
3-18-2017

Monitoring Institutions in Indefinitely Repeated Games


Gabriele Camera

Chapman University, camera@chapman.edu

Marco Casali

University of Bologna

Follow this and additional works at: https://digitalcommons.chapman.edu/esi_working_papers

 Part of the [Econometrics Commons](#), [Economic Theory Commons](#), and the [Other Economics Commons](#)

Recommended Citation

Camera, G., & Casali, M. (2017, Mar 18). Monitoring institutions in indefinitely repeated games. ESI Working Papers 17-12. Retrieved from http://digitalcommons.chapman.edu/esi_working_papers/224

This Article is brought to you for free and open access by the Economic Science Institute at Chapman University Digital Commons. It has been accepted for inclusion in ESI Working Papers by an authorized administrator of Chapman University Digital Commons. For more information, please contact laughtin@chapman.edu.

Monitoring Institutions in Indefinitely Repeated Games

Comments

Working Paper 17-12

Monitoring institutions in indefinitely repeated games*

Gabriele Camera
Chapman University
& University of Basel

and

Marco Casari
University of Bologna

18 March 2017

Abstract

Does monitoring past conduct facilitate intertemporal cooperation? We designed an experiment characterized by strategic uncertainty and multiple equilibria where coordinating on the efficient outcome is a challenge. Participants, interacting anonymously in a group, could pay a cost either to obtain information about their counterparts, or to create a freely available public record of individual conduct. Both monitoring institutions were actively employed. However, groups were unable to attain higher levels of cooperation compared to a treatment without monitoring. Information about past conduct alone thus appears to be ineffective in overcoming coordination challenges.

Keywords: coordination, information, equilibrium selection, conventions, social dilemmas
JEL codes: C70, C90, D80

1. Introduction

This paper presents an experimental study of monitoring institutions that are designed to facilitate long-run cooperation among strangers. In the experiment, efficiency is theoretically possible but is difficult to achieve in practice because interactions are anonymous, there are multiple equilibria and strategic uncertainty (Van Huyck et al., 1991).

To ameliorate these difficulties, we introduce means of monitoring past conduct and we

* We thank two referees for helpful comments, M. Bigoni for outstanding research assistance, N. Liu and T. Le for software coding, as well as seminar participants at several universities and conferences for comments on an earlier version of this paper. G. Camera acknowledges research support from the NSF grant DMS-0437210. Correspondence address: G. Camera, Economic Science Institute, Chapman University, Orange, CA 92866 (USA); camera@chapman.edu. M. Casari (marco.casari@unibo.it), University of Bologna, Department of Economics, Piazza Scaravilli 2, 40126 Bologna (Italy).

investigate whether or not this can help groups coordinating on the cooperative equilibrium. We consider two different monitoring institutions. In the Information Request treatment, participants can pay a small cost to receive accurate information about the counterpart; in the Information Provision treatment, participants can pay a small cost to help create an accurate public record of individual conduct. The Information Provision treatment introduces a monitoring institution that resembles a Better Business Bureau and, historically, it echoes the medieval Law Merchant institution for long-distance trade (Milgrom et al., 1990). The Information Request treatment introduces a monitoring institution similar to a credit bureau. By design, neither of these institutions expands the set of equilibrium payoffs, i.e., full cooperation is self-sustaining in *all* treatments through the aforementioned general convention. However, accurate monitoring can support conditional cooperation and, more generally, a variety of reciprocity-based strategies.

To model long-run cooperation, we study a helping game that is played in a stable group of four individuals (Camera et al., 2013a) and is indefinitely repeated (Palfrey and Rosenthal, 1994). In the group, participants interact in pairs with counterparts that change at random in each round of play. Relying on relational contracting is ruled out by design because participants' identities always remain hidden and participants cannot communicate. In every round, two pairs are formed, both with one buyer and one seller. The seller can either consume a good in her possession or transfer it to the buyer. Transferring the good is socially efficient because the buyer values it more than the seller does. However, the seller has an incentive to behave opportunistically and consume the good. A general convention of gift-exchange among group participants can sustain the efficient outcome in the long run, provided that it incorporates a decentralized punishment scheme capable of deterring defections. Theory suggests that if such convention is based on the threat of a permanent and irreversible switch from cooperation to

defection, then this is sufficient to remove opportunistic temptations in a group of homogeneous individuals (Kandori, 1992, Ellison, 1994).

Within this framework, full cooperation is self-sustaining and formal institutions that support communication, monitoring, punishment, etc. are typically considered theoretically unnecessary to reach efficiency. However, there is strategic uncertainty because individuals are unsure which equilibrium strategy others will use (Van Huyck et al., 1991, p. 234). It is interesting to note that a typical assumption in applications of the theory of infinitely repeated games is that individuals successfully coordinate on the best outcome. However, previous experiments have shown that the fact that cooperation is an equilibrium is not sufficient for cooperation to arise among strangers (Camera and Casari, 2009, 2014) or among partners (Dal Bó and Fréchette, 2011).

In our experiment, the availability of monitoring in the form of Information Request and Information Provision did not facilitate the successful coordination on long-run cooperation. Cooperative conventions struggled to emerge despite the heavy use of the monitoring institutions, something that in principle could have simplified coordination tasks. In fact, the data reveal that when participants had the option to contribute to creating a public record of past conduct, groups cooperated even less than in the absence of any monitoring institution. This evidence suggests that to overcome equilibrium coordination challenges, groups of strangers need institutions that are complementary to monitoring, such as enforcement institutions, for example.

The structure of the paper is the following. Section 2 describes the experimental design. Section 3 offers some theoretical considerations. Section 4 reports the results, Section 5 discusses related studies, and Section 6 offers some final considerations.

2. The design of the experiment

The experimental design is related to the model adopted in Camera and Casari (2014), where players interact as strangers, and can engage in an inter-temporal giving and receiving of goods. There are three treatments described in Table 1.

In all treatments, interaction consists of the following helping game. There is a seller who can consume a good in her possession or transfer it to a buyer who values it comparatively more. The seller can either consume the good (=defection) or transfer it (=cooperation), while the buyer has no action to take. The payoffs for seller and buyer are, respectively, (a, a) with $a > 0$ under defection, and (d, u) under cooperation, with $d \in (0, a)$ and $u > 2a - d$. In the experiment $d=2$, $a=8$, $u=20$. Although defection is a dominant action for the seller, cooperation maximizes total surplus to in the pair (6 points=22-16 points).

This helping game has some aspects in common with a two-stage *sequential* prisoner's dilemma (Clark and Sefton, 2001) where the player who moves in the first stage can be interpreted as the producer while the other is the consumer who moves in a second stage (when the players' roles are reversed). In that kind of sequential game, payoffs would be $u+d$ to each player under full cooperation; $2a$ to each player under full defection; instead, the initial producer would earn $d+a$ if she cooperates and the initial consumer later defects thus earning $u+a$. Using a helping game instead of a Prisoner's Dilemma with simultaneous moves focuses the subjects' attention on the intertemporal dimension of cooperation (more in Section 5).

In the experiment, players interact within a four-player group for an indefinite number of rounds. In each round, first two pairs are formed at random. Then, roles are randomly assigned so that each pair has one buyer and one seller. In this indefinitely repeated game, the *efficient* outcome is attained when cooperation occurs in every pair and in every round until the game

stops (= full cooperation).

We implement indefinite repetition following the technique in Roth and Murnighan (1978). A supergame or *cycle* is characterized by a random continuation rule. The rule specifies that at the end of each round an additional round of play takes place with probability $\delta = 0.93$. The duration of the interaction is therefore always uncertain. However, in each round the supergame is expected to go on for approximately 13 additional rounds, i.e., $1/(1-\delta)-1$, no matter how many rounds have been played. The continuation probability δ is interpreted as the discount factor of a risk-neutral participant. In our experiment, a computer randomly selected an integer between 1 and 100 from a uniform distribution, and the supergame terminated for all session participants, when the drawn number exceeded 93.

Each experimental session employed twenty participants who played five consecutive cycles. This amounts to twenty-five groups per session. We adopted a pre-determined matching protocol to arrange participants into groups across the five cycles. This was done to ensure no one could interact with anyone else for more than one cycle. In each cycle, participants were paired exclusively within their group. Each group included only participants who neither belonged to the same group in a past cycle, nor would belong to the same group in a subsequent cycle. Participants were informed about this matching protocol. Cycles terminated simultaneously for all groups.

Group participants were matched at random in each round of every cycle, and could never identify their counterpart. Each participant could meet one of the three other persons in her group, with equal probability. Hence, each participant had one-third probability of meeting any other group member in each round of a cycle. Once pairs were formed, in each pair a computer-determined coin flip assigned a seller role to one player, and a buyer role to the other. Hence, in

each round every group comprised two buyers and two sellers, and participants were equally likely to change or keep their role in the following round. Participants could not observe outcomes outside of their pair (private monitoring). This completes the description of what we call the Baseline treatment.

In the two additional treatments, called Information Provision and Information Request, we modified the Baseline design by adding a prototypical monitoring institution, or technology (Table 1). Each of these additional treatments introduces the possibility for participants to build a reputation through the creation of individual records while preserving anonymity of interaction. One institution (*Information Provision*) has function similar to the Better Business Bureau, where a buyer’s record is a public good and the seller can freely view it. The other institution (*Information Request*) is more similar to a credit agency, where the buyer’s record is a private good and the seller must pay to view it. Neither treatment can offer as much information as would be available with public monitoring.

Table 1: Experimental treatments

	<i>Baseline</i>	<i>Information Provision</i>	<i>Information Request</i>
Monitoring	No	Yes: buyer pays 1 point to report counterpart’s action	Yes: seller pays 1 point to see counterpart’s record
Information to seller	No information about counterpart’s past conduct	Summary of counterpart’s actions reported by buyers	Summary of counterpart’s actions taken in the past
Average no. of rounds	51.6	68.5	85.5

Notes: Seller was called Red in the experiment and buyer was called Blue. Conversion rate: 1 point = \$0.025. The sessions were run in Sep-Nov 2008. Two sessions for each treatment: one session was run at Purdue University and one at the University of Iowa.

Information Provision treatment. This treatment adds a post-exchange stage along the lines of the decentralized model of monitoring developed in Milgrom et al. (1990). After observing the

outcome of the helping game, the buyer can pay 1 point to report truthfully the seller's action. This information is added to her opponent's *record*, which is empty in round 1 of a cycle. Alternatively, the buyer can choose not to make a report. The seller never sees the buyer's choice in this post-exchange stage. In any given round of a cycle, the record of a participant spans (at most) the six preceding rounds in that same cycle.¹ The record excludes the participant's identity and displays a summary of her history that is based on voluntary reports. It includes the number of past rounds in which: the participant was a seller, her action was not reported, and her reported action was cooperation or defection. Before making a choice, the seller can review at no cost the record of the buyer and her own record. The buyer does not observe any record. Since records are anonymous (identities are excluded from records), random matching implies that sellers cannot directly identify a past opponent by simply looking at a record. Possible payoffs and outcomes are the same as in Baseline, with the exception that payoffs for buyers include the loss of 1 point if they report the seller's action. If no action is ever reported, then the Information Provision treatment is observationally equivalent to the Baseline treatment.

Information Request treatment. This treatment adds a pre-exchange stage. Before making a choice in the helping game, the seller can pay 1 point to view the buyer's record. Alternatively, the seller can choose not to view the buyer's record. The buyer never sees the seller's action in this pre-exchange stage. As in Information Provision, the participant's record spans (at most) the six preceding rounds in that cycle and does not include the participant's identity. Unlike Information Provision, the record is a summary based on an accurate and complete history of the roles and actions taken by the participant in those six previous rounds. The record displays the

¹ The expected duration of a cycle was about 14 rounds. Limiting records to 6 rounds allowed for some learning by insuring that an initial mistake would not permanently stain the reputation of a participant. In this sense, the design allows for a "fresh-start," which gives players a chance to re-coordinate on cooperation after experimenting with defection choices. Another experiment with a similar stage game (Bigoni et al., 2014) studies a monitoring system that is costless, and that keeps track of the entire history of the subject during the cycle.

number of rounds in which: the participant was a seller, and the number of cooperative and defection actions taken as a seller. Possible payoffs and outcomes are as in Baseline, with the exception that the seller pays 1 point to view the buyer's record.²

The Information Provision and Information Request treatments exhibit elements of commonality. First, a participant's record includes only information about her past actions as a seller, which helps in building a reputation but does not necessarily reveal all past outcomes (e.g., outcomes experienced as a buyer). This means that a participant's record cannot reveal a past defection unless that participant was the seller in *that* round and *she* defected. Second, a seller can only view the record of the buyer she is currently matched to, and the record does not reveal the buyer's identity (anonymity). Lastly, a participant's record includes neither the history of *his* opponents, nor the histories that the participant observed. For example, the record does not say if participants defected after observing a defection.

Considering all treatments, we recruited 120 participants through announcements in undergraduate classes, half at Purdue University and half at the University of Iowa. The experiment was programmed and conducted with the software z-Tree (Fischbacher, 2007). Given the random termination protocol, we recruited participants for three hours. Instructions were read aloud at the start of the experiment and left on the participants' desks (a copy is in the Appendix). No eye contact was possible among participants. Average earnings were about \$17 per participant. On average, a session lasted 71 rounds for a running time of 2 hours, including instruction reading and a quiz. Details about the sessions are in Table 1.

² The Provision and Request treatments can be interpreted as introducing an institution that processes, respectively, the information truthfully provided by individuals in the group and all the available information; see Kandori (1992) for a similar interpretation. The institution marks individuals who have defected and the mark is publicly observable at no cost in one case (Provision), but not in the other (Request).

3. Theoretical predictions

Here, we demonstrate that—under the assumption of identical players that are self-interested and risk-neutral—the equilibrium set consists of multiple equilibria *in all treatments*. The equilibrium set includes full defection and full cooperation, which is the efficient outcome. The following analysis is based on the folk theorem-like results established in Kandori (1992) and Ellison (1994).

Start by observing that defection is the dominant strategy if the helping game is played once or if it is finitely repeated. In an indefinitely repeated game, the payoff is the (ex-ante) expected discounted stream of payoffs attained in each of the one-shot interactions. In the *Baseline* treatment participants can neither observe the outcome in other pairs (private monitoring), nor can they identify or communicate with their counterparts (anonymity). In this scenario, “always defect” is always a sequential equilibrium because defect is a best response to everyone else defecting in every round in which they are sellers. In this case, the equilibrium payoff in the repeated game corresponds to the present discounted value $a/(1-\delta)$.

We next prove that the efficient outcome can also be sustained as a sequential equilibrium, as long as δ is sufficiently large. To prove it, following Kandori (1992) and Ellison (1994), we conjecture that *all* players adopt a rule of behavior called “grim trigger strategy”, consisting of a “desirable” action as sellers and of a sanction that is implemented as soon as the player finds out that some seller acted in an undesirable manner. We identify the desirable action with cooperation and the sanction with defection. The grim trigger strategy thus stipulates that the player should cooperate whenever she is a seller, as long as she has never experienced a defection; otherwise, she should always defect and never cooperate again.

This strategy relies on a form of community (or, decentralized) punishment that is used to

police defections. Community punishment implies an incremental defection process in the group. The first seller who defects triggers an unstoppable punishment process, which eventually leads to full defection. This threat is precisely what supports cooperation in the model.

Proposition 1. *In the indefinitely repeated helping game if $\delta > 0.808$, then a social norm based on the grim trigger strategy supports the efficient outcome as a sequential equilibrium.*

The sketch of the proof is as follows.³ In each round payoffs for (seller, buyer) are (u, d) if cooperation is the outcome, and (a, a) otherwise. If everyone adopts the grim trigger strategy, i.e., if such strategy is a social norm, then in equilibrium every seller cooperates so the payoff to a representative player is the expected discounted utility from buying or selling with equal probability, $(u+d)/[2(1-\delta)]$. One needs to check two incentive compatibility elements: first, the “grim” punishment threat must remove a seller’s temptation to defect in any equilibrium round; second, it must also ensure that a seller has an incentive to follow the punishment norm, off-equilibrium. In doing so, we rely on the unimprovability criterion and consider one-time deviations by a single seller.

Since surplus is lost if the group moves from full cooperation to full defection, then a one-time equilibrium defection is suboptimal as long as players foresee a sufficiently long interaction. Equivalently—as reported in Proposition 1—players must be sufficiently patient. Intuitively, the future reward from cooperating today must be greater than the extra utility provided by defecting today. If grim trigger is a social norm, then the initial defection will quickly lead to 100% defections because there are only four players in our groups. The rapid spread of punishment to the entire group also explains why sellers have no incentive to cooperate

³ Details of the proof are available in the Supplementary Materials to Camera and Casari (2014)

after observing a defection, such as in an attempt to stop the contagious punishment process in its tracks.

Proposition 2. *In the Baseline treatment, the equilibrium set includes full defection and the efficient outcome. In the Provision and Request treatments, the addition of monitoring institutions neither eliminates any of the equilibria available in the Baseline treatment, nor expands the efficiency frontier.*

To prove the first part of the statement, note that, due to indefinite repetition, if all participants play grim trigger, then the efficient outcome can be sustained as a sequential equilibrium. Given that in the experimental design the continuation probability is $\delta=0.93$, then according to Proposition 1 the efficient outcome is an equilibrium in every treatment and in every group. As noted above, full defection is also an equilibrium because defection in a round is always a best response if every other seller is expected to defect in every round.

The second part of the statement immediately follows because monitoring can always be ignored. Simply put, none of the strategies available in the *Baseline* treatment condition on information about counterparts, and all these strategies remain available in all other treatments. The central consequence is that the efficient outcome can be supported in all treatments and therefore the Information Provision and Information Request treatments do not expand the set of equilibrium payoffs.

It should be clear that additional strategies are available when participants can monitor past conduct, compared to when they cannot. Indeed, the monitoring institutions we introduced allow sellers to have accurate information of the counterpart's past behavior. In Information Request, a seller can always view an accurate record of the actions her counterpart took as a seller in the last six rounds. In Information Provision, accurate records can be created if buyers choose to report

the seller's action to a public repository. Although the additional strategies made possible through monitoring do not introduce Pareto-superior equilibria, they do expand the equilibrium set and therefore they might actually increase strategic uncertainty and coordination problems relative to Baseline. Yet, the expanded strategy set in Information Provision and Information Request treatments *neither constrains* participants to employ strategies that condition on the counterpart's past conduct, *nor precludes* the use of social norms based on decentralized enforcement.

Given this expanded strategy set, it is meaningful to quantify the efficiency that can be theoretically achieved through monitoring. We define the *efficiency loss* as 100% minus the realized surplus over the maximum surplus. Every strategy that uses monitoring generates a deadweight loss that lowers the efficiency frontier for any cooperation level achieved. In Information Provision and Information Request, the long-run efficiency loss from creating and viewing opponents' records is below 16.7%. The maximum loss occurs when all participants report and view the actions of all opponents, which costs 2 points out of a maximum surplus of 12 in each group.⁴

To summarize, the monitoring technologies considered in this study cannot expand the theoretical efficiency frontier relative to Baseline. In fact, their use would simply lower the theoretical efficiency frontier. However, we know from previous work that coordination on strategies that support the efficient outcome in groups of strangers is difficult because of the multiplicity of equilibria, including the inefficient outcome (Camera and Casari, 2014). It is therefore an open question whether or not the introduction of a monitoring institution (i) may alter cooperation rates relative to Baseline, or (ii) may facilitate coordination on efficient play.


⁴ With information provision, a buyer could report only the first defection observed and still generate an accurate record. Here we do not characterize the optimal strategy for providing information and for requesting information because it is beyond the scope of this study.

4. Results

There are four key results: Result 1 is on cooperation levels, Result 2 compares the strategies employed across treatments, Result 3 is about the deadweight loss from using the monitoring institutions, while Result 4 concerns the distribution of earnings. In this section, the empirical analysis adopts as unit of observation a group of four participants interacting in a cycle.⁵

Result 1. *In all treatments, the average cooperation rate was well below 100%. In Information Provision and Information Request, rates were similar or lower than in Baseline.*

Support for this result comes from Figure 1, Figure 2 and Tables 2, 3 and 4.



⁵ Data from the Baseline treatment are also analyzed in Camera and Casari (2014).





Figure 1. Average cooperation in a period, by cycle

Figure 1 gives an overview of average cooperation across the five cycles of each treatment (one observation is one group in a period of cycle, N=50 per treatment). Table 2 aggregates these data

at the cycle level and then at the treatment level. Overall, average cooperation in the Information Provision treatment was 37.5%, which is 10.7% lower than Baseline (Mann-Whitney test, p-values=0.031, n1=50, n2=50), and 11.9% lower than Information Request (Mann-Whitney test, p-values=0.026, n1=50, n2=50). Average cooperation in Information Request is not significantly different from Baseline (Mann-Whitney test, p-value=0.97, n1=50, n2=50). These observations confirm that cooperation in this environment is difficult to support even when participants have a monitoring technology at their disposal.




Figure 2. The distribution of average cooperation across groups

To gain more insight, we calculated the distributions of cooperation rates across groups, by treatment (see Figure 2; one observation is one group in a cycle, N=50 per treatment). In all treatments, few groups achieved very high and very low cooperation rates. However, the distribution in Information Provision is stochastically dominated by the distributions in Baseline and in Information Request. This evidence suggests that the higher cooperation rates in Baseline and Information Request are not due simply to some groups being very cooperative, but rather to

a generic increase in cooperativeness.⁶

When we focus on average cooperation in round 1 of each cycle, we find that the highest cooperation rate is 51.0%; see Table 3. Behavior in round 1 can reveal the existence of a focal point but we find little evidence of a focal point. Monitoring did not improve the ability of groups to coordinate on either cooperation or defection. The rates of coordination on either action in round 1 were 58% in Baseline, 30% in Information Provision, and 46% in Information Request. The treatment ranking in terms of coordination on cooperation or cooperation levels is identical (Table 3).

Table 2: Average cooperation rate in a group: all rounds

Cycle	Baseline	Information Provision	Information Request
1	0.475	0.410	0.387
2	0.441	0.342	0.526
3	0.563	0.448	0.485
4	0.487	0.425	0.543
5	0.446	0.249	0.530
Overall cooperation frequency	0.482	0.375	0.494
Net surplus (points)	5.78	4.01	5.69
Gross surplus (points)	5.78	4.50	5.93
Max theoretical surplus (points)	12	12	12

Notes: 1 obs. = 1 group (10 obs. per cycle, per treatment). Gross surplus in Information Provision and Information Request is the average cooperation rate multiplied by the maximum surplus (12 points). Net surplus is gross surplus minus the cost of the institution, i.e., 2 points multiplied by the frequency with which buyer (in Information Provision) or seller (in Information Request) used the institution.

⁶ Epps and Singleton tests reveal that the distributions in Baseline and Information Provisions are statistically different (Baseline vs. Information Request, p-value=0.513; Baseline vs. Information Provision, p-value=0.095; Information Provision vs. Information Request, p-value=0.167; n1=n2=50 for each pairwise comparison).

Table 3: Average cooperation rate in a group: round 1 of each cycle

Cycle	Baseline	Information Provision	Information Request
1	0.40	0.40	0.20
2	0.30	0.45	0.45
3	0.55	0.50	0.70
4	0.55	0.40	0.60
5	0.75	0.50	0.60
Overall frequency of cooperation	0.51	0.45	0.51
Fraction of groups with 100% cooperation	0.30	0.10	0.24
Fraction of groups with 100% defection	0.28	0.20	0.22

Notes: 1 obs. = 1 group (10 obs. per cycle, per treatment).

The probit regressions in Table 4 supply additional evidence about the poor performance of the monitoring treatments in terms of cooperation rates. The dependent variable is a seller's choice to cooperate (1) or not (0) in a round. When pooling all observations, cooperation in Information Provision is significantly lower than in Baseline (the marginal effect is of -9.8 percentage points, p-value 0.013, see column 2, Table 4). By contrast, cooperation in Information Request is statistically indistinguishable from Baseline (the marginal effect is of 0.1 percentage points, p-value 0.967, see column 2, Table 4). The regressions include controls for fixed effects (cycles, rounds within the cycle), for demographic characteristics, including gender and major, and for the duration of the previous cycle. In addition, the regressions trace the response of the representative participant in the rounds *following* an observed defection, as we report next.

Table 4: Probit regression on the seller’s choice to cooperate – marginal effects

	All treatments rounds 1 only	All treatments all rounds	Baseline	Information Provision	Information Request
Dependent variable: 1=cooperation 0=defection	(1)	(2)	(3)	(4)	(5)
<i>Treatment dummies:</i>					
Information Provision	-0.040 (0.054)	-0.098** (0.039)			
Information Request	0.033 (0.103)	0.001 (0.035)			
<i>Seller's record in Provision has:</i>					
at least one cooperative action				0.239*** (0.003)	
at least one defection action				-0.135*** (0.003)	
<i>Buyer's record in Provision has:</i>					
at least one cooperative action				0.157*** (0.010)	
at least one defection action				-0.096** (0.045)	
<i>Buyer's record in Request has:</i>					
cooperation rate > 50%					0.415*** (0.024)
defection rate > 50%					-0.020 (0.102)
<i>Strategy coding:</i>					
grim trigger		-0.488*** (0.057)	-0.416*** (0.045)	-0.229*** (0.009)	-0.535*** (0.042)
Lag 1		0.025 (0.034)	0.102 (0.115)	0.036*** (0.011)	0.005 (0.073)
Lag 2		-0.032* (0.016)	-0.007 (0.028)	-0.008 (0.021)	-0.027*** (0.007)
Risk neutral or low Risk aversion (questionnaire)	0.062 (0.083)	0.114 (0.072)	0.091 (0.114)	0.305*** (0.044)	-0.034 (0.078)
High Risk aversion (questionnaire)	-0.208** (0.093)	-0.090 (0.087)	-0.066 (0.149)	-0.207*** (0.068)	0.031 (0.059)
Controls	Yes	Yes	Yes	Yes	Yes
Pseudo-R2	0.111	0.114	0.077	0.265	0.173
Observations	300	3950	1010	1370	1570

Notes: Each observation refers to a seller in a pair. Marginal effects are computed at the mean value of regressors. 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. Round fixed effects are included (except in the first column) but not reported in the table (rounds 2-5, 6-10, 11-17, 18-25, >25). Controls are not reported and include cycles dummies 2, 3, 4, 5, duration of previous cycle (set to 14.3 rounds for cycle 1), Iowa location, Male, Business major, and Engineering, Science, and Mathematics major. Robust standard errors (in parentheses) are computed with a cluster on each session; * significant at 10 percent; ** significant at 5 percent; *** significant at 1 percent.

Result 2. *When a monitoring institution was available, the strategy of the representative seller conditioned on both the buyer's and her own record. Moreover, in Information Request, participants increased their reliance on grim trigger, while in Information Provision they reduced it in comparison to Baseline.*

Support is provided by the regression in Table 4, which traces the response over time of an individual after suffering an initial defection, as a buyer. We consider the individual's behavior in the first two instances in which she is a seller after suffering the defection. The econometric technique we employ allows us to identify the adoption of community punishment schemes such as grim trigger or T-round punishment.⁷ We construct three variables. The grim trigger variable takes value 1 in all rounds following the first defection suffered as a buyer, and 0 otherwise. The Lag regressors pertain to specific rounds in which the participant is a seller after experiencing her first defection as a buyer. The Lag 1 variable takes value 1 on the first round of the continuation game in which she is a seller, and 0 otherwise; the Lag 2 variable takes value 1 on the second round in which she is a seller.⁸ These two regressors allow us to study how the individual behaves on the first two occasions in which she has an opportunity to react to a defection.

Theoretically, if subjects support cooperation by following the grim trigger strategy, then the probability of choosing the cooperative action should permanently drop to zero for someone who suffers a defection because that subject will punish the deviation by defecting forever after. Empirically, subjects' behavior is heterogeneous, so—in the empirical estimation based on a representative subject—we do not see this sudden and permanent drop in cooperation. However,

⁷ Every participant observed a random number between 1 and 100 at the end of each period. A number below 94 meant that the cycle would continue into a new round. Theoretically, groups could have used this random number to coordinate on reverting to cooperation from a punishment phase triggered by a defection (Ellison, 1994).

⁸ For a detailed discussion on this econometric technique applied to a Prisoner's Dilemma game see Camera and Casari (2009); see Camera and Casari (2014) for an application to the same helping game as in this study.

in all treatments, we do observe a significant and persistent decline in the probability of cooperating after someone suffers a defection. The grim trigger marginal effects are all negative and highly significant (columns 3, 4, 5 in Table 4). This evidence is consistent with a subset of participants using the grim trigger strategy, while participants in the other subset behave unconditionally. Note that this negative response is more pronounced in Information Request compared to Baseline, and less pronounced in Information Provision. Moreover, the sum of the estimated marginal effects of grim trigger and Lag2 is -0.56 for Information Request, -0.24 for Information Provision, while for Baseline we have -0.42 (Table 4).

We now study how participants modified strategies in response to the possibility of accessing the history of actions of the opponent. Introducing the possibility of monitoring past actions of counterparts leaves untouched the set of theoretical discount factors supporting the efficient outcome because participants can always abstain from using the monitoring technology. However, if knowing the past behavior of an opponent is the key to raising cooperation levels in anonymous groups, then we should observe an extensive use of the monitoring technology. In the experiment, participants actively employed monitoring. Buyers in Information Provision paid a cost to report the anonymous seller's choice in 24.5% of cases; sellers in Information Request paid a cost to inspect the record of their anonymous buyer in 12.1% of cases.

This experiment offers an opportunity to study patterns of information flows because participants actively provided and requested information. In both treatments, there is evidence that the actions of the representative participant were significantly affected by their own record as well as their opponent's. The regressions in Table 4 show that in Information Provision and Information Request a seller was significantly more willing to cooperate with a buyer who was known to have cooperated in the past (see *Buyer's record* regressors in Table 4, columns 4-5). In

addition, in Information Provision, sellers were significantly less willing to cooperate with a buyer who was known to have defected at least once (see *Buyer's record* regressor in Table 4, column 4). This behavior is consistent with motivations related to indirect reciprocity or conditional cooperation (Fehr and Gaechter, 2000, Fischbacher et al., 2001).

In Information Provision, the use of the monitoring institutions was inefficient because of a lack of coordination. Participants would gain from coordinating on reporting either cooperation or defection, but not both. In particular, given a goal to sustain full cooperation, reporting only defection actions is the less costly strategy. Instead, buyers sometimes reported cooperation and sometimes reported defection. Buyers reported cooperative actions in 33.1% of cases and reported defection in 16.0% of cases (N=1370). In sum, there is evidence that Information Provision and Information Request, which can reduce informational frictions, were ineffective in increasing cooperation relative to the Baseline treatment, where the past behavior of opponents always remained hidden. The active use of monitoring had also implications for realized surplus.

Result 3. *The empirical deadweight loss from using the monitoring institutions was 2.0% of total surplus in Information Provision and 4.1% in Information Request.*

Support for Result 3 comes from Table 2. The direct costs of monitoring were small overall. In Information Request, the cost of monitoring derives from the 1 point paid by sellers to view the buyers' record. The average cost for the group was 0.49 points out of a total surplus of 12 points in a round. In Information Provision, the cost is 0.24 points. Given the realized cooperation rate and costs of the monitoring institution, one can calculate two measures of surplus. The *net surplus* for a group corresponds to the total points earned over and above the defection payoff. The *gross surplus* is the sum of the net surplus and the cost of the monitoring institution. The

Baseline treatment achieves the highest net surplus, which is 5.78 points out of a maximum of 12 points. The Information Provision treatment attains the minimum net surplus of 4.01 (Table 2).




Figure 3. Participants’ average profits by average frequency of cooperation

Notes: Next to each data point, we report the associated percentage of observations. Only observations where participants switch roles within the cycle are included (about 170 per treatment). Earnings were adjusted to account for the uneven frequency of a participant’s buyer and seller role: we separately computed average profits as buyer and as seller and then took their arithmetic average. Figure 3 does not qualitatively change when using raw average profits.

Result 4. *The addition of a monitoring institution redistributed surplus from frequent defectors to frequent cooperators to an extent insufficient to alter the incentives to coordinate on the efficient outcome.*

Support for this result comes from the illustration of the average per-capita profits in a round in Figure 3, which can range from a minimum of 8 $((8+8)/2)$ through a maximum of 11 $((20+2)/2)$.

We divide participants into five types according to their frequency of cooperation in a cycle. In

Information Provision there is no evidence that monitoring brings about a redistribution of surplus from frequent defectors $[0, 0.2)$ to frequent cooperators $[0.8, 1]$ with respect to Baseline, while there is some evidence in the Information Request treatment. Given that the availability of monitoring allows participants to identify free riders, one would have expected monitoring to reduce payoffs of for frequent defectors and to increase payoffs for frequent cooperators in comparison to Baseline. While in Information Request, one can notice a relative increase in the payoff of frequent cooperators and a relative decrease in payoffs for frequent defectors, no such redistributive pattern is evident in Information Provision.

5. Related experimental literature

The previous experimental literature on *indefinitely* repeated games mostly adopts a Prisoner's Dilemma as stage game, and follows a partners matching protocol (Dal Bó, 2005, Dal Bó and Fréchette, 2011, Palfrey and Rosenthal, 1994). By contrast, in our experiment the stage game consists of an individual decision problem, and all interaction takes place among strangers instead of partners (Camera and Casari, 2014). We now briefly discuss these two aspects.

The stage game that we adopt exhibits distinct features from a Prisoner's Dilemma. First, it is a cognitively simpler task. Second, the outcome depends on the choice of only one player because the counterpart is passive; hence, within the stage game, beliefs about others are not relevant. Third, our game removes motivations of positive and negative reciprocation compared to a sequential Prisoner's Dilemma (Clark and Sefton, 2001), because counterparts change at random in each round and their identities are unobservable. Third, the repetition of the task brings the focus onto the dynamic structure of incentives, as cooperation through a helping game essentially amounts to an intertemporal exchange of gifts.

Our design with strangers admits multiple equilibria similarly to what happens in indefinitely repeated games among partners. A partners setting is more conducive to reaching high levels of coordination on cooperation than a strangers setting because players can more easily resolve the inherent strategic uncertainty over the course of the game (for some evidence see Camera et al., 2013a). Hence, a partners setting is less suitable to study a possible role of monitoring institutions in resolving strategic uncertainty problems because in that setting players can directly observe the past conduct of counterparts.

This study is distinct from the branch of the literature that has studied *finitely* repeated social dilemmas, because in that setting coordination problems can emerge either in the presence of behavioral types of individuals (Cox et al., 2015) or when individuals have social preferences (Chen et al., 2014). By contrast, in our design individuals are assumed identical and self-interested, and coordination problems emerge because of the indefinite horizon of interaction.

Previous work on indefinitely repeated social dilemmas has illustrated the difficulties inherent to coordinating on the efficient outcome both in large and small groups (Camera and Casari, 2009, Camera et al., 2013a). Various remedies have been investigated to overcome coordination challenges. One branch of literature considers how reputation mechanisms may promote cooperation. For instance, Stahl (2009), Van Huyck et al. (1995), Ule et al. (2009), and Bolton et al. (2005) endow laboratory participants with a ready-made, costless reputational information-sharing technology. In an indefinitely repeated game, Stahl (2009) finds that a color-coded monitoring mechanism was not always effective at improving cooperation. In the field, however, monitoring institutions often require an effort or are endogenously created by the individual, which is an aspect that we study in the experiment. In our design, monitoring institutions are costly and monitoring past conduct involves an explicit action from the decision-maker. In the

labor market game of Gërkhani et al. (2013), employers can voluntarily share information about a job candidate's trustworthiness with other employers. They find that information sharing improves the recruiting of trustworthy workers and the earnings for employers and workers.

6. Final considerations

We have studied an indefinitely repeated social dilemma among strangers, where there is a social benefit from intertemporal cooperation and participants are prevented from relying on relational contracting. The game is characterized by multiple equilibria hence there is an issue of strategic uncertainty and equilibrium selection (Van Huyck et al., 1991). Presuming that participants aim at attaining the efficient outcome, they must be capable to coordinate on a suitable convention of decentralized punishment to remove opportunistic motivations. What sets this study apart from the bulk of the experimental literature on coordination is the presence of equilibrium multiplicity in the supergame and not in the stage game (Dal Bó and Fréchet, 2011).

We studied a Baseline scenario, where participants cannot obtain any information about the past conduct of their counterpart, and two treatments with two different institutions for monitoring past conduct. In the Information Request treatment, participants can receive accurate information about the counterpart, while in the Information Provision treatment participants can create an accurate public record of individual conduct. In the Baseline treatment, cooperation rates were about 48% for the average group, which is evidence that attaining full cooperation in groups of strangers is challenging even if the group is small. In experiments, cooperation rates could be substantially increased by changing the informational conditions or by introducing communication. Camera and Casari (2009) and Camera et al. (2013b) reports the results of

experiments in a similar setting but adopting a Prisoner's Dilemma game; revealing the laboratory identities of counterparts raised cooperation by 20 percentage points, and introducing free-form communication by means of chat boxes among groups of four raised cooperation by 40 percentage points.

Did better monitoring technologies allow participants to achieve similar or better coordination on the inter-temporal giving and receiving of goods, compared to the no-monitoring setup in Baseline? The short answer is no (Result 1), even if the availability of monitoring institutions helped in some respects to simplify the coordination task. To see this, note that in Information Request treatment having the option to pay to view the history of the opponent allows subsets of participants to coordinate on history-dependent strategies. For example, some participants may choose to cooperate only with those who have immaculate cooperation records even if not everyone in their group does so. With Information Provision, this strategy is possible only if everyone in the group coordinates on making reports and on using a history-dependent strategy, because information on past actions is not already available and must be created by the group. However, neither Information Request nor Information Provision helped in coordinating on sanctioning strategies that effectively removed the incentives to defect (Result 2), despite their relatively low cost (Result 3). What's more, neither Information Request nor Information Provision redistributed enough surplus from defectors to cooperators to influence the selection of a cooperative equilibrium (Result 4).

If improving knowledge of past behaviors through monitoring is the key to reducing the temptation to defect, then the Information Request and Information Provision treatments should exhibit a larger favorable impact on cooperation than the Baseline treatment. In practice, the data reveals that monitoring institutions available in those two treatments are at best ineffective in

increasing cooperation relative to Baseline. Additional experiments with variations in monitoring may further corroborate this conclusion and help uncover what additional elements are needed to enhance the effectiveness of monitoring institutions. Some specific factors might have weakened the usefulness of the monitoring institutions in supporting cooperation in our experiment. One factor is the heterogeneity in the use of the monitoring technologies, which might have magnified coordination difficulties, instead of reducing them. Another factor is the cost associated with using the monitoring institutions, which may have discouraged their use. Finally, the possible ambiguity in interpreting reports of defections may also have played a role: a defection in an individual's record could have been interpreted either as an act of free riding by that individual, or as a form of punishment if the individual had suffered a defection earlier in the game. These considerations suggest a few additional aspects that could be studied to understand better which types of monitoring institutions are empirically effective in promoting cooperation, and which ones are not.

References

- Bigoni, Maria, Gabriele Camera and Marco Casari (2014). Money is more than memory. Working Paper No. 14-17, Economic Science Institute, Chapman University.
- Bolton, G. E., Katok, E., and A. Ockenfels (2005). Cooperation among strangers with limited information about reputation. *Journal of Public Economics* 89(8), 1457-1468.
- Camera, Gabriele and Marco Casari (2009). Cooperation among strangers under the shadow of the future. *American Economic Review*, 99(3), 979–1005.
- Camera, Gabriele and Marco Casari (2014). The coordination value of monetary exchange: Experimental evidence. *American Economic Journal: Microeconomics*, 6 (1), 290-314.
- Camera, Gabriele, Marco Casari, and Maria Bigoni (2013a). Money and Trust among Strangers. *Proceedings of the National Academy of Science*, 110(37), 14889-14893.
- Camera, G., M. Casari, and M. Bigoni (2013b). Experimental Markets with Frictions. *Journal of Economic Surveys* 27(3), 536–553
- Chen, Yan, Sherry Xin Li, Tracy Xiao Liu and Margaret Shih (2014). Which Hat to Wear? Impact of Natural Identities on Coordination and Cooperation. *Games and Economic Behavior*, 84, 58-86.

- Clark, Kenneth and Martin Sefton (2001). The Sequential Prisoner's Dilemma: Evidence on Reciprocation. *Economic Journal*, 111(468), 51-68
- Cox, Caleb, Matthew Jones, Kevin Pflum, and Paul J. Healy (2015). Revealed Reputations in the Finitely-Repeated Prisoners' Dilemma. *Economic Theory*, 58(3), 441–484.
- Dal Bó, Pedro, and Guillaume Fréchette (2011). The Evolution of Cooperation in Infinitely Repeated Games: Experimental Evidence. *American Economic Review* 101(1), 411-429.
- Ellison, Glenn (1994). Cooperation in the prisoner's dilemma with anonymous random matching. *Review of Economic Studies*, 61, 567-88.
- Fehr, Ernst and Simon Gaechter (2000). Fairness and Retaliation: The Economics of Reciprocity. *Journal of Economic Perspectives*, 14(3), 159-181.
- Fischbacher, Urs, Simon Gaechter, and Ernst Fehr (2001). Are people conditionally cooperative? Evidence from a public goods experiment. *Economics Letters*, 71(3), 397-404.
- Klarita Gërzhani, Jordi Brandts, and Arthur Schram (2013). The emergence of employer information networks in an experimental labor market. *Social Networks*, 35(4), 541-560.
- Kandori, Michihiro (1992). Social norms and community enforcement. *Review of Economic Studies*, 59, 63-80.
- Milgrom, Paul, Douglass North, and Barry 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.
- Palfrey, Tom R., and Howard Rosenthal (1994). Repeated Play, Cooperation and Coordination: An Experimental Study. *Review of Economic Studies*, 61, 545-565.
- Roth, A.E., and Murnighan, K. (1978). Equilibrium behavior and repeated play of the prisoner's dilemma. *Journal of Mathematical Psychology*, 17, 189-98.
- Stahl, Dale O. (2013). An Experimental Test of the Efficacy of Simple Reputation Mechanisms to Solve Social Dilemmas. *Journal of Economic Behavior & Organization*, 94, 116–124.
- Ule, A., Schram, A., Riedl, A., and Cason, T. N. (2009). Indirect punishment and generosity toward strangers. *Science*, 326(5960), 1701–1704.
- Van Huyck, John B., Raymond C. Battalio and Richard O. Beil (1991). Strategic Uncertainty, Equilibrium Selection, and Coordination Failure in Average Opinion Games. *Quarterly Journal of Economics*, 106(3), 885-910.
- Van Huyck John, B., Battalio Raymond, C., Walters Mary, F. (1995). Commitment versus Discretion in the Peasant-Dictator Game. *Games and Economic Behavior* 10(1) 143-170.