KRANNERT GRADUATE SCHOOL OF MANAGEMENT

Purdue University
West Lafayette, Indiana

Cooperation among strangers: an experiment with in-definite interaction

By

Gabriele Camera
Marco Casari

Paper No. 1201
Date: June, 2007

Institute for Research in the
Behavioral, Economic, and
Management Sciences

Cooperation among strangers: an experiment with indefinite interaction [*]

Gabriele Camera
Purdue University
and
Marco Casari
Purdue University

June 2007

Abstract

We study the emergence of norms of cooperation in experimental economies populated by strangers interacting indefinitely and lacking formal enforcement institutions. In all treatments the efficient outcome is sustainable as an equilibrium. We address the following questions: can these economies achieve full efficiency? Which institutions for monitoring and enforcement promote cooperation? Finally, what classes of strategies are employed to achieve high efficiency? We find that, first, cooperation can be sustained even in anonymous settings; second, some type of monitoring and punishment institutions significantly promote cooperation; and, third, subjects dislike indiscriminate strategies and prefer selective strategies.

Keywords: experiments, repeated games, cooperation, equilibrium selection, prisoners' dilemma, random matching

JEL codes: C90, C70, D80

(*) Correspondence address: Camera (gcamera@purdue.edu); Casari (casari@purdue.edu), Krannert School of Management, Purdue University, 403 W. State St., West Lafayette, IN 47907, USA. Financial support for running the experiments was provided by Purdue's CIBER. Camera acknowledges research support from NSF grant.no. DMS-0437210. We thank Jingjing Zhang for valuable research assistance, and we thank for comments Michael Baye, John Duffy, Jason Abrevaya, Thomas Palfrey, as well as seminar participants at the Chicago ES meeting, Montreal ESA meeting, and Indiana University.

1 Introduction

There is a growing interest in economics for models of anonymous and decentralized interaction. A possible cause for this interest is that societies have become increasingly anonymous and the frequency of repeated interaction has declined. This interest is reflected in the adoption of trading environments populated by a large number of individuals who meet at random. Such frameworks are used, for example, in Diamond (1982) to model the existence of frictions in trading, in Kiyotaki and Wright (1989) to provide the foundations for monetary exchange, in Dixit (2003) to study economic governance, and in Shimer (2005) to analyze unemployment. When agents interact as strangers, as in the above settings, there exist frictions in cooperation and coordination among agents, hence achieving optimum outcomes is a challenge.

Economic theory has shown that even in anonymous groups, cooperation is theoretically possible as long as individuals are involved in a long-term interaction. The theoretical foundation can be traced back to the folk theorems for infinitely repeated games (supergames) of Friedman (1971) and the subsequent random-matching extensions in Kandori (1992) and Ellison (1994). The basic theoretical result is that cooperation is an equilibrium if agents are sufficiently patient. There exists very limited empirical evidence, however, regarding the above environments.

This paper studies matching economies in an experiment where pairs of strangers "infinitely" play a prisoners' dilemma. Strangers are anonymous subjects who are randomly matched in each period, and their histories are private information. In these economies the Pareto efficient outcome is not an equilibrium in the one-shot game, but, for an appropriate choice of parameters, it is one of the equilibria if the horizon is infinite. Kandori (1992) and Ellison (1994) proved that the Pareto efficient outcome can be achieved by adopting social norms of cooperation that rely on the threat of a "grim-trigger" punishment scheme, i.e., economy-wide defection. Basically, a subject cooperates unless someone has been caught defecting, in which case the subject should forever defect.

In practice, however, achieving the Pareto efficient outcome may be problematic because subjects are not in a stable partnership, cannot communicate their intentions to others, and can neither commit to nor enforce cooperation. One also wonders whether the subject perceives the grim-trigger punishment as a plausible threat. Given these frictions, subjects face a double challenge: not only must they be able to coordinate on the Pareto efficient outcome, but also coordinate on a credible threat that can support continuous cooperation. Our goal is to identify behavioral elements and institutional characteristics that are associated to the emergence, sustainability, and breakdown of cooperation.

This paper reports the experimental results from four treatments of matching economies where interaction is indefinitely repeated, based on a probabilistic continuation rule. Treatments differ in two dimensions: the level of information about action histories and the punishment technology. Under *private monitoring*, subjects observed only their own history and under *public monitoring*, they observed the history of the whole economy. In some treatments subjects could only punish by defecting, while in the *personal punishment* treatment, they could pay a cost to inflict a loss on their opponent.

Our study addresses the following research questions: can strangers who interact indefinitely achieve substantial levels of cooperation and efficiency? Which institutions for monitoring and enforcement promote cooperation? What classes of strategies are adopted in economies that achieve high efficiency? We obtained the following results. First, efficiency levels in our experimental economies are high and increasing with experience, even under private monitoring; this result provides empirical support for the theoretical findings in Kandori (1992) and Ellison (1994). Second, costly personal punishment significantly promotes cooperation; however, not all monitoring institutions promote cooperation. We report high cooperation levels in situations where subjects know identities and histories of opponents (*non-anonymous public monitoring*) but not when identities are unobservable (*anonymous public monitoring*). Finally, subjects appear to have preferences for certain strategies. In particular, the average subject: (a) avoids indiscriminate strategies; (b) shows a strong tendency to defect with opponents who have “cheated” her in the past; and (c) tends to disregard information on the opponent’s behavior in other matches.

One can identify three main contributions. Our study of indefinite interactions among strangers complements and extends the experimental literature on indefinitely repeated games, which has mostly focused on interactions among partners (recent examples include Palfrey and Rosenthal, 1994; Aoyagi and Frechette, 2003; Dal Bó, 2005; Duffy and Ochs, 2006). Second, our experimental findings can help define an empirically-relevant criterion for equilibrium selection, based on behavioral considerations. This is important from a practical standpoint because random matching models often display multiple equilibria with various levels of efficiency, but an unambiguous equilibrium selection criterion is missing (e.g., see the monetary equilibria in Aliprantis et al., 2006, 2007). Our laboratory findings shed light on what type of economic institutions may facilitate the emergence of norms of cooperation in anonymous societies, complementing a growing literature devoted to uncover theoretical links between the availability of enforcement and punishment institutions on one side, and patterns of exchange and cooperation on the other (e.g., Krasa and Villamil, 2000; Dixit, 2003).

The paper proceeds as follows: Section 2 discusses the related literature; Section 3 presents the experimental design; Section 4 provides a theoretical analysis; results are reported in Section 5; and Section 6 concludes.

2 Related experimental literature

Our paper builds on the experimental literature on infinitely repeated games (supergames), whose theoretical foundation can be traced back to Friedman (1971). Roth and Murnighan (1978) were the first to implement *infinitely* repeated games in an experiment by employing a probabilistic continuation rule, which transforms it into an *indefinitely* repeated game. For risk-neutral subjects, a constant continuation probability is theoretically equivalent to assuming a constant discount rate and an infinite horizon.

A number of experiments have adopted probabilistic continuation rules to study the empirical validity of folk theorems for supergames. A basic result is that subjects perceive the differences in the incentive structure of a finitely repeated versus an indefinitely repeated interaction, and react in the expected direction. For example, Palfrey and Rosenthal (1994) and Dal Bó (2005) report lower cooperation for finite duration experiments in comparison to indefinite duration experiment with a the same expected length. Moreover, the higher the discount rate the lower the cooperation. For a recent discussion see Normann and Wallace (2006).

In order to place our contribution within the existing literature, and given our focus on the models in Kandori (1992) and Ellison (1994), we will discuss indefinitely repeated experiments whose stage game is a prisoner's dilemma (for experiments with other games see Cason and Khan, 1999, Engle-Warnick and Slonim, 2004, 2006, Engle-Warnick, 2007). It is helpful to classify experiments with indefinite interaction according to two aspects, the matching protocols and the availability of information supplied about other subjects. The protocol to match subjects *within a supergame* is an empirically relevant and theoretically interesting parameter. Furthermore, all experiments we surveyed include several supergames within a session, and hence need an additional protocol to match the subjects after each supergame. We will come back later to this matching across supergames and for now focus on matching within a supergame.

The most common matching protocol within a supergame is *fixed matching*. For instance see Palfrey and Rosenthal (1994), Aoyagi and Frechette (2003), or Dal Bó (2005). Under this design, which we refer to as "partner", subjects always interact with the same person and generally support a significant level of cooperation, sometimes full cooperation. The present study employs instead a *random matching* protocol within a supergame as, for instance, in Schwartz et al. (1999) and Duffy and Ochs (2006). In any given period subjects still meet in pairs but after each period new pairs are randomly formed drawing among subjects from a larger economy with N>2 people.

A comparison of fixed matching (partner) versus a random matching (stranger) in finitely repeated games can be found in Andreoni and Croson (2002) and indefinitely repeated games can be found in Duffy and Ochs (2006). This latter study has a random matching treatment with private monitoring and the parameters were set in a way that full cooperation was an equilibrium outcome. The study finds remarkably higher cooperation in fixed than in random matching economies. Therefore, despite the theoretical viability of cooperative equilibria with random matching and private monitoring, it seems that they are empirically difficult to attain.

A key parameter when comparing cooperation rates between fixed and random matching treatments is the expected number of encounters with any given person. This number is higher under fixed matching than random matching for economies of equal size and identical continuation probability. As a consequence the deck is stacked in favor of observing higher cooperation in fixed matching. To avoid this bias, we do not use a partner treatment. Instead we introduce a novel design that, on one hand, equalizes the expected number of encounters with any given person across treatments, while on the other hand, provides as much information as in the partner treatment. The new design has random matching and public monitoring, as we provide the complete history for each agent in the economy. A subject knows the identity of their opponent (non-anonymous) as well as what their opponent chose when meeting other participants.[1]

A second novel feature of our study is to understand which one of the several available strategies that support a given equilibrium outcome have been employed.[2] This issue has been largely unexplored in the experimental literature on supergames, as it has mostly focused on measuring the levels of cooperation. As we will later clarify, we develop a design where we can exploit differences in information across treatments in order to change the strategy set and hence identify the type of strategies employed.

We also relate the choice of punishment strategy in an indefinitely repeated setting to the literature on costly personal punishment in one-shot settings. Experimental studies of finitely repeated social dilemmas have evidenced a surprising tendency of subjects to engage in costly personal punishment of others, in particular defectors. Although, this

[1] Schwartz et al. (1999) and Duffy and Ochs (2006) also consider treatments when subjects receive some information about the reputation of their current opponent while preserving anonymity of the opponent. Our non-anonymous public monitoring gives all individual histories and so provides information about reputation and it also reveals the identity of the opponent.

[2] These strategies include off-equilibrium threats that on the equilibrium path will never be employed. The features of such threats are largely irrelevant as long as they are credible and they generate a sufficiently low continuation payoff.

behavior is inconsistent with personal income maximization, it has been shown to be remarkably robust (Ostrom et al, 1992; Fehr and Gaechter, 2002; Casari and Plott, 2003). A third novel feature of our study is to examine how this behavioral trait may be employed in supporting the cooperative equilibrium in an infinitely repeated game, where there does already exist a punishment technology. This design may be useful in isolating possible elements or economic institutions that can facilitate selecting the cooperative equilibrium in a more general setting.

As noted earlier, the matching protocol across supergames is also important because of possible contagion effects across supergames. It is therefore helpful to mention the various protocols adopted in previous experiments. To play a supergame in a session with N participants, subjects can be partitioned into K economies. The way we ran multiple supergames is to ensure that any two subjects were never assigned to the same economy for more than one supergame. A more rigorous partitioning procedure in the experimental literature is to rule out that anyone may share a common past opponent.[3] Both procedures control for contagion effects.[4] This contrasts with randomly matching the same set of subjects after each period and after each supergame (for instance, Schwartz et al., 1999, and random pairing in Duffy and Ochs, 2006).

3 Experimental design

This experiment has four treatments (Table 1). While the stage game (Table 2), the continuation probability, and matching protocols were identical across treatments, we manipulated the amount of information and the punishment options available to subjects. The efficient outcome can be supported as an equilibrium in all treatments.

[3] In Dal Bó (2005) each subject plays three supergames (treatment). In the “Dice” sessions, in each supergame participants are partitioned into K=(N/2) two-person economies. The partitioning across supergames is such that the decisions one subject made in one supergame could not affect, in any way, the decisions of subjects he or she would meet in the future. Ensuring the absence of contagion effects in this manner requires very large session sizes. For a theoretical discussion of matching procedures see Aliprantis et al. (2006, 2007).

[4] In our study each subject played for five supergames. Subjects may have shared a common past opponent in supergames three or later. Aoyagi and Frechette (2003) used a different in between matching protocol; each agent plays G>10 supergames. In the first 10 supergames they partition agents as in the former way described in the main text above and in the last (G-10) supergames the randomly rematch participants.

Table 1: Comparison of experimental designs[(*)]

Matching protocol (within the economy)	*Anonymity (identity of opponent is unknown)*	*Designs*	*Strategy type* *Global*	*Reactive*	*Targeted*
Fixed matching	No	Partner-*not in this study*	(i)	(i)	(i)
Random matching	No	Public monitoring (non-anonymous)	Yes	Yes	Yes
	Yes	Anonymous public monitoring	Yes	Yes	
	Yes	Private monitoring		Yes	
	Yes	Private monitoring with punishment		Yes	(ii)
			High power Not selective	*Medium power Moderately selective*	*Low power Highly selective*

Table 2: The stage game

(A) Notation in the theoretical analysis

Player 1/ player 2	*Y*	*Z*
Y	y, y	l, h
Z	h, l	z, z

(B) Parameterization of the experiment

Player 1/ player 2	*Y*	*Z*
Y	25, 25	5, 30
Z	30, 5	10, 10

The stage game. The stage game is a standard prisoners' dilemma with payoffs determined according to Table 2 (payoffs to column and row players, respectively).[5] We call action Y *cooperate* and action Z *defect*. So, we say that there is *cooperation* in the

[(*)] (i) In partner, the distinction among targeted, reactive, and global strategies is irrelevant because of the fixed matching. (ii) One could interpret the possibility of personal punishment as a form of targeted strategy, although the personal punishment reduces the continuation payoffs for the punisher more than with the reactive strategy. Personal punishment expands the set of strategies. In particular it allows for a targeted strategy because an agent can punish his opponent after observing the choice of his opponent.

[5] We selected this parameterization as it scores high on the indexes proposed by Rapoport and Chammah (1965), Roth and Murnighan (1978), and Murnighan and Roth (1983) that correlate with the level of cooperation in the indefinitely repeated prisoners' dilemma in a partner protocol. Also, in Table 2 we have $0 \leq l < z < y < h$ and $2z < l+h < 2y$.

pair only if both subjects choose Y. Consequently, we will define the degree of cooperation in the economy according to how many pairs cooperate.

The supergame. A supergame (or *cycle*, as we will call it) consists of an indefinite interaction among subjects achieved by a random continuation rule, as first introduced by Roth and Mangham (1978). A supergame that has reached period t continues into t + 1 with a probability $\delta \in (0, 1)$, so the interaction is of finite but uncertain duration. We interpret the continuation probability δ as the discount factor of a risk-neutral subject. The expected duration of a supergame is 1/(1−δ) periods, and we set δ = 0.95, so in each period the supergame is expected to go on for 20 (additional) periods.[6] In our experiment the computer drew a random integer between 1 and 100, using a uniform distribution, and the supergame terminated with a draw of 96 or of a higher number. All session participants observed the same number, and so it could have also served as a public randomization device.

*The experimental session***.** Each experimental session involved twenty subjects and exactly five cycles. We built twenty-five economies in each session by creating five groups of four subjects in each of the five cycles. This matching protocol across supergames was applied in a predetermined, round-robin fashion. More precisely, in each cycle each economy included only subjects who had neither been part of the same economy in previous cycles nor were part of the same economy in future cycles. Subjects did not know how groups were created but were informed that no two participants ever interacted together for more than one cycle.

Participants in an economy interacted in pairs according to the following matching protocol within a supergame. At the beginning of each period of a cycle, the economy was randomly divided into two pairs. There are three ways to pair the four subjects and each one was equally likely. So, a subject had one third probability of meeting any other subject in each period of a cycle. For the whole duration of a cycle a subject interacted

[6] With continuation probability δ, the expected number of periods is $S = \sum_{n=1}^{\infty}(1-\delta)\delta^{n-1}n = 1/(1-\delta)$.

exclusively with the members of her economy. By design, cycles for all economies terminated simultaneously.

Treatments. The experiment consisted of four different treatments that differed in the availability of information and punishment options (Table 3). All treatments maintained the same continuation probability, stage game parameters, and matching protocols. Two treatments were characterized by *private monitoring,* i.e., subjects could observe actions and outcomes in their pair, but not the identity of their opponent. One, denoted *private monitoring*, was the benchmark case as in Kandori (1992). The other, denoted *private monitoring with punishment*, added the possibility of personal punishment. Subjects could lower the earnings of their opponent, at a cost, after having observed their opponent's action. In order to do so, we added a second stage to the one-shot game. The first stage was the prisoners' dilemma in Table 2B. In the second stage actions were revealed, and subjects had the opportunity to pay 5 points to reduce the opponent's earnings by 10 points. No one could observe any of the actions outside their pair, including the personal punishment. The remaining two treatments were characterized by *public monitoring*, which simply means that every subject could observe the actions taken in every pair. In one treatment, denoted *non-anonymous public monitoring*, histories were associated with identities of subjects. In the remaining treatment, denoted *anonymous public monitoring*, subjects observed histories but not identities.

To summarize, the availability of information about actions in the economy was set at one of three different degrees. First, subjects could be aware only of their own history (private monitoring, private monitoring with punishment) or of the history of the entire economy. Second, the history of the economy could be made available at an aggregate (anonymous public monitoring) or individual level (non-anonymous public monitoring). The history of the economy was provided at the aggregate level by listing everyone's actions in random order and without identifiers. On the contrary in the non-anonymous public monitoring treatment, individual histories were listed with the person's ID as label. This allowed a subject to inspect the opponent's actions in previous encounters with her as well as the opponent's behavior with others.

We recruited 160 subjects through announcements in undergraduate classes at Purdue University and signed up online. The sessions were run at the Vernon Smith Experimental Economics laboratory at Purdue University. No eye contact was possible among subjects, and copies of the instructions were on all desks. Instructions were read aloud.[7] Average earnings were $29.50 per subject. A session lasted on average 110 periods for a running time of 2.5 hours, including instruction reading and a quiz. Details about the number and length of sessions are provided in Table 3 (each session had 20 participants and 5 cycles).

Table 3: Four experimental treatments

	Private monitoring		*Anonymous public monitoring*		*Private monitoring with punishment*		*Public monitoring (non-anonymous)*	
Information	No subject IDs; own and current opponent's action		No subject IDs; list of all 4 group members' actions in random order		No subject IDs; own and current opponent's actions		Subject IDs are public; individual histories of all 4 group members	
Punishment	Subjects can only punish by defecting		Subjects can only punish by defecting		Subject can pay 5 points to reduce opponent's payoff by 10 points		Subjects can only punish by defecting	
Session date	21.4.05	7.9.05	27.4.05	1.9.05	28.4.05	6.9.05	12.4.05	8.9.05
Show-up fee	$5	$5	0	0	$5	0	0	0
Periods	71	104	129	125	139	99	86	128

4 Theoretical predictions

We first introduce a theoretical framework for the private monitoring treatment based on Kandori (1992) and then discuss the other treatments, in particular the private monitoring with punishment and public monitoring. The analysis is based on the assumption of identical players, who are self-regarding and risk-neutral, in the absence of commitment and enforcement.[8]

[7] A copy of the instructions can be found at http://www.mgmt.purdue.edu/faculty/casari/anonymous.htm

[8] The theoretical framework is one of a homogeneous population. An alternative approach is to consider subjects of different types in the experiment as, for example, in Costa- Gomes, Crawford, and Broseta (2001) and Healy (forthcoming).

An "economy" is composed of four players *a*, *b*, *c*, and *d* who interact for an indefinite number of periods denoted t = 1, 2, ... Participants are randomly paired to play the prisoners' dilemma of Table 2. There are three ways to pair participants in an economy, {*ab*, *cd*}, {*ac*, *bd*}, or {*ad*, *cb*}, and in each period one pairing was randomly chosen with equal probability.[9]

4.1 Equilibrium in the stage game

Consider the stage game described in Table 2A, which is a prisoners' dilemma. The players simultaneously and independently select an action from the set $\{Y, Z\}$. We allow for mixed-strategies. Let $\pi \in [0,1]$ denote the probability that the representative player selects Y, and $1-\pi$ the probability that he selects Z. We use $\Pi \in [0,1]$ to denote the given selection of the opponent.

The unique Nash equilibrium is defection. In equilibrium both players choose *Z*, the minmax action, and earn *z*, the minmax payoff. The representative player's payoff is simply his expected utility, denoted *U*. This can be rearranged as:

$$U = z + \Pi(h-z) - \pi[\Pi(h-y) + (1-\Pi)(z-l)].$$

The player maximizes *U* by choosing π, so can assure himself payoff *z*, independent of Π. Notice that *U* is linear in π, and we have assumed $y < h$ and $l < z$. It follows that the player's best response is to set $\pi = 0$, for any $\Pi \geq 0$.

Since $2z < l + h < 2y$, total surplus in the economy is maximized when each pair cooperates. Thus, we refer to the outcome where both players in both pairs select *Y* as the (Pareto) efficient or fully cooperative outcome. If both pairs in the economy select $\{Z, Z\}$, then we say that the outcome is inefficient. A Nash equilibrium is a fixed point in the players' aggregate best response, so $\pi = \Pi = 0$ is the unique equilibrium.

4.2 Equilibrium in the indefinitely repeated game with private monitoring

With private monitoring indefinite repetition of the stage game with randomly selected opponents can expand the set of equilibrium outcomes. In this section we

---

[9] Strictly speaking, we are dealing with a game with varying opponents, since players are paired randomly at each point in time. However, action sets and payoff functions are unchanging. Thus, we refer to it as a supergame, following the experimental literature.

provide sufficient conditions so that the equilibrium set includes the efficient outcome, following the work of Kandori (1992) and Ellison (1994). This is achieved when everyone cooperates in every match and in every period.

The *private monitoring* treatment is characterized by two informational frictions. Players cannot observe identities of opponents, so we say that players are *strangers*. Second, players can neither communicate with each other nor observe action histories of others; they can only observe the outcome resulting from actions taken in their pair.

Clearly, the inefficient outcome can be supported as a sequential equilibrium through the strategy "defect forever." Because repeated play does not decrease the set of equilibrium payoffs, *Z* is always a best response to play of *Z* by any randomly selected opponent. In this case the players' payoff in the indefinitely repeated game is the present discounted value of the minmax payoff forever $z/(1-\delta)$.

If δ is sufficiently high, however, then the efficient outcome can be sustained as a sequential equilibrium. Formally, we have the following result.

Lemma 1. *Let* $\delta^* \in (0,1)$ *be the unique value of* δ *that satisfies*

$$\delta^2(h-z) + \delta(2h-y-z) - 3(h-y) = 0\,.$$

If $\delta \geq \delta^*$*, then the efficient outcome can be sustained as a sequential equilibrium. In an economy with full cooperation, every player receives payoff* $y/(1-\delta)$.

The proof is in Appendix A and follows that found in Kandori (1992). Here, we provide intuition. Conjecture that players behave according to actions prescribed by a social norm; a social norm is simply a rule of behavior that identifies "desirable" play and a sanction to be selected if a departure from the desirable action is observed. We identify the desirable action by *Y* and the sanction by *Z*. Thus, every player must cooperate as long as she has never played *Z* or has seen anyone select *Z*. However, as soon as a player observes *Z*, then she must select *Z* forever after. This is known as a *grim trigger strategy*. In our experiments, this strategy is equivalent to what we call a *reactive strategy* (i.e., a player will choose *Z* if and only if one of his opponents has chosen *Z*).

Given this social norm, on the equilibrium path everyone cooperates so the payoff to everyone is the present discounted value of *y* forever: $y/(1-\delta)$. A complication arises when a player might want to defect since $h > y$. Hence, since $z < y$, it must be the threat of minmax forever that deters a player from defecting. Notice that a player deviates in several instances—first, in equilibrium, if she has not observed play of *Z* in the past but chooses *Z* currently, and second, off-equilibrium, if she has observed play of *Z* in the past but plays *Y* currently.

Consider one-time deviations by a single player (unimprovability criterion). It should be clear that cooperating when no defection has ever been observed is optimal only if the agent is sufficiently patient. The future reward from cooperating today must be greater than the extra utility generated by defecting today. Instead, if a defection occurs and everyone plays according to the social norm, then everyone will end up defecting since the initial defection will spread by contagion. Given that the economy has only four players, this contagion in our experimental economies should occur very quickly. This is illustrated in Figure 1, by the line labeled reactive strategy.

Cooperating after observing a defection should also be suboptimal. Choosing *Y* in this instance can delay the contagion but cannot stop it. To see why, suppose a player has observed *Z*. If he meets a cooperator in the following period, then choosing *Y* generates a current loss to the player because he earns y (instead of h). If he meets a deviator, choosing *Y* also generates a current loss because he earns l rather than z. Therefore, the player must be sufficiently impatient to prefer play of *Z* to *Y*. The smaller are l and y, the greater is the incentive to play *Z*. Our parameterization ensures this incentive exists for all $\delta \in (0, 1)$ so it is a dominant strategy to play Z after observing (or selecting) Z.

Assuming a homogenous population in our experimental economies, the preceding discussion has two immediate predictions, which are put forward below.

Proposition 2. *In our experimental economies with private monitoring, the efficient outcome can be sustained as an equilibrium.*

Proposition 2 follows directly from Lemma 1. For the efficient outcome to be a feasible, we need $\delta \geq \delta^*$. In our experimental design $\delta = 0.95$ and $\delta^* = 0.443$, a value that solves the condition in Lemma 1 for the parameterization given in Table 2B. [10]

Proposition 3. *In our experimental economies with private monitoring, the efficient outcome can be sustained as an equilibrium without using personal punishment. Instead, this outcome cannot be sustained as an equilibrium by relying exclusively on personal punishment.*

Recall that with personal punishment an agent has the option, at a cost, to lower the current earnings of his opponent only after observing the outcome of the prisoners' dilemma. In a one-shot interaction, choosing personal punishment is a dominated action because it is costly for the punisher. In our design the interaction is indefinitely repeated, but personal punishment is still individually suboptimal for the same reason it is in the one-shot game.

Personal punishment is dominated for two reasons. First, it does not trigger a faster contagion to the state of economy-wide defection. In our design agents are anonymous, randomly matched in each period, and can only observe actions and outcomes in their pair. Hence, to someone outside the match, a choice of personal punishment is no more visible than a choice of defection. Because of private monitoring, personal punishment is no more efficient than a "grim trigger" defection strategy, and in addition, it is costly.

Second, the sole use of personal punishment cannot sustain cooperation, even with public monitoring. The reason is that personal punishment is not a credible threat because after observing a defection, it is never individually optimal to pay the cost for personal punishment.[11] For instance, a strategy where agents always cooperate and respond to a defection only with personal punishment for the period cannot sustain cooperation. After the opponent defects, an agent has no incentive to inflict personal punishment because it simply adds a further loss. Additionally, the incentive to defect in following periods

[10] Contagion equilibria as in Kandori (1992) are not robust to adding a small amount of noise in the observation of individual behavior. With noise, equilibria arise similar to those in the continuum limit where individual behavior is unobservable (e.g., see Al-Najar and Smorodinsky, 2001, Fudenberg, Levine, and Pesendorfer, 1998, Levine and Pesendorfer, 1995). One can suppose that the larger the population, the greater the instances of noise in observability. To lessen such instances in our experimental economies, we work with four-agent economies, the smallest possible number that allow pairwise anonymous matching.

[11] On the contrary, defecting after having observed a defection is an optimal strategy.

remains because defection is the unique best response in the one-shot game. In conclusion, though personal punishment is a sufficient threat to sustain cooperation, it is not a credible one.

4.3 Equilibrium in the indefinitely repeated game with public monitoring

In this section we specify that the efficient outcome can also be sustained as a sequential equilibrium in the treatments in which the history of actions taken in the economy is public information. Of course, with more information the possible strategies that sustain the efficient outcome are expanded.

Proposition 4. *In our experimental economies with public monitoring, the efficient outcome can be sustained as an equilibrium.*

When we allow for public monitoring, instead, the value of δ^* can only fall. It is now 0.25 since according to the grim trigger strategy, a current defection implies a sure defection by any future partner. This is illustrated in Figure 1 by the line denoted *global strategy*, representing a grim trigger strategy in which permanent defection occurs as soon as a defection is detected *anywhere* in the economy (in or outside the pair).

Figure 1: Dynamic reaction to a defection in an economy

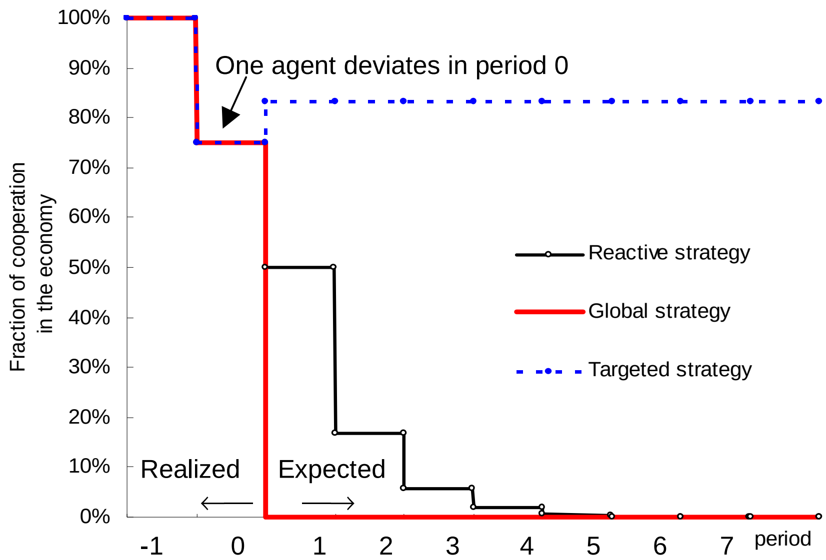

The important aspect of public monitoring is that giving more information about actions is beneficial to cooperators in several different respects. First, a player who observes a deviation might have the option to defect in the future only with a subset of players (for instance, those known to have deviated). This can only increase the frequency of cooperation in the economy because it allows players to cooperate with those known to cooperate. Second, a player is less likely to experience a defection as a result of a past defection by someone else. In addition more information is detrimental to deviators, since they can be targeted more effectively. All of these elements serve to increase the payoff for a cooperator and decrease it for a deviator, which generates incentives to cooperate for even lower discount factors.

Below we identify three broad classes of strategies. They do not exhaust all possible behaviors but are indicative of three intuitive ways of behaving. First, players could switch from a cooperative mode to a punishment mode when they observe a defection, no matter if coming from an opponent or someone else in the economy. We have already called it a *global strategy*. Conversely, players could switch to a punishment mode when they observe an opponent defect, but stay in cooperative mode if a defection is observed elsewhere in the economy, what we refer to as a *reactive strategy*. Finally, an even more selective strategy would involve a player switching to a punishment mode after observing an opponent defect, limiting defections only to future encounters with the same opponent, while staying in a cooperative mode with anyone else. We refer to this as a *targeted strategy*. It is easily demonstrated that, with a targeted strategy, the efficient outcome is optimal as long as δ is greater than 0.5.

In random matching with non-anonymous public monitoring all classes of strategies are available. On the contrary, with private monitoring reactive strategies are available, but global and targeted strategies are not. Hence, variations in cooperation level between treatments could suggest what class of strategies enhances cooperation (see Table 1).

One can classify strategies also using a “power” and a “selectivity” score. Power relates to the incentives to keep cooperating, while selectivity relates to the incentives to punish following a defection. The *power* of a strategy is the maximum punishment that can be inflicted on a defector, which depends on the immediacy and frequency of punishment. Global strategies have the most power because punishment can take place the following period and applies to everyone (Figure 1)—all else being equal, the greater the power, the lower the continuation payoff for someone who defects when everyone else is cooperating. Hence, a strategy with greater power reduces the incentives to defect. Global strategies provide the largest possible threat since punishment is immediate and indiscriminate. Targeted strategies have the least power, while reactive strategies occupy a space in-between the two. For example, with public monitoring, the lower bound for δ falls by about 40% when we move from a reactive strategy to a global strategy, and by about 50% when we move from a targeted strategy to a global strategy.

The *selectivity* of a strategy is linked to the incentives to punish defectors. Targeted strategies are the most selective and allow agents to punish at the lowest cost. We use the term *selectivity* because it is related to the richness of information needed to support cooperation. A more selective strategy requires a finer partition of the information set.

5 Results

We first present results on the aggregate outcome (Results 1-6) and then on the strategies employed to sustain those outcomes (Results 7-11). The section is broken up by treatment so that the discussion can be more focused.

Result 1. *The introduction of public monitoring in the non-anonymous treatment increased cooperation over private monitoring.*

Figures 2 and 3 and Table 4 provide support for Result 1. In the non-anonymous public monitoring average cooperation across economies is 81.5%. For an economy k=1,..,50 we define the action a_{it}^k of an agent j=1,..,4 in period t=1,..,T^k of economy k to be an element $a_{it}^k \in \{0,1\} \equiv \{Z, Y\}$; a cooperative action is coded as 1 and a defection is coded as 0. Therefore, the average cooperation in an economy k is $c_k = \frac{1}{4}\sum_{t=1}^{T^k}\sum_{i=1}^{4} a_{ti}^k$ and across economies is $c = \frac{1}{50}\sum_{k=1}^{50} c_k$. Although economies have different length T^k, they are given equal weight in our measure *c* of average cooperation across economies.

Figure 2: Average cooperation across treatments[(*)]

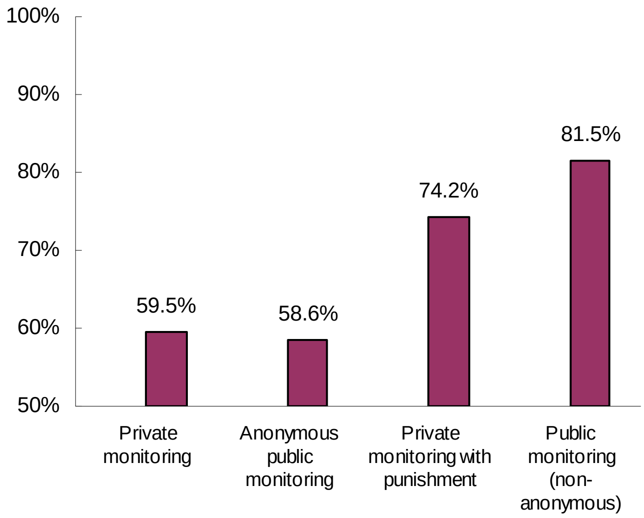

A Mann-Whitney test conducted on cooperation in non-anonymous public monitoring shows significant difference with private monitoring (59.5%, p-value 0.0001) and with anonymous public monitoring (58.6%, p-value 0.0000). Result 1 is consistent with data reported in the literature of high levels of cooperation in the partner treatment. Similar to a partner design, participants interact in pairs and know the whole individual history of interaction, but unlike it, the match for the period is randomly picked from a group of three other individuals.

We also report the distribution of the fifty economies by average cooperation level, which is illustrated in Figure 3. About 38% of the economies have cooperation rates above 98%. The superiority of non-anonymous public monitoring is clear also from the average cooperation in the initial period across economies, shown in Table 4.

[(*)] We aggregated economies from all cycles by treatment and carried out Mann-Whitney tests of pairwise differences in cooperation between treatments. Differences are statistically significant at 1% level with two exceptions: private monitoring vs. anonymous public monitoring and private monitoring with punishment vs. non-anonymous public monitoring. One economy is one observation; in each comparison $n_1=n_2=50$.

Figure 3: Empirical distribution of average cooperation by economy(*)

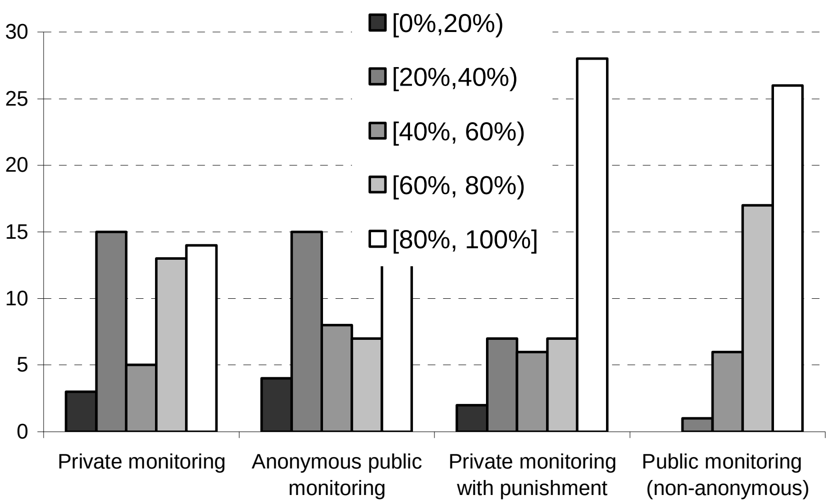

Table 4: Cooperation in the first period of an economy(**)

Number of cooperative actions	Private monitoring	Anonymous public monitoring	Private monitoring with punishment	Public monitoring (non-anonymous)
		Average cooperation		
	73.5%	70.5%	84.5%	87.0%
		Frequency of cooperation in an economy		
4	36%	26%	50%	54%
3	30%	42%	38%	40%
2	28%	22%	12%	6%
1	4%	8%	0%	0%
0	2%	2%	0%	0%
		Frequency of cooperation in a match		
2	58%	51%	71%	75%
1	31%	39%	27%	24%
0	11%	10%	2%	1%

(*) Kolmogorov-Smirnov two-tail two-sample test on distributions confirms results from the Mann-Whitney tests on the differences between averages. On one hand private monitoring and anonymous public monitoring are not statistically different (10% confidence level, n_1=n_2=50). Conversely, private monitoring with punishment and non-anonymous public monitoring are not statistically different. Treatments from the two groups are instead statistically different at least at a 5% level.
(**) In each treatment the number of observations is 50 for "average" and "frequency of cooperation in an economy" and 100 for "frequency of cooperation in a match."

Result 2. *Cooperation did emerge in economies with private monitoring. Cooperation in later cycles is higher than in earlier cycles.*

Figures 2 and 3 and Table 4 provide support for Result 2. In the private monitoring treatment average cooperation across economies was 59.5% for all periods and 73.5% for just the first periods, which are remarkably high given results in previous studies (Schwartz et al., 1999; Ochs and Duffy, 2006). This provides support to the empirical relevance of the theoretical results of Kandori (1992) and Ellison (1994).

The average cooperation had an increasing trend across cycles, as seen in Figure 4.

Figure 4: Average cooperation across cycles

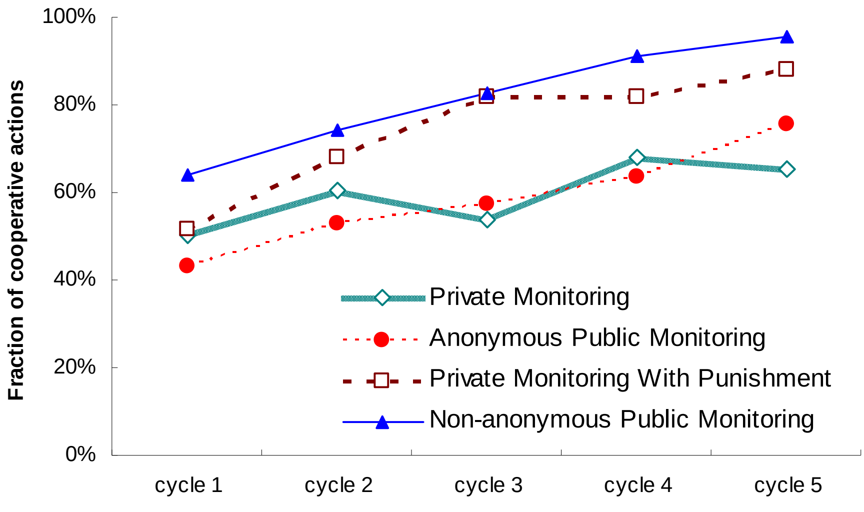

This figure suggests that as subjects became familiar with the incentive structure of the indefinite repetition, they responded by increasing cooperation level.[12] This finding is in line with previous studies (Aoyagi and Frechette, 2003; Dal Bó and Frechette, 2006) and marks a difference with finitely repeated prisoners' dilemmas and voluntary public

[12] There is statistically significant learning across cycles. We aggregated economies from all treatments by cycle and carried out Mann-Whitney tests of pairwise differences in cooperation between cycles. The increment between cycle 1 and 5 is significant at 1% level. The most significant jump in cooperation level is from cycle 1 to cycle 2 (5% significance) while the difference between cycles 4 and 5 is not significant at a 10% level (one economy is one observation; in each comparison $n_1=n_2=40$).

good games, where average cooperation in experiments is generally positive but declining over time (Palfrey and Rosenthal, 1994). Even if one expects some degree of cooperation in our private monitoring economies given the salience of the payoffs, the increase in cooperation from one cycle to the next displayed in Figure 4 is in sharp contrast with the evidence from finitely repeated games.

Result 3. *In the anonymous treatments, the introduction of public monitoring did not improve cooperation over private monitoring.*

Subjects in public monitoring possess information about the choices of others that is unavailable in private monitoring. Figure 2 shows that when this information is anonymous, it does not foster cooperation. Average cooperation is around 59% in both treatments, and the difference is statistically insignificant (Mann-Whitney test, p-value 0.418, n_1=n_2=50). First period averages lead to the same conclusion (Table 4).

Result 4. *The introduction of personal punishment in the anonymous treatments increased cooperation.*

Figures 2 and 3 and Table 4 provide support for Result 4. When we add personal punishment to economies with private monitoring, average cooperation jumps from 59.5% to 74.2%. This difference is statistically significant at a 1% level (Mann-Whitney test, p-value 0.0067). This difference is also evident when comparing average cooperation in the first period of each cycle (73.5% vs. 84.5%, Table 4). Surprisingly, average cooperation is statistically indistinguishable from the non-anonymous public monitoring treatment (Mann-Whitney test, p-value 0.154).

Result 5 (Realized efficiency). *In the anonymous treatments, the introduction of personal punishment increased realized efficiency over private monitoring and over public monitoring.*

The comparison among treatments in terms of realized efficiency substantially confirms the conclusions drawn in Results 1-4 in terms of average cooperation. We define realized efficiency in an economy k as $e_k = \frac{1}{T^k}\sum_{t=1}^{T^k}\sum_{i=1}^{4}\left(\pi_{ti}^k - 10\right)/\left(25-10\right)$.

Where the payoffs of the static-game are given Table 1 and range from minimum earnings of 10 to maximum earnings of 25 points. Realized efficiency e_k ranges from 0 to 1. In particular, e_k=0 when everyone in the economy always defect, and e_k=1 when everyone in the economy always cooperates. With personal punishment realized efficiency can be negative, with a minimum of -1 when everyone always defects and always punishes. The realized efficiencies for the four treatments in the experiment were 59.5%, 58.6%, 65.2%, and 81.5%, respectively.

Results 6 (selection of equilibrium). *In all treatments, period 1 cooperation is significantly different than zero. Hence, there is no evidence of coordination on the inefficient outcome.*

Table 4 provides evidence for Result 6. Choices in the first period of each economy suggest whether some equilibrium among the many possible had a particularly strong drawing power. One can examine how subjects coordinated in the initial period by looking either at agreement of choices in the economy or in the pairwise match; see Table 5 in Appendix B. Either way, we can rule out that subjects attempted to coordinate on defection. In particular, at least half of the economies started with full cooperation in two treatments, public monitoring (non-anonymous) and private monitoring with punishment. If we consider matches as the relevant unit of observation, both subjects cooperated in more than 50% of the matches in every treatment.

Result 7 (strategy employed). *There is evidence of use of reactive strategies in the private monitoring treatment. Subjects who observed a defection by their opponent switched from a cooperative mode to a punishment mode.*

Table 5 (in Appendix B) and Figure 4 provide support for Result 7. Table 5 reports the results from a probit regression that explains the individual choice to cooperate (1) or not (0) using as regressors dummies that control for fixed effect (cycles, periods within the cycle, individuals), as well as the duration of the previous cycle, and a set of six regressors to trace the response of the representative subject to an observed defection. In this manner we can understand the type of strategy employed by the representative subject. The specific choice of regressors is one among many possible ways to trace strategies. One advantage is to detect whether subjects followed theoretically well-known strategies such as grim trigger or tit-for-tat (Axelrod, 1984). For simplicity our strategies embed a maximum delay of five periods. The grim trigger regressor[13] has a value of 1 in *all* periods following an observed defection and 0 otherwise. The five tit-for-tat regressors have a value of 1 only in *one* period following an observed defection and 0 otherwise. The first takes value 1 in the period immediately following the defection. The second takes value 1 in the second period following a defection, and so on.

If subjects switched from a cooperative to a punishment mode following an observed defection, we expect at least one of the strategy regressors to be negative. For example, if subjects punished for just two periods following a defection, we expect the sum of the estimated coefficients of the grim trigger regressor and the tit-for-tat regressors to be negative for the first and second period following a defection, and zero afterwards.

Figure 5: Strategies of the representative subject in private monitoring

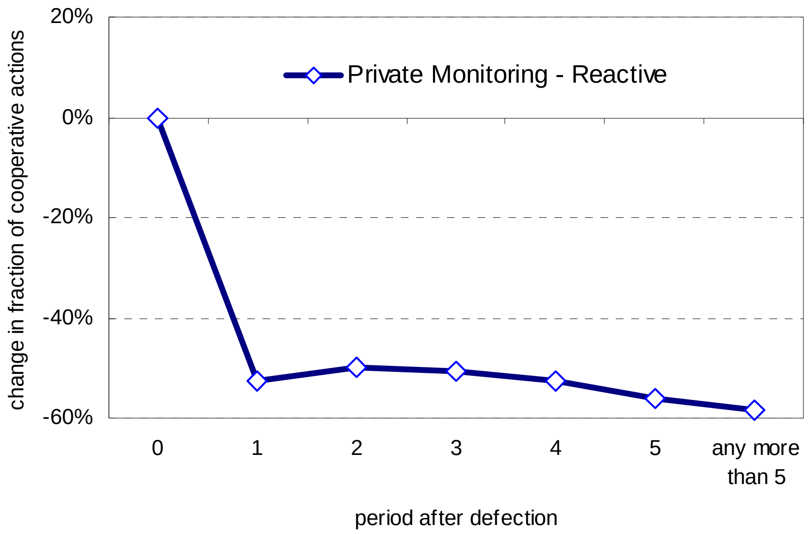

[13] We label a regressor "grim trigger" because it reminds us of the well-known grim trigger strategy, which specifies a permanent shift to punishment following a defection.

Figure 5 illustrates the marginal effect on the frequency of cooperation in the periods that followed an observed defection.[14] The focus on the first five periods is for convenience in showing patterns in the results. The representation for "any more than five" periods is based on the marginal effect of the grim trigger regressor only. The representation for periods 1 though 5 is based on the sum of the marginal effects of the grim trigger regressor and the relevant tit-for-tat regressor. The L-shaped pattern of response suggests a downward shift in cooperation levels immediately after a defection. This shift appears to be persistent. The grim trigger regressor is significant at a 1% level, and all other strategy regressors are significant at 10% level or more (Table 5).[15]

Result 8 (strategy employed). *Two types of strategies were used in the private monitoring with punishment treatment. Subjects who observed a defection by their opponent sometimes employed personal punishment while staying in a cooperative mode. Other times, they switched from a cooperative mode to a punishment mode.*

Table 5 and Figure 6 provide support for the first part of Result 8. Similar to the private monitoring treatment, Figure 6 (see footnote 13) suggests a downward shift in cooperation levels following a defection, which is immediate and persistent. The grim trigger regressor is significant at a 1% level (Table 5). In contrast to the private monitoring treatment, the magnitude of the downward shift in cooperation levels is now substantially lower.

One may conjecture that subjects sometimes continued cooperating but sanctioned through personal punishment. Table 6 shows that most of the personal punishment was given by cooperators when their opponent defected. In about 58% of such encounters, the cooperator requested personal punishment be inflicted on the opponent.

---

[14] Figure 5 is based on Table 5 using the coefficient estimates coding reactive strategies. Period 0 is exogenously set at 0%. The point for "any more than 5" is the marginal effect on the frequency of cooperation of the grim trigger regressor. Periods 1 through 5 are the sum of two marginal effects on the frequency of cooperation, the effect of the grim trigger regressor plus the proper tit-for-tat regressor (i.e. coding reaction one period after the observed defection for period 1, coding reaction two periods after the observed defection for period 2, etc.). Marginal effects for the tit-for-tat regressors are computed for grim trigger regressor set at 1 (i.e. defection)

[15] Table 6 reports that the actual length of the previous cycle influenced the propensity of participants to cooperate—the longer the previous cycle, the higher the current cooperation level. This confirms the finding reported in Aoyagi and Frechette (2003) and Engle-Warnick and Slonim (2004).

Figure 6: Strategies of the representative subject in private monitoring with punishment

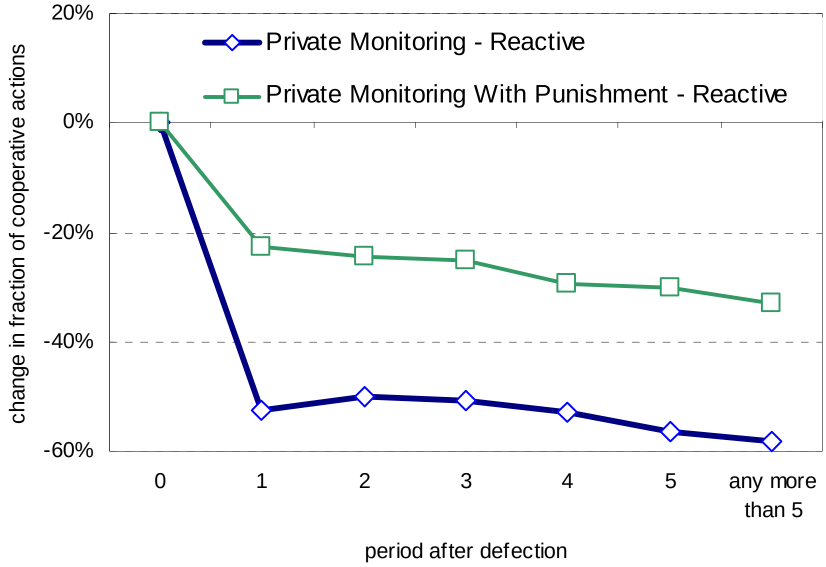

Table 6: Frequency of personal punishment(*)

		Opponent receiving punishment	
		Cooperate	Defect
Subject requesting punishment	Cooperate	0.1%	58.3%
	Defect	5.4%	5.2%

Whenever a subject requested personal punishment, she was more likely to continue cooperating. Table 7A suggests that a cooperator encountering a defector was much more likely to cooperate in the following period when she requested personal punishment than when she did not; there was an astounding 29% point difference (75.5% vs. 46.7%).[16]

In interpreting Result 8, recall that our theoretical framework is one of a homogeneous population, as in Kandori (1992) and Ellison (1994). Within this framework the punishment behavior we have observed seems at odds with equilibrium

(*) The unconditional frequency of personal punishment is 9.1%. Each cell indicates the frequency of personal punishment inflicted on the opponent conditional on the outcome in the match in stage one (there are four possible outcomes). The outcome (Cooperate, Defect) occurred 509 times.

[16] Table 7B suggests that a defector who had been punished by a cooperator was more likely to cooperate in the following period (34.5% vs. 24.1%). Once we controlled for all other factors, however, the evidence is not so clear-cut (Table 5).

predictions, since subjects could only achieve cooperation by triggering to permanent defection. As already noted, for our parameterization the set of outcomes does not expand with the addition of personal punishment. There are, however, behavioral reasons why subjects may have preferred personal punishment to community punishment. In the concluding section of the paper we will put forward some conjectures.

Result 9 (strategy employed). *Subjects preferred reactive strategies over global strategies. In the anonymous public monitoring treatment, a defection by an opponent generated a stronger response than a defection elsewhere in the economy.*

In anonymous public monitoring subjects observed whether a defection had occurred in the match or elsewhere in the economy. Hence, both reactive and global strategies were available. A subject using a reactive strategy punished everyone once a defection in the match had been experienced, but kept cooperating if a defection was observed outside the match. In contrast a subject using a global strategy started punishing everyone once she had observed a defection, no matter if it came from an opponent or someone else.

Figure 7: Strategies of the representative subject in anonymous public monitoring

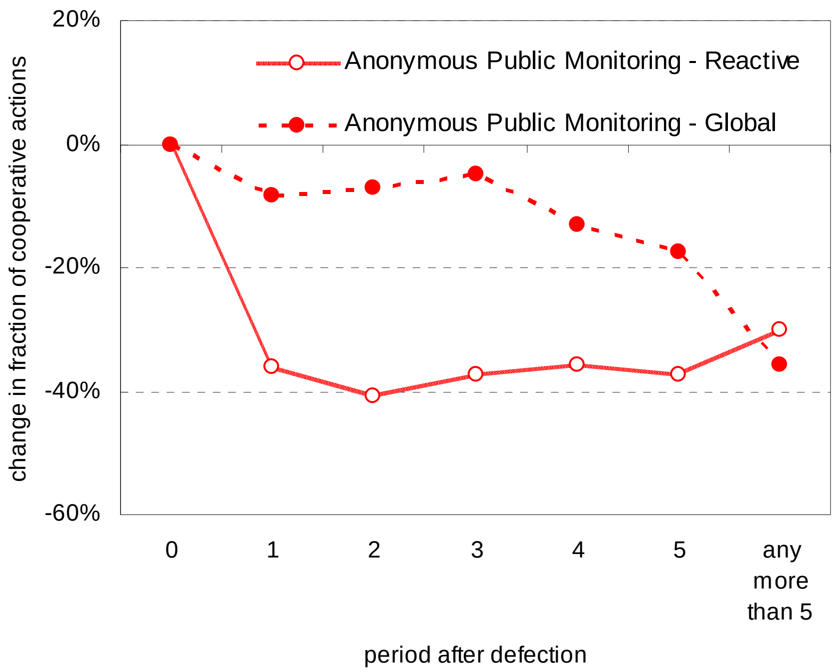

Figure 7 and Table 5 provide evidence for Result 9. Figure 7 is based on the marginal effects estimated using regressions in Table 5.[17] In addition to what has already been explained above in relation to Figure 5, the cooperation choices for anonymous public monitoring includes additional six strategy regressors to trace global strategies (Table 5). The representative subject that experienced a defection displayed a strong and persistent decrease in future cooperation levels (solid line in Figure 7). Conversely, when the representative subject observed a defection outside the match but did not experience it, the response was much weaker (dashed line in Figure 7).[18]

Table 7: Transitional matrices in private monitoring with punishment

(A) Choice after a subject cooperated and opponent defected

Did the subject *request* personal punishment?	Subject choice in the following period: Cooperate	Defect
Yes	75.5%	24.5%
No	46.7%	53.3%

(B) Choice after a subject defected and opponent cooperated

Did the subject *receive* personal punishment?	Subject choice in the following period: Cooperate	Defect
Yes	34.5%	65.5%
No	24.1%	75.9%

Result 10 (strategy employed). *Subjects preferred targeted strategies over reactive and global strategies. In the non-anonymous public monitoring treatment, a defection by an opponent generated a strong response in future encounters with the same opponent, while defections outside the match were ignored.*

In non-anonymous public monitoring subjects observed all individual histories. Hence targeted, reactive and global strategies were all available. Recall that a subject using a targeted strategy punished only opponents who had defected in previous encounters but cooperated with everyone else, even if they had defected in the past with someone else.

[17] The graph uses the coefficient estimates coding reactive and global strategies, respectively. See footnote for Figure 5. Marginal effects for the reactive strategies were computed for the average values of global strategies regressors. Marginal effects for the global strategies were computed for the average values of reactive strategies regressors.
[18] The two lines in Figure 7 overlap for periods "any more than 5" because of how reactive and global strategy regressors are defined (see Figure 1).

Table 5 and Figure 8 provide evidence for Result 10. Figure 8 reports the marginal effects estimated using regressions in Table 5.[19] In addition to what has already been referenced in relation to Figure 7, the cooperation choices for non-anonymous public monitoring includes six additional strategy regressors to trace targeted strategies (Table 5). The representative subject that experienced a defection displayed a strong and persistent decrease in cooperation levels when future encounters involved the same opponent (dark solid line in Figure 8). In contrast, there is little support for the use of either reactive or global strategies (light solid and dashed lines in Figure 8).

Figure 8: Strategies of the representative subject in non-anonymous public monitoring

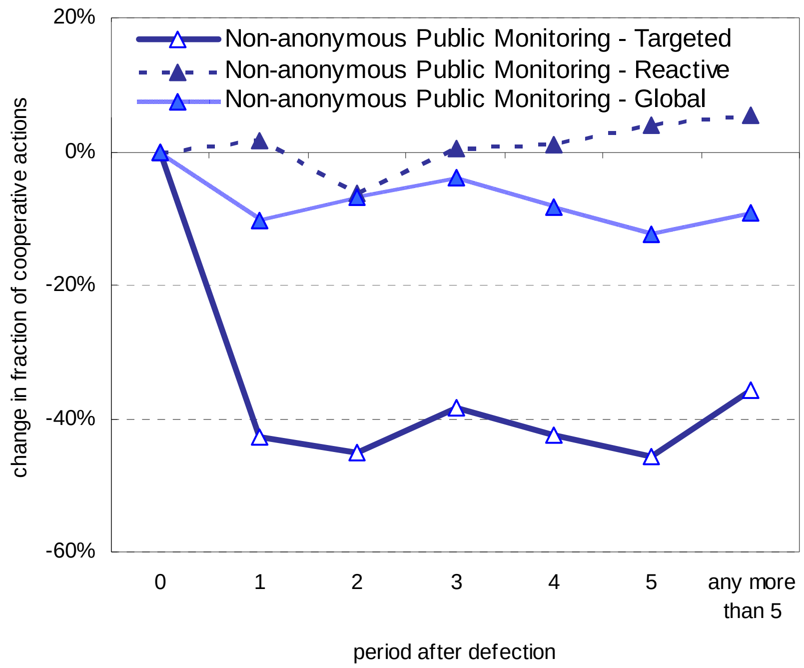

[19] The graph uses the coefficient estimates coding targeted, reactive and global strategies, respectively. See notes on Figure 5. Marginal effects for the targeted strategies were computed for the average values of reactive and global strategies regressors. Marginal effects for the reactive strategies were computed for the average values of targeted and global strategies regressors. Marginal effects for the global strategies were computed for the average values of targeted and reactive strategies regressors.

Result 11 (strategy employed). *Subjects employed grim trigger strategies and did not revert to a cooperative mode. In all treatments a defection of an opponent triggered a persistent decrease in cooperation. In particular, following a defection, economies in public monitoring treatments did not appear to revert to a cooperative mode.*

While in private monitoring treatments, cooperation could be supported only through grim trigger strategies; in public monitoring treatments cooperation could also be supported through T-period trigger strategies.[20] Regression results from Table 5 allow us to detect if such type of strategies were actually employed. In that circumstance one should see after T period from a defection a full "recovery" to pre-defection cooperation levels. However, no such recovery can be detected from Figures 7 and 8. This is consistent with previous findings in different settings (Mason and Phillips, 2002). After an initial drop, one period after the defection, one should observe an upward trend in the marginal effect curves of Figure 7 and 8. Instead, the curves look generally flat.

6 Final Remarks

We studied long-run equilibria in experimental economies composed by *strangers* who play indefinitely a prisoners' dilemma in pairs. Subjects are randomly matched and cannot directly communicate, and their identities and histories are private information. Achieving cooperation in this setting is difficult because subjects can neither commit to cooperation nor enforce it, especially because opponents vary randomly over time. Contrary to our expectations, we found that subjects did overcome these hurdles and cooperated at high and increasing rates (private monitoring treatment). This result provides empirical support to the well-known theoretical results of Kandori (1992) and Ellison (1994), who specify conditions under which cooperation is an equilibrium of infinitely repeated games among strangers. Our empirical finding is a novel contribution given the weak evidence provided on this point by previous experimental studies (Schwartz et al., 1999; Duffy and Ochs, 2006).

[20] At the end of each period, everyone observes the random draw. That number can be used as a coordination device. In particular, even in private monitoring subjects could coordinate a reversion to cooperation using that publicly observed number.

We then built on this initial finding by studying if and how the introduction of some prototypical institutions, capable of reducing either informational or enforcement frictions, would impact the emergence of cooperation (private monitoring with punishment, anonymous public monitoring, non-anonymous public monitoring treatments). According to theory, none of these institutions alters the lower or upper bound of cooperation possible in equilibrium. Yet, they had a remarkable impact on cooperation levels observed in the experiment.

In some treatments we increased the available information by displaying the histories of actions of everyone in the economy (public monitoring). Such information sometimes had no effect on aggregate cooperation levels and sometimes had startling effects. It turns out that unless histories could be traced back to a specific individual, this additional information was not used. In the anonymous public monitoring treatment, subjects received aggregate information about histories in the economy, but failed to exploit the information to increase cooperation above the private monitoring treatment. Instead, when details about identities were added to this aggregate information (non-anonymous public monitoring), cooperation was considerably higher. Second, in some treatments subjects had the costly option to lower the opponent's payoff. In this personal punishment treatment cooperation levels increased so dramatically that they are statistically indistinguishable from the non-anonymous public monitoring treatment.

Another main contribution of the paper is to shed light on the classes of strategies employed by subjects who indefinitely play a prisoners' dilemma. The subjects' behavior in our experimental economies suggests a strong preference for strategies that are *selective* in punishment (i.e., a preference for narrowing down the sets of targets of punishment). Indeed, when strategies with different levels of selectivity were available, subjects invariably chose the one with the most selective punishment. For example, when subjects remained anonymous but could see all histories in the economy, the representative subject mostly defected only after having directly experienced a defection (reactive strategy). When subjects could also see individual identities, then the representative subject essentially targeted her punishment toward those who directly cheated her in previous encounters, but cooperated with everyone else. This is remarkable because the power of a targeted strategy (punish the culprit only) is lower than that of a global strategy (punish everyone as soon as one sees a defection); the latter strategy immediately triggers an economy-wide defection, and as a result incorporates a bigger threat.[21] In fact our data suggest that the threat of economy-wide defection has low credibility. For instance, when economy-wide defection was the *only* available threat to support a cooperative outcome (private monitoring treatment), we observed the lowest levels of cooperation in all treatments in period 1. This result indicates that subjects may doubt that a single defection will trigger an economy-wide punishment.

We derived some possible reasons for the frequent use of some classes of strategies. First, subjects may have other-regarding preferences,[22] in which case they would prefer punishment schemes that decreased the harm to cooperators while raising it for defectors. This attitude would suggest a strong preference for targeted strategies over reactive or global strategies, and therefore, a reluctance to engage in economy-wide defection. Second, subjects may prefer simpler strategies because of cognitive costs.[23] The results obtained provide mixed evidence on this point. A grim-trigger reactive strategy may be the simplest choice available because it requires knowledge of the outcome only in the current period and only in the subject's match. Other strategies may involve a higher cognitive cost because they require the monitoring of identities, as when strategies are targeted, or of outcomes in other matches. Another dimension of complexity could be time-dependence as in t-period punishment strategies, which are not observed. In public monitoring treatments t-period punishment strategies are feasible and deliver higher continuation payoff. Self-regarding agents, and even more so other-regarding agents, should prefer t-period punishment to grim trigger strategies. Yet, punishment following a

[21] If power is a criterion to select strategies, then in the anonymous public monitoring everyone should use a global strategy, which is not observed. In the non-anonymous public monitoring one should observe that a defector is punished by everyone in every future match, which is not observed.

[22] There is an experimental literature that validates this conjecture (e.g., see Fehr and Gaecther, 2002) and several models of other-regarding preferences exists that alternatively focus on: altruism,inequality aversion or reciprocity (see Sobel, 2005 for a review).

[23] The economies included just four subjects, and information was clearly displayed and easily accessible. So, one can hardly argue that monitoring identities and histories was a demanding task.

defection appears to have no reversal trend (i.e., we see little evidence of time-dependent strategies). Although this observation may suggest that simplicity plays a role in the selection of strategies, we also observe the use of more complex strategies that involve several contingencies, such as targeted strategies.

The widespread use of personal punishment also deserves some discussion. Through personal punishment, a subject can directly and immediately lower the earnings of her opponent, which is not a best response for a self-regarding, rational agent (proposition 3). In the experiment, however, availability of personal punishment remarkably increased aggregate cooperation from the very first period. One can think of several reasons for the use of personal punishment. One is reciprocity because a subject may be happy to pay a cost to lower her opponent's earnings in order to reciprocate for her defection. In this manner she avoids harming cooperators through punishing only those who have been unkind. Under private monitoring, a reciprocator had no other equilibrium strategy with comparable selectivity in punishing defectors.[24] Another reason is simplicity because personal punishment neither requires knowledge of others' strategies nor coordination on some informal punishment scheme. Moreover, personal punishment is unavoidable. When using a reactive strategy, instead, punishing by defecting is uncertain because the interaction could suddenly end. A final reason for using personal punishment involves using a channel of costly communication, which may have helped in coordinating (e.g., see Cooper et al., 1996, Crawford, 1998, Van Huyck et. al, 2002).

A tentative conclusion is thus that cognitive costs may play a minor role in driving strategy choice, while other-regarding preferences may be more relevant. We plan to tackle this issue in future work.

[24] If the subject uses a reactive strategy, she will punish the defector in future periods. Moreover, others in the economy will eventually punish that defector.

Appendix A

Proof of Lemma 1

In this section we develop the proof of the Proposition. We start by discussing payoffs, given a deviation. A player deviates from desired play in two instances: In- or off-equilibrium, if she has not observed a deviation in the past but chooses *Z*, currently. Off-equilibrium, if he has observed a deviation in the past but plays Y, currently. Since the environment is stationary, by the unimprovability criterion we restrict attention to one-time deviations. We also consider only single-player deviations. While this simplifies the analysis, deriving off-equilibrium payoffs still requires a bit of work (which is why we include the proof in the appendix). The problem is that players observe only the actions in their pair; in order to calculate expected values, we must know how uncooperative behavior spreads to the economy after a defection is observed.

A1 The diffusion of sanctions in the economy

Consider a representative period *t* and recall that there are 3 possible ways to pair four players. Thus, if $d=1,\cdots,4$ is the number of players who choose Z currently, then $d'=d,\cdots,4$ is the number of deviators tomorrow, which depends on the current realization of the random pair.[25] As noted above, the central concern of a player is the likelihood that her/his opponent does not cooperate. Thus, we report the probabilities ρ_d that a player who selects *Z* today will meet a cooperator today, given that d players choose Z today. We also calculate $\Pr[d'|d]$ (i.e., the probability that tomorrow there are d' individuals who play Z, given that today there are d).

The first set of probabilities is needed to determine the expected current utility to a player who is aware of a deviation or that deviates, selecting *Z*. The second set of probabilities is needed to calculate the continuation payoff for a player who is aware of a deviation or that deviates, selecting *Z*. Indeed, they will give us *transition matrices*, allowing us to calculate the various probabilities that the sanction spreads to the rest of the economy. Notice that there will be two contingencies to consider. In one case we calculate

[25] Clearly, if $D_t=0$, then $D_{t+1}=0$ with certainty.

probabilities under the conjecture that every player follows the sanctioning behavior specified by the social norm; the other is derived under the conjecture that one player deviates from the sanctioning behavior, once.

Case 1. Off-equilibrium, everyone sanctions. Consider a player who currently selects *Z*. Let $\rho=(\rho_1,\rho_2,\rho_3,\rho_4)$ with ρ_d being the probability that he meets a cooperator given that $d=1,\cdots,4$ players currently select Z. Clearly, the probability that he meets someone who chooses *Z* is 1 − ρ. Recall that each player can be paired to three other players, with equal probability. Therefore, we have

$$\rho = (1, \frac{2}{3}, \frac{1}{3}, 0)$$

Here, the transition matrix is:

$$\mathbf{A} = \begin{bmatrix} & \mathbf{1} & \mathbf{2} & \mathbf{3} & \mathbf{4} \\ \mathbf{1} & 0 & 1 & 0 & 0 \\ \mathbf{2} & 0 & 1/3 & 0 & 2/3 \\ \mathbf{3} & 0 & 0 & 0 & 1 \\ \mathbf{4} & 0 & 0 & 0 & 1 \end{bmatrix}$$

The bold numbers in the rows (columns) indicate the number D_t (D_{t+1}) of players who currently (next period) play *Z*. Each cell represents the corresponding conditional probability $\Pr[D_{t+1} | D_t]$. Clearly, Pr [2|1] = Pr [4|3] = 1 since if an odd number of players plays *Z* today, then at least one of them is paired to a cooperator. The latter will choose *Z* in $t+1$. Also, Pr [4|4] = 1, since the social norm does not specify reversion to cooperation. To see why Pr [2|2] = 1/3 and Pr [4|2] = 2/3 recall that there are three possible pairings. One of those involves the two players who currently choose *Z*. So, with probability 1/3 the sanctioning behavior does not spread further. If that pairing is not realized, then *Z* will be necessarily seen by the remaining two cooperators. So, with probability 2/3, next period everyone will choose the sanction, Z.

Case 2. Off equilibrium, one player does not use the prescribed sanction.

Suppose, off-equilibrium, a player who observed a deviation in the past chooses to deviate from the sanctioning rule and plays *Y* this period. Instead, everybody else follows the social norm. Consider this player. Again, ρ_d is the probability that he meets a cooperator given that $d = 1, \cdots, 4$ players have observed Z in the past. The probability that this player meets someone who chooses Z is $1-\rho$.

However, since this player chooses *Y* instead of *Z*, the transition matrix is now different:

$$\tilde{A} = \begin{bmatrix} & \mathbf{1} & \mathbf{2} & \mathbf{3} & \mathbf{4} \\ \mathbf{1} & 0 & 1 & 0 & 0 \\ \mathbf{2} & 0 & 1/3 & 2/3 & 0 \\ \mathbf{3} & 0 & 0 & 1/3 & 2/3 \\ \mathbf{4} & 0 & 0 & 0 & 1 \end{bmatrix}$$

Again, Pr [2|1] = 1, since this is the case when no one observed a deviation in the past but someone chooses to play *Z* today. Also, Pr [4|4] = 1 since the player who deviates today by choosing cooperation will revert back to playing *Z*, in the next period (onetime deviation). Now, consider the second row, i.e., the case when two players observed a deviation, but only one of them plays *Z* today (reverting to playing *Z*, tomorrow). Here, only one of the three possible pairings includes both players who observed a deviation. In this case the sanctioning behavior does not spread further. In the other two pairings, it spreads only to one more player, since only one player plays *Z* today. Hence, we have Pr [2|2] = 1/3 and Pr [3|2] = 2/3. The third line is similarly explained.

A2 Off-equilibrium payoffs

Using the matrices above, we can now construct off-equilibrium payoffs in two contingencies. The first is in equilibrium, when the player deviates for the first time, choosing *Z*. The second is off equilibrium, when the player has observed uncooperative behavior in the previous date (i.e., has seen *Z* for the first time) and now deviates by cooperating, choosing *Y*.

Payoff from a deviation, when everyone follows the sanctioning rule. Suppose that every player follows the social norm. Define the column vector:

$$\boldsymbol{V} = \begin{bmatrix} V_1 \\ V_2 \\ V_3 \\ V_4 \end{bmatrix}$$

where V_d denotes the expected lifetime utility at the start of a period, before pairing takes place, to a player who selects Z currently, given that $d = 1, \cdots, 4$ players choose Z currently.[26] Using the vector of probabilities ρ and the transition matrix $\boldsymbol{A}$, where we denote $\boldsymbol{A}_d$ the d^{th} row of matrix $\boldsymbol{A}$, we have:

$$V_d = z + \rho_d (h - z) + \delta \boldsymbol{A}_d \boldsymbol{V}$$

That is, the expected current utility depends on the probability of encountering a cooperator. When he meets a cooperator the player earns h , and otherwise he earns z . The continuation payoff is 0 with probability $1-\delta$, and it is $\boldsymbol{A}_d \boldsymbol{V}$ with probability δ. The latter component tells us that current play may lead to different numbers of cooperators tomorrow, depending on the outcome of the pairing process. Specifically, we have:

$$\begin{aligned} V_1 &= h + \delta V_2 \\ V_2 &= z + \frac{2}{3}(h - z) + \delta \left(\frac{1}{3} V_2 + \frac{2}{3} V_4 \right) \\ V_3 &= z + \frac{1}{3}(h - z) + \delta V_4 \\ V_4 &= z + \delta V_4 \end{aligned} \qquad (1)$$

To see how we derive them, we discuss the first two lines. Consider the first line. If a player is the initial deviator, then he is certainly paired to a cooperator, i.e., $\rho_1 = 1$ and earns current payoff h. The current cooperator will choose *Z* in the future. Thus, the current deviator's continuation payoff is $\delta \mathrm{V}_2$. Consider the second line. Since the player chooses Z currently, he earns z if he meets the other only player who chooses *Z* (with probability $\frac{1}{3}$) and earns h if he meets a cooperator (with probability $\frac{2}{3}$). This gives expected current utility $z + \frac{2}{3}(h - z)$. The continuation payoff depends on which one of these pairs took place. If he met the other deviator, no cooperator observes *Z* today, so tomorrow there will still be two players who select Z. Otherwise, both cooperators

[26] Clearly, the agent selects Z as a deviation from equilibrium when d = 1. In this case the agent is the initial deviator. If d = 2, instead, the agent may select Z simply because he observed Z in the past and now follows the sanctioning rule.

observe Z today and tomorrow select Z as well. Simple manipulations of (1) give:

$$
\begin{aligned}
V_1 &= \frac{h(3+\delta)}{3-\delta} + \frac{z\delta(1+\delta)}{(3-\delta)(1+\delta)} \\
V_2 &= \frac{2h}{3-\delta} + \frac{z(1+\delta)}{(3-\delta)(1-\delta)} \\
V_3 &= z + \frac{1}{3}(h-z) + \delta\frac{z}{1-\delta} \\
V_4 &= \frac{z}{1-\delta}
\end{aligned}
\qquad (2)
$$

Payoff from a deviation, when a player does not follow the sanctioning rule. Suppose d players have observed a deviation in the past, and everyone follows the social norm except one of these players. This player defects from the sanctioning rule and cooperates. Let $\tilde{V}_d$ denote the expected lifetime utility at the start of a period, before pairing takes place, to the player that has observed a deviation in the past but selects Y currently, given d. Using the vector of probabilities ρ and the transition matrix $\tilde{\boldsymbol{A}}$, where we denote $\tilde{\boldsymbol{A}}_d$ its d^{th} row, we have:

$$
\tilde{V}_d = \begin{cases} V_1 & if \;\; d = 1 \\ l + \rho_d (y-l) + \delta \tilde{\boldsymbol{A}}_d \boldsymbol{V} & if \;\; d = 2 \end{cases}
$$

When d = 1, this means that no deviation was observed previously but someone chooses to deviate today. Therefore $\tilde{V}_1 = V_1$, since it is the first period in which a deviation is observed. For the case $d \geq 2$ notice that only $d-1$ players choose Z currently, the remaining one choosing Y. Therefore we must use the matrix $\tilde{\boldsymbol{A}}$.

In that case, we see that $l + \rho_d (y-l)$ is the expected current utility from meeting either a cooperator or not. Since the player cooperates, when he meets a cooperator, he earns y, and otherwise he earns l. Again, the continuation payoff is 0 with probability $1-\delta$, and it is $\tilde{\boldsymbol{A}}_d \boldsymbol{V}$ with probability δ. We use $\boldsymbol{V}$ and not $\tilde{\boldsymbol{V}}$ in the continuation payoff, since everyone reverts to the sanctioning rule specified by the social norm, from tomorrow on.

As done for the case above, simple calculations generate:

$$\tilde{V}_2 = l + \frac{2}{3}(y-l) + \delta(\frac{1}{3}V_2 + \frac{2}{3}V_3)$$
$$\tilde{V}_3 = l + \frac{1}{3}(y-l) + \delta(\frac{1}{3}V_3 + \frac{2}{3}V_4) \quad (3)$$
$$\tilde{V}_4 = l + \delta V_4$$

A3 Requirements for individual optimality

In this section we check that the actions recommended by the social norm are best responses after any history of play. To do so we consider two issues. First, we derive a condition ensuring that choosing *Z* is not a best response on the equilibrium path. Second, we check that playing *Y* instead of *Z*, after having observed a deviation, is never optimal.

Suboptimality of a deviation, in equilibrium. We must check that deviating by choosing *Z* is suboptimal, relative to cooperating. That is:

$$\frac{y}{1-\delta} \geq V_1 = \frac{h(3+\delta)}{3-\delta} + \frac{z\delta(1+\delta)}{(3-\delta)(1-\delta)}$$

an inequality that is rearranged as:

$$\delta^2(h-z) + \delta(2h-y-z) - 3(h-y) \geq 0$$

Let $f(\delta)$ define the expression on the RHS of the inequality. Notice that since $h > y > z$ and $\delta \in (0,1)$, then $f(\delta) \geq 0$ for all $\delta \geq \delta^*$ where $\delta^* \in (0,1)$ is the unique value of δ that solves $f(\delta) = 0$. We have $\delta^* > 0$ since $f(0) < 0$ and $f'(\delta) > 0$. Also, $\delta^* < 1$ since $f'(\delta) > 0$ for $\delta > 0$ and $f(1) = 2(y-z) > 0$. The parameterization of our experiment implies $\delta^* = 0.443$.

Suboptimality of a deviation, off-equilibrium. Here, we check that if a player has observed *Z* in the past, then *Y* today is suboptimal. That is, since we have shown that choosing *Z* is never optimal, when $d = 1$, we must find conditions such that $V_d \geq \tilde{V}_d$ for all $d \geq 2$. To do so, use (1), (2) and (3). Clearly, $V_4 \geq \tilde{V}_4$ since $z \geq l$. Now consider the inequality $V_3 \geq \tilde{V}_3$. Rearranging:

$$\tilde{V}_3 = l + \frac{1}{3}(y-l) + \delta\frac{1}{3}(V_3 - V_4) + \delta V_4,$$

we have $V_3 \geq \tilde{V}_3$, if:

$$z + \frac{1}{3}(h - z) - [l + \frac{1}{3}(y - l)] \geq \delta \frac{1}{3}(V_3 - V_4)$$
$$\Rightarrow \delta \leq 3 \times \frac{2(z - l) + (h - y)}{h - z}$$

since from (1) we have $V_3 - V_4 = \frac{1}{3}(h - z)$. It is easy to see that for our parameterization this is satisfied by all $\delta \in (0,1)$ since $\frac{2(z - l) + (h - y)}{h - z} = 34$.

Finally, consider $V_2 \geq \tilde{V}_2$. This inequality is immediately rewritten as:

$$z + \frac{2}{3}(h - z) - [l + \frac{2}{3}(y - l)] \geq \delta \frac{2}{3}(V_3 - V_4)$$
$$\Rightarrow \delta \leq \frac{3}{2} \times \frac{z - l + 2(h - y)}{h - z}$$

which always holds for all parameters since $\frac{z - l + 2(h - y)}{h - z} = 34$.

The intuition is simple. Cooperating instead of sanctioning after observing a defection may be helpful to the player, since it delays the spread of the sanction. However, doing so generates a current loss to the player since he earns y (instead of h) if he meets a cooperator, and l (instead of z) if he meets a deviator. Therefore, the player must be sufficiently impatient to prefer play of Z to Y—clearly, the smaller l and y, the greater the incentive to follow with the sanction. Our parameterization insures that this incentive exists for all $\delta \in (0,1)$.

Appendix B

Table 5: Probit regression on individual choice to cooperate – marginal effects(*)

Dependent variable: 1=cooperation 0=defection	Private Monitoring	Anonymous Public Monitoring	Private Monitoring With punishment	Public Monitoring (non-anonymous)	All treatments	All treatments, first periods only
Treatment dummies:						
Anonymous Public Monitoring					-0.04566*	-0.02924
					(0.02450)	(0.07313)
Private Monitoring With punishment					0.99813***	0.09188
					(0.00010)	(0.06737)
Public Monitoring (non-anonymous)					0.94668***	0.11747*
					(0.00951)	(0.06082)
Cycle dummies: Cycle 2	0.03938	0.05750	0.08336***	-0.00304*	0.06190***	-0.03677
	(0.10412)	(0.03787)	(0.02618)	(0.00179)	(0.02336)	(0.02852)
Cycle 3	0.07621	0.05045	0.11154***	0.02014***	0.09321***	0.00613
	(0.06949)	(0.05102)	(0.01972)	(0.00188)	(0.02688)	(0.02911)
Cycle 4	0.13560***	0.18856***	0.14888***	0.12585***	0.17362***	0.04925
	(0.00770)	(0.02515)	(0.03042)	(0.02706)	(0.02205)	(0.03559)
Cycle 5	-0.16040***	0.28969***	0.13911***	0.13946***	0.21443***	0.08162***
	(0.04322)	(0.03196)	(0.03295)	(0.00420)	(0.02114)	(0.03126)
Duration of previous cycle	0.00133*	0.00272***	0.00249***	0.00413***	0.00362***	0.00327***
	(0.00076)	(0.00015)	(0.00033)	(0.00058)	(0.00082)	(0.00073)
Reactive strategies:						
Grim trigger	-0.55054***	-0.26642***	-0.38185***	0.07527	-0.38854***	
	(0.01458)	(0.07434)	(0.10006)	(0.05562)	(0.04064)	
Tit-for-tat with lag 1	0.08814**	-0.04836**	0.05643*	-0.06063	0.01797	
	(0.04303)	(0.02415)	(0.02986)	(0.03933)	(0.02747)	
Tit-for-tat with lag 2	0.11608***	-0.09505***	0.04598*	-0.14011***	-0.02680	
	(0.03620)	(0.01776)	(0.02692)	(0.03122)	(0.03870)	
Tit-for-tat with lag 3	0.10324**	-0.07310*	0.04044	-0.06265***	-0.00994	
	(0.04237)	(0.04174)	(0.03453)	(0.00683)	(0.02686)	
Tit-for-tat with lag 4	0.08009***	-0.05838	0.01518	-0.05327	-0.03325	
	(0.00524)	(0.04672)	(0.04463)	(0.06054)	(0.02879)	
Tit-for-tat with lag 5	0.02974**	-0.07094***	0.01432	-0.01770	-0.04448*	
	(0.01385)	(0.00729)	(0.03021)	(0.04159)	(0.02311)	

(*) Marginal effects are computed at the mean value of regressors. Robust standard errors for the marginal effects are in parentheses computed with a cluster on each session; * significant at 10%; ** significant at 5%; *** significant at 1%. For a continuous variable the marginal effect measures the change in the likelihood to cooperate for an infinitesimal change of the independent variable. For a dummy variable the marginal effect measures the change in the likelihood to cooperate for a discrete change of the dummy variable.
First periods of each cycle are excluded with the exception of the last column. Individual fixed effects and period fixed effects are included but not reported in the table (individual dummies: s2-s30 s32-s37 s39 s41-s60 s62-s97 s99-s159; period dummies: 3, 4, 5, 6-10, 11-20, 21-30, >30). A *grim trigger regressor* has value 1 in *all* periods following a defection and 0 otherwise. Five *tit-for-tat regressors* have value of 1 only in *one* period following a defection and 0 otherwise; we trace response up to a five-period delay. Details on strategy coding are in the text after Results 7. Duration of previous cycle was set to 20 for cycle 1.

Dependent variable: 1=cooperation, 0=defection	Private Monitoring	Anonymous Public Monitoring	Private Monitoring With punishment	Public Monitoring (non-anonymous)	All treatments	All treatments, first periods only
Global strategies:						
Grim trigger		-0.31089** (0.13126)		-0.11638*** (0.00182)		
Tit-for-tat with lag 1		0.22700*** (0.01652)		0.02296 (0.05930)		
Tit-for-tat with lag 2		0.22948*** (0.06277)		0.02812 (0.04358)		
Tit-for-tat with lag 3		0.24272*** (0.01007)		0.04799** (0.02376)		
Tit-for-tat with lag 4		0.17551*** (0.03108)		0.00537 (0.02106)		
Tit-for-tat with lag 5		0.15494*** (0.01241)		-0.03226 (0.05412)		
Targeted strategies:						
Grim trigger				-0.36340*** (0.04707)		
Tit-for-tat with lag 1				-0.04395*** (0.00521)		
Tit-for-tat with lag 2				-0.05715*** (0.01385)		
Tit-for-tat with lag 3				-0.01781 (0.03301)		
Tit-for-tat with lag 4				-0.04346*** (0.00326)		
Tit-for-tat with lag 5				-0.06268*** (0.01560)		
Personal punishment						
Requested (lag)			-0.07641 (0.08505)			
Requested (lag) × opponent defected (lag)			0.02845 (0.02884)			
Received (lag)			0.06715* (0.03858)			
Received (lag) × subject defected (lag)			-0.32918*** (0.09712)			
Observations	3320	4880	4400	4280	16680	800

References

Aliprantis, C.D., Camera, G., and Puzzello, D. (2006). Matching and anonymity. *Economic Theory* 29, 415-432

Aliprantis, C.D., Camera, G., and Puzzello, D. (2007). A random matching theory. *Games and Economic Behavior.*

Al-Najar, N., and R. Smorodinsky, R. (2001). Large nonanonymous repeated games. *Games and Economic Behavior, 37,* 26-39.

Andreoni, J., and R. Croson, R. (2002). Partners versus strangers: The effect of random rematching in public goods experiments. To appear in C. Plott, and V. Smith, (Eds.), *Handbook of Results in Experimental Economics Results.*

Aoyagi, M., and G. Frechette, G. (2003). Collusion in repeated games with imperfect monitoring. unpublished working paper, Osaka University and Harvard University.

Axelrod, R. (1984). *The evolution of cooperation*. New York: Basic Books.

Camerer, C., and K. Weigelt, K. (1988). Experimental tests of the sequential equilibrium reputation model. *Econometrica, 56,* 1-36.

Casari, M. and Plott, C. 2003. Decentralized management of common property resources: experiments with a centuries-old institution. *Journal of Economic Behavior and Organization* 51, 2, 217-247

Cason, T. N., and F. U. Khan, F.U. (1999). A laboratory study of voluntary public goods provision with imperfect monitoring and communication. *Journal of Development Economics, 58,* 533-552.

Cooper, R., D. DeJong, D. and R. Forsythe, R. (1996). Cooperation Without Reputation: Experimental Evidence from Prisoner's Dilemma Games. *Games and Economic Behavior, 12*, 187-218.

Costa- Gomes, M., Crawford, V., and Broseta B. (2001). Cognition and behavior in normal-form games: An experimental study. *Econometrica 69*, 1193-1235

Crawford, V. (1998). A survey of experiments on communication via cheap talk. *Journal of Economic Theory 78,* 286-298.

Dal Bó, P. (2005). Cooperation under the shadow of the future: Experimental evidence from infinitely repeated games. *American Economic Review 95* (5), 1591-1604.

Dal Bó, P. and Frechette, G. (2006). The evolution of cooperation in infinitely repeated games: Experimental evidence. Unpublished manuscript. Brown University and New York University.

Diamond, P. (1982). Aggregate Demand Management in Search Equilibrium. *Journal of Political Economy* 90, 881-894.

Dixit, A. (2003). On modes of economic governance. *Econometrica* 71(2), 449-481.

Duffy, J., and Ochs, J. (2006). Cooperative behavior and the frequency of social interaction. Unpublished working paper. University of Pittsburgh.

Ellison, G. (1994). Cooperation in the prisoner's dilemma with anonymous random matching.Prisoners' Dilemma with Anonymous Random Matching. *Review of Economic Studies 61*, 567-88.

Engle-Warnick, J. (2007). Five indefinitely repeated games in the laboratory. CIRANO, Working Paper 2007s-11.

Engle-Warnick, J., and Slonim R. L. (2004). The evolution of strategies in a repeated trust game. *Journal of Economic Behavior and Organization 55* (4): 553-573.

Engle-Warnick, J., and R. L. Slonim R.L. (2006). Learning to trust in indefinitely repeated games. *Games and Economic Behavior 54* (1): 95-114.

Fehr, E. and S. Gaechter, S. (2002). Altruistic Punishment in Humans, *Nature* 415, 137-140.

Friedman, J. W. (1971). Non-cooperative equilibrium for supergames. *Review of Economic Studies* 38, 113, 1-12.

Fudenberg, D., Levine, D.K., and Pesendorfer, W. (1998). When are nonanonymous players negligible? *Journal of Economic Theory 79*, 46-71.

Healy, P.J. Group reputations, stereotypes, and cooperation in a repeated labor market, *American Economic Review* forthcoming

Kandori, M. (1992). Social norms and community enforcement. *Review of Economic Studies 59*, 63-80.

Krasa, S., and A. Villamil (2000). Optimal contracts when enforcement is a decision variable. *Econometrica* 64(1), 119-138.

Levine, D. K., and Pesendorfer, W. (1995). When are agents negligible? *American Economic Review 85*, 1160-1170.

Mason, C. F., and Phillips, O.R. (2002). In support of trigger strategies: Experimental evidence from two-person non-cooperative games. *Journal of Economics and Management Strategy 11* (4), 685-716.

Murnighan, J. K., and Roth, A.E. (1983). Expecting continued play in prisoners-dilemma games – A test of several models. *Journal of Conflict Resolution 27* (2): 279-300.

Normann, H.-T. and Wallace, B. (2006). The Impact of the Termination Rule on Cooperation in a Prisoner's Dilemma Experiment. working paper, Royal Holloway.

Ostrom, E., Walker, J., and Gardner, R., (1992). Covenants with and without a sword: self governance is possible. *American Political Science Review* 86, 404-417.

Palfrey, T.R., and Rosenthal, H. (1994). Repeated play, cooperation and coordination: An experimental study. *Review of Economic Studies 61*, 545-565.

Rapoport, A., and Chammah (1965). *Prisoner's dilemma*: *A study in conflict and cooperation*. Ann Arbor: The University of Michigan Press.

Roth, A.E., and Murnighan, K. (1978). Equilibrium behavior and repeated play of the prisoner's dilemma. *Journal of Mathematical Psychology 17*, 189-98.

Schwartz, S., Young, R. and Zvinakis, K. (2000). Reputation without repeated interaction: A role for public disclosures. *Review of Accounting Studies 5*, 351-375.

Shimer, R. (2005). The Assignment of Workers to Jobs in an Economy with Coordination Frictions. *Journal of Political Economy* 113(5): 996-1025.

Sobel, J. (2005). Interdependent Preferences and Reciprocity, *Journal of Economic Literature* 43, 392-436.

Van Huyck, J. B., Wildenthal, J.M., and Battalio, R.C. (2002). Tacit cooperation, strategic uncertainty, and coordination failure: Evidence from repeated dominance solvable games. *Games and Economic Behavior 38*, 156-175.

(NOT FOR PUBLICATION)

Table 1A: Example of strategy coding

Period	ID	choice	**opponent choice (*)**	Reactive strategies are based on (*)			**opponent and other pair choices (**)**	Global strategies are based on (**)			**opponent's ID (***)**	Targeted strategies are based on (*) and on (***)		
				Grim trig.	TfT lag 1	TfT lag 2		Grim trig.	TfT lag 1	TfT lag 2		Grim trig.	TfT lag 1	TfT lag 2
1	7	0	**1**	0	0	0	**1**	0	0	0	**16**	0	0	0
2	7	1	**0**	0	0	0	**0**	0	0	0	**18**	0	0	0
3	7	1	**1**	1	1	0	**1**	1	1	0	**9**	0	0	0
4	7	1	**0**	1	0	1	**0**	1	0	1	**16**	0	0	0
5	7	0	**1**	1	1	0	**1**	1	1	0	**16**	1	1	0
6	7	1	**1**	1	0	1	**1**	1	0	1	**18**	1	1	0
7	7	1	**1**	1	0	0	**0**	1	0	0	**18**	1	0	1
8	7	1	**1**	1	0	0	**0**	1	1	0	**9**	0	0	0
9	7	1	**1**	1	0	0	**1**	1	0	1	**9**	0	0	0
10	7	1	**1**	1	0	0	**0**	1	0	0	**18**	1	0	0
11	7	1	**1**	1	0	0	**1**	1	1	0	**9**	0	0	0
12	7	1	**1**	1	0	0	**1**	1	0	1	**18**	1	0	0
13	7	1	**0**	1	0	0	**0**	1	0	0	**16**	1	0	1
14	7	1	**1**	1	1	0	**0**	1	1	0	**18**	1	0	0
15	7	1	**1**	1	0	1	**0**	1	0	1	**9**	0	0	0
16	7	0	**0**	1	0	0	**0**	1	0	0	**16**	1	1	0
17	7	0	**0**	1	1	0	**0**	1	0	0	**16**	1	0	1
18	7	0	**1**	1	0	1	**1**	1	0	0	**16**	1	0	0
19	7	1	**1**	1	0	0	**1**	1	0	0	**9**	0	0	0
20	7	1	**1**	1	0	0	**1**	1	0	0	**18**	1	0	0
21	7	1	**1**	1	0	0	**1**	1	0	0	**18**	1	0	0
22	7	1	**0**	1	0	0	**0**	1	0	0	**16**	1	0	0
23	7	0	**0**	1	1	0	**0**	1	1	0	**16**	1	1	0
24	7	1	**1**	1	0	1	**0**	1	0	1	**18**	1	0	0
25	7	0	**1**	1	0	0	**0**	1	0	0	**9**	0	0	0

Notes: (*) 1=cooperation, 0=defection, (**) 1=3 persons cooperated, 0=less than 3 persons cooperated; TfT=tit-for-tat regressor. Experimental data from session 8, cycle 1, periods 1-25 (non-anonymous public monitoring treatment).