms we consider to simplify the strategies that agents are following. There is a second point at which there would be some difficulty in embedding in a social norm an optimal penal code of the sort that Abreu considers. In Abreu's optimal penal codes, if a person deviates, the other players are to play the strategies that are part of the subgame perfect equilibrium that yields the lowest payoff to the deviator. In a social norm this could be accomplished by the transition function assigning the deviating player a particular status,

say B, and having the social standard of behavior prescribing that any person matched with a B choose the strategy that would correspond to the play in an optimal penal code. The difficulty would arise if there were a second deviator. In Abreu's optimal penal codes, for such a history, the strategies have players choosing those actions that are part of the subgame perfect equilibrium yielding the lowest payoff to the second deviator. For our symmetric games, this would be equivalent to assigning the second deviator the status B, but simultaneously reverting the previous deviator back to "normal" status. This isn't possible within a social norm, as a person's status is to depend only on his status and that of his current opponent and the play of their stage game; it is not allowed to depend on the play of others' stage games. Of course, we could have altered the transition function to allow this, but this would violate the informational decentralization which has motivated our study.

The first difficulty mentioned above in embedding Abreu's penal codes into social norms had to do with the number of status levels necessary to do so. This difficulty - that one would need as many status levels as periods in which a deviator was to be punished - is reminiscent of the questions of complexity of strategies in modelling play by automata (see, e.g., the survey by Kalai [1987]). There is a great deal of similarity both between the motivations behind modelling play by finite automata and behind social norms and in the problems the two studies raise. The status levels in our model have a close relationship with the states of an automaton. The main differences formally between those models and ours is that a social norm prescribes an action as a function of both his and his opponent's status. This would be equivalent to having an automaton-like machine whose action was a function of both its state and the state of the opposing automaton.

This discussion highlights a feature of our approach that is worth discussing a bit further. A social standard of behavior prescribes an action that is a function of the status levels of the matched pair. This allows for a comparison of a particular strategy across different sized games, that is, across random matching games with different numbers of players. This allowed us to investigate how well the continuum model approximated large finite societies keeping fixed the stage game to be played. This question is of interest in its own right.

We should point out that the relative ease of constructing social norms that sustain cooperation comes at a cost. The strategies embodied in a social standard of behavior need not (and will not, in general) be subgame perfect. Massive deviations from the social standard of behavior may cause the status distribution to change so dramatically that it is no longer optimal to follow the standard of behavior. As emphasized above, however, we think the "local perfectness" of norm equilibria serves much the same purpose as subgame perfection.

The last comment we will make about the model is our choice of a random matching game. In reality, social status is likely to be linked to the organization for which a player works and the job status he holds. In this sense, we think that it is appropriate to apply norm equilibrium in the case in which the

stage game a player will play in each period may depend upon his status as well. A simple version of this idea is exploited in Okuno-Fujiwara [1989]. We plan to extend our analysis to this more general setting in the future.

7.1 Related Work

Bendor and Mookherjee [1990] present a model in which they compare third party sanctions with direct sanctions. In their model a group of n people is to play a game repeatedly. In each period, each player is to choose a distinct action toward each of the other $n-1$ players. They examine when third party sanctions can be useful, that is, they ask when higher payoffs can be supported by equilibria in which a given player punishes "bad" behavior of other players, even when the bad behavior was not directed toward the given player. There is a similarity between our work and that of Bendor and Mookherjee in that one can think of our social standard of behavior calling upon players to carry out what might be termed third party sanctions. Our model and motivation differs from theirs, however, in that our main concern is in the informational role that what we call status plays in the coordination of play among a large number of people who are randomly matched.

Akerlof [1980] presents a very interesting discussion of caste behavior that bears a resemblance to some of the norm equilibrium examples in this paper. As in the Bendor and Mookherjee paper, the Akerlof model is one of fixed matching and, hence, can play no informational role.

REFERENCES

- Abreu, D. [1988]
"On the Theory of Infinitely Repeated Games with Discounting," Econometrica, Vol.56, pp.383-82.
- Akerlof, G.A. [1980]
"A Theory of Social Custom, of which Unemployment may be One Consequence," Quarterly Journal of Economics, Vol.94, pp.749-75.
- Aumann, R. [1987]
"Correlated Equilibrium as an Expression of Bayesian Rationality," Econometrica, Vol.55, pp.1-18.
- Bendor, J. and D. Mookherjee [1990]
"Norms, Third Party Sanctions and Cooperation," Journal of Law, Economics and Organization, 6, pp. 33-64.
- Bertsekas, D.P. [1976]
Dynamic Programming and Stochastic Control, Academic Press, New York.
- Binmore, K.G. [1987]
"Modeling Rational Players," Economics and Philosophy, Vol.3, pp.179-214.
- Derman, C. [1970]
Finite State Markovian Decision Processes, Academic Press, New York.
- Fudenberg, D. and E. Maskin [1986]
"The Folk Theorem in Repeated Games with Discounting and with Incomplete Information", Econometrica, Vol.54, pp.533-54.
- Howard, R. [1960]
Dynamic Programming and Markov Processes, John Wiley and Sons, New York.
- Kalai, E. [1987]
"Artificial Decisions and Strategic Complexity in Repeated Games," Northwestern University Discussion paper.
- Kalai, E. and A. Neme [1988]
"The Strength of a Little Perfection," Technical Report No.544, Institute for Mathematical Studies in the Social Sciences, Stanford University.
- Kandori, M. [1988]
"Social Norms and Community Enforcement," mimeo, Stanford University.
- Kaneko, M. [1987]
"The Conventionally Stable Sets in Noncooperative Games with Limited Observations I: Definitions and Introductory Arguments," Mathematical Social Sciences, Vol.13, pp.93-128.
- Kreps, D. [1989]
"Corporate Culture and Economic Theory," in Alt and Shepsle (eds.) Rational Perspectives in Political Sciences, Harvard University Press, Cambridge.

- Kreps, D. and R. Wilson [1982]
"Reputation and Imperfect Equilibrium," Journal of Economic Theory, Vol.27, pp.253-79.
- Kreps, D., P. Milgrom, J. Roberts and R. Wilson [1982]
"Rational Cooperation in the Finitely Repeated Prisoners' Dilemma," Journal of Economic Theory, Vol.27, pp.245-52.
- Milgrom P., D. North and B. Weingast [1988]
"Third Party Enforcement of Norms and Contracts: A Theoretical-Historical Analysis," mimeo.
- Okuno, M. [1984]
"Corporate Loyalty and Bonus Payments: An Analysis of Work Incentives in Japan," in M. Aoki ed. The Economic Analysis of the Japanese Firm, North-Holland, Amsterdam.
- Okuno-Fujiwara, M. [1987]
"Monitoring Cost, Agency Relationships, and Equilibrium Modes of Labor Contracts," Journal of the Japanese and International Economies, Vol.1, pp.147-167.
- Okuno-Fujiwara, M. [1989]
"On Labor Incentives and Work Norm in Japanese Firms," Journal of the Japanese and International Economies, Vol.4, pp. 367-84.
- Rosenthal, R.W. [1979]
"Sequences of Games with Varying Opponents," Econometrica, Vol.47, pp.1353-66.
- Rosenthal, R.W. and H.J. Landau [1979]
"A Game-Theoretic Analysis of Bargaining with Reputations," Journal of Mathematical Psychology, Vol.20, pp.233-55.
- Whittle, P. [1983]
Optimization over Time, John Wiley and Sons, New York.