SOCIAL NORMS, COOPERATION AND INEQUALITY

PEDRO DAL BÓ

Working Paper Number 802
Department of Economics
University of California, Los Angeles
Bunche 2263
Los Angeles, CA 90095-1477
April 13, 2001

Social Norms, Cooperation and Inequality

Pedro Dal Bó*

Abstract

Social norms can influence behavior by establishing reward and punishment schemes that operate through community enforcement and in this way support outcomes that cannot be supported by personal enforcement alone. This paper analyzes the outcomes that can be supported through social norms in a society. I consider a society of infinitely long-lived and very patient agents that are randomly matched in pairs every period to play a given game and, unlike previous work that considers a society divided in two groups, I only restrict this matching procedure to be independent of history and time. I find that any mutually beneficial outcome can be supported by a self-enforcing social norm under both perfect information and a simple local information system. These Folk Theorem results explain not only how social norms can provide incentives to forestall opportunistic behavior and support cooperation in a community but also how they can support outcomes characterized by inequality. For example, self-enforcing social norms can support outcomes such as a "kingdom" (a society in which one member reaches the highest possible payoff) or a caste system (in which some groups of the population obtain higher payoffs than others).

Key words: social norms, random matching, repeated games, folk theorem, inequality, discrimination.

Journal of Economic Literature Classification Numbers: C72, J7, Z13.

*I am grateful to David Levine for invaluable guidance and ideas. I want to thank Anna Aizer, Hongbin Cai, Walter Cont, Ernesto Dal Bó, Jean-Laurent Rosenthal, Federico Weinschelbaun, William Zame and seminar participant at Universidad de Buenos Aires and UCLA for very useful comments and discussions. Any remaining errors are mine.

1 Introduction

In a stable society, codes of conduct or social norms guide interaction among people. A social norm functions as an implicit rule of behavior that guides the actions of society members under different circumstances. In this sense, social norms organize interaction among the members of a society, thereby reducing the level of uncertainty.

An important feature of social norms is their ability to curtail opportunistic behavior through the establishment of reward and punishment schemes that operate through social reaction. In this way social norms can support high levels of cooperation among the members of a society. Greif [6] presents an example of this found among the Maghribi traders in the Mediterranean during the 11th century. With very little support from formal institutions to enforce contracts governing overseas trading, the Maghribi traders followed a simple social norm: no Maghribi trader would trade with another Maghribi trader who had cheated a Maghribi trader before. In this way the Maghribi punished deviations and reduced the incentives to cheat, hence, supporting efficient levels of commerce that could not have been achieved if they had relied upon personal retaliation alone.

But social norms not only can lead to efficient outcomes, they can also lead to inefficient outcomes or outcomes with inequality. Akerlof [2] presents several examples of social norms that lead to inefficient and unequal outcomes. Among those he considers the traditional Indian caste system, in which a mutually beneficial transaction -like marriage- between two members of different castes may be forestalled by the expected reaction of third parties.

This paper analyzes the outcomes (efficient or inefficient, equal or unequal) that can be supported by social norms. I consider a simple society consisting of a fixed number of infinitely long-lived members. Every period these members are randomly matched in pairs to play a given game. The only restriction on these matches is that the probability of each match is fixed and independent of time and past behavior. I assume that the members are selfish in the sense that they do not care about others or the social norm per se. Hence, no members follow the social norm because they derive any direct utility from doing so. In this paper, social norms are followed because no member can ever profit from deviating from it on his or her own. In other words, I study social norms that are a subgame perfect equilibrium or a sequential equilibrium, depending on the information requirements.

I consider two societies under different information requirements: one characterized by perfect information and another with a more restrictive information process. For the former, I consider the case in which it is possible for every member to observe each other's behavior. In this situation of perfect information, I prove that if the agents are sufficiently patient, any feasible and individually rational outcome can be supported by a social norm that is a subgame perfect equilibrium. Therefore, a society can achieve in equilibrium any feasible and individually rational outcome through community enforcement, that is, with a social norm that specifies rewards and punishments.

To understand the relevance of this result it is important to note the difference between community enforcement and personal enforcement. Under personal enforcement, a cheater will only face retaliation by the victim. On the contrary, under community enforcement all the members of the society react to a deviation. Hence, community enforcement can offer a punishment that is more swift and severe than personal enforcement can. This implies that any outcome that can be supported by personal enforcement can also be supported by community enforcement. But most interestingly, there are outcomes that can not be supported by personal enforcement that can be supported by community enforcement. In fact, there are highly unequal outcomes that can only be supported by community enforcement, as the next example shows.

Consider a society with four members that are matched every period to play the following prisoner's dilemma game:

	c	d
c	2 , 2	-2 , 4
d	4 , -2	0 , 0

Assume that each pair has the same probability of being matched and consider the following social norm: Player 1, who is called "king", always plays *d*, players 2, 3 and 4, who are called "serfs", play c if no one has deviated and play d if someone has ever deviated. Under this social norm, the king receives a payoff of 4 every period while the serfs receive an average payoff of $\frac{2}{3}$ in equilibrium. It is easy to verify that this social norm is a subgame perfect equilibrium for discount factors greater than $\frac{3}{4}$.[1] It is interesting to note that with personal enforcement the king would only be able to obtain payoffs lower than 3.[2] Therefore, community enforcement allows the king to obtain a higher payoff than what he could obtain under personal enforcement. The reason for this is that under community enforcement each serf knows that if he deviates when playing with the king, no other serf will cooperate with him. Hence, serfs accept the negative payoff they receive every time they are matched with the king because this enables them to reap the benefits of full cooperation among the serfs. With personal enforcement the king can not use this threat and therefore he can not achieve the high level of payoffs he receives under community enforcement.

[1]The King has clearly no incentives to deviate. Seeing that the serfs do not have incentives to deviate requires some calculations. If no one has deviated, a serf facing the King would get an expected payoff of $-2(1-\delta)+\delta\frac{2}{3}$ for following the social norm, where δ is the discount factor, and he would get 0 for deviating. Hence, if no one has deviated, the serfs will not deviate when playing against the King if $\delta \geqslant \frac{3}{4}$. Similar calculations show that for those discount factors a serf would not deviate when playing against another serf. Therefore, serfs have no incentives to deviate if no one has deviated before. If someone has deviated, they do not have incentives to deviate since the prescribed actions correspond with the stage game Nash equilibrium. Therefore, if $\delta \geqslant \frac{3}{4}$, no player has an incentive to deviate and, then, the social norm is a subgame perfect equilibrium .

[2]Under personal enforcement, a payoff higher than 3 for the King would result in a negative expected payoff for the serf in the matches with the King. Under personal enforcement the serf would not accept that, since he can secure for himself a minimum of zero by playing *d*.

In the social norm used in the previous example all the players are punished for the deviation of one of the members. While societies may use this kind of punishment schemes to reduce opportunistic behavior, schemes that only punish deviators seem to be more realistic. Using social norms that only punish deviators, I show that any feasible and individually rational outcome can still be supported in a subgame perfect equilibrium under some restrictions on the stage game.

But the requirement of perfect information is unrealistic when the size of the society is large. Therefore, I also study environments with less demanding information requirements. As an alternative to perfect information I consider the existence of a local information processing system, following the seminal papers of Kandori [8] and Okuno-Fujiwara and Postlewaite [11]. In this case, in addition to knowledge gained from their own experience, players have access to information from a system that assigns status levels to players depending on their past behavior.

These status levels enable the social norm to establish punishments and rewards. For example, in the case of the king and the serfs presented before there can be two status levels: "good" and "bad" serfs. If failing to cooperate with the king results in becoming a bad serf and nobody cooperates with a bad serf, a good serf may be willing to cooperate with the king, even when the latter never cooperates.

With local information systems added to social norms, I prove that, if players are sufficiently patient, any feasible and individually rational outcome can be supported by a social norm that is a sequential equilibrium. Therefore, even under very limited information, a society can achieve in equilibrium any feasible and individually rational outcome with a social norm that specifies rewards and punishments based on the information provided by a local information system. I also show a similar result, under some restrictions on the stage game, for social norms that only punish deviators.

These Folk Theorem results explain not only how social norms can provide incentives to forestall opportunistic behavior and support cooperation in a community but also how they can support outcomes characterized by inequality.

The following section presents the relevant literature, comparing previous findings with my own. Section 3 presents the model. Section 4 presents the perfect information Folk Theorem results and Section 5 presents the local information system Folk Theorem results. Section 6 concludes.

2 Relevant literature

Game theorists have long recognized that repeated playing and the possibility of future retaliation modifies current behavior, for example see Luce and Raiffa [10]. In the case in which the same set of players play the same game repeatedly, other studies have found the conditions under which any feasible and individually rational outcome can be supported in equilibrium[3]. However, in many interesting cases the same players do not meet repeatedly but rather switch partners over time. For example see the cases of the already mentioned medieval trade coalitions studied in Greif [6]. While the changing of partners might seem to make cooperation impossible by reducing the possibility of personal retaliation, that is not necessarily the case. As Kandori [8] and Okuno-Fujiwara and Postlewaite [11] show, social norms can create incentives for players to punish deviators even if the deviation occurred against another player, since failing to punish can be itself punishable.

Both papers consider a society divided in two groups and every period the players in one group are randomly matched with the players in the other group. In addition, the authors restrict players in the same group to have the same equilibrium payoff[4].

[3]See Aumann and Shapley [3] for the case without discounting, Fudenberg and Maskin [5] for the case of discounting with perfect information and Fudenberg, Levine and Maskin [4] for the case of discounting with imperfect public information.

[4]This characterization of the game has the appealing graphical property that we can represent an equilibrium outcome of the game in $\Re^2$. Since all the members of a group receive the same payoff in equilibrium and there are only two groups, equilibrium payoffs can be written as the pair (v_1, v_2), where v_1 and v_2 denote the utility received by each player in group 1 and 2, respectively.

Kandori [8] shows that with perfect information any feasible and individually rational payoff pair can be supported as a subgame perfect equilibrium. Therefore, any outcome that can be reached in the long-term relationship of two agents can also be reached by long-term relationship of two groups. In the case of two groups it is possible to construct credible group retaliations that mimic the ones needed for the Folk Theorem with only two players. This paper studies the set of equilibrium payoffs when we abandon those restrictions and we only require the matching procedure to be independent of history and time. In Section 3 I show that under perfect information, community enforcement can support not only those outcomes that can be supported by personal enforcement, but it can support other outcomes as well (as the king example in the introduction). In fact, I show that under perfect information any feasible and individually rational outcome can be supported by a social norm if the players are patient enough.

Since the requirement of perfect information is unrealistic when the size of the society is large, Kandori [8] and Okuno-Fujiwara and Postlewaite [11] study the existence of Folk theorem results with a local information processing system. Okuno-Fujiwara and Postlewaite [11] prove a Folk Theorem for a weak (non-perfect) equilibrium concept (Norm equilibrium) and Kandori [8] proves a Folk Theorem for sequential equilibrium under certain assumptions of the stage game. While Kandori [8] shows that infrequent transactions and limited information can still be overcome to achieve a Folk Theorem, he does so in the restrictive environment of two groups in which all the members of a group receive the same payoff in equilibrium. Without these restrictions, Section 4 explains not only how social norms can support cooperation in a community but also how they can support outcomes characterized by inequality[5].

The idea that social norms may support inequality is not new in the literature. Akerlof [2] presents several examples of social norms that support unequal (and inefficient) payoffs in (non-perfect) equilibrium. While those examples show that social norms can

[5]In addition, the proof of Theorem 4 shows how theorem 2 in Kandori [8] could be proved without restrictions on the stage game.

support inequality in equilibrium, the punishments specified in them are not credible because the equilibria studied are not perfect. In contrast, the equilibria presented in this paper are perfect. In addition my results are general to any stage game and not limited to particular examples.

3 The Matching Game

The society consists of N players, where N is an even number. In each stage, each of the players is matched with another player to play the stage game Γ. I assume that the matching of players is independent of past actions or time: the probability that player i is matched with player j is α_{ij}, $0 \leq \alpha_{ij} \leq 1$, for every period and history. This definition allows for partitions of players as in Kandori [8] and Okuno-Fujiwara and Postlewaite [11].

The stage game Γ is a symmetric game played by two players, with actions $a \in A$, for both players and payoffs $g : A^2 \rightarrow \Re^2$, with the property that $g_{row}(a', a) = g_{col}(a, a')$, from the symmetry of the game. Given that the stage game is symmetric and that the row or column positions are not important I simplify notation writing $g(a, a') = g_{row}(a, a') = g_{col}(a', a)$. Therefore $g(a, a')$ denotes the payoff for the player that is playing a when the other is playing a'.

To minimax the other player the prescribed strategy is $m = \arg\min_{a \in A} \left(\max_{a' \in A} g(a, a') \right)$. I normalize the payoffs to have the minimax payoffs, not $g(m, m)$, equal to zero.[6] If both players play m, each obtains a payoff of $g_m = g(m, m)$. Since m may not be the best response to m, it is the case that $g_m \leq 0$. The maximum payoff that can be obtained in the stage game is $\overline{g} = \max_{a, a' \in A} g(a, a')$ and the minimum payoff is $\underline{g} = \min_{a, a' \in A} g(a, a')$. I assume that players can condition their actions on public randomization devices, that is, they can play correlated strategies. Define $\triangle(A^2)$ as the set of possible correlated strategies in the stage game. Abusing notations I denote an element of that set for

[6] I assume for convenience that m is not a mixed action.

player i and j as $(a_{ij}, a_{ji}) \in \triangle(A^2)$.

Now I proceed to define the set of feasible payoffs of the random matching game. I define first the "play" of the stage game: the play of the stage game describes what profile of actions would be played by each possible matching of players in each period, that is $play : pair \times \{0, 1, 2....\} \to \triangle(A^2)$.[7] The play indicates what should be played by each possible pair in each period, for example a_{ij}^t denotes what i should play when matched with j in period t. Therefore, the expected stage payoff for i in period t is $\sum_{j \neq i} \alpha_{ij} g(a_{ij}^t, a_{ji}^t)$. If $\left(a_{ij}^t\right)$ is the "play" for every period and δ is the discount factor, the average expected payoff of player i is $v_i = (1-\delta) \sum_{t=0}^{\infty} \delta^t \sum_{j \neq i} \alpha_{ij} g(a_{ij}^t, a_{ji}^t)$. Each play defines an expected payoff for every player $v = (v_1, v_2, ..., v_N)$. Therefore the set of feasible payoffs are defined by the payoffs that result from every possible play, denote this set as $V \subset \Re^N$.[8]

4 Perfect Information

In this section I consider societies in which it is possible for every member to observe each other's behavior. In this situation of perfect information, I show that if the agents are sufficiently patient, any feasible and individually rational outcome $(v \in V : v \gg 0)$ can be supported by a social norm that is a subgame perfect equilibrium.

Theorem 1 *(Folk theorem with perfect information) With perfect information any feasible and individually rational payoff $(v \in V : v \gg 0)$ can be supported by a subgame perfect equilibrium for δ large enough.*

[7]If there are N players the number of possible pairs is $\begin{pmatrix} N \\ 2 \end{pmatrix} = \frac{N!}{2*(N-2)!}$.

[8]Note that the set of feasible payoffs V is different from the set that arises in a two-player repeated game or in a repeated random matching game with two groups with all the members of each group receiving the same payoff as in Kandori [8] and Okuno-Fujiwara and Postlewaite [11]. In fact, the set V has a different dimension.

Proof. Consider the following social norm to support v: if no one has deviated in the last T periods follow the "play" that yields v, if someone has deviated in the last T periods, play m.

First I check that no player has incentives to deviate if no one deviated in the last T periods. For player i the expected utility of conforming with the equilibrium is at least $(1-\delta)\underline{g}+\delta v_i$ while he would get at most $(1-\delta)\overline{g}+\delta v_i^p$ by deviating, where $v_i^p=(1-\delta^T)g_m+\delta^T v_i$. Choose T so that, no matter the value of δ, δ^T is always equal to some fixed number $d\in(0,1)$. Then, given that $v_i>0\geq g_m$, it is true that $v_i>v_i^p$ independently of δ. Therefore for δ large enough it is true that $(1-\delta)\underline{g}+\delta v_i>(1-\delta)\overline{g}+\delta v_i^p$.

Now I consider the incentives to deviate if someone has deviated in the last T periods. In this case the incentives to deviate are the greatest when the future reward for facing the present punishment is as far in the future as possible, that is when all the players have to face T periods of punishment. If player i plays m, as the strategy prescribes, he receives a payoff of $v_i^p=(1-\delta^T)g_m+\delta^T v_i$. If he deviates he receives at most δv_i^p. Choosing $\delta^T=d$ large enough for v_i^p to be positive (it is here that v_i strictly greater than zero is required), the player has no incentive to deviate during the punishment stage for $\delta<1$.

Note that there is no contradiction in the requirements made on δ and T in the two parts of the proof, as it is only required that $\delta^T=d$ and δ are large enough. ■

The social norm used in the proof is a subgame perfect equilibrium given that regardless of the history of the game no player has incentives to deviate. In particular note that during the punishment stage no player has incentives to deviate since doing so only restarts the punishment stage. Therefore all the members of the society have incentives to enforce the social norm when someone deviates and, then, the punishments are credible.

While it may seem that Theorem 1 is a consequence of Fudemberg and Maskin [5] result for games with N players, this is not the case. While their results apply to games in which N players play the same stage game in all periods, in this case players may change partners every period. In addition, Theorem 1 can not be derived from their results by defining "meta-actions," that is functions from possible matches to actions for each player, given that Fudemberg and Maskin [5] assumes observable actions. However, by defining those "meta-actions" it can be shown that Theorem 1 is a consequence of Fudenberg, Levine and Maskin [4] folk theorem under imperfect public information. Nevertheless, the proof presented here has two advantages: one, its simplicity, and two, the fact that the implicit lower bound on the discount factor to support an individually rational and feasible outcome does not depend on the size of the population as may be the case in Fudenberg, Levine and Maskin [4].

Note that during the punishment stage in the proof of Theorem 1 all the players receive the same low payoff (g_m) regardless of who has deviated. In this social norm, then, all the players are punished for the deviation of one of the members. While societies may use this kind of punishment schemes to reduce opportunistic behavior, punishment schemes that only punish deviators seem to be more appealing.

Definition 2 *A social norm displays personal punishment*[9] *if the prescribed actions for two players who have not deviated in the past can not depend on the past actions of the rest of the players.*

In this way, under personal punishment, only the deviators can be punished. Of course punishing may impose a cost to the player facing the deviator even under personal punishment

But personal punishment introduces new problems to the design of social norms that support a folk theorem. In fact, the social norm used in the proof of Theorem 1 may not be an equilibrium. Given the possible inequality of payoffs with personal punishment

[9] Personal punishment should not be confused with personal enforcement. Personal enforcement means that cheaters are only punished by the player that was cheated, while personal punishment means that only cheaters are punished.

some players may have an expected payoff lower than g_m during the punishment stage, and therefore may have incentives not to punish, breaking in that way the credibility of the social norm[10]

Fortunately under some conditions it is possible to support any feasible and individually rational payoff with social norms with personal punishment.

Assumption 1: $\exists r \in A$ such that $g(m,r) > g_m \geq g(r,m) = \underline{g}$.

Under this assumption there is an action r that allows a deviator to "ask for forgiveness" by taking the lowest possible payoff in the game $\underline{g}$, and giving the punisher a payoff higher than g_m while playing m. As such, it is possible to create punishment schemes in which the punisher earns a higher payoff than the deviator and in which the deviator by refusing to take the punishment can at most obtain a payoff of zero. This allows me to construct a social norm that ensures that players have incentives to follow it even when some of the other players have deviated[11]. This assumption is satisfied, for example, by the prisoner's dilemma game, in which m stands for d and r stands for c.

Theorem 3 *(Folk theorem with perfect information and personal punishment) Under perfect information, Assumption 1 and personal punishment, any feasible and individually rational payoff ($v \in V : v \gg 0$) can be supported by a subgame perfect equilibrium for δ large enough.*

Proof. Consider the following social norm that yields v in equilibrium: if player i meets player j and neither has been the last player to deviate in the last T periods, they

[10]For example consider the case of the kingdom presented in the introduction but with personal punishment instead: only the players that have deviated are punished. Imagine now that two of the serfs have deviated and have to be punished for τ periods. In this case the third serf will earn zero each time he meets with one of the other serfs and -2 each time that he meets with the king. Then, he would get an expected payoff of $-\frac{2}{3}$ each period during the punishment stage. By deviating he can get a zero payoff during T periods and, then, he may be willing to deviate just to avoid the negative payoff he earns if he does not deviate.

[11]Without Assumption 1, Theorem 2 is still true under alternative conditions like uniform random matching games and N large.

play (a_{ij}, a_{ji}) (which yields v in equilibrium); the last player to have deviated in the last T periods (simultaneous deviations are ignored) plays r and his match plays m.

First, I check that no player has incentives to deviate if no one has deviated in the last T periods. For player i the expected utility of conforming with the equilibrium is at least $(1-\delta)\underline{g} + \delta v_i$ while he would get at most $(1-\delta)\overline{g} + \delta v_i^p$ by deviating, where $v_i^p = (1-\delta^T)g\,(r,m) + \delta^T v_i$. Choose T to have $\delta^T = d \in (0,1)$. Then, given that $v_i > 0 \geq g\,(r,m)$, it is true that $v_i > v_i^p$ independently of δ. Therefore for δ large enough it is true that $(1-\delta)\underline{g} + \delta v_i > (1-\delta)\overline{g} + \delta v_i^p$.

Second, I consider the case in which player i has been the last player to deviate in the last T periods. In this case the incentives to deviate for i are the greatest when the future reward for taking the present punishment is as far in the future as possible, that is when he has to face T periods of punishment. If player i plays r, as the strategy prescribes, he receives a payoff of $v_i^p = (1-\delta^T)g\,(r,m) + \delta^T v_i$. If he deviates he receives δv_i^p. Choosing $\delta^T = d$ large enough for v_i^p to be positive, the player has no incentive to deviate during the punishment stage since $\delta < 1$.

Third, I consider the case in which player j has been the last player to deviate in the last T periods. In this case the incentives for i to deviate depend on the payoff for i when meeting j when the latter has not deviated, say g_{ij}. If $g_{ij} > g(m,r)$ it can be easily shown that the incentives for i to deviate are higher when j has to be punished for T periods. In this case, the expected utility of conforming with the equilibrium for i is at least $(1-\delta)\underline{g} + \delta v_i'$, where $v_i' = (1-\delta^{T-1})\left(\alpha_{ij} g(m,r) + (1-\alpha_{ij})\underline{g}\right) + \delta^{T-1} v_i$, while he would get at most $(1-\delta)\overline{g} + \delta v_i^p$ by deviating, where $v_i^p = (1-\delta^T)g\,(r,m) + \delta^T v_i$. Choose T to have $\delta^T = d \in (0,1)$. Then, given that $v_i > 0 \geq g\,(r,m) = \underline{g}$ and $g(m,r) \geq g\,(r,m)$, it is easy to see that $v_i' > v_i^p$ independently of δ. Therefore for δ large enough it is true that $(1-\delta)\underline{g} + \delta v_i' > (1-\delta)\overline{g} + \delta v_i^p$. If on the contrary, $g_{ij} > g(m,r)$, the incentives for i to deviate are higher when j has not deviated. But we have already proven in the first part of the proof that in that case i does not want to deviate.

Note that there is no contradiction in the requirements made on δ and T in the different parts of the proof, as it is only required that $\delta^T = d$ and δ are large enough. ∎

As I mentioned before, with personal punishment a player, say player 1, may earn a very low payoff during the punishment stage (it may be the case that the player being punished, say player 2, is the only player that gives player 1 a positive payoff on the equilibrium path). Assumption 1 allows me to construct a social norm that makes sure that by deviating during the punishment stage, player 1 can only reduce his payoff even more and, then, has no incentives to deviate and the punishment is credible.

The social norms used in the proofs of the former theorems, in which players punish deviators since failing to do so is itself punished, resemble, in some ways, the enforcement of castes in India. When describing marriage customs in India, Akerlof [2] says: "The caste rules dictate not only the code of behavior, but also the punishment for infractions: violators will be outcasted; furthermore, those who fail to treat outcastes as dictated by caste code will themselves be outcasted."

While perfect information may be plausible in a small community it will certainly be implausible in a large one: it would be difficult for each player to know what every other player has done in the past if the number of players is very large. Therefore, Theorem 1 and 3 would not apply to the study of social norms and their impact on cooperation and inequality in large communities. In the next section I study the outcomes that can be supported by social norms under lower information requirements.

5 Local information

Even though in particular cases it is possible to forestall opportunistic behavior without the players having more information than their own experience, as the "contagious equilibrium" in Kandori [8], in more general cases extra information is necessary to provide the needed structure of punishments. In this section I assume that, in addition to their own experience, players have access to a local information processing system that gives players some information regarding their opponent's past behavior. Following Okuno-Fujiwara and Postlewaite [11], the local information processing system has the following structure: 1) in period t agent i has a "status" or "flag" $z_i(t) \in Z_i$, where Z_i is a finite set (without loss of generality I can assume that $Z_i \subset N_+$); 2) if players i and j are matched in period t and play $a_i(t)$ and $a_j(t)$, the update of status follows a transition mapping $(z_i(t+1), z_j(t+1)) = \tau_{ij}(a_i(t), a_j(t), z_i(t), z_j(t))$; 3) if at time t player i is matched with player j, the former only knows his own history and $(z_i(t), z_j(t))$.

Based on the local information processing system, a social norm prescribes the behavior for each player as a function of his past history, his status and the status of the matched player. I show that any feasible and individually rational outcome can be supported by a social norm in equilibrium.

Theorem 4 *(Folk theorem with local information): With local information any feasible and individually rational payoff ($v \in V : v \gg 0$) can be supported by a sequential equilibrium for δ large enough.*

Proof. Consider the following social norm that yields v in equilibrium: if player i meets player j, and both are "nice" ($z_i = z_j = 0$), they play (a_{ij}, a_{ji}) (which yields v in equilibrium); if any of the two is "guilty" ($z_i \neq 0$ or $z_j \neq 0$) they both play m. The local information system, which assigns the "nice" and "guilty" labels, works as follows: if a nice player conforms he keeps the $z = 0$ flag; if a nice player meets a guilty player with $z = \tau$, he gets the $\tau - 1$ flag next period; if a player deviates, he and his match get a flag $z = T$; if a guilty player conforms, he has the flag reduced one unit. Hence, $z_i = 0$ denotes that i has not deviated, has not seen a deviation in the last T periods and has not met someone (that met someone that met someone....) that deviated in the last T periods. Instead, $z_i > 0$ denotes that i has deviated in the last T periods or is aware that someone has deviated in the last T periods. Summarizing, the social norm is:

$$s_i(t) = \begin{cases} a_{ij} \ if \ z_i(t) = z_j(t) = 0 \\ m \ if \ z_i(t) \ or \ z_j(t) \neq 0 \end{cases} \quad \text{and}$$

$$z_i(t+1) = \begin{cases} T \ if \ a_i(t) \neq s_i(t) \ or \ a_j(t) \neq s_j(t) \\ \tau - 1 \ if \ z_i(t) \ or \ z_j(t) = \tau \ and \ a_{i,j}(t) = s_{i,j}(t) \\ 0 \ if \ z_{i,j}(t) = 0 \ and \ a_{i,j}(t) = s_{i,j}(t) \end{cases}$$

To prove that this social norm is a sequential equilibrium I show that no player has incentives to deviate in any possible information set if he believes the rest of the players will follow the social norm. I assume that outside the equilibrium path players believe that there are no more "guilty" players than those they have evidence there are. That is, if player i has been matched with player j with $z_j = \tau > 0$, and no other guilty player, he believes that there are no more players with "guilty" flags than those that played with j in the last $T + 1 - \tau$ periods (that is, all the players that have been matched with j when j obtained a "guilty" flag or after). It is easy to see that this beliefs are consistent (as defined in Kreps and Wilson [9]): if the probabilities of trembles are converging to zero the probability that other deviations happened (besides the ones observed by the player) also converges to zero.

Next I check that in every information node no player has incentives to deviate. First consider the case in which player i has the $z_i = 0$ flag (he has not deviated and he has not seen a deviation or a "guilty" flag in the last T periods. Then he believes that all players have "nice" flags. In this situation the player i expects to earn at least $(1-\delta)\underline{g} + \delta v_i$ by conforming, and at most $(1-\delta)\overline{g}+\delta v_i^p$ by deviating, where $v_i^p = (1-\delta^T)g_m+\delta^T v_i$. Choose T to have $\delta^T = d \in (0,1)$. Then, given that $v_i > g_m$, it is true that $v_i > v_i^p$ independently of δ. Therefore for δ large enough it is true that $(1-\delta)\underline{g} + \delta v_i > (1-\delta)\overline{g} + \delta v_i^p$.

Second consider the case in which player i has the $z_i = \tau$ flag. This could be because either i has deviated in the past or because he has been matched with someone with a guilty flag. Then, player i believes that in τ periods all players will be nice and he will earn v_i every period. The incentives for i to deviate and try to avoid the punishment are larger the farther away the end of the punishment phase is, that is when $\tau = T$. In this case i obtains $(1-\delta^T)g_m + \delta^T v_i = v_i^p$ by conforming, and δv_i^p by deviating. Given $v_i > 0$, I can choose $\delta^T = d$ large enough for v_i^p to be positive and, then, he has no incentive to deviate during the punishment stage since $\delta < 1$.

Third consider the case of a player i with $z_i = 0$ who is matched with a player j with $z_j = \tau$. The analysis of this case coincides with the case above and i has no incentives to deviate.

Note that there is no contradiction in the requirements made on δ and T in the different parts of the proof, as it is only required that $\delta^T = d$ and δ are large enough. ■

Note that any player that deviates, sees a deviation or knows that a deviation occurred will be punished as the deviator until the end of the punishment stage. In this way, when a player knows that there has been a deviation his incentives to enforce the social norm do not depend on who has deviated. Whoever has deviated, once out-of-the-path beliefs are specified, it is easy to check that every player will enforce the punishment. The lack of personal punishment in this social norm allows me to prove the folk theorem without restrictions on the stage game. A similar social norm could be used to proof Theorem 2 of Kandori [8] without the restriction imposed in that paper on the stage game.

But the social norm in the proof of Theorem 4 may be criticized, precisely, because not only the deviator is punished. The next result shows that with local information and Assumption 1 any feasible and strictly individually rational outcome can be supported in a sequential equilibrium with personal punishment.

Theorem 5 *(Folk theorem with local information and personal punishment) Under local information, Assumption 1 and personal punishment any feasible and individually rational payoff ($v \in V : v \gg 0$) can be supported by a sequential equilibrium for δ large enough.*

Proof. In Appendix ■

The local information systems needed in the previous two proofs in this section are "simple", in the sense that the number of flags needed is finite and does not increase with time or the number of deviations. Since the punishment stage consists of T periods an information mechanism with at least $T+1$ flags per player is needed: one for each period of punishment and one for when the player is not in the punishment stage.

While the number of flags needed in Theorem 4 and 5 is finite, it can be very large. A way to drastically reduce the number of needed flags is to allow for a random transition rule of flags. In that case we can have two types of flags per player: guilty and nice, and,every period, all the guilty players that have conformed with the punishment are forgiven with probability $p \in (0,1)$ and they become nice. In this way, p can be use to establish the severity of punishment, as T was doing before, with the need of only two flags per player. Random forgiveness eliminates the need of counting the number of periods of punishment. As the next proposition shows, under Assumption 1 and personal punishment, any payoff vector that is strictly individually rational and feasible can be supported by a sequential equilibrium with only two flags (nice and guilty) if a random transition rule is allowed.

Proposition 6 *(Folk theorem with local information, personal punishment and random transition rule) Under local information processing, Assumption 1, personal punishment and random transition rule, any feasible and individually rational payoff* $(v \in V : v \gg 0)$ *can be supported by a sequential equilibrium for* δ *large enough.*

Proof. In Appendix ■

The equilibria described in this paper present some characteristics that are worth mentioning. First, in the equilibria in this section, the best response of any player in any situation depends only on his own and his match's flags. Any other information that players may have is irrelevant for making decisions: the best response is to follow the social norm, which tells players what to do under every combination of flags. In this way, the flags are sufficient statistics for the players decision making since they summarize all the relevant information.[12]

[12]This property of equilibria is called "straightforward" in Kandori [8].

Second, the long run behavior of the community is not affected by any finite sequence of deviations. Contrary to some proofs of the Folk Theorem for N players[13], if there have been deviations the prescribed actions revert to the original ones after T periods of punishment in the equilibria of this paper. Hence, the actions on the equilibrium path are globally stable: regardless of how many deviations have been up today, in the future the play of the game will return to the equilibrium play. This property is of special importance when studying societies with a large number of members. If a single deviation may take the community out of the equilibrium path for ever, it would be difficult to observe the equilibrium behavior in a large community in which each member has a small probability of making mistakes.

Third, the equilibria described in this paper are robust to small perturbations of the payoffs matrix of the stage game (of course this perturbations can not violate Assumption 1 in the cases in which this assumption is needed). Given that in the proofs of this paper all the inequalities are strict, if a social norm is an equilibrium under a given payoff matrix, it will also be an equilibrium with a payoff matrix that is arbitrarily close to the original one. Therefore, the equilibria presented here do not depend on a precise characterization of the players payoffs.

6 Conclusions

This paper analyzes the outcomes that can be supported by social norms in a society of infinitely long-lived and very patient agents that are randomly matched in pairs every period to play a given game. Unlike previous work that considered a society divided in two groups and all the members of each group receiving the same payoff, I only restrict this matching procedure to be independent of history and time. I find that any feasible and individually rational outcome can be supported by a self-enforcing social norm under both perfect information and a simple local information system. I also find

[13]See Theorem 2 in Fudenberg and Maskin [5] or Theorem 1 in Abreu, Dutta and Smith [1].

that the same result holds, under some restrictions on the stage game, if the social norms can only punish deviators.

To show the richness of the equilibria analyzed in this paper I present here several outcomes that can be supported in equilibrium by social norms in a simple community. I consider a community of ten members that are matched uniformly to play the following prisoner's dilemma:

	c	*d*
c	*2 , 2*	*-1 , 4*
d	*4 , -1*	*0 , 0*

I present first a society in which social norms support an equal and efficient outcome.

Optimal egalitarian society: In equilibrium all the players play c and receive a payoff of 2. This outcome is feasible and individually rational and, hence, can be supported by a self-enforcing social norm under either perfect information or local information. Therefore, a social norm, with its promise of punishment to deviators (and the consequent inequality under personal punishment) can support an egalitarian outcome that pareto dominates the inefficient egalitarian equilibrium of the one shot game.

But the results in this paper explain not only how social norms can provide incentives to curtail opportunistic behavior and support cooperation in a community, but also how they can support outcomes characterized by inequality as the next two examples illustrate.

Kingdom: As in the example in the Introduction, consider a society in which in equilibrium one player, the "king", always plays d and the rest of the players, the "serfs", play c. In equilibrium the king receives a payoff of 4 and each of the serfs receive $\frac{5}{3}$. This outcome is feasible and, since both payoffs are positive, it is also individually rational and, then, can be supported by a self-enforcing social norm under either perfect information or local information. The king gets the maximum payoff of the game since each serf prefers to be exploited by the king instead of rebelling and suffering the future punishment.

Caste (or Class) society: Consider a society divided in three castes: one player belongs to the high caste and in equilibrium he always plays d; three players belong to the middle caste and they play c when matched with a member of the same of higher cast and d otherwise; and the remaining six players belong to the lower cast and they always play c in equilibrium. Then, the high caste member receives 4, the middle cast members receive 3 and the low caste members receive $\frac{2}{3}$. This outcome is feasible and individually rational and, then, can be supported by a self-enforcing social norm under either perfect information or local information.

These examples show that social norms can support unequal outcomes even when all the members of the community are basically equal. While in these examples the division of members among the different groups is arbitrary, in reality it may correspond to differences in race, religion or gender. In this way, the results in this paper show how self-enforcing social norms may perpetuate discrimination among members of society even when all of them are intrinsically equal. These results show that discrimination and inequality can exist even when there are no differences in human capital or productivity and no taste for discrimination.

7 Apendix

Proof of Theorem 5: Consider the following social norm that yields v in equilibrium: if player i meets player j, and both are nice ($z_i = z_j = 0$), they play (a_{ij}, a_{ji}); if a nice player i meets a guilty player j ($z_i = 0$, $z_j \neq 0$) the former plays m and the latter plays r, (that is, the nice player punishes the guilty one, and this "asks" for forgiveness); and if two guilty players ($z_i \neq 0$, $z_j \neq 0$) meet they both play m. The local information system works as follows: if a player deviates he gets a flag $z = T$, denoting that he has to be punished for T periods; if a guilty player conforms he has his flag reduced one unit and if a nice player conforms he keeps the $z = 0$ flag.

To prove that the former is a sequential equilibrium I show that no player has incentives to deviate after any possible history if the rest of the players follow the social norm. As in the previous proposition, this makes beliefs unimportant, since regardless of what has happened in the past no player has incentives to deviate. First I check that a nice player has no incentives to deviate and second, I check that a guilty player has no incentives to deviate.

When a nice player studies if he should conform or deviate he must not only consider today's profit from the deviation but also the loss in the future T periods of punishment. This loss depends on whether the other players are nice or guilty and the payoff that the player receives upon meeting each nice player. Define $g_{ij} = g(a_{ij}, a_{ji})$ as the payoff that a nice player i receives for playing with a nice player j. If j is nice, player i would face a future loss for deviating today equal to $g_{ij} - \underline{g}$ every time he meets player j in the next T periods (remember that $g(r,m) = \underline{g}$). If j is guilty the future loss by deviation for i, when he meets j, would be $g(m,r) - g_m$. Hence for a nice player i the future loss for deviating today each time he meets player j would be greater when the latter is nice if $g_{ij} - \underline{g} > g(m,r) - g_m$. Therefore when $g_{ij} > g(m,r) - g_m + \underline{g}$ player i has more incentives to deviate when j is guilty than when j is nice.

If the former inequality holds for every j with $\alpha_{ij} > 0$, a nice player i would have the greatest incentive to deviate when the other players are T-guilty. In that case, if he conforms with the prescribed strategy he would get at least $(1-\delta)g(m,r) + \delta v_i'$ where $v_i' = \left(1-\delta^{T-1}\right)g(m,r) + \delta^{T-1}v_i$. By deviating he would get at most $(1-\delta)\overline{g} + \delta v_i^{p1}$, where $v_i^{p1} = \left(1-\delta^{T-1}\right)g_m + \delta^{T-1}(1-\delta)\underline{g} + \delta^T v_i$. Choose T so that $\delta^T = d \in (0,1)$. Then, given that $g(m,r) > g_m \geq \underline{g}$, it is true that $v_i' > v_i^{p1}$ independently of δ. Therefore, for δ large enough it is true that $(1-\delta)g(m,r) + \delta v_i' > (1-\delta)\overline{g} + \delta v_i^{p1}$.

If $g_{ij} < g(m,r) - g_m + \underline{g}$ for every j with $\alpha_{ij} > 0$, player i has more incentives to deviate when all the other players are nice. In this case he would get at least $(1-\delta)\underline{g} + \delta v_i$ by conforming and at most $(1-\delta)\overline{g} + \delta\left[\left(1-\delta^T\right)\underline{g} + \delta^T v_i\right]$ by deviating. Since $v_i > \underline{g}$, choosing T so that $\delta^T = d \in (0,1)$, the former is greater than the latter for δ large enough.

It could also be the case that player i faces players with whom the inequality is satisfied with probability α, and with probability $(1-\alpha)$ faces players with whom it is not satisfied. In this case the incentives for i to deviate will be the highest when the players belonging to the first group are T-guilty and the rest are nice. By conforming he would get at least $(1-\delta)\underline{g} + \delta v_i''$, where $v_i'' = \left(1-\delta^{T-1}\right)\left(\alpha g(m,r) + (1-\alpha)\underline{g}\right) + \delta^{T-1}v_i$ and by deviating he would get at most $(1-\delta)\overline{g} + \delta v_i^{p2}$, where $v_i^{p2} = \left(1-\delta^{T-1}\right)\left(\alpha g_m + (1-\alpha)\underline{g}\right) + \delta^{T-1}(1-\delta)\underline{g} + \delta^T v_i$.. Choose T so that $\delta^T = d \in (0,1)$. Then, given that $g(m,r) > g_m$, it is true that $v_i'' > v_i^{p2}$ independently of δ. Therefore, for δ large enough it is true that $(1-\delta)\underline{g}+\delta v_i'' > (1-\delta)\overline{g}+\delta v_i^{p2}$.

Consider the case of a $\tau-$guilty player i. Define α_t as the probability that i will be matched with a guilty player at time t. The least he can make by conforming with the prescribed strategy is:

$$(1-\delta)\left[\underline{g} + \sum_{t=2}^{\tau}\delta^{t-1}\left(\alpha_t g_m + (1-\alpha_t)\underline{g}\right) + \sum_{t=\tau+1}^{T}\delta^{t-1}\left(\alpha_t g(m,r) + (1-\alpha_t)\underline{g}\right)\right] + \delta^T v_i$$

(remember that $g(r,m) = \underline{g}$) and by deviating he can get at most:

$$(1-\delta)\left[\sum_{t=2}^{\tau}\delta^{t-1}\left(\alpha_t g_m + (1-\alpha_t)\underline{g}\right) + \sum_{t=\tau+1}^{T}\delta^{t-1}\left(\alpha_t g_m + (1-\alpha_t)\underline{g}\right) + \delta^T\underline{g}\right] + \delta^{T+1}v_i$$

Therefore the gains from conforming are at least:

$(1-\delta)\left[\delta^T v_i + \sum_{t=\tau+1}^{T} \delta^{t-1}\alpha_t \left(g(m,r) - g_m\right) + \left(1-\delta^T\right)\underline{g}\right]$.[14] By Assumption 1 the second term is non negative for any sequence of α_t. Hence, it is enough to have $\delta^T v_i + \left(1-\delta^T\right)\underline{g} > 0$ for the gains from conforming to be positive. This can be done by choosing $\delta^T = d$ large enough.

Note that there is no contradiction in the requirements made on δ and T in the two parts of the proof, it is only required that $\delta^T = d$ and δ are large enough. ■

Proof of Proposition 6: Consider the following social norm that yields v in equilibrium: if player i meets player j, and both are nice ($z_i = z_j = 0$), they play (a_{ij}, a_{ji}); if a nice player i meets a guilty player j ($z_i = 0,\ z_j = 1$) the former plays m and the latter plays r; and if two guilty players ($z_i = z_j = 1$) meet they minmax each other. The local information system works as follows: if a player deviates he gets a flag $z = 1$, denoting that he has to be punished; with probability $p \in (0,1)$ all the guilty players that have conformed last period are forgiven and get a flag $z = 0$ or remain guilty ($z = 1$) with probability $(1-p)$; and if a nice player conforms he keeps the $z = 0$ flag.

First I check that a nice player has no incentives to deviate and second, I check that a guilty player has no incentives to deviate.

As in the proof of Theorem 5, for a nice player i the future loss by deviating today each time he meets player j would be greater when the latter is nice if $g_{ij} - \underline{g} > g(m,r) - g_m$. Therefore if $g_{ij} > g(m,r) - g_m + \underline{g}$ player i has more incentives to deviate when j is guilty than when j is nice.

If the former inequality holds for every j with $\alpha_{ij} > 0$, a nice player i would have the greatest incentive to deviate when the other players are guilty. Using the recursiveness of the problem it can be found that the expected utility player i receives by conforming is $\frac{1}{1-\delta p}\left[(1-\delta)g(m,r) + \delta(1-p)v_i\right]$ while by deviating he receives at most $(1-\delta)\overline{g} + \frac{\delta(1-\delta)}{1-\delta p}\left[pg_m + (1-p)\underline{g}\right] + \frac{\delta^2(1-p)}{1-\delta p}v_i$.

Calculating the difference between them and simplifying, for player i not to have incentives to deviate it must be the case that $\frac{1}{1-\delta p}\left[g(m,r) - \delta\left(pg_m + (1-p)\underline{g}\right) + \delta(1-p)v_i\right] \geq \overline{g}$. From Assumption 1 and $v_i > 0$, it must be true that the term in brackets is positive, therefore, the inequality is satisfied for δp large enough.

If $g_{ij} < g(m,r) - g_m + \underline{g}$ for every j with $\alpha_{ij} > 0$, player i has more incentives to deviate when all the other players are nice than when they are guilty. In this case he would get at least $(1-\delta)\underline{g} + \delta v_i$ by conforming and (using the recursiveness of the problem) at most $(1-\delta)\overline{g} + \delta(1-\delta)\underline{g} + \frac{\delta^2}{1-\delta p}\left[p(1-\delta)\underline{g} + (1-p)v_i\right]$ by deviating. Then, for player i not to have incentives to deviate it must be the case that $\frac{\delta}{1-\delta p}\left(v_i - \underline{g}\right) \geq \overline{g} - \underline{g}$. Given that $v_i > \underline{g}$, the inequality is satisfied for δp large enough.

[14] This formulas are only valid for the case of $\tau < T$. For the case in which $\tau = T$, a similar formula for the gains from conforming results. The only difference is that in the latter case the second term disappears.

It could also be the case that player i faces players with whom the inequality is satisfied with probability α, and with probability $(1-\alpha)$ faces players with whom it is not satisfied. In this case the incentives for i to deviate will be the highest when the players belonging to the first group are guilty and the rest are nice. In this case it can be shown that player i has no incentives to deviate if $\frac{\delta}{1-\delta p}\left[p\alpha(g(m,r)-g_m)-(1-p)\underline{g}+\delta(1-p)v_i\right] \geq \overline{g}-\underline{g}$. From Assumption 1 and $v_i > 0$, it must be true that the term in brackets is positive (remember that $\underline{g} \leqslant 0$), therefore, the inequality is satisfied for δp large enough.

If player i is guilty it can be shown that, regardless of the other players flags, he does not have incentives to deviate if $\delta(1-p)v_i + (1-\delta)\underline{g} \geq 0$. This inequality is satisfied for δp large enough.

Note that there is no contradiction in the requirements made on δ and p in the four different parts of the proof: it is only required that $\delta, p \in (0,1)$ with δ and p large enough. ■

References

[1] Abreu, D., Dutta, P.K. and Smith, L.: "The Folk Theorem for Repeated Games: A NEU Condition", *Econometrica*, 62(1994), 4, 939-948.

[2] Akerlof, G.: "The Economics of Caste and of the Rat Race and other Woeful Tales", *Quarterly Journal of Economics*, 90(1976), 4, 599-617.

[3] Aumann, R. and Shapley, L.: "Long Term Competition: A Game Theoretic Analysis", mimeo, Hebrew University, (1976).

[4] Fudenberg, D., Levine, D. and Maskin, E.: "The Folk Theorem with Imperfect Information", *Econometrica, 62(1994),* 5, 997-1039.

[5] Fudenberg, D. and Maskin, E.: "The Folk Theorem in Repeated Games with Discounting or with Incomplete Information", *Econometrica,* 54(1986), 3, 533-554.

[6] Greif, A.: "Contract Enforceability and Economic Institutions in Early Trade- the Maghribi traders coalition", *American Economic Review*, 83(1993), 3, 525-548.

[7] Greif, A., Milgrom, P. and Weingast, B.: "Coordination, Commitment, and Enforcement - the case of the merchant guild", *Journal of Political Economy*, 102(1994), 4, 745-776.

[8] Kandori, M.: "Social Norms and Community Enforcement", *Review of Economic Studies*, 59(1992), 1, 63-80.

[9] Kreps, D.M. and Wilson, R.: "Sequential Equilibrium", *Econometrica,* 50(1982), 4, 863-894.

[10] Luce, R.D. and Raiffa, H.: *Games and Decisions.* New York: Wiley, (1957).

[11] Okuno-Fujiwara, M. and Postlewaite, A.: "Social Norms and Random Matching Games", *Games and Economic Behavior*, 9(1995), 1, 79-109.