Social Norms and Community Enforcement

MICHIHIRO KANDORI
Princeton University

The present paper extends the theory of self-enforcing agreements in a long-term relationship (the Folk Theorem in repeated games) to the situation where agents change their partners over time. Cooperation is sustained because defection against one agent causes sanction by others, and the paper shows how such a "social norm" is sustained by self-interested agents under various degrees of observability. Two main results are presented. The first one is an example where a community can sustain cooperation even when each agent knows nothing more than his personal experience. The second shows a Folk Theorem that the community can realize any mutually beneficial outcomes when each agent carries a label such as reputation, membership, or licence, which are revised in a systematic way.

1. INTRODUCTION

It is widely recognized that in many economic transactions, informal means are employed to execute mutually beneficial agreements. As S. Macaulay (1963) points out, "social pressure" and "reputation" are perhaps more widely used than formal contracts and filing suits. In many cases, people behave honestly because honesty is rewarded and/or defection is punished in future transactions.

Such informal enforcement mechanisms can be classified into two categories. One is *personal enforcement*, in which cheating triggers retaliation by the victim. These mechanisms are effective only if quick and substantial retaliations are available—that is, they work best in frequent and long-term relationships. The Folk Theorem in the repeated game literature (Rubinstein (1979) and Fudenberg and Maskin (1986)) provides a formal model of personal enforcement, showing that any mutually beneficial outcome can be sustained as a subgame-perfect equilibrium if the same set of agents frequently play the same stage game *ad infinitum*. However, many important transactions are *in*frequent in nature. As economic historians argue, the division of labour and specialization are important driving forces of economic progress. Potential gains are larger in diverse transactions with different specialists than with fixed partners. Therefore, the control of incentives in such an infrequent trade is of vital importance to understand the organization of economic transactions. This observation leads to the second category of informal enforcement mechanisms, *community enforcement*, where agents change their partners over time and dishonest behaviour against one partner causes sanctions by other members in the society. The present paper is devoted to the study of such mechanisms.

The specification of desirable behaviour together with sanction rules in a community may be regarded as a *social norm*, and we analyze how such social norms work to support efficient outcomes in infrequent transactions. Our approach to this problem assumes the standard axiom of economics that agents only care about their own utility; that is, we do *not* assume that people follow a social norm for its own sake, but we investigate how such a rule is sustained by self-interested community members. For a social norm to be

sustainable, it must provide proper incentives to the members in every respect. Thus, not only are deviators from the desired behaviour punished, but a person who fails to punish is in turn punished. In other words, we will investigate the (subgame-perfect or sequential) equilibria of the game played by community members.

In small communities where members can observe each other's behaviour, community enforcement works in much the same way as personal enforcement. We can easily modify the usual Folk Theorem to show that any efficient and individually rational outcome can be sustained when there are frequent interactions among the community members as a whole, and this is true even if the transactions among any given pair of agents are infrequent. This observation shows that changing partners itself is unimportant and the crux of the matter is information transmission among the community members. Given this, we propose the following research programme: What is the minimal information transmission necessary to sustain efficient outcomes by community enforcement? As a first step to answer this question, the present paper shows what a community can achieve under various degrees of observability.

When the observability is not perfect, each agent typically possesses a piece of information about the history of trades in the community which may not be known to others. The presence of such private information characterizes the main theoretical difference between our models and the models of standard repeated games, where all relevant information is assumed to be common knowledge among all agents in the game.[1] We encounter new problems and the analysis of such games turns out to be a non-trivial extension of the usual repeated games. The source of difficulty is illustrated by the following observation. Suppose the community can somehow "mark" deviators, say, by putting dark spots on their foreheads, and suppose that the community norm requires that an agent should cooperate if and only if the partner is unspotted. Now consider the agents' incentives to follow the norm. Clearly, no one wishes to deviate from the equilibrium path if the punishment is severe enough. To show that this social norm is an equilibrium, however, we must also show that everybody has incentives to follow the norm *after any history*. Consider two unspotted players matched to each other, and suppose that one of them is likely to encounter many spotted partners in the future. If the punishment is costly to carry out, this may destroy the incentives for them to cooperate, because one of them doesn't have enough stake in the future. And just as in any games with private information, their incentives depend on how this information is known to them, which is determined by their personal experiences (for example, how many deviators each of them has seen and when they were observed). Checking incentives and specifying equilibrium behaviour in those cases can potentially be very complicated, because (i) private information (what the players have observed) and its distribution will be increasingly complicated over time, and (ii) the players' private information do not come from a common prior distribution after deviations occur. Those points show that the game does not have the usual recursive structure which is possessed by the standard models of repeated games with public information. The analysis of models with public information is greatly simplified because after any history the continuation strategies correspond to an equilibrium in the original game (see Abreu, Pearce and Stacchetti (1990)). In contrast, in our model continuation strategies correspond to *a part* of (partially) correlated

equilibrium where the correlation device is the private observation of the past history. Point (i) above says that the distribution of the correlation device is non-stationary and point (ii) says that the players do not share the same prior of the correlation device. Hence it is not obvious if the seemingly straightforward social norm described above is in fact an equilibrium; players' behaviour off the equilibrium may need to be modified, and the modifications can be highly complicated.

The paper presents two main results. The first one is an example where a community can sustain cooperation even when each agent only observes the results of the trades he engaged in. An agent knows what his partners did to him, but he does not observe their identity nor what they have done to other agents. We show that even in that case, cooperative behaviour can sometimes be sustained. The social norm supporting cooperation in such a situation uses a contagious process of defection, and the scope and limitations of such a social norm will be discussed. The second results gives the Folk Theorem under mild assumptions, assuming the existence of a mechanism or institution which processes a certain class of information honestly. Being interested in the minimal information dissemination, we put the following restrictions on the form of information transmission.

1. A label is attached to each agent.
2. Before executing trade each agent observes his and his partner's label.
3. A player and his partner's actions and labels today determine their labels tomorrow.

Requirement 2 says that each agent has only local knowledge of the current state of the community, and requirement 3 indicates that the revision of labels does not require knowledge of the entire society, and allows the information processing to be potentially decentralized. Reputation, membership, citizenship, social status, and credit cards can be regarded as examples of this class of information transmission mechanisms.[2] We show how such an information processing device, which does not have any enforcement power of its own, facilitates community enforcement of efficient trade. These mechanisms are studied in the pioneering work by Okuno-Fujiwara and Postlewaite (1989) and they show that the Folk Theorem holds in a random matching game when such information transmission is available. However, as we have seen, the presence of private information complicates the analysis, and to overcome this difficulty they either assume infinite population and uniform random matching (where no one expects to meet a defector after any finite number of deviations), or employ an equilibrium concept which is weaker than sequential equilibrium. Under mild assumptions on the payoff functions, however, the present paper proposes a simple strategy profile which defines an exact (i.e. sequential) equilibrium for any population size and matching rule, but still is able to support any feasible and individually rational payoffs. The equilibrium reveals the importance of such notions as repentance and forgiveness in social norms.

The role of information in matching games is also analyzed by Rubinstein and Wolinsky (1990) in a different context. In their model, buyers and sellers are randomly matched, and once a seller and a buyer agree to trade, they leave the market. So the structure of their model is different from repeated games, but they derive results which are similar in flavour to ours. Their analysis shows that if each player has enough

2. The following may be helpful to motivate the reader who is familiar with the dynamic programming decomposition of repeated games (See Abreu, Pearce and Stacchetti (1990)). If we treat the whole matching game as a huge repeated game with all community members, continuation payoffs serve as labels summarizing all relevant information about the past history, if each player can observe all players' labels and the revision of each label can depend on the actions and labels of all players. The question is whether we can find something analogous which works in a decentralized fashion.

information about other players' history, there exist many equilibria which are not sustained when such information is unavailable.

The paper is organized as follows. Section 2 introduces the formal model of repeated matching games. Section 3 shows that community enforcement works in much the same way as personal enforcement under public observability. Section 4 provides an example where a community can sustain an efficient outcome even when each agent observes the history of his transactions only. Section 5 deals with the Folk Theorem when labels are attached to agents. The final section discusses the related literature and possible extensions of our model.

2. REPEATED MATCHING GAMES

The basic structure of the repeated matching game is described as follows. The set of players $N = \{1, \ldots, 2n\}$ is partitioned into two sets of equal size, $N_1 = \{1, \ldots, n\}$ and $N_2 = \{n+1, \ldots, 2n\}$, where N_k is the set of type-k players ($k = 1, 2$). In each stage, each type-1 player is matched with a type-2 player according to some matching rule, and they play a two-player stage game. This procedure is repeated infinitely and each player's total payoff is the expected sum of his stage payoffs discounted by $\delta \in (0, 1)$. Most of the results in this paper do not depend on the way players are matched. For example, they can be endogenous and history-dependent. Let $\mu(i, t)$ be player i's match at time t. For some cases, we assume *uniform random matching* in which

$$\text{Prob} \{\mu(i, t) = j\} = 1/n \quad \text{for all } i \in N_1 \text{ and } j \in N_2 \text{ and for all } t$$

and the matching in each stage is independent. The stage game is described by a payoff function $g : A \to \mathbb{R}^2$, where $A = A_1 \times A_2$ and A_k is the finite set of actions for type-k players ($k = 1, 2$). The minimax point $M^1 \in A$ for type-1 players is defined by

$$M_2^1 \in \arg\min_{a_2 \in A_2} (\max_{a_1 \in A_1} g_1(a_1, a_2))$$

$$M_1^1 \in \arg\max_{a_1 \in A_1} g_1(a_1, M_2^1)$$

and M^2 is defined similarly. For simplicity, the "mutual minimaxing point" (M_1^2, M_2^1) is denoted by $m = (m_1, m_2)$. We normalize the payoffs in such a way that the minimax payoffs are equal to zero $(g_1(M^1) = g_2(M^2) = 0)$. Finally, the set of feasible and individually rational payoffs in the stage game is defined by

$$V = \{v \in \text{cog}(A) \mid v \gg 0\},$$

where $\text{cog}(A)$ is the convex hull of the set $g(A)$.

3. FOLK THEOREMS UNDER PUBLIC OBSERVABILITY

We first point out that, if all agents' past actions are publicly observable, every point $v \in V$ is sustained by a perfect equilibrium, even when agents change their partners over time. Perhaps the simplest way to prove the assertion is to utilize the same strategies as the Folk Theorem for two-player repeated game.

Proposition 1. *If $v \in V$ is supported by an equilibrium in the two-player repeated game for some δ, then it is also supported by an equilibrium in the matching game for the same δ with arbitrary population size and matching rule.*

The point of this assertion is that the same δ works for any population and matching rule; under public observability, each agent has as strong an incentive to cooperate as if he faced the same partner in each period. This is true even when the chance of meeting the same partner in the future is very small, or even zero.

Proof. The players in the matching game start by playing the equilibrium path of the two-player game. If the type-l player deviates $(l = 1, 2)$, *all* type-l's are punished by *all* of the other type, using the same punishment strategies as in the two-player game. The same principle applies to any further deviations, and simultaneous deviations are ignored. Then, it is clear that the incentives of each player are identical with those in the two-player game because each player encounters the same sequence of action profiles as in the two–player repeated game, only the opponents being changed over time. Hence the prescribed strategies are in fact a perfect equilibrium.   ‖

In the above proposition, an equilibrium is constructed where a large number of innocent players are held responsible for the deviation of a defector. Since punishment never occurs on the equilibrium path, this does not entail actual welfare loss, but such an equilibrium is intuitively unappealing. If only the deviator is to be punished and innocent pairs are to play the originally prescribed actions, however, more conditions are necessary to achieve subgame perfection. As was discussed in the introduction, difficulties may arise in providing incentives for innocent players after some players' defections. To maintain the faith of innocent agents, we have to avoid the situation where too much burden of punishment is imposed on them. Such a situation arises when the community is highly populated with "guilty" agents, or the matching is highly non-uniform. Introducing some "forgiveness" in the social norm, we can avoid the first problem.

Proposition 2. *Under uniform random matching, any payoff point $v \in V$ is sustained by an equilibrium where only defectors are punished, if $\delta \in (\delta^*, 1)$, where δ^* can be chosen independently of the population size.*

Remark. Without the independence of population, the statement would be trivial because personal retaliations may be effective when players are very patient.

Proof. Let $a^* \in A$ be the action profile achieving the payoff v. On the equilibrium path, action profile a^* is played.[3] If a player deviates from the equilibrium path, he and his opponents will play mutual minimax m for T periods and then revert to the original actions a^*. Meanwhile, all the other players keep on playing a^*. When the deviator defects while he is punished, the T-period punishment is restarted. If another player deviates when a player is punished, the former deviator is forgiven and *only the latest deviator is punished.* Simultaneous deviations are ignored.

Let $x = g(m)$ and define the average punishment payoffs

$$V = (1 - \delta^T)x + \delta^T v.$$

Choose $\delta^T \in (0, 1)$ so that $V \gg 0$. Then, no player deviates when he is punished, because he can earn at most $0 + \delta V_k$ $(k = 1, 2)$ by a deviation, which is less than the original payoff V_k.

---

3. Or use a correlated strategy if necessary, and assume that the corelation devices are observable. The same remark will apply to all propositions below.

When a type-k player $(k = 1, 2)$ is innocent, his average payoff is at least

$$\Pi = (1-\delta)x_k + (\delta - \delta^T)[(1 - 1/n)v_k + (1/n)x_k] + \delta^T v_k,$$

while by a deviation he receives at most

$$\Pi' = (1-\delta)v_k^* + \delta V_k,$$

where $v_k^* = \max_{a \in A} g_k(a)$. As $\delta \uparrow 1$ with δ^T held constant,

$$\Pi - \Pi' \to (1 - \delta^T)(1 - 1/n)(v_k - x_k) > 0,$$

and therefore the deviation is unprofitable. Since the difference $\Pi - \Pi'$ is smallest for $n = 2$, the lower bound on δ for $n = 2$ works for all larger n.   ‖

Those assertions show that under public observability it is unimportant who punishes the defector; community enforcement works just as well as personal retaliation in the usual repeated games. In this sense, observability in the community is a substitute for having a long-term frequent relationship with a fixed partners.

4. COOPERATION WITHOUT INFORMATION PROCESSING

This section analyzes an example where very limited observability is svailable, and illustrates what kind of difficulty a community confronts in such a case. The example shows the ability and limitations of social norms with very little information, which helps to explain why a community needs a certain amount of information to sustain its norms. In this section, the stage game is assumed to be the prisoner's dilemma game described by Table 1.

	c	d
c	1,1	$-\ell$, 1+g
d	1+g, $-\ell$	0,0

TABLE 1

The numbers g (the gain from defection) and l (the loss when cheated) in the above table are positive, and action c stands for cooperation and d stands for defection. The first number in each entry indicates the row player's payoff and the second is the column player's. Since the game is symmetric, assume n pairs are formed randomly out of $2n$ players without any distinction between type-1 and type-2. For notational convenience, let $M = 2n$ denote the population size.

We consider the case in which each player observes only the history of action profiles in the stage games which he has played. In this situation, a player knows nothing about the identity of players or what has gone on in the rest of the community. Direct communication among the players is assumed to be non-existent. We call this case *no information processing*. In this situation, there is no way to implement the equilibria in the Folk Theorems under perfect information described in Section 3, and personal

retaliation is not available either. One may further conjecture that the only equilibrium is the repetition of the Nash equilibrium in the stage game. However, we will show that the community may sometimes still sustain cooperation by what may be called the "contagious equilibrium".

The idea of the "contagious equilibrium" is that trust is attached to the community as a whole, not to each individual. Therefore, a *single* defection by a member means the end of the *whole* community trust, and a player who sees dishonest behaviour starts cheating all of his opponents. As a result, defection spreads like an epidemic and cooperation in the whole community breaks down. We are going to show that this somewhat peculiar social norm is indeed justified as a sequential equilibrium of the game.[4]

To describe the equilibrium conditions, let us introduce the "types" of players. A player is of *type c* if nobody has deviated in the past history of his stage games, and otherwise he becomes *type d* . This means that, once a player cheats or is cheated, he becomes type *d* for good. The equilibrium is described by a simple rule; each player chooses the action which is equal to his type.

Theorem 1. *Under uniformly random matching, the contagious strategy described above constitutes a sequential equilibrium strategy for any given g and M if δ and l are sufficiently large.*

Remark. If a player can believe that a large number of players have already been cheated whenever he seees a defection, then it is in his best interest to defect forever, as is prescribed by the contagious strategy. This system of beliefs, however, is *not consistent* in the sense of Kreps–Wilson (1982)[5] so that it is not compatible with sequential equilibrium. For example, if a player sees a defection in the first period, he must believe that everybody else is cooperating, because the definition of consistent beliefs requires that defections by different players are statistically independent. By the same token, when a player defects when he has seen no defection, he must believe that he is the first person to defect in the society. Therefore, checking the incentive to follow contagious process of defection is a non-trivial problem.

The intuition of Theorem 1 is straightforward. When a player defects, his future payoffs are destroyed by the contagious process of defection. The loss of future payoffs outweighs the instantaneous gain if the player does not discount the future payoffs much. Hence a player never deviates from the equilibrium path if δ is close to unity. On the other hand, once the contagious process has started, a player faces the following trade-off. If a player chooses *c* rather than *d*, he can slow down the contagious process to enjoy high payoffs in the future. However, playing *c* is costly in terms of instantaneous gain. In particular, there is always a possibility of meeting a *d*-type, in which case playing *c* induces loss *l*. Therefore, if *l* is large enough, it is in a player's best interest to follow the contagious process once it has started.

To derive the formal equilibrium conditions, let us introduce some notation. Let X_t be the number of type *d* players at time *t*, and define the $M \times M$ transition probability

---

4. Harrington (1989) independently discovered a similar equilibrium in the case of Poisson matching in continuous time.

5. A system of beliefs is consistent if it is derived by the following procedure. First, perturb each player's strategy so that all actions are taken with positive probabilities. The perturbation of strategies is required to be independent across players and over time. Then we can unambiguously define the beliefs by Bayes' rule, because all information sets are reached with positive probabilities. Finally, take the limit of the beliefs as the perturbation tends to zero.

matrix $A = (a_{ij})$ where $a_{ij} = \text{Prob}(X_{t+1} = j \mid X_t = i)$. Also define $B = (b_{ij})$ by $b_{ij} = \text{Prob}(X_{t+1} = j \mid X_t = i$ and one of the d-types deviates to play c at time $t)$, and let $H = B - A$. Matrix H indicates how the diffusion of d-types is delayed by unilateral deviation of one of d-types.[6] The probability that a d-type player encounters a c-type is summarized by a column vector,

$$\rho = \frac{1}{M-1}(M-1, M-2, \ldots, 1, 0)^\top$$

where the ith element of ρ is the conditional probability that a d-type player meets a c-type when there are i d-types. Finally, e_i stands for $1 \times M$ vector whose ith element is 1 with zeros everywhere else. With this preparation, the formal equilibrium conditions are given as follows.

Lemma. *The contagious strategy constitutes a sequential equilibrium if*

$$\frac{1}{1+g} \geq (1-\delta)e_1(I - \delta A)^{-1}\rho \tag{1}$$

and

$$\frac{\left(\dfrac{M-k}{M-1}\right)g + \left(\dfrac{k-1}{M-1}\right)l}{1+g} \geq \delta e_k H (I - \delta A)^{-1}\rho \quad \text{for} \quad k = 2, 3, \ldots, M. \tag{2}$$

The proofs of Theorem 1 and the Lemma are given in the Appendix.

Note that the contagious equilibrium is more effective at deterring defection than personal retaliation, where a defector is punished only by the person he cheated. This is because under the contagious equilibrium a defector not only loses cooperation with the victim, but also some of the future payoffs with other players are destroyed by the contagious process. More precisely, if (g, δ) is such that a player is just indifferent between cooperating and cheating under personal retaliation, he strictly prefers to cooperate under the contagious equilibrium. Since the incentive to start the contagious process (the inequality (2)) can be maintained when being cheated is costly (when l is large), there are situations in which cooperation is sustained by the contagious equilibrium, but not by personal retaliation.

There are two things going on in this equilibrium; information transmission and punishment. A player who was cheated starts cheating, which has the effect of telling the other members that someone has deviated, and of punishing the cheater who cannot be identified. The players have positive incentives to follow such behaviour because the community trust will completely break down sooner or later irrespective of his action, so that there is no point in sticking to honest behaviour (see condition (2)). As a result, quick breakdown of community trust occurs after a deviation in a self-enforcing way.

Three points should be noted about the limitations of social norms under no information processing:

(1) The above social norm does not have as much cheating deterrence power as the community enforcement under perfect information, because a cheater is not punished immediately. And the conditions for sustaining cooperation under

6. For the closed forms of A and H, see the original version of this paper (CARESS working paper No. 89–14, University of Pennsylvania, 1989).

perfect information are much weaker, so that there exists a case in which cooperation is sustained under perfect information but not under no information processing.

(2) In a large population, it is difficult to sustain cooperation under no information processing, because it takes a long time for a defector to meet players who have met someone who have met ... the player who was cheated by the defector. Given this observation, it is straightforward to show the following impossibility result.

Proposition 3. *Consider a general matching game defined in Section 2 with uniform random matching and no information processing. Suppose $v \in V$ is not the value of a Nash equilibrium of the stage game. Then, for any δ, v cannot be sustained by a sequential equilibrium if the population size is large enough.*

The proof is contained in the Appendix.

(3) Even when cooperation is sustained, the norm under no information transmission has the unfortunate feature that innocent players will necessarily be punished. Also, it is fragile in the sense that a little bit of noise ("trembling hands") causes complete breakdown of cooperation in the community. This may be the main reason why we do not observe such a norm very often. Given this point, we feel that a norm should be evaluated not only by its equilibrium payoffs, but also by its (suitably defined) "robustness".

Given those limitations of social norms under no information processing, we will now turn to the case where information is transmitted in a systematic way. Unlike the case of no information transmission, we will show that we can construct equilibria which work for any population size and are robust in a suitably defined sense.

5. LOCAL INFORMATION PROCESSING

This section is devoted to the study of social norms when there is a mechanism or institution which systematically processes some information among community members. The processing of information is treated as exogenous, and it is assumed that information is transmitted honestly. Assuming the existence of such a mechanism our main concern here is to identify a class of simple and "robust" equilibria which sustain any feasible and individually rational outcomes by community enforcement with as little information as possible. To this end, we require a series of desirable properties and show that there exists a class of equilibria which satisfy all of them.

The first requirement is informational decentralization. We require that both community members' decision making and the update of information can be made without the knowledge of the entire society. Although computer networks in advanced societies may facilitate instant dissemination of information about all the community members, the presumption here is that such a central coordination device is costly. We restrict our attention to the situations where each agent carries a label and the necessary information is transmitted by agents' labels. Reputation, membership, citizenship, social status, and credit cards can be regarded as examples of the labels. When two players are matched, they observe each other's label first and then take some actions. After that, their labels are updated depending only on *their* original labels and actions by a given rule. Thus, the choice of actions and the revision of labels in this setting are based only on the local

information, and this information structure is referred to as *local information processing*. Such a mechanism is first introduced by Okuno-Fujiwara and Postlewaite (1989). It is formally defined as follows.

Definition 1.   A matching game with *local information processing* has the following information structure.

1. A state $z_i(t) \in Z_k$ is assigned to player $i \in N_k$ $(k = 1, 2)$ at time t.
2. When player i and j meet at time t and take actions $(a_i(t), a_j(t))$, their next states are determined by

$$(z_i(t+1), z_j(t+1)) = \tau(z_i(t), z_j(t), a_i(t), a_j(t)).$$

3. At time t, player i can observe at least $(z_i(t), z_{\mu(i, t)}(t))$ before choosing his action.

Although this information structure potentially can convey large amounts of information when the community can utilize a large number of labels, we are interested in finding as simple a process as possible in this class. The minimal number of labels will be discussed later in this section.

The second requirement is the simplicity of decision making and the robustness with respect to the information structure. In potential applications of the present model, it is highly likely that the players may be able to observe something other than the labels of their partners, so it is desirable that the equilibrium does not depend on the fine detail of the information structure. In other words, the labels should be sufficient statistics and it should summarize all relevant information for agents' decision making. An equilibrium in which each labels are sufficient statistics is called *straightforward*.

Definition 2.   A sequential equilibrium in a matching game with local information processing is *straightforward* if given that all other players' choice of actions depends only on their and their partners' labels, a player's best response also depends only on his and his partner's labels, even if he had more information than those;

$$a_i(t) = \sigma_k(z_i(t), z_{\mu(i, t)}(t)) \quad \text{for } i \in N_k, k = 1, 2.$$

In a straightforward equilibrium, each agent does not need any information other than his and his partner's labels, and this, together with the local information processing, represents a strong form of informational decentralization.

Third, we also require that the equilibrium should be independent of such fine details of the game as the matching rule and the size of population. Since matching in the real world is determined in a rather complicated way, the applicability of the model would be limited if it depended too much on the specification of the matching mechanism. Note that this requirement drastically reduces the information an agent needs to know. To check if the norm is actually an equilibrium, one does not have to know the number of people in the community, or how people are matched.

The next requirement is the stability of the equilibrium. As the example in the previous section shows, a community may be able to sustain an efficient outcome under limited observability by using harsh punishments, but the equilibrium may be fragile in the sense that a small amount of noise (deviation) causes the breakdown of the honest behaviour in the whole community. On the contrary, if the equilibrium always goes back to the original payoff point, it is robust to the mistakes of players and it also allows players to test various actions in order to learn the social norm. A strong notion of dynamic stability is defined as follows.

Definition 3.   An equilibrium sustaining payoffs $v \in V$ is *globally stable* if for any given finite history of actions h,

$$\lim_{t \to \infty} E(v_i(t) \mid h) = v_k \quad \text{for all } i \in N.$$

where $v_i(t)$ is player i's continuation payoffs at t and $E(\cdot \mid h)$ is the conditional expectation.

Finally, equilibria should be simple. For a social norm to be effective, the rule must be easy to understand. Especially when there are newcomers to the community, simplicity is crucial for them to learn and follow the community norm. An important measure of the simplicity of an equilibrium is the number of different actions prescribed to the players on and off the equilibrium path. Another measure would be the number of the states ($\# Z_k$, $k = 1, 2$). Those numbers should be small for an equilibrium to be appealing.

In the rest of this section, it will be shown that those requirements can be achieved for some important classes of stage games. The first class of games is those in which only one type of player has some incentive problems. A canonical example is a credit market, where the lender has a choice between lending money or not, and the borrower can either pay back or not. An efficient choice of action is to lend and to pay back, where only the borrower can gain by unilateral deviation. In such cases, it is well known that a straightforward and globally stable equilibrium can be constructed to sustain the efficient point. For completeness, let us define this class of games and briefly state the result.

Definition 4.   A stage game g has a *one-sided incentive problem* for the type-1 player at $a^* \in A$ if a_2^* is a best response to a_1^* and there is a Nash equilibrium in a^0 in the stage game such that $g_1(a^0) < g_1(a^*)$. A symmetric definition applies to the type-2 player.

Then, the following assertion is rather obvious; cooperation is sustained because if a player with the incentive problem cheats, he will simply be punished by his partners for a finite time by the one-shot Nash Equilibrium.

Proposition 4.   *If g has a one-sided incentive problem at $a^* \in A$, then, under local information processing, there exists a straightforward and globally stable equilibrium supporting a^*, if $\delta \in (\delta, 1)$ for some δ^*, and the statement is independent of the matching rule and the population size.*

Proof.   See Appendix.   ‖

The credit bureau is an example of the possible application of the above proposition. Credit bureaus are private institutions which keep track of consumers' credit history and sell the information to potential lenders, who decide whether to give credit to a consumer depending on the record . A recent paper by D. Klein (1989) analyzes a model of credit bureaus. Another example is the reputation of a firm which can potentially cheat its consumers by supplying low-quality goods. This was analyzed by B. Klein and K. B. Leffler (1981). Avner Greif's (1989) remarkable analysis of medieval trade reveals the potential applicability of the idea of community enforcement to the problems of economic history before the creation of law enforcement agencies. He shows some evidence that medieval Jewish traders controlled their agents' behaviour through their reputation, and again the structure of the model is characterized by the one-sidedness of the incentive problem (only the agents have the opportunity to cheat). The point is that the simple form of community enforcement in those examples is due to the one-sidedness of the

incentive problem, and is not readily extended to the general case of two-sided incentive problems unless we impose some assumptions.

If both types of players have incentive problems, a more sophisticated treatment is necessary. The source of the difficulty is to check the incentives of "honest" players matched together when there are many "guilty" players in the society.[7] The second class of the games we consider cover those cases and are characterized by the following property. (Recall that m represents the mutual minimax point.)

Assumption (A1).   There exists $r \in A$ such that

$$g_1(m_1, r_2) > g_1(m) \geqq g_1(r_1, m_2)$$

$$g_2(r_1, m_2) > g_2(m) \geqq g_2(m_1, r_2).$$

This property has the following implications. We will introduce an equilibrium in which two different punishments are used. If a guilty player meets another guilty player, they mutually minimax each other. If a guilty player and an innocent player are matched, the innocent minimaxes the guilty but the latter is not supposed to minimax the former; instead, he "repents" by choosing an action r, which is less harmful for the opponent (the strict inequalities in (A1)) but more costly for himself (the weak inequalities). By those dual punishments, honest behaviour is enforced for innocent players even though they expect to encounter guilty players for a long time in the near future. The property (A1) is a natural assumption and satisfied by many important games. Examples in this class include the prisoner's dilemma game, where m is defection and r is cooperation, and bilateral trades where the minimax point is no trade and r is interpreted as paying a fine to innocent players. With those two punishments, we can construct a "robust" equilibrium which sustains any mutually beneficial outcomes.

Theorem 2.   *Under the Assumption (A1), every point $v \in V$ is sustained by a straight-forward and globally stable equilibrium with local information processing, if $\delta \in (\delta^*, 1)$ for some δ^*, which is independent of the matching rule and the population size. Furthermore, only three actions are prescribed to each player.*

Proof.   The state spaces are finite sets $Z_1 = Z_2 = Z = \{0, 1, \ldots, T\}$. State 0 indicates that the player is innocent and otherwise he is guilty. The T states for guilty players count the number of punishments which last for T periods. Let $a^* \in A$ be the action profile to achieve the designated payoffs v. The equilibrium utilizes only three actions, a^*, m, and r. If two innocent players are matched, they choose the designated action a^*. If two guilty players meet, they mutually minimax each other. If an innocent player encounters a guilty player, the former minimaxes the latter but the latter chooses the "repenting" action r defined in (A1). That is, for $z \in Z \times Z$,

$$\sigma(z) = \begin{cases} a^* & \text{if } z = (0, 0) \\ (m_1, r_2) & \text{if } z_1 = 0, z_2 \neq 0 \\ (r_1, m_2) & \text{if } z_1 \neq 0, z_2 = 0 \\ m & \text{if } z_1, z_2 \neq 0. \end{cases}$$

7. Recall the discussion in the introduction about the difficulty of the social norm with "spotted" deviators.

The state transition obeys a simple rule; any deviation starts a T-period punishment. For $z \in Z \times Z$,

$$\tau_1(z, a) = \begin{cases} 0 & \text{if } z_1 = 0 \text{ and } a_1 = \sigma_1(z) \\ z_1 + 1 \pmod{T+1} & \text{if } z_1 \neq 0 \text{ and } a_1 = \sigma_1(z) \\ 1 & \text{if } a_1 \neq \sigma_1(z), \end{cases}$$

and the symmetric formula applies to type-2 players.

Then the following conditions ensure that type-1 players cannot profitably deviate from the prescribed strategies. Symmetric conditions apply to type-2 players.

$$(1 - \delta^T) g_1(r_1, m_2) + \delta^T v_1 \geq 0 \tag{3}$$

$$(1 - \delta^T) \min \{g_1(m_1, r_2), v_1\} + \delta^T v_1 \geq (1 - \delta) v_1^* + \delta[(1 - \delta^T) g_1(m) + \delta^T v_1], \tag{4}$$

where $v_1^* = \max_{a \in A} g_1(a)$. Inequality (3) guarantees that a type-1 player follows the designated action when he is guilty. This is seen as follows. If a player is guilty, his payoff is at least

$$x(1) + \delta x(2) + \cdots + \delta^{T-1} x(T) + (\delta^T + \delta^{T+1} + \cdots) v_1, \tag{5}$$

where $x(t)$ is either $g_1(m)$ or $g_1(r_1, m_2)$ depending on the opponent. Note that if all other players conform to the equilibrium strategies, every player will be "forgiven" and will become innocent T periods ahead, so that the player will surely be matched with honest players and will receive v_1 after that. If he deviates, on the other hand, he can earn at most

$$0 + \delta x(2) + \cdots \delta^T x(T) + \delta^T g_1(r_1, m_2) + (\delta^{T+1} + \delta^{T+2} \cdots) v_1. \tag{6}$$

Note that at $T + 1$ periods ahead he will get $g_1(r_1, m_2)$ because his partner will be innocent. Thus the difference between (5) and (6) is

$$x(1) - \delta^T g_1(r_1, m_2) + \delta^T v_1 \geq (1 - \delta^T) g_1(r_1, m_2) + \delta^T v_1,$$

which is non-negative by (3).

Inequality (4) assures innocent player's actions to be in their best interest. The left-hand side of that inequality is the lower bound of an innocent player's average payoff when he conforms to the equilibrium strategy, and the right-hand side is the upper bound when he deviates.

Next, we will show that (3) and (4) can be satisfied for some δ and T. Equation (3) is satisfied if δ^T is close to one, because $v \gg 0$. If we increase δ holding δ^T constant, the right-hand side of equation (4) approaches

$$(1 - \delta^T) g_1(m) + \delta^T v_1,$$

which is strictly less than the left-hand side of (4), because $v_1 > 0 \geq g_1(m)$ and $g_1(m_1, r_2) > g_1(m)$ (by (A1)).

Finally, the equilibrium is globally stable because given any history, all players will be forgiven and go back to the original payoff point v after at most T periods. $\parallel$

Finally, let us examine the minimum number of states ($\#Z_k$, $k = 1, 2$) to obtain the Folk Theorem. When random state transitions are allowed, the informational requirement may be further weakened because in that case counting the number of punishments may be unnecessary. Suppose each guilty player has an independent chance of being forgiven in each period. Note that varying the probability of such chances is similar to varying

the number of deterministic punishments. In the standard repeated game, equilibrium can be constructed in either way. In the matching game, however, stochastic punishment requires more conditions. This is because unlike in the deterministic case, *a player may encounter a guilty player when he is forgiven*, which may destroy the incentives of guilty players. Therefore, an innocent player's payoff when he is matched with a guilty player should not be too small. If this is guaranteed, we can reduce the number of the states to just two. In a sense, this is the minimal information necessary to sustain any efficient outcomes. Formally, we need the following assumption.

Assumption (A2)   There exists $r \in A$ such that

$$g_1(m_1, r_2)g_1(m) - g_1(r_1, m_2) \geqq 0$$
$$g_2(r_1, m_2) > g_2(m) - g_2(m_1, r_2) \geqq 0.$$

For example, this assumption is satisfied by bilateral trades in which exchange of commodities is mutually beneficial; in that case, m stands for retaining the endowments and r for exchange.

Theorem 3.   *When assumption (A2) is satisfied and stochastic state transition is allowed, any payoff $v \in V$ is sustained by a straightforward and globally stable equilibrium under local information processing with two states for each player, if $\delta \in (\delta^*, 1)$, and δ^* can be chosen independently of the matching rule and the population size.*

Proof.   See Appendix.   ‖

Remark. If there is a random event which is observable to all the members in the community, then forgiving guilty players can be coordinated. That is, all guilty players can be forgiven at once based on the outcome of the publicly observable event. When this is possible, a guilty player always expects to be matched with innocent players whenever he is forgiven, and therefore assumption (A1) is sufficient to establish the same result as Theorem 2.

6. CONCLUDING REMARKS

In this paper, we have shown how informal sanctions by community members can induce desirable behaviour in infrequent trades. Specifically, it is shown that a simple action rule and local information transmission are sufficient to sustain any mutually beneficial outcomes under weak conditions.

Technically, our model is an extension of the theory of repeated games to the case of matching games. Early literature on the study of repeated matching games includes R. Rosenthal (1979), and R. Rosenthal and H. Landau (1979). The former deals with a model with adverse selection and proves the existence of Markov equilibrium when players observe their partners' last actions. The latter analyzes examples of equilibria in a bargaining game with a specific local information transmission mechanism, which they call reputation. The attempt to generalize the Folk Theorem of repeated games to the case of matching games was initiated by P. Milgrom, D. North and B. Weingast (1990), and M. Okuno-Fujiwara and A. Postlewaite (1989). The former analyzed concrete examples of information transmission mechanisms and the latter introduced the notion

of local information processing. Both of them, however, mainly deal with the infinite population case to avoid potentially complicated problems of incentives on off-equilibrium paths. Our paper shows that such problems can be resolved in a simple way if the stage game satisfies a certain weak condition. Equilibria constructed in our paper work for any population size and any matching rule, and are robust to changes in information structures.

Perhaps the most important question which is unanswered by the present paper concerns the way in which the information transmission postulated in our model is implemented. Three important problems arise in that respect.

1. What is the cost of information transmission?
2. Who pays the cost?
3. How are proper incentives maintained in the information transmission mechanism?

Papers by Milgrom *et al.* (1990) and Klein (1988) analyse some of those problems for such institutions as "law merchants" during the early middle ages and credit bureaus in modern times, but general theoretical examination of those problems is yet to be done.

APPENDIX

Proof of Theorem 1 *and Lemma.* By the Principle of Dynamic Programming, we have only to check that one-shot deviations from the strategy are unprofitable after any history (see Abreu (1988)). The condition that a one-shot deviation from the equilibrium path is unprofitable is

$$\frac{1}{1-\delta} \geq \sum_{t=0}^{\infty} \delta^t e_1 A^t \rho (1+g). \qquad (*)$$

The left-hand side is the payoff from cooperating forever, and the right-hand side is the payoff when the player defects forever. The expression $e_1 A^t \rho$ is the probability of meeting a c-type at time t given that the player was the first person to deviate at time 0. Since the equilibrium requires that a player play defection after he has defected, he receives payoff $(1+g)$ if he is matched with a c-type, and gets zero payoff otherwise. Summing up the geometric sequence shows that condition $(*)$ is equivalent to inequality (1).

To show the unprofitability of deviation from an off-equilibrium path (i.e. playing c after observing any defection including his own), we need in principle to specify a player's belief on X_t (the number of type-d players), because we are interested in a sequential equilibrium. Rather than following this procedure, we identify a sufficient condition for a one-shot deviation from an off-equilibrium path to be unprofitable under any consistent beliefs. This not only avoids complication but also guarantees that the equilibrium is robust to refinements of sequential equilibria. The condition we are going to use is that a one-shot deviation from playing d forever is unprofitable for a d-type player given any number of d-types (i.e. $X_t = k$ for all $k = 2, 3, \ldots, M$);

$$\sum_{t=0}^{\infty} \delta^t e_k A^t \rho (1+g) \geq \left(\frac{M-k}{M-1} \right) - \left(\frac{k-1}{M-1} \right) l + \delta \sum_{t=0}^{\infty} \delta^t e_k B A^t \rho (1+g).$$

The left-hand side is the payoff from playing d forever when there are k d-type players including the player himself, and right-hand side is what the player receives when he plays c today and then plays d forever; $(M-k)/(M-1)$ and $(k-1)/(M-1)$ are the probabilities of meeting type-c and type-d players respectively, and $e_k B$ is the distribution of the number of type d players tomorrow given that there are k d-types and one of them (the player) deviates to play c today. This can be manipulated to get inequality (2).

Given M and g, the equilibrium conditions (1) and (2) can be satisfied for large δ and l because $\lim_{\delta \to 1} (I - \delta A)^{-1} \rho < \infty$. This is seen as follows. Notice that

$$(I - \delta A)^{-1} \rho = \sum_{t=0}^{\infty} \delta^t A^t \rho = \sum_{t=0}^{\infty} \delta^t \tilde{A}^t \rho = (I - \delta \tilde{A})^{-1} \rho,$$

where $\tilde{A}$ is a matrix obtained by replacing the last column of A by zeros. Such replacement is justified because $X_t = M$ is the absorbing state and the Mth element of ρ is zero. Given this, we have only to show the existence

of $(I - \tilde{A})^{-1}$. Since the number of d-types never declines, $\tilde{A}$ is upper-triangular and so is $(I - \tilde{A})$. The determinant of an upper-triangular matrix is the products of its diagonal elements, which are all strictly positive for $(I - \tilde{A})$. ‖

Proof of Proposition 3.   Without loss of generality, suppose the type-1 player has an incentive to deviate from v, and suppose the gain from defecting is $d_1 > 0$. Let $\Delta_1 = \max_{a \in A} g_1(a) - \min_{a \in A} g_1(a)$. Given a strategy, suppose V and V' be the continuation payoffs for a type-1 player when he conforms to a strategy and deviates from it respectively (a symmetric argument applies to type-2 players). Then, for sufficiently large n, $V - V' < d_1$ and therefore the only equilibrium is the repetition of a one-shot Nash equilibrium. This is because

$$V - V' \leq \sum_{t=1}^{\infty} \delta^* t \frac{\min\{2^{t-1}, n\}}{n} \Delta_1 \to 0 \quad \text{as } n \to \infty.$$

The inequality is explained as follows. Let I_t be the set of players whose behaviour at $T + t$ is affected by the defection at time T. If player i is the defector at T, $I_1 = \{i, \mu(i, T)\}$ and for $t > 1$, $I_t = \{j \mid j \in I_{t-1} \text{ or } j = \mu(k, t-1) \text{ for some } k \in I_{t-1}\}$. The number of such players is maximal when all members of I_t are matched with players in $N \backslash I_t$, so that an upper bound of $\# I_t$ is $\min\{2^t, 2n\}$. Hence $\min\{2^{t-1}, n\}$ is an upper bound of the number of type-2 players in I_t, so that $\min\{2^{t-1}, n\}/n$ is an upper bound of the probability for the type-1 defector to meet a type-2 player in I_t at $T + t$. ‖

Proof of Proposition 4.   Suppose the type-1 players have an incentive problem, and let $Z_1 = \{0, 1, \ldots, T\}$ be the state space for type-1 and let type-2 players' state be identically equal to zero. The state transition rule is the following.

$$\tau(0, 0, a) = \begin{cases} (0, 0) & \text{if } a_1 = a_1^* \\ (1, 0) & \text{otherwise} \end{cases}$$

and $\tau(k, 0) = k + 1 \pmod{T + 1}$ for $k = 1, \ldots, T$ and for all $a \in A$. The action choice is given by

$$\sigma(k, 0) = \begin{cases} a^* & \text{if } k = 0^* \\ a^0 & \text{otherwise} \end{cases}$$

Thus, after a deviation, a type-1 player is punished for T periods by one-shot Nash equilibrium. Type-2 players have no incentive to deviate because their future payoffs are independent of their actions and they are always taking one-shot best responses. If T and δ are sufficiently large, the loss caused by the T periods punishment becomes greater than the gain from a one-shot deviation, and therefore the type-2 conforms to the prescribed strategy. ‖

Proof of Theorem 3.   The state spaces are $Z_1 = Z_2 = Z = \{0, 1\}$, where 0 indicates innocence and 1 means that the player is guilty. Let $a^* \in A$ be the action profile which achieves v. The action rule is, for $z \in Z$,

$$\sigma(z) = \begin{cases} a^* & \text{if } z = (0, 0) \\ (m_1, r_2) & \text{if } z_1 = 0, z_2 = 1 \\ (r_1, m_2) & \text{if } z_1 = 1, z_2 = 0 \\ m & \text{if } z_1 = z_2 = 1 \end{cases}$$

The state transition is determined as follows. If $z_i(t) = 0$ and player i does not deviate, then $z_i(t+1) = 0$, and if he deviates, $z_i(t+1) = 1$. When $z_i(t) = 1$, $z_i(t+1) = 1$ if he deviates, and if he conforms,

$$z_i(t+1) = \begin{cases} 1 & \text{with probability } p \\ 0 & \text{with probability } 1 - p. \end{cases}$$

We will check the incentives of type-1 players. Symmetric arguments apply to type-2 players. When a type-1 player is guilty, his total payoff is

$$\Pi = x(1) + \sum_{k=2}^{\infty} (\delta p)^{k-1} (x(k) - u(k)) + \sum_{k=2}^{\infty} \delta^{k-1} u(k),$$

where

$$u(k) = \theta(k) g_1(m_1, r_2) + (1 - \theta(k)) v_1,$$
$$x(k) = \theta(k) g_1(m) + (1 - \theta(k)) g_1(r_1, m_2),$$

and $\theta(k)$ is the probability of meeting a guilty opponent after k periods. Note that when the player is forgiven, his payoff changes from $x(k)$ to $u(k)$. If he deviates, he earns at most

$$\Pi' = 0 + \delta x(2) + \delta^2 (p x(3) + (1 - p) u(3)) + \cdots$$
$$= \delta \sum_{k=2}^{\infty} (\delta p)^{k-2} (x(k) - u(k)) + \sum_{k=2}^{\infty} \delta^{k-1} u(k).$$

The difference is

$$\Pi - \Pi' = x(1) + (p\delta - \delta) \sum_{k=2}^{\infty} (\delta p)^{k-2} (x(k) - u(k)).$$

Note that $x(k) - u(k) = \theta(k)(g_1(m)) - g_1(m_1, r_2)) + (1 - \theta(k))(g_1(r_1, m_2) - v_1) < g_1(r_1, m_2)$ by $v_1 > 0$ and Assumption (A2). Therefore,

$$\Pi - \Pi' > x(1) + \frac{\delta p - \delta}{1 - \delta p} g_1(r_1, m_2) \to x(1) - g_1(r_1, m_2) \geq 0$$

as $\delta \uparrow 1$ with $\delta p \in (0, 1)$ held constant. Thus we have shown that any deviation is unprofitable for guilty players. When a player is innocent, his total payoff is

$$U = \sum_{k=1}^{\infty} \delta^{k-1} u(k),$$

and one-shot deviation earns at most $U' = \max_{a \in A} g_1(a) + \Pi'$. Therefore,

$$U - U' = (u(1) - \max_{a \in A} g_1(a) + \delta \sum_{k=2}^{\infty} (\delta p)^{k-2} (u(k) - x(k)).$$

Since $u(k) - x(k) > 0$, this is positive if δ and δp are close to 1. ‖

REFERENCES

ABREU, D. (1988), "On the Theory of Infinitely Repeated Games with Discounting", *Econometrica*, **56**, 383–96.
ABREU, D., PEARCE, D. and STACCHETTI, E. (1990), "Toward a Theory of Discounted Repeated Games with Imperfect Monitoring", *Econometrica*, **58**, 1041–1063.
BEN-PORATH, Y. (1980), "The *F*-Connection: Families, Friends, and Firms and the Organization of Exchange", *Population and Development Review*, **6**, 1–30.
FUDENBERG, D. and LEVINE, D. (1991), "An Approximate Folk Theorem with Imperfect Private Information", *Journal of Economic Theory*, **54**, 26–47.
FUDENBERG, D., LEVINE, D. and MASKIN, E. (1989), "The Folk Theorem with Imperfect Public Information" (mimeo).
FUDENBERG, D. and MASKIN, E. (1986), "The Folk Theorem in Repeated Games with Discounting or with Incomplete Information", *Econometrica*, **50**, 533–554.
GREIF, A. (1989), "Reputation and Coalitions in Medieval Trade", *Journal of Economic History*, **49**, 857–882.
HARRINGTON, J. (1989), "Cooperation in Social Settings" (mimeo, Johns Hopkins University).
KANDORI, M. (1989) *Information and Coordination in Strategic Interaction over Time* (Ph. D. thesis, Stanford University).
KANDORI, M. (1989), "Repeated Games Played by Overlapping Generations of Players", *Review of Economic Studies* (forthcoming).
KANDORI, M. (1989), "The Use of Information in Repeated Games with Imperfect Monitoring", *Review of Economic Studies* (forthcoming).
KLEIN, B. and LEFFLER, K. B. (1981), "The Role of Market Forces in Assuring Contractual Performance", *Journal of Political Economy*, **89**, 615–641.
KLEIN, D. (1989), "Cooperation through Collective Enforcement: A Game-Theoretic Model of Credit Bureaus" (mimeo, U. C. Irvine).
KREPS, D. M. and WILSON, R. (1982), "Sequential Equilibria", *Econometrica*, **50**, 863–894.
MACAULAY, S. (1963), "Non-Contractual Relations in Business: A Preliminary Study", *American Sociological Review*, **28**, 55–67.
MILGROM, P., NORTH, D. and WEINGAST, B. (1990), "The Role of Institutions in the Revival of Trade: The Law Merchant, Private Judges, and the Champagne Fairs", *Economics and Politics*, **2**, 1–23.
NORTH, D. (1987), "Institutions, Transaction Costs and Economic Growth", *Economic Inquiry*, **25**, 419–428.
OKUNO-FUJIWARA, M. and POSTLEWAITE, A. (1989), "Social Norms in Random Matching Game" (mimeo., University of Tokyo and University of Pennsylvania).
ROSENTHAL, R. W. (1979), "Sequences of Games with Varying Opponents", *Econometrica*, **47**, 1353–1366.

ROSENTHAL, R. W. and LANDAU, H. J. (1979), "A Game-Theoretic Analysis of Bargaining with Reputations", *Journal of Mathematical Psychology*, **20**, 233–255.

RUBINSTEIN, A. (1979), "Equilibrium in Supergames with the Overtaking Criterion", *Journal of Economic Theory*, **21**, 1–9.

RUBINSTEIN, A. and WOLINSKY, A. (1990), "Decentralized Trading, Strategic Behavior and the Walrasian Outcome", *Review of Economic Studies*, **57**, 63–78.

WILLIAMSON, O. (1985) *Economic Institution of Capitalism* (New York: Free Press).