

Cooperation among strangers under the shadow of the future

Gabriele Camera and Marco Casari *

Abstract

We study the emergence of norms of cooperation in experimental economies populated by strangers interacting indefinitely. Can these economies achieve full efficiency even without formal enforcement institutions? Which institutions for monitoring and enforcement facilitate cooperation? Finally, what classes of strategies do subjects employ? 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 mostly employ strategies that are selective in punishment.

Keywords: experiments, repeated games, cooperation, equilibrium selection, prisoners' dilemma, random matching. *JEL codes:* C90, C70, D80

(*) Corresponding author: G.Camera, University of Iowa, Department of Economics, W 210 PBB, 21 E. Market Street, Iowa City, IA 52242-1994, Phone: 319-335 3125, Fax: 319-335 1956, email: gabriele-camera@uiowa.edu; M.Casari, Purdue University and University of Bologna, Piazza Scaravilli 2, 40126 Bologna, Italy, Phone: ++39-051-209 8662, Fax: ++39-051-209 8493, marco.casari@unibo.it.

Financial support for running the experiments was provided by Purdue's CIBER. G. Camera acknowledges research support from NSF grant DMS-0437210. Jingjing Zhang provided valuable research assistance. We thank for comments three anonymous referees, Masaki Aoyagi, Roko Aliprantis, Michael Baye, John Duffy, Jason Abrevaya, Thomas Palfrey, seminar participants at the ASSET conference in Padova, ES 2007 meeting in Chicago, ESA 2007 meeting in Montreal, Indiana University, University of Bologna, University of Pittsburgh, Osaka University, City University of Hong Kong, and University of Torino.

Despite its relevance to macro and microeconomics, there are few experimental studies on strategic behavior in long-term relationships with uncertain endings (e.g., Tom R. Palfrey and Howard Rosenthal, 1994; Masaki Aoyagi and Guillaume Frechette, 2005; Pedro Dal Bó, 2005; John Duffy and Jack Ochs, 2006). This paper fills the gap by studying economies where subjects repeatedly interact in pairs formed at random, and the economy has an indefinite duration, i.e., infinitely repeated (matching) games. Such an approach is of general interest for two reasons.

First, the underlying theoretical platform is widely used in economics. Infinitely repeated games with random matching of anonymous agents have been employed in macroeconomics to model trading frictions (Peter Diamond, 1982), to analyze labor markets and equilibrium unemployment (Dale Mortensen and Christopher Pissarides, 1994), and in monetary economics to make explicit obstacles to credit arrangements (Nobuhiro Kiyotaki and Randall Wright, 1989). In microeconomics, similar models have been used to study the emergence of social norms in anonymous societies (Michihiro Kandori, 1992, Glenn Ellison, 1994), the organization of commerce (Paul R. Milgrom, Douglass C. North and Barry R. Weingast, 1990), and economic governance (Avinash Dixit, 2003). Empirical studies of infinitely repeated games focus overwhelmingly on interactions in stable pairs of partners (Aoyagi and Frechette, 2005; Dal Bó, 2005, Jim Engle-Warnick and Robert L. Slonim, 2004, 2006) and not on interactions among randomly matched, anonymous agents. Instead, we investigate which institutions are behaviorally associated to the emergence, sustainability, and breakdown of cooperation in anonymous economies.

Second, random matching models are richer than fixed matching model in terms of the set of strategies that can be adopted. In general, models based on indefinitely repeated games have multiple equilibria. Agents wanting to support a cooperative outcome face a double challenge:

not only must they be able to coordinate on an outcome, but must also coordinate on a credible threat that can support uninterrupted long-run cooperation. The above models assume, often implicitly, that self-interested agents will select the most efficient among the available equilibria. While convenient, equilibrium selection criteria based on efficiency have no solid foundations either in theoretical or in empirical arguments. To identify an equilibrium selection criterion it is crucial to understand what strategies agents adopt. While the vast majority of experimental studies on indefinitely repeated games concern fixed pairs of subjects, a random matching design allows for a richer set of strategies and it is better suited to isolate behavioral components in strategy selection. The data from our experimental economies demonstrate, for instance, that subjects have “preferences” over strategies, and this crucially influences the outcome selected. In designs that are identical except for the classes of strategies available, we find that outcomes can differ considerably in terms of cooperation level because of a reluctance to use some classes of strategies, despite their theoretical effectiveness.

In our experiment we simplify as much as possible the coordination task by designing economies of four agents. Each period they are randomly paired to play a prisoners’ dilemma. The economy has an indefinite duration, based on a probabilistic continuation rule. The theoretical foundation for this design can be traced back to the folk theorems for infinitely repeated games (supergames) of James W. Friedman (1971), and the subsequent random-matching extensions in Kandori (1992) and Ellison (1994). The basic theoretical result is that cooperation is an equilibrium when agents are involved in a long-term interaction, are sufficiently patient, and have sufficient information on the actions of others. This result is very powerful and it extends even to anonymous economies where action histories are private information. Parameters in the experiment are set to ensure that the efficient outcome can be

sustained as one of the possible equilibria, when agents adopt the following simple social norm. Every agent cooperates unless someone has been caught defecting, in which case those who see the defection should forever defect (“grim trigger” strategy). In practice, however, achieving the efficient outcome may be problematic because subjects in the experiment are not in a stable partnership, cannot communicate their intentions to others, and can neither commit to nor enforce cooperation. We study the effect of various levels of information about action histories and the punishment technologies that are made available to subjects.

Our study revolves around the following research questions: first, can strangers who interact indefinitely achieve substantial levels of cooperation and efficiency? Strangers are anonymous subjects who are randomly matched in each period and their histories are private information. Second, which institutions for monitoring and enforcement promote cooperation? And, finally, what classes of strategies are adopted in economies that achieve high efficiency?

Our results bring new insights in understanding long-term relationships in anonymous economies. First, cooperation levels in our experimental economies are high and increasing with experience, even when action histories are private information. The result is novel. Second, this study sheds some light on the type of economic institutions that may facilitate the emergence of norms of cooperation in experimental anonymous societies. For instance, not all monitoring institutions promote cooperation. We report high cooperation levels in situations where subjects know identities and histories of opponents, but not when they see aggregate outcomes without observing identities. Moreover, costly personal punishment significantly promotes cooperation. Under this institution subjects can pay a cost to inflict a loss on their opponent. The effect of this institution has been studied in settings with finitely repeated interaction (Elinor Ostrom, James Walker, and Roy Gardner, 1992, Ernst Fehr and Simon Gaechter, 2000), but not when interaction

is indefinitely repeated, which is when many *informal* equilibrium punishment strategies are also available. Our work complements a growing economics literature devoted to uncover theoretical links between the (un)availability of enforcement and punishment institutions on one side, and patterns of exchange and cooperation on the other (e.g., Stefan Krasa and Anne Villamil, 2000, Dixit, 2003, Charalambos D. Aliprantis, Gabriele Camera and Daniela Puzzello, 2007). Finally, subjects appear to have preferences for some classes of strategies. The average subject avoids indiscriminate strategies, shows a strong tendency to defect with opponents who have “cheated” her in the past, and tends to disregard information on the opponent’s behavior in other matches. These findings help define an empirically-relevant criterion for equilibrium selection—one of the unsolved questions of the theory of repeated games—using behavioral considerations.

The paper proceeds as follows: Section I discusses the related literature; Section II presents the experimental design; Section III provides a theoretical analysis; results are reported in Section IV; and Section V concludes.

I. Related experimental literature

Our work builds on the experimental literature on infinitely repeated games (supergames). Alvin E. Roth and Keith Murnighan (1978) were the first to implement infinitely repeated games in an experiment by employing a probabilistic continuation rule, thus transforming it into an *indefinitely* repeated game. Many experiments have followed this design (e.g., Duffy and Ochs, 1999, Aldo Rustichini and Anne Villamil, 2000) because for risk-neutral subjects a constant continuation probability is theoretically equivalent to a constant time-discount rate and an infinite horizon.

Several 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 finitely repeated versus indefinitely repeated interaction, and react in the expected direction. For example, Dal Bó (2005) reports lower cooperation for finite duration experiments in comparison to indefinite duration with the same expected length; the higher the discount rate the lower the cooperation. See also Hans-Theo Normann and Brian Wallace (2006).

The closest literature considers indefinitely repeated experiments whose stage game is a prisoner's dilemma (for other games see Tim Cason and Feisal Khan, 1999, Engle-Warnick and Slonim, 2004, 2006, Engle-Warnick, 2007). Two aspects of these experiments are important: the matching protocols and the availability of information about other subjects. *Within a supergame*, subjects are matched either using a fixed or a random protocol. Since all experiments surveyed include several supergames within a session, they also specify an additional protocol to match the subjects after each supergame. We will come back to this point, later.

Most studies use a *fixed matching* protocol within a supergame; see Palfrey and Rosenthal (1994), Aoyagi and Frechette (2005), or Dal Bó (2005). Under this design, referred to as “partner”, subjects always interact with the same person and generally support a significant level of cooperation, sometimes full cooperation. Instead, our study employs a *random matching* protocol within a supergame as, for instance, in Steven Schwartz, Richard Young, and Kristina Zvinakis (2000) and Duffy and Ochs (2006). In any given period subjects meet in pairs but after each period matches are destroyed and new pairs are formed drawing subjects at random from the entire economy. Duffy and Ochs (2006) found 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 they are empirically difficult to attain.

A novel feature of our study is that it helps us understand which one of the several available strategies that support a given equilibrium outcome have been employed.¹ This issue has been largely unexplored in the experimental literature on supergames, as it has mostly focused on measuring the levels of cooperation; an exception is Engle-Warnick and Slonim (2006). Our experimental design allows us to 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. Subjects in experimental studies of finitely repeated social dilemmas have shown a surprising tendency to engage in costly personal punishment of others, especially defectors. Though this behavior is inconsistent with personal income maximization, it has been shown to be remarkably robust (Ostrom, Walker, and Gardner, 1992; Fehr and Gaechter, 2000, Marco Casari and Charles Plott, 2003).² A third novel feature of our study is to examine how this behavioral trait may be employed in sustaining the cooperative equilibrium in an infinitely repeated game, where there does already exist an (informal) 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.

The matching protocol across supergames is also important because of possible contagion effects. Indeed, to play a supergame in a session, there are several ways to partition a pool of subjects into several 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

¹ The strategies include off-equilibrium threats that are not carried out on the theoretical equilibrium path. The features of these threats are irrelevant as long as they are credible and generate a sufficiently low continuation payoff.

² For example, Fehr and Gaechter (2000) consider a stranger matching model (a finite sequence of one-shot interactions), and the intensity of punishment does not show any end game effect; punishment seems to be driven by strong negative emotions.

rigorous partitioning procedure is to avoid that anyone shares a common past opponent. Both procedures control for contagion effects.³ This contrasts with randomly re-matching the same set of subjects in each period and after each supergame (e.g., Schwartz, Young and Zvinakis, 2000).

	Private monitoring		Private monitoring with punishment		Anonymous public monitoring		Public monitoring (non-anonymous)	
Matching protocol within an economy	Random		Random		Random		Random	
Anonymity	No subject IDs		No subject IDs		No subject IDs		Subject IDs are public	
Information	Action of current opponent		Action of current opponent		History of all actions taken in the economy without IDs (no individual histories)		Individual histories of everyone in the economy	
Ways to punish	Only by defecting		Pay 5 (points) to reduce opponent's payoff by 10		Only by defecting		Only by defecting	
<i>Available strategies:</i>								
- Global (not selective)					Yes		Yes	
- Reactive (moderately selective)	Yes		Yes		Yes		Yes	
- Targeted (highly selective)			(°)				Yes	
Session dates	21.4.05	7.9.05	28.4.05	6.9.05	27.4.05	1.9.05	12.4.05	8.9.05
Show-up fee	\$5	\$5	\$5	0	0	0	0	0
No. of periods	71	104	139	99	129	125	86	128

Table 1: Experimental treatments⁴

³ In Dal Bó (2005) each subject plays three supergames. In each supergame of the “Dice” sessions N participants are partitioned into N/2 two-person economies. The partitioning across supergames is such that a subject’s decisions in a supergame could not affect the decisions of subjects met in future supergames. Ensuring the absence of contagion effects in this manner requires very large session sizes (see the theory of anonymous matching procedures in Aliprantis, Camera and Puzello 2006, 2007). In our study each subject played five supergames. Subjects may have shared a common past opponent in supergames three or later. Aoyagi and Frechette (2005) use a different in between matching protocol; each agent plays G>10 supergames. In the first ten supergames they partition agents in a round robin fashion and in the last (G-10) supergames they randomly re-matched participants.

(°) 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.

⁴ For comparison purposes, note that a “partner” treatment (e.g., as in Dal Bó, 2005 or Duffy and Ochs, 2006) differs from our treatments in the matching protocol (fixed pairings instead of random), may differ in Anonymity (subject IDs may be public or not), and is otherwise identical to the Private Monitoring treatment in Information and Ways to punish. Of course, with fixed pairings the distinction among targeted, reactive, and global strategies is irrelevant.

II. Experimental design

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

The stage game. The stage game is a standard prisoners' dilemma with payoffs determined according to Table 2.⁵ We call action Y *cooperate* and action Z *defect*. We say that there is coordination on *cooperation* in the pair only if both subjects choose Y . So, we will define the degree of coordination on cooperation in the economy according to how many pairs cooperate.

(A) Notation in the theoretical analysis			(B) Parameterization of the experiment		
Player 1/ Player 2	Y	Z	Player 1/ Player 2	<i>Cooperate</i> (Y)	<i>Defect</i> (Z)
Y	y, y	l, h	<i>Cooperate</i> (Y)	25, 25	5, 30
Z	h, l	z, z	<i>Defect</i> (Z)	30, 5	10, 10

Table 2: The stage game

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 with probability one 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 $1/(1-\delta)$ periods, and we set $\delta = 0.95$, so in each period the

⁵ We selected this parameterization as it scores high on the indexes proposed by Anatol Rapoport and Albert 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. In Table 2 we have $0 \leq l < z < y < h$ and $2z < l+h < 2y$.

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 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 1).⁷ 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

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

⁷ Following a referee's suggestion, we ran a fifth treatment under private monitoring with economies of 14 subjects interacting for only one cycle. We ran 4 sessions at Purdue University drawing subjects from the same pool. On average, a session lasted 40 minutes and paid \$11.70 per person, including an \$8 or \$10 show-up fee.

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 current actions taken in every pair. In one treatment, denoted *non-anonymous public monitoring*, histories were associated with identities of subjects.⁸ 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 with others.

⁸ In a finitely-repeated trust game experiment Iris Bohnet and Steffen Huck (2004) inform the trustor about her trustee's past behavior in each period. In our *non-anonymous public monitoring* treatment we provide information about identities, actions and matching histories of everyone in the economy, not only of the current opponent.

We recruited 160 subjects through announcements in undergraduate classes at Purdue University and signed up online. The experiment was programmed and conducted with the software z-Tree (Fischbacher 2007) at Purdue University. No eye contact was possible among subjects, and copies of the instructions were on all desks. Instructions were read aloud. A copy of the instructions is in the online supplementary material. 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 1 (each session had 20 participants and 5 cycles).

III. 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 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.⁹

An “economy” is composed of four players a , b , c , and d who interact for an indefinite number of periods denoted $t = 1, 2, \dots$. 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.¹⁰

III.A. Equilibrium in the stage game

Consider the stage game described in Table 2A, which is a prisoners’ dilemma. The players

⁹ The theoretical framework is one of a homogeneous population. An alternative is to consider subjects of different types in the experiment as, for example, in Miguel Costa-Gomes, Vincent Crawford, and Bruno Broseta (2001) or Paul J. Healy (2007). The assumption of identical risk-neutral players is, of course, open to question but it has been retained since it is a useful abstraction

¹⁰ 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.

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, i.e. 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:

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

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 everyone cooperates, i.e. all players choose Y . Thus, we refer to the outcome where every player in the economy selects Y as the *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.

III.B. Equilibrium in the indefinitely repeated game with private monitoring

Under private monitoring indefinite repetition of the stage game with randomly selected opponents can expand the set of equilibrium outcomes. Following the work in Kandori (1992) and Ellison (1994), here we present sufficient conditions so that the equilibrium set includes the efficient outcome, which is achieved when everyone cooperates in every match and all periods.

The *private monitoring* treatment is characterized by two informational frictions. First, 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.

The inefficient outcome can be supported as a sequential equilibrium using the strategy “defect forever.” Since repeated play does not decrease the set of equilibrium payoffs, Z is always a best response to play of Z by any randomly chosen opponent. In this case the payoff in the indefinitely repeated game is the present discounted value of the minmax payoff, $z/(1-\delta)$.

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

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

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

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 Camera and Casari (2007) and in the online supplementary material, 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 called a *reactive strategy*, i.e., a player will choose Z if and only if his opponent 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.

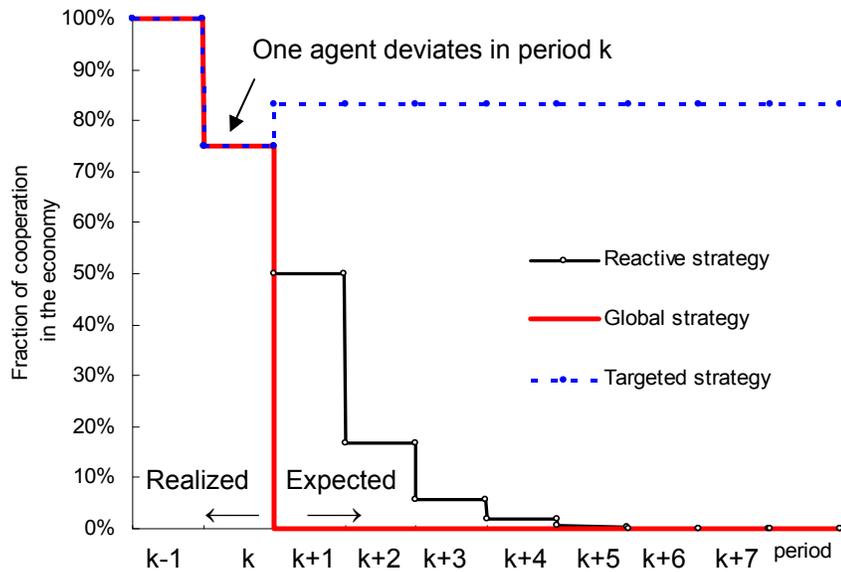


Figure 1: Dynamic reaction to a defection in an economy

Consider one-time deviations by a single player. 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 (unimprovability criterion). Instead, if a defection occurs and everyone plays according to the social norm, then everyone ends 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 observes Z . If he meets a cooperator in the next period, then choosing Y produces a current loss to the player because he earns y (instead of h). If he meets a deviator, choosing Y also causes a current loss because he earns l rather than z . Hence, 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 optimal 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 Proposition 1. For the efficient outcome to be feasible, we need $\delta \geq \delta^*$. In our experimental design $\delta = 0.95$ and $\delta^* = 0.443$, a value that solves the condition in Proposition 1 for the parameterization given in Table 2B.¹¹

Proposition 3. *In our experimental economies with private monitoring, the use of personal punishment is neither necessary nor sufficient to sustain the efficient outcome as an equilibrium.*

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. To sustain the efficient outcome in private monitoring, subjects can use a grim trigger strategy, hence the use of personal punishment is theoretically not necessary. In a one-shot interaction, choosing

¹¹ 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 David K. Levine and Wolfgang 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 study, we work with four-agent economies, the smallest possible number that allow pairwise anonymous matching.

personal punishment is a dominated action because it is costly for the punisher. Given indefinite repetition, personal punishment in our design is not a credible threat and cannot be part of any sequential equilibrium, on or off the equilibrium path. Personal punishment is not an optimal strategy 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, personal punishment is not theoretically sufficient to sustain the efficient outcome because the threat of personal punishment alone 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. On the contrary, defecting after having observed a defection is an optimal strategy. 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. Moreover, the incentive to defect in following periods remains because defection is the unique best response in the one-shot game. In conclusion, though personal punishment is a big enough threat to sustain cooperation, it is not a credible one.

III.C. 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).

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, if players cooperate with those known to have cooperated in the past, then, loosely speaking, 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, off the equilibrium path, which generates incentives to cooperate for even lower discount factors.

Below we identify three broad classes of strategies. 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*. Second, players could switch to a punishment mode when they observe an opponent defect, but

stay in a cooperative mode if a defection is observed elsewhere in the economy, what we refer to as a *reactive strategy*. Third, 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*, because the subject punishes only those who have defected in a match with her. It is easily demonstrated that, with a targeted strategy, the efficient outcome is optimal as long as δ is greater than 0.5. Of course, these three classes of strategies do not exhaust all possible behaviors; for instance, players can punish anyone who has ever defected. However, they are indicative of three intuitive ways of behaving and, as we will see, can explain a great deal of subjects' behavior.

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—global, reactive, targeted (or something else that we do not characterize)—enhances cooperation (see Table 1).

One can classify strategies also using “power” and “selectivity” scores. The *power* of a strategy is the maximum loss that can be inflicted on a defector, which depends on the immediacy and frequency of punishment. The greater the power, the lower is the defector's continuation payoff, and the greater is the incentive to cooperate. Among the three strategies considered, global strategies have the most power as they provide the largest possible threat: everyone defects right after a deviation (Figure 1). Targeted strategies have the least power.¹² A strategy's *selectivity* depends on who gets punished. Targeted strategies are the most selective

¹² For example, with public monitoring, the lower bound for δ falls by about 40 percent when we move from a reactive to a global strategy and by about 50 percent when we move from a targeted to a global strategy.

and imply the lowest cost of punishment: only opponents who defected are punished.

IV. Results

We first present results on the aggregate outcome (Results 1-5) and then on the strategies employed to sustain those outcomes (Results 6-10).

Result 1. *In economies with private monitoring, cooperation emerged and was higher in later cycles than in earlier cycles.*

Kandori (1992) and Ellison (1994) proved the theoretical possibility of cooperation with private monitoring. In our data the rate of (coordination on) cooperation was remarkably high, 59.5 percent when averaging across all periods (Figure 2). In addition, the data displays an increasing cooperation trend across cycles (Figure 3). Both aspects are novel.

Consider an economy $k=1,\dots,50$ as a unit of observation. For an economy k we define the action a_{it}^k of an agent $i=1,\dots,4$ in period $t=1,\dots,T^k$ of the economy as 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

$$(3) \quad c_k = \frac{1}{4T^k} \sum_{t=1}^{T^k} \sum_{i=1}^4 a_{it}^k,$$

and across economies is $c = \frac{1}{50} \sum_{k=1}^{50} c_k$. Thus, although economies have different length T^k , they are given equal weight in our measure c of average cooperation, since we consider each economy a unit of observation.

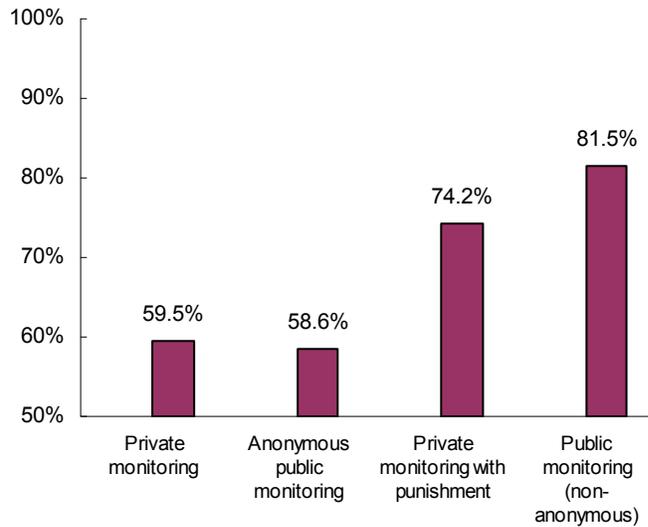


Figure 2: Average cooperation across treatments¹³

It is instructive to compare our results to those from related experiments with a random matching protocol. Duffy and Ochs (2006) report a low average rate of cooperation of 6.3 percent and, most importantly, a declining trend across supergames; cooperation declined from 8.7 to 3.9 percent (first half of the sessions). Schwartz, Young and Zvinakis (2000) report a 19.2 percent cooperation rate and also a declining trend across supergames from 25.1 percent (cycles 1-2) to 10.8 percent (cycles 5-7). Trends are important because they point to a direction of learning. In our study, instead, cooperation in the first two cycles was 55.2 percent and grew to 66.5 percent in the last two cycles. Why do we see these differences? In all these experimental designs the incentives satisfy the *necessary* conditions to support full-cooperation as a theoretical equilibrium. However, such incentives are not theoretically *sufficient* to achieve full cooperation because subjects may coordinate on a less efficient equilibrium outcome. Behavioral components

¹³ 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 percent 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 $n_1=n_2=50$.

may thus influence the outcome. In particular, three elements may generate lower cooperation rates in Duffy and Ochs (2006) than in our experiment: their stage game payoffs reward cooperation less ($l=0$, $h=30$, $y=10$, $z=20$); they have lower continuation probability (0.90 vs. 0.95); and they have larger group sizes (6-14 vs. 4). All these features generate stronger incentives towards cooperation in our design.¹⁴ Consequently, we surmise that both our as well as the other related studies point to some missing elements in the theory of social norms. To assess the impact of the economy size with private monitoring we ran four additional sessions with economies of 14 subjects interacting for one cycle. Average cooperation was 23.6 percent, which is in line with previous studies.¹⁵ A systematic exploration of the behavioral conditions for the breaking down of cooperation with a random matching protocol is left to future work.

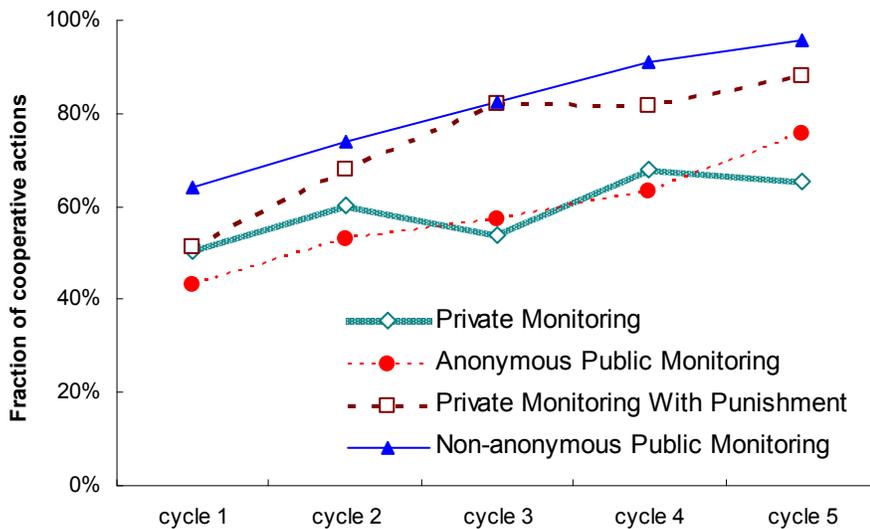


Figure 3: Average cooperation across cycles

¹⁴ Similar considerations hold for the design in Schwartz, Young and Zvinakis (2000). In their stage game subjects first chose between a safe, outside option of 60 or to play a prisoner's dilemma ($l=10$, $h=170$, $y=120$, $z=60$). The continuation probability was about 0.89 and economy sizes 10-14.

¹⁵ Period 1 cooperation was 55.3 percent versus 42.9 percent in Duffy and Ochs (2006) in a comparable treatment (Random $I=0$).

More generally, the literature reports increasing cooperation rates over cycles in indefinitely repeated experiments among partners (Aoyagi and Frechette, 2005; Dal Bó and Frechette, 2006; Duffy and Ochs, 2006). Instead, results are sharply different than in finitely repeated prisoners' dilemmas and voluntary public good experiments, where cooperation is generally declining over time (Palfrey and Rosenthal, 1994).

Result 2. *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 if this information is anonymous, then it does not foster cooperation. Average cooperation for all periods is around 59 percent in both treatments; the difference is statistically insignificant (Mann-Whitney test, p-value 0.418, $n_1=n_2=50$). If we consider average cooperation in first periods only, then we reach the same conclusion (see Table 3). Considering first-period choices of an economy is important because they supply a complementary measure of cooperation, which is independent from the choices of other subjects in the economy. In particular, as is shown later, first-period behavior suggests whether some equilibrium among the many possible had a particularly strong drawing power.

Number of cooperative Actions	Private monitoring	Anonymous public monitoring	Private monitoring with punishment	Public monitoring (non-anonymous)
Average cooperation (percentages)				
	73.5	70.5	84.5	87.0
Frequency of cooperation in an economy (percentages)				
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 (percentages)				
2	58	51	71	75
1	31	39	27	24
0	11	10	2	1

Table 3: Cooperation in the first period of an economy¹⁶

Result 3. *The introduction of personal punishment in the anonymous treatments increased both cooperation and realized efficiency.*

Figures 2 and Table 3 provide support for Result 3. When we add personal punishment to economies with private monitoring, average cooperation jumps from 59.5 to 74.2 percent. This difference is statistically significant at a 1 percent level (Mann-Whitney test, p-value 0.007). This difference is also evident when comparing average cooperation in the first period of each cycle (73.5 vs. 84.5 percent, Table 3). As shown in Result 4, average cooperation is statistically indistinguishable from the non-anonymous public monitoring treatment (Mann-Whitney test, p-value 0.154). The observed pattern of high cooperation in the first period not only is an indicator of subjects' preferences for the efficient outcome, but it also suggests that subjects might have anticipated that the efficient outcome would be enforced by personal punishment. This anticipation was correct because, as shown in Result 7, personal punishment was indeed administered to defectors.

¹⁶ 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.”

The comparison among treatments in terms of realized efficiency substantially confirms the conclusions drawn in Results 1-5, in terms of average cooperation. Ranking among treatments may have been different because personal punishment is a deadweight loss. We define *realized efficiency* in an economy k by

$$(4) \quad e_k = \frac{1}{4T^k} \sum_{t=1}^{T^k} \sum_{i=1}^4 \frac{\pi_{it}^k - 10}{25 - 10}$$

The payoff to subject i , in economy k , of period t is denoted π_{it}^k (and given in Table 2). The denominator reports the average payoff in a match, which ranges from a minimum of 10 to a maximum of 25 points. Realized efficiency e_k ranges from 0 to 1. In particular, $e_k=0$ when everyone in the economy always defects 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. Only 2 out of 50 economies had a negative realized efficiency (the minimum was $e_k = -0.074$). Average realized efficiencies e_k for the four treatments in the experiment were 0.595, 0.586, 0.652, and 0.815 (ordering as in Table 3). Given that private monitoring with punishment displayed average cooperation levels comparable to non-anonymous public monitoring (Table 3, footnote 18), the difference in efficiency between these two treatments originates from the deadweight loss due to personal punishment.

Result 4. *The introduction of public monitoring in the non-anonymous treatment increased cooperation over private monitoring.*

Figure 2 and Table 3 provide support for Result 4. In the non-anonymous public monitoring treatment average cooperation across economies was 81.5 percent. A Mann-Whitney test conducted on cooperation in non-anonymous public monitoring shows significant difference with private monitoring (59.5 percent, p-value 0.0001, $N_1 = N_2 = 50$) and with anonymous public

monitoring (58.6 percent, p-value 0.0000).¹⁷ Result 4 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 analyzed the distribution of average cooperation levels across the fifty economies.¹⁸ About 38 percent of the economies have cooperation rates above 98 percent. The superiority of non-anonymous public monitoring is clear also from the average cooperation in the initial period across economies, as shown in Table 3. In this treatment the full past record of the opponent is available, hence each participant could develop a reputation over time, which, as will be shown in Result 9, was the key to the observed success in coordinating on the fully efficient equilibrium.

The remaining results report about the strategies adopted by the representative subject by considering three elements: (1) how subjects played the first period of each cycle, (2) how subjects reacted after seeing a defection, and (3) whether, after seeing a defection, subjects eventually reverted to cooperation. This allows us to establish the empirical relevance of several available strategies, which may or may not be consistent with equilibrium.

Results 5. *In all treatments, period 1 cooperation was significantly different than zero. Hence, there is no evidence of coordination on the inefficient outcome.*

Table 3 provides evidence for Result 5. As noted earlier, choices in the first period of each economy suggest whether some equilibrium among the many possible had a particularly strong

¹⁷ The unit of analysis used in all tests is an economy. Strictly speaking, all observations are independent only if we focus on the first cycle. The results of the test rely on all observations being independent.

¹⁸ Kolmogorov-Smirnov two-tail two-sample tests on distributions confirm results from the Mann-Whitney tests on the differences between averages. On one hand private monitoring and anonymous public monitoring are not statistically different treatments (10 percent confidence level, $n_1=n_2=50$). On the other hand, private monitoring with punishment and non-anonymous public monitoring are not statistically different either. Instead, treatments from the two groups are statistically different at least at a 5 percent level.

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 3. 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, in period 1 both subjects cooperated in more than 50 percent of the matches in every treatment.¹⁹ Furthermore, Table 5 includes an analysis of coordination on cooperation to all periods. Coordination on cooperation in an economy ranges from 28 percent in private monitoring to 50 percent in private monitoring with punishment.

Table 4: Probit regression on individual choice to cooperate – marginal effects²⁰

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
<i>Treatment dummies:</i>						
Anonymous Public Monitoring					-0.046*	-0.029 (0.073)
Private Monitoring With punishment					0.998*** (0.000)	0.092 (0.067)
Public Monitoring (non-anonymous)					0.947*** (0.009)	0.117* (0.061)
<i>Cycle dummies:</i>						
Cycle 2	0.039 (0.104)	0.057 (0.038)	0.083*** (0.026)	-0.003* (0.002)	0.062*** (0.023)	-0.037 (0.028)
Cycle 3	0.076 (0.069)	0.050 (0.051)	0.111*** (0.020)	0.020*** (0.002)	0.093*** (0.027)	0.006 (0.029)
Cycle 4	0.136*** (0.008)	0.188*** (0.025)	0.149*** (0.030)	0.126*** (0.027)	0.174*** (0.022)	0.049 (0.035)
Cycle 5	-0.160***	0.290***	0.139***	0.139***	0.214***	0.082***

¹⁹ Of course, there is variation in subjects' period one behavior. Consider all cycles; the variance of average cooperation across all subjects is 0.136, 0.117, 0.056, and 0.059 (treatments ordered as in Table 3). The percentages of subjects who cooperated in period one of all cycles are 44, 36, 46, and 52, respectively.

²⁰ 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 percent; ** significant at 5 percent; *** significant at 1 percent. 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 (except the last column). Individual fixed effects and period fixed effects are included (except in the last column) 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). Duration of previous cycle was set to 20 for cycle 1.

	(0.043)	(0.032)	(0.033)	(0.004)	(0.021)	(0.031)
Duration of previous cycle	0.001*	0.003***	0.002***	0.004***	0.004***	0.003***
	(0.001)	(0.000)	(0.000)	(0.001)	(0.001)	(0.001)
<i>Reactive strategies:</i>						
Grim trigger	-0.550***	-0.266***	-0.382***	0.075	-0.388***	
	(0.014)	(0.074)	(0.100)	(0.055)	(0.041)	
lag 1	0.088**	-0.048**	0.056*	-0.061	0.018	
	(0.043)	(0.024)	(0.030)	(0.039)	(0.027)	
lag 2	0.116***	-0.095***	0.046*	-0.140***	-0.027	
	(0.036)	(0.018)	(0.027)	(0.031)	(0.039)	
lag 3	0.103**	-0.073*	0.040	-0.063***	-0.010	
	(0.042)	(0.042)	(0.034)	(0.007)	(0.027)	
lag 4	0.080**	-0.058	0.0152	-0.053	-0.033	
	(0.005)	(0.047)	(0.045)	(0.060)	(0.029)	
lag 5	0.030**	-0.071***	0.014	-0.018	-0.044*	
	(0.014)	(0.014)	(0.030)	(0.041)	(0.023)	
<i>Global strategies:</i>						
Grim trigger		-0.311**		-0.116***		
		(0.131)		(0.002)		
lag 1		0.227***		0.023		
		(0.016)		(0.059)		
lag 2		0.229***		0.028		
		(0.063)		(0.043)		
lag 3		0.243***		0.048**		
		(0.010)		(0.024)		
lag 4		0.175***		0.005		
		(0.031)		(0.021)		
lag 5		0.155***		-0.032		
		(0.012)		(0.054)		
<i>Targeted strategies:</i>						
Grim trigger				-0.363***		
				(0.047)		
lag 1				-0.044***		
				(0.005)		
lag 2				-0.057***		
				(0.014)		
lag 3				-0.018		
				(0.033)		
lag 4				-0.043***		
				(0.003)		
lag 5				-0.063***		
				(0.016)		
<i>Personal punishment</i>						
Requested (lag)			-0.076			
			(0.085)			
Requested (lag)			0.028			
× opponent defected (lag)						
			(0.029)			
Received (lag)			0.067*			
			(0.038)			
Received (lag)			-0.329***			
× subject defected (lag)						
			(0.097)			
Observations	3320	4880	4400	4280	16680	800

Result 6. *In the private monitoring treatment, the representative subject who observed her opponent defect switched from a cooperative mode to a punishment mode. Hence, there is evidence of use of reactive strategies.*

Table 4 and Figure 4 provide support for Result 6. Recall that a reactive strategy involves a shift to a punishment mode following a defection of the opponent. A grim trigger strategy lies in this class and can theoretically sustain an equilibrium with full cooperation in our setting.

Table 4 reports the results from a probit regression that explains the individual choice to cooperate (1) or not (0) using two groups of regressors. First, we introduce several dummy variables that control for fixed effect (cycles, periods within the cycle, individuals), as well as for the duration of the previous cycle. Second, we include a set of regressors used to trace the response of the representative subject in the periods *following* an observed defection. For simplicity, we limit our focus to the five periods following an observed defection. This specification is more general than tracing behavior in periods 1-5 only, and it allows us to shed light on the type of strategy employed by the representative subject. Of course, there are several ways to choose regressors in order to trace strategies. Our specification has the advantage to detect whether subjects followed theoretically well-known strategies, such as grim trigger or tit-for-tat (Robert Axelrod, 1984). Indeed, we include a “grim trigger” regressor, which has a value of 1 in *all* periods following an observed defection and 0 otherwise. We also include five “lag” regressors, which have a value of 1 only in *one* period following an observed defection and 0 otherwise. For example, the “lag 1” regressor takes value 1 exclusively in the period after the defection (0 otherwise). The “lag 2” regressor takes value 1 exclusively in the second period following a defection (0 otherwise). And so on.

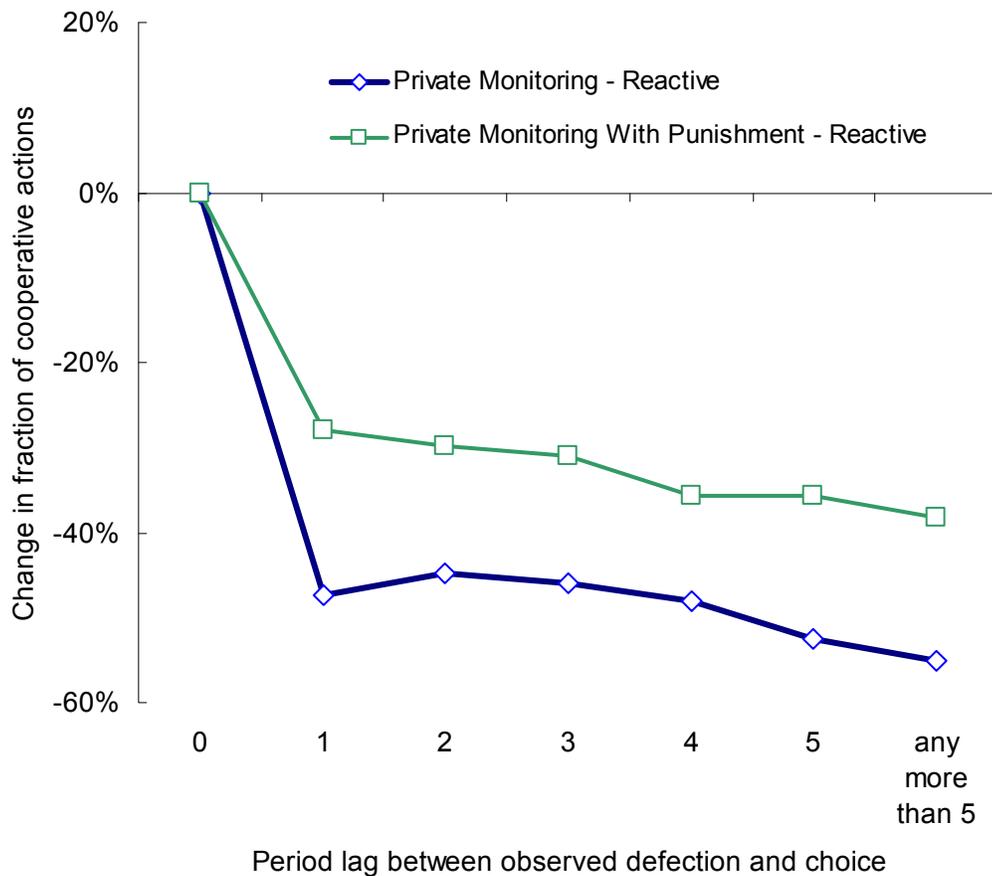


Figure 4: Strategies of the representative subject in private monitoring with and without personal punishment

If the representative subject switched from a cooperative to a punishment mode after seeing a defection, then the estimated coefficient of at least one of the six strategy regressors should be negative. For example, if subjects punished for just two periods following a defection, then the sum of the estimated coefficients of the grim trigger regressor and the lag regressors should be negative for the first and second period following a defection, and zero afterwards.

Figure 4 illustrates the marginal effect on the frequency of cooperation in the periods that followed an observed defection.²¹ The focus on the five-period lags is for convenience in

²¹ Figure 4 is based on Table 4 using the coefficient estimates coding for reactive strategies. Zero-period lag is exogenously set at 0 percent. The point for “any more than 5” is the marginal effect on the frequency of cooperation

showing patterns in the results. The representation for “any more than five” period lags is based on the marginal effect of the grim trigger regressor only. The representation for period lags 1 through 5 is based on the sum of the marginal effects of the grim trigger regressor and the lag regressor with the appropriate lag. The L-shaped pattern of response to an observed defection suggests a persistent downward shift in cooperation levels immediately after a defection. The grim trigger coefficient estimate is significantly different than zero at a 1 percent level. All other strategy regressors are significant at 10 percent level or more (Table 4).²² While there is evidence that the representative subject employed a reactive strategy, not all observed actions fit this type of strategy. Indeed, the transitional matrices displayed in Table 5A indicate that about 40 percent of individual actions are not compatible with reactive strategies.

of the grim trigger regressor. Lags 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 lag 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 lag regressors are computed for grim trigger regressor set at 1 (i.e. defection)

²² Table 4 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 (2005) and Engle-Warnick and Slonim (2004).

A - Private Monitoring

C - Private Monitoring with Punishment

No. cooperative actions in the current period	No.cooperative actions in next period						No.cooperative actions in next period					
	0	1	2	3	4	totals	0	1	2	3	4	totals
	(percentages)						(percentages)					
0	10	4#	#	#	#	14	8	3#	#	#	#	10
1	4	9#	3#	#	#	16	3	5#	2#	1#	#	10
2	1	4#	15	3#	#	23	1	3#	5	3#	1#	12
3	#	#	5	11#	3#	19	#	1#	4	10#	3#	17
4	#	#	#	3#	24	28	#	#	#	3#	46	50
						100						100

B - Anonymous Public Monitoring D - Non-Anonymous Public Monitoring

No. cooperative actions in the current period	No.cooperative actions in next period						No.cooperative actions in next period					
	0	1	2	3	4	totals	0	1	2	3	4	totals
	(percentages)						(percentages)					
0	7	4	1			13	1	1	1			3
1	5	9	4	1		19	1	1	2	1		5
2	1	5	8	3		17	1	2	13	7	3	25
3		1	4	10	4	19		1	8	11	5	26
4				4	28	32			3	5	33	41
						100						100

Note: When everyone uses only reactive strategies (grim trigger) the cells with the # sign should be empty. A blank cell indicates no observation or a frequency below 0.5 per cent. Frequencies are rounded to the nearest integer percentage point. All periods included except the last one of each cycle. No. of observations is 3320 (A), 4400 (B), 4880 (C), and 4080 (D).

Table 5: Transitional matrices in an economy

Result 7. *In the private monitoring with punishment treatment, the representative subject who observed her opponent defect sometimes employed personal punishment while staying in a cooperative mode.*

Tables 4-7 along with Figure 4 provide support for Result 7. The L-shaped pattern of response to an observed defection in Figure 4 (lighter line) suggests a persistent downward shift in cooperation levels immediately after a defection. The estimated coefficient for the grim trigger regressor is significant at a 1 percent level (Table 4).

As already noted, for our parameterization the addition of personal punishment does not expand the set of equilibrium outcomes. However—in contrast to the private monitoring treatment without punishment (Result 6)—we do find behavioral differences. First, the magnitude of the downward shift in cooperation levels is now substantially smaller (compare the darker and lighter lines in Figure 4). Second, subjects employed personal punishment in 9.1 percent of the matches. In particular, Table 6 shows that personal punishment was mostly used by cooperators against an opponent who defected. In about 58 percent of such cooperator-defector encounters, the cooperator requested that personal punishment be inflicted on the opponent.

		Action of opponent receiving punishment (percentages)	
		Cooperate	Defect
Action of subject requesting punishment	Cooperate	0.1	58.3
	Defect	5.4	10.4

Table 6: Frequency of personal punishment²³

These two changes in observed behavior are correlated. When observing a defection, subjects at times switched from a cooperative mode to a punishment mode. However, subjects often continued cooperating but sanctioned through personal punishment. That is, subjects sometimes treated personal punishment as a substitute for *informal* punishment, i.e., defecting in following periods. Table 7A supports this interpretation. In particular, a cooperator encountering a defector subsequently cooperated 75.5 percent of the times if she requested personal punishment, but only 46.7 percent of the times, if she did not punish the defector. Reversing the viewpoint, 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 percent). Once we control for all other factors, however, the

²³ 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.

evidence on this point is mixed (Table 4). Personal punishment seems to boost cooperation levels only in small part by deterring defection and in large part by avoiding that cooperators switch to defection after punishing. This finding is interesting because the existing literature mostly places emphasis on the former aspect, though recent studies on peer punishment find the latter aspect is very important, even in finitely repeated interaction (Marco Casari and Luigi Luini, 2007).

(A) Choice after a subject cooperated and the opponent defected			(B) Choice after a subject defected and the opponent cooperated		
Did the subject <i>request</i> personal punishment?	Subject choice in the following period (<i>percentages</i>)		Did the subject <i>receive</i> personal punishment?	Subject choice in the following period (<i>percentages</i>)	
	Cooperate	Defect		Cooperate	Defect
Yes	75.5	24.5	Yes	34.5	65.5
No	46.7	53.3	No	24.1	75.9

Table 7: Transitional matrices in private monitoring with punishment

To interpret Result 7, recall that our theoretical framework presumes a homogeneous population, as in Kandori (1992) and Ellison (1994). Within this framework, the observed punishment behavior seems at odds with equilibrium predictions. Subjects should theoretically achieve cooperation only by threatening and eventually triggering to permanent defection. In finitely repeated experiments, subjects employ personal punishment for behavioral reasons—for instance distributional justice or revenge—and for lack of any alternative equilibrium punishment strategy (Ostrom, Walker and Gardner, 1992, Casari and Luini, 2007). In our study, instead, subjects show a preference for personal punishment over (equilibrium) informal punishment schemes. This gives even stronger support to the notion of the usefulness of personal punishment in sustaining cooperation. In the concluding section of the paper we will put forward various conjectures to explain this behavior.

To complement the evidence for Result 7 we calculated individual average profit and punishment points given (data not reported). When considering the individual averages within a cycle, greater profits are associated to less punishment being given. However, there is significant variability in punishment: subjects with the lowest average profit tend to punish more, but not all of them engage in punishing. Thus, costly personal punishment seems to be a public good. On the one hand it significantly increases cooperation as well as realized efficiency (Result 3). On the other hand the subjects who benefit the most are cooperators who punish little or not at all.

Result 8. *In the anonymous public monitoring treatment, the representative subject selected reactive strategies over global strategies.*

In anonymous public monitoring subjects observed whether a defection had occurred in the match or elsewhere in the economy. In the experiment, a defection by an opponent generated a stronger response than a defection elsewhere in the economy. This conclusion is based on the estimated coefficients for reactive and global strategies. Both strategies were available in the anonymous public monitoring treatment (Table 4, Figure 5). A subject using a reactive strategy punished everyone after seeing a defection in the match, but kept cooperating after seeing a defection outside the match. In contrast a subject using a global strategy started punishing everyone after observing a defection, no matter if it came from an opponent or someone else.

Figure 5 is based on the marginal effects estimated using regressions in Table 4.²⁴ In addition to what has been explained after Result 6, the probit regression for anonymous public monitoring includes six additional strategy regressors, which are used to trace global strategies. The representative subject who experienced a defection displayed a strong and persistent decrease in

²⁴ Figure 5 uses the coefficient estimates coding reactive and global strategies, respectively. 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.

future cooperation levels (reactive strategy: solid line in Figure 5). Conversely, the response was much weaker when the representative subject observed a defection *outside* her match (global strategy: dashed line in Figure 5).

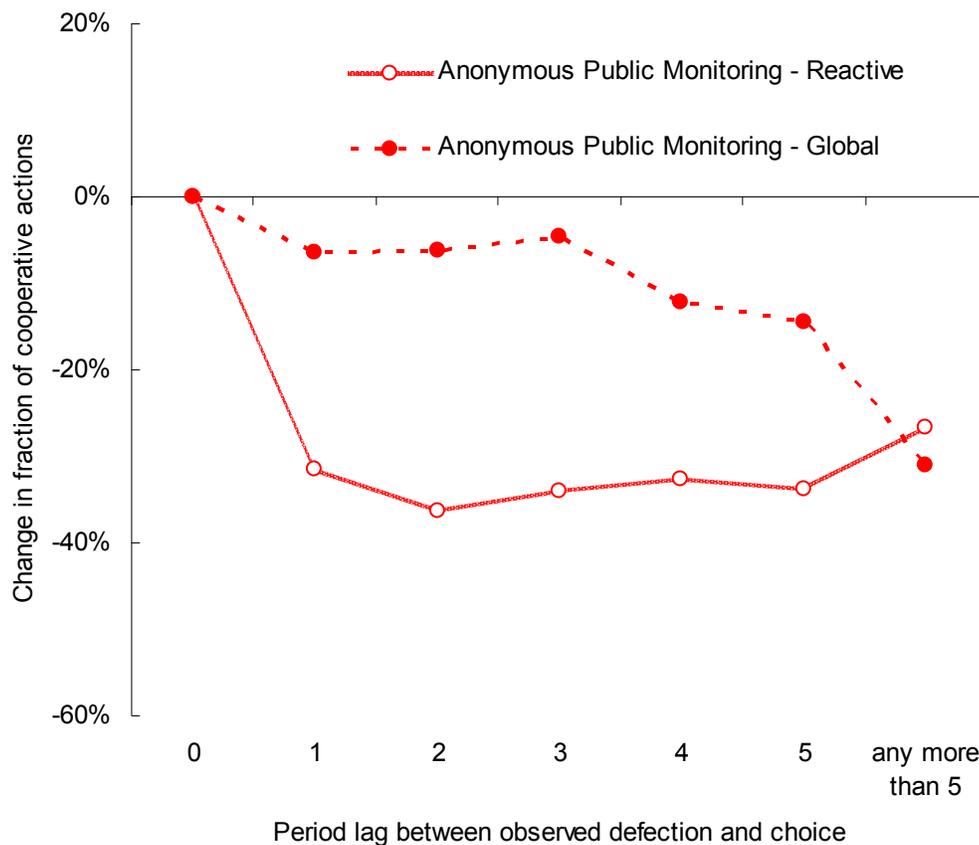


Figure 5: Strategies of the representative subject in anonymous public monitoring²⁵

Result 9. *In the non-anonymous public monitoring treatment, the representative subject selected targeted strategies over reactive and global strategies.*

In non-anonymous public monitoring subjects observed all individual histories. In the experiment, a defection by an opponent generated a strong response in future encounters with the

²⁵ The two lines overlap for periods “any more than 5” because of how reactive and global strategy regressors are defined (see Figure 1).

same opponent. However, defections outside the match were largely ignored. This conclusion is based on the estimated coefficients for targeted, reactive and global strategies. These three strategies were all available with non-anonymous public monitoring (Table 4 and Figure 6). Recall that a subject using a targeted strategy punished only opponents who defected in previous encounters but cooperated with everyone else, even if they defected with someone else.

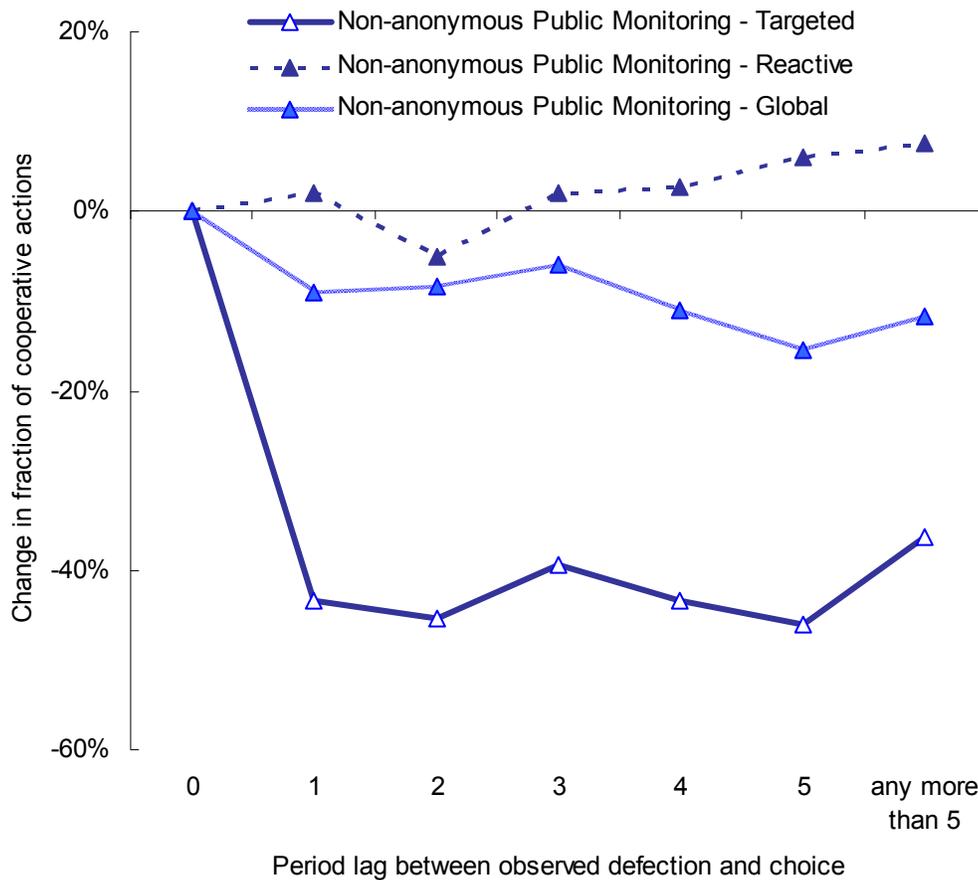


Figure 6: Strategies of the representative subject in non-anonymous public monitoring

Figure 6 reports the marginal effects estimated using regressions in Table 4.²⁶ In addition to what has already been discussed in relation to Figures 4 and 5, the (cooperation) choices for non-anonymous public monitoring include six additional strategy regressors, which we used to trace targeted strategies. Figure 6 is interpreted as follows. The dark solid line indicates that a subject who experienced a defection displayed a strong and persistent decrease in cooperation levels when future encounters involved the same opponent. In contrast, the light solid and the dashed lines in Figure 6 reveal that there is little support for the use of either reactive or global strategies. We draw the following lesson: individual-specific information appears to be much more effective than aggregate information in promoting cooperation.

Result 10. *In all treatments a defection of an opponent triggered a persistent decrease in cooperation, and the representative subject did not 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.²⁷ Regression results from Table 4 allow us to detect if such type of strategies were actually employed. In all treatments, including economies with public monitoring, the defection of an opponent triggered a persistent decrease in cooperation with very little reversion to a cooperative mode. If the representative subject employed a T-period trigger strategy, one should detect a U-shape pattern and not an L-shape pattern in the marginal effect

²⁶ Figure 6 uses the coefficient estimates coding targeted, reactive and global strategies, respectively. Marginal effects for targeted strategies were computed for the average values of reactive and global strategies regressors. Marginal effects for reactive strategies were computed for the average values of targeted and global strategies regressors. Marginal effects for global strategies were computed for the average values of targeted and reactive strategies regressors.

²⁷ At the end of each period, everyone observes the same random draw concerning the continuation of the cycle. So, subjects could coordinate a reversion to cooperation using that random draw, even with private monitoring.

curves of Figures 4-6. Instead, after an initial drop, the curves look generally flat, and no recovery to pre-defection cooperation levels after five periods can be detected.²⁸

V. 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). We find that strangers can achieve remarkably high rates of cooperation in small economies, and cooperation increases with experience. This empirical finding is novel. The theoretical work of Kandori (1992) and Ellison (1994) ensures that full cooperation is an equilibrium in our experimental design, but previous studies of similar indefinitely repeated prisoner's dilemmas report that subjects selected a different equilibrium. Yet, our experimental results suggest that these theories of cooperation among strangers seem to lack some fundamental element to describe human behavior, because subjects appear to strongly focus over some classes of the strategies to support the cooperative equilibrium.

We have 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

²⁸ Wald tests reveal lag regressors are often jointly significant at 5 percent level (except in private monitoring with punishment, and global strategies in non-anonymous public monitoring), which suggests the use of other strategies in addition to a permanent punishment. We did not expand on this because their magnitude is small.

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. Unless histories could be traced back to a specific individual, then 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, cooperation was considerably higher when details about identities were added to this aggregate information (non-anonymous public monitoring). The lesson that we draw is that information must be linked to a particular individual, in order to have an effect on cooperation. This result suggests that reputation-tracking institutions, such as personal credit history in financial markets, play an important role in sustaining compliance without relying frequently or exclusively on costly enforcement institutions, such as courts of law. 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. However, it is important to realize that though adding either individual histories or personal punishment increased cooperation to similar levels, the use of personal punishment generates a deadweight loss.

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 focus on strategies that are *selective* in punishment (i.e., strategies that narrow 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, which of course comes at a higher efficiency cost.²⁹ 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 put forward several possible reasons for the frequent use of some classes of strategies. First, subjects may have other-regarding preferences. Indeed, there is an experimental literature that validates this conjecture and several models of other-regarding preferences exist that alternatively focus on: altruism, inequality aversion or reciprocity (see Sobel, 2005 for a review). Subjects with a reciprocity or “punishment that fits the crime” norm, for instance, may prefer punishment schemes that decrease 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.

²⁹ 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.

Second, subjects may prefer simpler strategies because of cognitive costs. The results reported 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. However, 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. 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 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. In fact, subjects using a reactive strategy must punish everyone in order to eventually punish the defector. 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., Russell Cooper, Douglas DeJong, and Robert Forsythe, 1996, Crawford, 1998, John B. Van Huyck, John M. Wildenthal, and Raymond C. Battalio, 2002).

A tentative conclusion is 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.

References

- Aliprantis, Charalambos D., Gabriele Camera, and Daniela Puzzello.** 2006. "Matching and Anonymity." *Economic Theory*, 29: 415-432.
- Aliprantis, Charalambos D., Gabriele Camera, and Daniela Puzzello.** 2007. "Anonymous Markets and Monetary Trading." *Journal of Monetary Economics*, 54(7): 1905–1928.
- Andreoni, Jim, and Rachel Croson.** 2002. "Partners versus strangers: The Effect of Random Rematching in Public Goods Experiments." In *Handbook of Experimental Economics Results* ed. Charles Plott and Vernon Smith.
- Aoyagi, Masaki, and Guillaume Frechette.** 2005. "Collusion as Public Monitoring Becomes Noisy: Experimental Evidence." Unpublished.
- Axelrod, Robert.** 1984. *The evolution of cooperation*. New York: Basic Books.
- Bohnet, Iris and Steffen Huck.** 2004. "Repetition and Reputation: Implications for Trust and Trustworthiness When Institutions Change." *American Economic Review Papers and Proceedings*: 362-366.
- Camera, Gabriele and Marco Casari.** 2007. "Cooperation Among Strangers: An Experiment with Indefinite Interaction," Purdue University, Krannert Working Paper no. 1201.
- Camerer, Colin and Keith K. Weigelt.** 1988. "Experimental Tests of the Sequential Equilibrium Reputation Model." *Econometrica* 56: 1-36.
- Casari, Marco and Luini, Luigi.** 2007. "Group Cooperation under Alternative Peer Punishment Technologies: An Experiment." *Journal of Economic Behavior and Organization*. Forthcoming.
- Casari, Marco and Charles Plott.** 2003. "Decentralized Management of Common Property Resources: Experiments with a Centuries-Old Institution." *Journal of Economic Behavior and Organization* 51(2): 217-247.
- Cason, Tim N., and Feisal U. Khan.** 1999. "A Laboratory Study of Voluntary Public Goods Provision with Imperfect Monitoring and Communication." *Journal of Development Economics* 58: 533-552.
- Cooper, Russell, Douglas DeJong, and Robert Forsythe.** 1996. "Cooperation without Reputation: Experimental Evidence from Prisoner's Dilemma Games." *Games and Economic Behavior*, 12: 187-218.
- Costa- Gomes, Miguel, Vincent Crawford, and Bruno Broseta.** 2001. "Cognition and Behavior in Normal-Form Games: An Experimental Study." *Econometrica*, 69: 1193-1235.
- Crawford, Vincent.** 1998. "A Survey of Experiments on Communication via Cheap Talk." *Journal of Economic Theory* 78: 286-298.
- Dal Bó, Pedro.** 2005. "Cooperation under the Shadow of the Future: Experimental Evidence from Infinitely Repeated Games." *American Economic Review*, 95(5): 1591-1604.
- Dal Bó, Pedro, and Guillaume Frechette.** 2006. "The Evolution of Cooperation in Infinitely Repeated Games: Experimental Evidence." Unpublished.
- Diamond, Peter.** 1982. "Aggregate Demand Management in Search Equilibrium." *Journal of Political Economy*, 90: 881-894.
- Dixit, Avinash.** 2003. "On Modes of Economic Governance." *Econometrica*, 71(2): 449-481.
- Duffy, John, and Jack Ochs.** 1999. "Emergence of Money as a Medium of Exchange: An Experimental Study." *American Economic Review*, 89(4): 847-877.

- Duffy, John, and Jack Ochs.** 2006. "Cooperative Behavior and the Frequency of Social Interaction." Unpublished.
- Ellison, Glenn.** 1994. "Cooperation in the Prisoner's Dilemma with Anonymous Random Matching." *Review of Economic Studies*, 61: 567-88.
- Engle-Warnick, Jim.** 2007. "Five Indefinitely Repeated Games in the Laboratory." CIRANO, Working Paper 2007s-11.
- Engle-Warnick, Jim, and Robert L. Slonim.** 2004. "The Evolution of Strategies in a Repeated Trust Game." *Journal of Economic Behavior and Organization*, 55(4): 553-573.
- Engle-Warnick, Jim, and Robert L. Slonim.** 2006. "Learning to Trust in Indefinitely Repeated Games." *Games and Economic Behavior*, 54 (1): 95-114.
- Engle-Warnick, Jim, and Robert L. Slonim.** 2006. "Inferring Repeated-Game Strategies from Actions: Evidence from Trust Game Experiments." *Economic Theory*, 28: 603-632.
- Fehr, Ernst, and Simon Gaechter.** 2000. "Cooperation and Punishment in Public Goods Experiments." *American Economic Review*, 90(4): 980-994.
- Friedman, James W.** 1971. "Non-Cooperative Equilibrium for Supergames." *Review of Economic Studies*, 38(113): 1-12.
- Healy, Paul J.** 2007. "Group Reputations, Stereotypes, and Cooperation in a Repeated Labor Market." *American Economic Review*, 97(5): 1751-1773.
- Kandori, Michihiro.** 1992. "Social Norms and Community Enforcement." *Review of Economic Studies*, 59: 63-80.
- Kiyotaki, Nobuhiro, and Randall Wright.** 1989. "On Money as a Medium of Exchange." *Journal of Political Economy*, 97: 927-954.
- Krasa, Stefan, and Anne Villamil.** 2000. "Optimal Contracts When Enforcement Is a Decision Variable." *Econometrica*, 68(1): 119-138.
- Levine, David K., and Wolfgang Pesendorfer.** 1995. "When Are Agents Negligible?" *American Economic Review*, 85: 1160-1170.
- Milgrom, Paul R., Douglass C. North, and Barry R. Weingast.** 1990. "The Role of Institutions in the Revival of Trade: The Law Merchant, Private Judges, and the Champagne Fairs." *Economics and Politics*, 2: 1-23.
- Mortensen, Dale and Christopher Pissarides.** 1994. "Job Creation and Job Destruction in the Theory of Unemployment." *Review of Economic Studies*, 61: 397-415.
- Murnighan, Keith, and Alvin E. Roth.** 1983. "Expecting Continued Play In Prisoners-Dilemma Games – A Test Of Several Models." *Journal of Conflict Resolution*, 27(2): 279-300.
- Normann, Hans-Theo, and Brian Wallace.** 2006. "The Impact of the Termination Rule on Cooperation in a Prisoner's Dilemma Experiment." Unpublished.
- Ostrom, Elinor, James Walker, and Roy Gardner.** 1992. "Covenants With and Without a Sword: Self Governance Is Possible." *American Political Science Review*, 86: 404-417.
- Palfrey, Tom R., and Howard Rosenthal.** 1994. "Repeated Play, Cooperation and Coordination: An Experimental Study." *Review of Economic Studies*, 61: 545-565.
- Rapoport, Anatol, and Albert Chammah.** 1965. *Prisoner's dilemma: A study in conflict and cooperation*. Ann Arbor: The University of Michigan Press.

- Roth, Alvin E., and Keith Murnighan.** 1978. "Equilibrium Behavior and Repeated Play of The Prisoner's Dilemma." *Journal of Mathematical Psychology*, 17: 189-98.
- Rustichini, Aldo and Anne Villamil.** 2000. "Intertemporal Pricing in Laboratory Posted Offer Markets with Differential Information." *Economic Theory*, 16: 613-637.
- Schwartz, Steven, Richard Young, and Kristina Zvinakis.** 2000. "Reputation without Repeated Interaction: A Role for Public Disclosures." *Review of Accounting Studies*, 5: 351-375.
- Sobel, Joel.** 2005. "Interdependent Preferences and Reciprocity." *Journal of Economic Literature*, 43: 392-436.
- Van Huyck, John B., John M. Wildenthal, and Raymond C. Battalio.** 2002. "Tacit Cooperation, Strategic Uncertainty, and Coordination Failure: Evidence from Repeated Dominance Solvable Games." *Games and Economic Behavior*, 38: 156-175.