

Aoyagi, Masaki; Bhaskar, V.; Frechette, Guillaume R.

Working Paper

The impact of monitoring in infinitely repeated games: Perfect, public, and private

ISER Discussion Paper, No. 942

Provided in Cooperation with:

The Institute of Social and Economic Research (ISER), Osaka University

Suggested Citation: Aoyagi, Masaki; Bhaskar, V.; Frechette, Guillaume R. (2015) : The impact of monitoring in infinitely repeated games: Perfect, public, and private, ISER Discussion Paper, No. 942, Osaka University, Institute of Social and Economic Research (ISER), Osaka

This Version is available at:

<https://hdl.handle.net/10419/127071>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

**THE IMPACT OF MONITORING
IN INFINITELY REPEATED GAMES:
PERFECT, PUBLIC, AND PRIVATE**

Masaki Aoyagi
V. Bhaskar
Guillaume R. Frechette

July 2015

The Institute of Social and Economic Research
Osaka University
6-1 Mihogaoka, Ibaraki, Osaka 567-0047, Japan

The Impact of Monitoring in Infinitely Repeated Games: Perfect, Public, and Private *

Masaki Aoyagi
Osaka University

V. Bhaskar
University of Texas at Austin

Guillaume R. Fréchette
New York University

July 27, 2015

Abstract

This paper uses a laboratory experiment to study the effect of a monitoring structure on the play of the infinitely repeated prisoner's dilemma. Keeping the stage game fixed, we examine the behavior of subjects when information about past actions is perfect (*perfect* monitoring), noisy but public (*public* monitoring), and noisy and private (*private* monitoring). We find that the subjects sustain cooperation in every treatment, but that their strategies differ substantially in the three treatments. Specifically, we observe that the strategies are more complex under public and private monitoring than under perfect monitoring. We also find that the strategies under private monitoring are more lenient than under perfect monitoring, and less forgiving than under public monitoring.

JEL classification: C72, C73, C92

Key words: Infinitely repeated games, monitoring, perfect, public, private, experiments

*We wish to thank Matthew Embrey and Sevgi Yuksel as well as seminar participants at IAS Princeton, Southern Methodist University, WZB, Stanford, Toulouse, Science Po, Harvard University, the University of Montréal, University of Heidelberg, the University of Munich, the Max-Planck Institute, the University of Indiana, and conference participants at the North American Meeting of the ESA, the Social Dilemmas conference, and the SITE conference for helpful comments and suggestions that improved the paper. We would further like to thank the National Science Foundation (grant SES-SES-1225779), the CV Starr Center and the Center for Experimental Social Science (CESS) for research support. We are responsible for all errors.

1 Introduction

Many economic situations involve repeated interactions among players who do not know for sure what actions other players have chosen, or what information they have observed about those actions. In a pioneering work, Stigler (1964) studies a model of secret price cutting and points out the difficulty faced by the firms who attempt to collude under (*imperfect*) *private monitoring*, where information about the past actions (price choices) of other players is noisy and privately observed. As Kandori (1992) notes, and by now is well recognized, private monitoring implies the absence of a coordination device for the players, and hence the lack of a recursive structure in the repeated game. In this sense, there is a fundamental difference between private monitoring and the more traditional environments of *perfect monitoring*, where information about past actions is perfect (and public), and (*imperfect*) *public monitoring*, where the information is noisy but public. Theory shows that any equilibrium that sustains a positive degree of cooperation under private monitoring must entail the play of an intricate mixed strategy.

A natural question then is whether people's behavior under private monitoring is indeed significantly different from that under perfect or public monitoring. Answering this question in the field is difficult because of data restriction: Private information is unlikely to be replicated in publicly available data sets, and so is information about other critical parameters such as the discount factor and the conditional distribution of signals given actions. In this light, a laboratory experiment offers a valuable alternative to field research, and the objective of this paper is to create such environments in a laboratory and explore the behavior of subjects.

The first question this paper intends to explore concerns the ability of experimental subjects to sustain cooperation under private monitoring. Given the complexity of strategic interaction, it is only recently that theorists began to understand the possibility of cooperation in a private monitoring environment. Hence, it would not be surprising if the subjects fail to cooperate in a laboratory. The difficulty of cooperation under private monitoring can be measured most accurately if we vary the monitoring structures but keep the other aspects of the game fixed.

Our next question concerns if and how the subjects' behavior under private monitoring differs from that under perfect or public monitoring. We will answer this question in two different ways. First, we examine if their behavior after certain histories is different under different monitoring structures. Second, we estimate their strategies and check whether the most popular strategies are different across treatments. These analyses together help us answer the following questions among others: How long back in history does a strategy look when choosing actions? Is it *lenient* in the sense that it does not revert to a punishment after a single bad signal, or *forgiving* in the sense that it returns to cooperation after punishing the opponent?

Our analysis is also an indirect test of the theory of private monitoring that implies the use of extremely complex strategies required to sustain cooperation.

Our experiments use the repeated prisoners' dilemma (PD) with a random termination probability of 0.1. There are three treatments corresponding to three different monitoring structures as follows. In the *perfect* monitoring treatment, player i observes player j 's action choice a_j^t in each round t . In the two imperfect monitoring treatments, on the other hand, player j 's action a_j^t in each round t generates signal ω_i^t , which is correct and equals a_j^t with probability 0.9, but is incorrect with probability 0.1: In each round t , player i observes a pair of noisy signals (ω_i^t, ω_j^t) in the *public* monitoring treatment, whereas he only observes ω_i^t in the *private* monitoring treatment.¹ We keep fixed other elements of the game as much as possible: The three treatments have the same expected stage-payoffs, and the payoff in the perfect monitoring treatment is determined randomly by the same probability distribution as in the imperfect monitoring treatments to control for the effect of uncertainty.

Our findings can be summarized as follows. First, we find that the rate of the cooperative action under private monitoring is comparable to those under perfect and public monitoring.² Furthermore, the rate of coordination (either on (C, C) or (D, D)) is slightly lower under private monitoring than under perfect and public monitoring, but significantly higher than implied by independent action choices. These positive results on cooperation and coordination are remarkable in view of the theoretical difficulties associated with private monitoring. Second, our strategy estimation suggests that the subjects play substantially different strategies in the three treatments. In particular, when we focus on the cooperative strategies that are found in the most significant proportions in each treatment, none of them is lenient under perfect monitoring, but all of them are lenient under public and private monitoring. As for forgiveness, the strategies used under private monitoring are not as forgiving as those under public monitoring. Furthermore, when we represent the strategies by finite automata and measure their complexity by the number of states in this representation, we find the strategies under private monitoring more complex than those under perfect monitoring. These findings suggest that the subjects find ways to cooperate and coordinate through a substantially different mechanism under each monitoring structure.

The organization of the paper is as follows: In the next section, we give a brief review of the literature. Section 3 formulates a model of repeated PD, Section 4

¹Hence, a player's signal takes one of two values in the private monitoring treatment whereas it takes one of four values in the public monitoring treatment.

²In addition, we establish that their cooperation levels are higher than those in one-shot PD in the literature.

provides a theoretical background, and Section 5 describes the experimental design. The questions our analysis attempts to answer are listed in Section 6, and the results are presented in Section 7. Section 8 concludes with a discussion.

2 Related Literature

There is only indirect evidence from observational data as to whether repeated interactions under private monitoring lead to cooperation. A meta study by Levenstein & Suslow (2006) identifies the existence of a joint sales agency or industry associations as mechanisms that help cartels through the collection and dissemination of information. Harrington & Skrzypacz (2007, 2011) find that cartels for such products as citric acid, lysine, and vitamins went to great lengths to make sales public information amongst members and also used inter-firm sales as a way to transfer profits to sustain collusion. Their theory demonstrates that supporting collusion requires side transfers in environments where information about prices is private. These findings indirectly attest to the difficulty of sustaining collusion just with the help of private information.

As mentioned in the Introduction, the primary objective of this paper is to identify the pure effect of the monitoring structure while keeping other aspects of the game fixed as much as possible. Although there is now a growing literature on repeated game experiments, we are aware of no work that makes cross comparison of different monitoring structures including private monitoring.³

Early experimental studies find some cooperation when subjects engage in repeated interactions under perfect monitoring.⁴ Further evidence of cooperation in repeated games was provided by Engle-Warnick & Slonim (2004, 2006*b,a*), Dal Bó (2005), Aoyagi & Fréchette (2009), and Duffy & Ochs (2009) in various settings (these subsequent studies differ from the earlier ones in that they allow subjects to play multiple repeated games). Dal Bó & Fréchette (2011) find in perfect monitoring games that cooperation rates by experienced subjects are 1) very low when cooperation is theoretically infeasible, and 2) higher when it is theoretically feasible, and very high for certain parameter values. Furthermore, Dal Bó & Fréchette

³Experiments on infinitely repeated games address a number of different questions. They include, to mention a few, Schwartz et al. (2000), Dreber et al. (2008) on modified PD, Cason & Mui (2015) on a collective resistance game. Cooper & Kühn (2014) on the role of communication and renegotiation, Fudenberg et al. (2014) on the relationship between behavior in the dictator game and that in an infinitely repeated game, Cabral et al. (2014) on reciprocity, and Bernard et al. (2014) on a gift exchange game. Other forms of dynamic games are studied by Battaglini et al. (2015), and Vespa (2015).

⁴See Roth (1995). Early studies include Roth & Murnighan (1978), Murnighan & Roth (1983), Feinberg & Husted (1993), Holt (1985), and Palfrey & Rosenthal (1994).

(2013) find that in the repeated PD with perfect monitoring, the strategies used by the majority of subjects are simple, and can be classified into one of 1) Always D (defect), 2) Grim Trigger, which begins with C (cooperate) but switches to D forever following a defection, and 3) Tit-For-Tat, which begins with C and thereafter mimics the other player's action in the previous round.

On games with imperfect public monitoring, Aoyagi & Fréchette (2009) find that subjects support cooperation in a repeated PD with a noisy continuous signal, and that their payoff is a decreasing function of noise in the public signal in line with the theoretical prediction on the maximal symmetric PPE payoff.⁵ Fudenberg et al. (2012) study a model of repeated PD under imperfect public monitoring that closely resembles our public monitoring treatment. They formulate a model in which each player observes a profile of implemented actions that differ from the intended actions with positive probability, and examine the effects of stage payoffs and noise levels on subjects' behavior.⁶ Fudenberg et al. (2012) find no clear difference in the levels of cooperation caused by the monitoring structure, but that the subjects' strategies under public monitoring are more *lenient* and more *forgiving* than under perfect monitoring. Using the strategy frequency estimation method proposed in Dal Bó & Fréchette (2011), they support this finding by showing that the estimated strategies of the subjects choose actions based on the signals of the past few rounds instead of the most recent one. In light of these results, it is particularly interesting to see what happens to leniency and forgiveness when monitoring becomes private.⁷

To the best of our knowledge, the paper by Matsushima & Toyama (2013) is the only other laboratory study of private monitoring in repeated games.⁸ They use a repeated PD with private monitoring, and focus on how player behavior changes in two treatments that change the noise level in the signal.⁹ As one would expect,

⁵Some, including Cason & Khan (1999), study repeated games with imperfect monitoring but do not use random termination, which has become the standard procedure for implementing infinitely repeated games in a laboratory since Roth & Murnighan (1978). See Fréchette & Yüksel (2013) for some alternative termination methods and Sherstyuk et al. (2013) for alternative payment methods.

⁶The payoff depends on one's own choice and signal in our model whereas it depends on the two implemented actions in Fudenberg et al. (2012).

⁷Our analysis also provides a robustness check of Fudenberg et al. (2012), who use different stage-payoff tables for different noise levels. In contrast, we keep the expected payoff table fixed across treatments.

⁸Holcomb & Nelson (1997) observe in a repeated duopoly model (without random termination) that the experimenter's manipulation of information about a subject's quantity choice "does significantly affect market outcomes" (p.79). Feinberg & Snyder (2002) also study the effect of occasional manipulation of payoff numbers in a version of repeated PD, and find less collusive behavior when such manipulation is ex post not revealed than when it is.

⁹The noise here refers to the probability that the signal differs from the actual action choice. In our notation, see page 8, they have $g = \ell = 5/17$ in their 0.9 treatment and $g = \ell = 5/19$ in their 0.6

they find that a player is more likely to cooperate when he observes a c (cooperative or good) signal rather than a d (defective or bad) signal in both treatments, and that the cooperation rate is higher when the noise is smaller. The most interesting is their finding that players are more responsive to the signal when the noise is smaller. That is, the difference in the rates of choosing C after a c signal and after a d signal is larger when the noise is smaller. This last finding is at odds with the theoretical prediction based on a memory-one belief-free equilibrium.¹⁰ Our empirical findings point in the same direction with respect to responsiveness: players are more responsive to the opponent's signal under perfect monitoring than under public or private monitoring. We emphasize however that unlike in Matsushima & Toyama (2013), we don't limit ourselves to memory-one strategies that condition behavior only on the most recent signal.

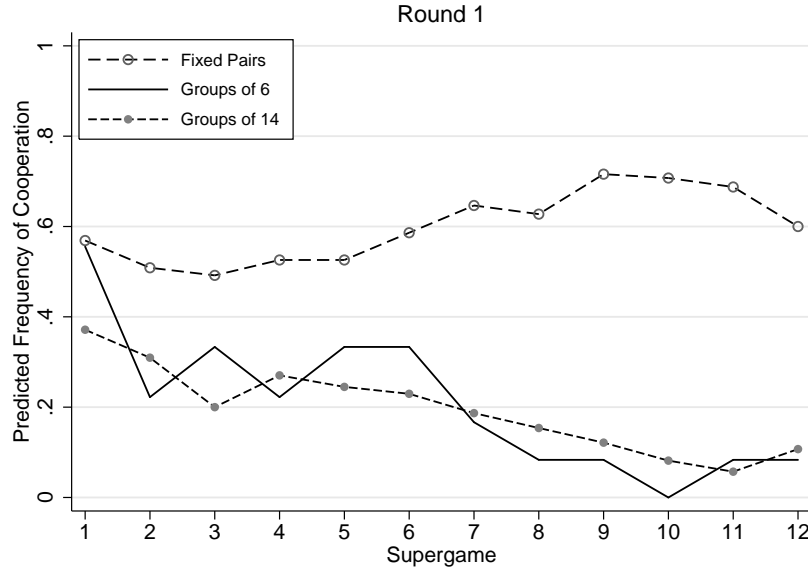


Figure 1: Duffy & Ochs (2009): random re-matching in small groups

Closely related to repeated games with private monitoring are models of random matching within a group where a group of players are matched in pairs to a treatment.

¹⁰See Section 4 for the discussion of the belief-free equilibrium. Matsushima & Toyama (2013) interpret these findings as a result of their subjects using a behavioral tit-for-tat strategy that incorporates psychological costs of cooperating and defecting.

different partner every round. Monitoring is private – although a player perfectly observes the action of his opponent in the current supergame, he does not observe the actions taken in other pairs. Theoretically, regardless of the group size, cooperation can be sustained in equilibrium if δ is large enough through a contagious grim-trigger strategy which cooperates as long as all past interactions have resulted in (C, C) , but defects otherwise. Thus a single defection results in the breakdown of cooperation through the contagion process. Duffy & Ochs (2009) find that even when the group size admits cooperation in theory, their subjects cannot sustain cooperation in the random matching environment. Their results are replicated in Figure 1. Subsequent experiments by Camera & Casari (2009), Camera et al. (2012), and Camera & Casari (2011) also confirm that cooperation in the random rematching environment is fragile and is possible only for very small groups. The difficulty to support cooperation with random rematching in small groups naturally poses a question as to whether players can sustain cooperation in bilateral interactions with private monitoring.

3 Models of Repeated Prisoners' Dilemma

Two players play a symmetric 2×2 stage-game infinitely often. The set of actions for each player i is denoted $A_i = \{C, D\}$. Player i 's action $a_i \in A_i$ generates a signal $\omega_j \in \{c, d\}$ with noise $\varepsilon = 0.1$. The probability distribution of ω_j conditional on a_i is given by $\Pr(\omega_j = c \mid a_i = C) = \Pr(\omega_j = d \mid a_i = D) = 1 - \varepsilon$. The two signals ω_1 and ω_2 are independent conditional on the action profile $a = (a_1, a_2)$: $\Pr(\omega_1, \omega_2 \mid a) = \Pr(\omega_1 \mid a_2) \Pr(\omega_2 \mid a_1)$.

Under *perfect monitoring*, player i observes j 's action a_j . Under (*imperfect*) *public monitoring*, player i observes the signal profile $\omega \equiv (\omega_1, \omega_2)$. Under (*imperfect*) *private monitoring*, player i only observes ω_i .¹¹ With the identification $\omega_i = a_j$ under perfect monitoring, player i 's stage-payoff is a function of his own action a_i and the signal ω_i about player j 's action, and denoted by $g_i(a_i, \omega_i)$. Player i 's expected stage-payoff u_i is a function of the action profile a and is given by

$$u_i(a) = \sum_{\omega_i \in A_j} \Pr(\omega_i \mid a_j) g_i(a_i, \omega_i). \quad (1)$$

We specify the function g_i so that the expected stage payoffs (u_1, u_2) form a PD as

¹¹Note that under private monitoring, there is no common knowledge between the two players about each other's action or signal. Furthermore, independence of ω_i and ω_j implies that i 's private signal is uninformative about ω_j under the unconditional choice of actions.

follows:

$$\begin{array}{cc|cc}
& & C & D \\
C & 1 & 1 & -\ell & 1+g \\
D & 1+g & -\ell & 0 & 0
\end{array} \tag{2}$$

In our experiments, the parameters g and $\ell > 0$ are chosen to satisfy $g = \ell$.¹² Let x_i^t denote player i 's information at the end of round t regarding events in the round as described above.¹³ Player i 's history up to round t is the sequence $h_i^t = (x_i^1, \dots, x_i^t)$. Let H_i^t be the set of i 's histories up to t and let $H_i = \cup_{t=1}^{\infty} H_i^t$. Player i 's (behavioral) strategy σ_i is a collection $(\sigma_i^t)_{t=1}^{\infty}$ such that $\sigma_i^1 \in \Delta A_i$ and for $t \geq 2$, $\sigma_i^t : H_i^{t-1} \rightarrow \Delta A_i$, where ΔA_i is the set of probability distributions over A_i . Denote by $\delta \in (0, 1)$ the common discount factor of the players, and let $\pi_i(\sigma)$ be player i 's expected payoff in the repeated game under the strategy profile $\sigma = (\sigma_1, \sigma_2)$. Likewise, let $\pi_i(\sigma \mid h_i)$ be i 's expected continuation payoff under σ following history $h_i \in H_i$.¹⁴ A strategy profile $\sigma = (\sigma_1, \sigma_2)$ is a *perfect Bayesian equilibrium* (PBE, or simply an *equilibrium*) of the repeated game if for $i = 1, 2$,

$$\pi_i(\sigma \mid h_i) \geq \pi_i(\sigma'_i, \sigma_j \mid h_i)$$

for any alternative strategy σ'_i and any private history $h_i \in H_i$.¹⁵ Under perfect monitoring, σ is a PBE if and only if it is a subgame perfect equilibrium (SPE). Under public monitoring, a strategy σ_i is *public* if σ_i^t is a function only of $(\omega^1, \dots, \omega^t)$ and not that of (a_i^1, \dots, a_i^t) . A PBE σ is a *perfect public equilibrium* (PPE) if each σ_i is public.

4 Theoretical Background

This section collects some background material that is well recognized in the theoretical literature but is useful for the interpretation of our experimental results.

¹²The equality $g = \ell$ implies that the expected payoff table has a benefit-cost form à la Fudenberg et al. (2012). Namely, a player choosing C incurs cost g but gives benefit $1 + g$ to the other player, whereas action D entails no cost or benefit. The same condition is referred to as separability in Matsushima & Toyama (2013). The benefit-cost (b/c) ratio is given by $\frac{1+g}{g}$.

¹³As mentioned earlier, the payoffs in the perfect monitoring treatment are randomly generated according to the same distribution as in the imperfect monitoring treatments. As such, x_i^t formally includes the realization of the random payoff, which is by design independent of any player's action choice. See Section 5 for details of the actual implementation.

¹⁴Throughout, we consider the average discounted payoff, which equals the sum of discounted stage payoffs multiplied by $1 - \delta$.

¹⁵Under public and private monitoring, the definition of a PBE omits reference to the belief system given the full-support assumption on the signal distribution \Pr .

One essential observation concerns the players' preference for efficiency/cooperation and the severity of a punishment. Under perfect monitoring, cooperation in every period on the equilibrium path can be enforced by non-lenient and non-forgiving strategies such as the grim-trigger strategy. Since no bad signal is observed on the equilibrium path, leniency or forgiveness is immaterial for the efficiency of an outcome. Under imperfect public monitoring, on the other hand, bad signals arise even when both players cooperate. To achieve efficiency, hence, the strategy must be lenient in the sense that a punishment is started either only after the consecutive occurrence of bad signals, or with a small probability after each occurrence of such a signal. Furthermore, if the players are concerned with efficiency *a posteriori* after the punishment is triggered, then the strategy must be forgiving so that the cooperative phase can be restored after a fixed number of rounds or after the occurrence of a good signal during the punishment phase. The situation is significantly more complex with private monitoring. If player i believes that his opponent j is playing a strategy that chooses C with probability one today but is not lenient, then i 's strategy must be lenient: If i observes a bad signal today and responds with D , then it will likely cause j to observe a bad signal and hence revert to a punishment. On the other hand, if i is lenient and plays C instead, it will likely keep j in the cooperative phase. After all, j does not know that i has observed a bad signal, and since it is caused by the noise in monitoring, i might as well ignore it. This reasoning excludes the possibility of an equilibrium that entails the unconditional play of C on the path along with a non-lenient response to a bad signal. Theory of private monitoring suggests that j 's strategy must be finely adjusted in the level of leniency and forgiveness so that i has an incentive to play C after a good signal and D after a bad signal.

A more specific description of an equilibrium in each case is as follows. Under perfect monitoring, mutual cooperation is an SPE outcome if $\delta \geq \frac{g}{1+g}$. For example, if we denote by CC the action-signal pair $(a_i, a_j) = (C, C)$, the grim-trigger strategy σ_G that begins with $a_i = C$ and plays C if $h_i^t = (CC, \dots, CC)$ but plays D otherwise, is an equilibrium strategy and generates the maximum symmetric average discounted payoff of 1.

Under public monitoring, a pair (σ_G, σ_G) of grim-trigger strategies that revert to punishment when the history $h_i^t \neq (Ccc, \dots, Ccc)$ is also a PPE for δ sufficiently large and ε sufficiently small.¹⁶ However, such an equilibrium entails a significant efficiency loss since permanent defection is triggered with probability $1 - (1 - \varepsilon)^2$ in every round.¹⁷ We can verify that the use of a lenient strategy that triggers a

¹⁶As before, $(a_i, \omega_i, \omega_j) = (C, c, c)$ is abbreviated as Ccc . (σ_G, σ_G) is a PPE if $(1 - 2\varepsilon)(1 - \varepsilon) - \varepsilon(2 - \varepsilon)g \geq \frac{1-\delta}{\delta}g$.

¹⁷The equilibrium expected payoff under this grim-trigger strategy equals $\frac{1-\delta}{1-\delta+\delta\varepsilon(2-\varepsilon)}$, which equals $0.369 \ll 1$ under our parameter values ($\delta = 0.9$ and $\varepsilon = 0.1$).

punishment with a positive but small probability after a bad signal yields a more efficient equilibrium.

In the case of private monitoring, the lack of common knowledge of histories becomes a major obstacle for cooperation. In particular, an equilibrium should be in mixed strategies if it supports any level of cooperation. As mentioned in the Introduction, two approaches to the problem have been developed in the literature as illustrated below.

The belief-based approach attempts to provide a proper incentive after each history by considering a mixture of repeated game strategies. Specifically, consider a mixture between the grim-trigger strategy σ_G and the strategy σ_D of choosing D always. Note that the continuation strategy of such a mixed strategy after each history is either σ_D or again a mixture of σ_G and σ_D . The initial probability weights on σ_G and σ_D are chosen so that after every history, it is incentive compatible to revert to σ_D if and only if a player observes a d signal.¹⁸ In one interpretation, when players are randomly matched to play the repeated game as in our experimental setting, a mixed strategy played by a single opponent corresponds to the population of opponents playing different pure strategies.

Another approach to private monitoring is belief-free equilibria in which each player plays a behavior strategy that makes the other player indifferent between C and D after every history. Specifically, player i 's strategy makes player j indifferent between his actions independent of the history observed by player j . This makes player j 's belief about player i 's (private) history irrelevant, and substantially simplifies the equilibrium analysis. A belief-free equilibrium constructed by Ely & Välimäki (2002) can be illustrated as follows:¹⁹ Let p be the probability that player i plays C when he observes signal $\omega_i = c$ in the previous round, and q be the same probability when he observes $\omega_i = d$. Furthermore, let $W(c)$ denote i 's continuation payoff when his opponent j observes $\omega_j = c$, and $W(d)$ denote his continuation payoff when j observes $\omega_j = d$. Since player i is assumed indifferent between playing C and D , we should have

$$(1 - \delta)g = \delta(1 - 2\varepsilon) [W(c) - W(d)], \quad (3)$$

where the left-hand side is i 's payoff gain in the current round from playing D rather than C , and the right-hand side is the increase in continuation payoff from playing C rather than D , which increases the probability of player j observing $\omega_j = c$ by

¹⁸It is typically the case that with high discount factors, the players do not have an incentive to switch to σ_D when observing $\omega_i = d$. This is the case with our specification of $\delta = 0.9$, and it is necessary to lower the effective discount factor by partitioning the supergame into several segments so that each segment is played only once in several rounds. See for example Sekiguchi (1997).

¹⁹See also Piccione (2002) for an alternative formulation of belief-free equilibria.

$1 - 2\varepsilon$.²⁰ Next, if player j observed $\omega_j = c$ in round $t - 1$ and player i plays D in round t , then player i 's continuation payoff from round t on is given by

$$W(c) = (1 - \delta) \{p(1 + g) + (1 - p) \cdot 0\} + \delta \{(1 - \varepsilon)W(d) + \varepsilon W(c)\}.$$

On the other hand, if player j observed $\omega_j = d$ in round $t - 1$ and player i plays D in round t , then player i 's continuation payoff from round t on is given by

$$W(d) = (1 - \delta) \{q(1 + g) + (1 - q) \cdot 0\} + \delta \{(1 - \varepsilon)W(d) + \varepsilon W(c)\}.$$

These equations together imply

$$W(c) - W(d) = (1 - \delta)(p - q)(1 + g). \quad (4)$$

Combining (3) and (4), we obtain

$$p - q = \frac{g}{\delta(1 - 2\varepsilon)(1 + g)}. \quad (5)$$

(5) provides one behavioral prediction when the subjects play belief-free equilibrium that has *memory-one* in the sense that mixed actions are determined by the signal realization of the previous round. Let the *responsiveness* of a strategy be defined by

$$\Pr(a_i^{t+1} = C \mid \omega_i^t = c) - \Pr(a_i^{t+1} = C \mid \omega_i^t = d). \quad (6)$$

When the subjects play the memory-one belief-free equilibrium described above, this quantity equals $p - q$ and is expressed in terms of the underlying parameters as in (5).

It can also be verified that the above belief-free strategy profile is an equilibrium not only under private monitoring but also under perfect and public monitoring. Consequently, if the subjects play the memory-one belief-free equilibrium in every monitoring treatment, then they should exhibit the same responsiveness in both the public and private monitoring treatments where $\varepsilon = 0.1$, and a lower responsiveness value in the perfect monitoring treatment where ε can be interpreted as 0.

5 Experimental Design

The experiment has three treatments corresponding to the three monitoring structures described above. The public and private monitoring treatments use the payoff

²⁰Note that the gain from playing D does not depend on j 's action when $g = \ell$ as assumed.

Table 1: Summary statistics

Treatments	Subjects	Sessions	Subjects per session	Supergames per session	Rounds per Supergame			Subject earnings (US\$)		
					avg	min	max	avg	min	max
Perfect	66	4	16, 18, 16, 16	11, 12, 19, 11	10.3	1	37	32.91	19.78	43.67
Public	68	4	18, 20, 14, 16	11, 11, 22, 11	10.1	1	37	34.87	23.40	48.12
Private	72	4	20, 18, 14, 20	12, 9, 19, 11	10.4	1	37	31.92	22.34	44.12

function $g_i(a_i, \omega_i)$ given by

$a_i \backslash \omega_i$	c	d
C	46	8
D	54	16

(7)

The expected payoffs are then generated according to (1). Our perfect monitoring treatment introduces the same random relationship between the payoffs and the action profile as follows: For each action profile (a_i, a_j) , player i 's payoff in the perfect monitoring treatment is generated by the lottery that yields $g_i(a_i, \omega_i = a_j)$ with probability $1 - \varepsilon$ and $g_i(a_i, \omega_i \neq a_j)$ with probability ε . For example, when the action profile is (C, C) , each subject (independently) receives 46 with probability $1 - \varepsilon$ and 8 with probability ε so that

$$u_i(C, C) = (1 - \varepsilon) g_i(C, c) + \varepsilon g_i(C, d),$$

just like in the other two treatments. It follows that our three treatments have exactly the same expected stage-payoff table. With our choice of $\varepsilon = 0.1$, it is given by

$a_1 \backslash a_2$	C	D
C	42.2, 42.2	11.8, 50.2
D	50.2, 11.8	19.8, 19.8

(8)

Note that the payoff matrix (8) is strategically equivalent to (2) for

$$g = \ell = \frac{5}{14} \approx 0.357.^{21}$$
(9)

In each of the three treatments, these parameter values ensure that the type of strategy discussed in the previous section is an equilibrium in which the players cooperate with strictly positive probability at least initially.

The experiments use the between-subject design so that each subject participates in one and only one treatment. Sessions were conducted at the CESS lab at

²¹Simply apply the affine transformation $22.4 u_i(a) + 19.8$. The benefit-cost ratio mentioned in Footnote 12 hence equals $\frac{1+g}{g} = 3.8$.

NYU.²² In each session, after the instructions are read aloud, subjects are randomly and anonymously paired via computer with another subject to play a supergame.²³ All supergames in a session are simultaneously terminated after every round with probability 0.1, and subjects are randomly rematched to play another supergame.²⁴ After every supergame, subjects are informed of the complete history of choices and signals by both players to ensure that feedback is the same for all treatments and that the only difference among them is the information structure *within* a supergame. This process repeats itself until 75 minutes of play have elapsed; the first supergame to end after that marks the end of a session. Four sessions of each treatment were conducted. The supergames lasted between 1 and 37 rounds, and averaged 10.3 rounds (close to the expected value of 10).²⁵ The sessions were approximately 1 hour and 40 minutes, and subjects earned between \$19.78 and \$48.12 with an average earning of \$33.21.²⁶ These and other summary statistics are provided in Table 1.²⁷

6 Directions of Analysis

As discussed in the Introduction, the primary focus of our research is on (1) whether subjects can support cooperation under private monitoring and how the level of cooperation compares with those under public and perfect monitoring, and (2) whether or not their behavior is different in the three treatments. While our investigations are more of exploratory nature rather than hypothesis-testing of theoretical predictions, we will relate the results to insights and predictions provided by the

²²Subjects who had participated in previous experiments with randomly terminated games or a PD as a stage-game were excluded.

²³In the experimental instructions, the term “match” is used in place of “supergame.”

²⁴The length of a supergame in each session of the perfect monitoring treatment was determined by a random number generator. Each session in the other treatments then used the same sequence of supergames as the corresponding session in the perfect monitoring treatment to control for the effect of the length of supergames on the evolution of play. Dal Bó & Fréchette (2011) and Engle-Warnick & Slonim (2006b) both document the impact of the length of supergames on behavior.

²⁵The difference in the average number of rounds results from the variation in number of supergames between sessions.

²⁶Points are converted to dollars at a pre-announced exchanged rate. Since the earnings for the first session of each treatment were slightly lower than expected (between \$19.78 and \$33.52 with an average of \$28.64), the minimum time of play was increased from 60 to 75 minutes and the exchange rate was decreased from 0.01 to 0.0075 for the subsequent sessions.

²⁷Given the difference in the number of supergames across treatments, the analysis uses data from only the first k_n supergames in session n of each treatment, where k_n is the minimal number of supergames in session n across treatments. For instance, since the second sessions of the three treatments have 9, 11 and 12 supergames, only the first 9 supergames are used in the analysis. Recall that the length of the k th supergame in session n is the same regardless of the treatments.

theory to the extent possible. In what follows we provide a set of questions that provide a guide for our analysis.

Question 1 (Cooperation and coordination) Is the level of cooperation and coordination lower under private monitoring than under perfect or public monitoring?

Theory clearly suggests that it is significantly more difficult to sustain cooperation under private monitoring since it requires the use of intricate mixed strategies as seen in Section 4. The lack of a coordination device under private monitoring also makes it difficult for the subjects to coordinate their actions beyond round 1. These considerations suggest an affirmative answer to question 1.

We examine our second question on the constancy of behavior across the three treatments from several different perspectives. Following the literature and as described earlier, we say that strategies are *lenient* if they do not prescribe sure defection following a single bad signal, and *forgiving* if they return to cooperation after having played defect.

Question 2 (Leniency and forgiveness) Are strategies more lenient and forgiving under public and private monitoring than under perfect monitoring?

As mentioned in Section 4, a player's preference for efficiency implies that his strategy under public monitoring should be more lenient and forgiving than under perfect monitoring. Indeed, previous work confirms this view. In the perfect monitoring environment, Dal Bó & Fréchette (2013) find grim trigger, which is not lenient, among one of the three most frequently observed strategies. To the contrary, both Fudenberg et al. (2012) and Embrey et al. (2013) find that the subjects' strategies are more lenient and more forgiving in the imperfect public monitoring environments. A similar observation can be made on the strategies that best describe the subjects' behavior in Aoyagi & Fréchette (2009): as noise in public information increases, the range of a "bad" signal which causes transition from the cooperation phase to the punishment phase shrinks, and the range of a "good" signal which causes transition from the punishment phase to the cooperation phase widens.²⁸

Under private monitoring, on the other hand, a strategy may be more lenient than under perfect monitoring if a player believes that his opponent plays C with high probability: Because of the noise in monitoring, he will give a benefit of doubt when he observes a single bad signal, and will also be reluctant to punish it since it

²⁸The signal space is continuous and the estimated strategies shift between the cooperation and punishment phases based on a threshold on the public signal. Data of Aoyagi & Fréchette (2009) show that this threshold decreases as noise increases.

will more likely invoke an adverse reaction from the opponent. The answer to the question on forgiveness under private monitoring is more difficult to predict since forgiveness, if any, takes place further down the history where it is difficult to infer if the opponent is cooperative or punitive.

Question 3 (Memory length) Do strategies have longer memory under public and private monitoring than under perfect monitoring?

An affirmative answer to this question is partially implied by the affirmative answers to Question 2 on leniency since leniency requires the examination of history over the past few rounds rather than just one.²⁹

Question 4 (Responsiveness) Is the level of responsiveness lowest under perfect monitoring and the same under public and private monitoring?

A belief-free strategy profile as described in Section 4 is an equilibrium in every monitoring treatment we consider. The memory-one belief-free equilibrium is the simplest of them, and if the subjects indeed play such an equilibrium, the responsiveness defined in (6) should be as described above.

7 Results

We present our results in two parts. The first part is a direct analysis of cooperation and coordination rates as well as action choices conditional on some histories. The second part is an analysis based on the estimation of strategies.

7.1 Cooperation Rates

Cooperation rates in the three treatments can be assessed visually in the left panel of Figure 2. In light of the variation in the number of supergames across sessions, the figure presents data in three categories: the first four supergames to the left, the last four supergames to the right, and a single point in the middle (labeled “other”) that corresponds to the average of the rates in all other supergames. As such, every point in Figure 2 (with the exception of the middle point) represents the average of four supergames, one from each session.

Observation 1 Subjects support cooperation under perfect, public and private monitoring.

²⁹Leniency can be defined in terms of either the number of bad signals before action *D* is chosen, or the probability with which action *C* is chosen after each bad signal. In the latter case, more leniency does not necessarily imply longer memory.

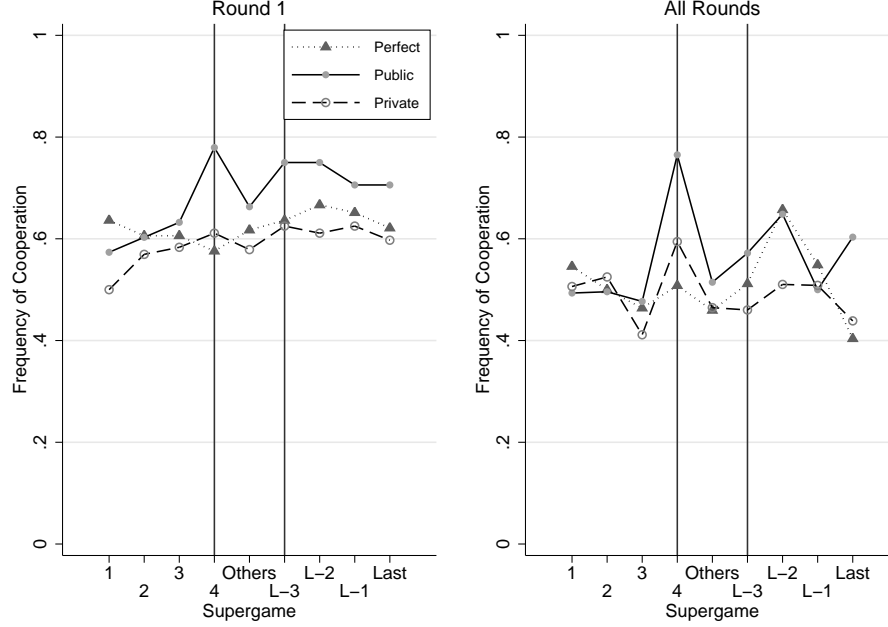


Figure 2: Cooperation rates by supergame

Observation 1 on perfect monitoring replicates earlier results in the literature and extends them to the environment with random payoffs. The round 1 cooperation rate in the last four supergames is 65%, which is statistically different from 0 at the 1% level.³⁰

Although there is no direct comparison available, the level of cooperation observed under perfect monitoring is considered high even when we use as a benchmark the level of cooperation in one-shot PD. In fact, the finding of Dal Bó & Fréchette (2013) suggests that the level of cooperation observed here can be sustained only if the discount factor (continuation probability) is higher than the critical level required for the existence of a cooperative equilibrium.³¹

³⁰Throughout the paper, unless stated otherwise, statistical tests are obtained by t-tests clustering the standard errors by session using only the last four supergames. The clustering is to account for potential session-effects. The interested reader is referred to Fréchette (2012). When results are referred to as not statistically significant, it implies a p-value greater than 10%.

³¹See the Appendix A.1, which replicates the finding of Dal Bó & Fréchette (2013). In their model, cooperation is a subgame perfect equilibrium outcome if and only if $\delta \geq 0.72$. As seen in Figure 6, cooperation rates for $\delta = 0.9$ and $\delta = 0.5$ diverge as the subjects accumulate experience.

Recall that in the present perfect monitoring treatment, payoffs are randomly generated conditional on the pair of action choices. In contrast, in the perfect monitoring repeated PD of Rand et al. (2015), action choices are implemented with error although the intended action choice is observed. The cooperation rates observed here are similar to those found in Rand et al. (2015), suggesting the robustness of subject behavior with respect to the way randomness is introduced.

Turning now to public monitoring, we see in the left panel of Figure 2 that the round 1 cooperation rate in the last four supergames is 73%, which is again statistically different from 0 at the 1% level. Positive cooperation in our experiments is in line with the findings in the literature on various forms of public monitoring: Fudenberg et al. (2012) introduces implementation errors in the subjects' action choices unlike in our experiment where noise is in the observation of the other player's action choice. In Aoyagi & Fréchette (2009), the subjects' action choices give rise to a one-dimensional continuous signal with infinite support that does not statistically identify the deviator, for example, from the action profile (C, C) . In Embrey et al. (2013), outcomes depend probabilistically on the subjects' action choices as in Fudenberg et al. (2012), but the public signal is binary.

Between perfect and public monitoring, there is no statistical difference in the round 1 cooperation rates. As depicted in the right panel of Figure 2, the cooperation rates over all rounds in the last four supergames are 46% under perfect monitoring and 58% under public monitoring. The rates are both statistically different from zero ($p < 0.01$), and statistically different from each other ($p < 0.01$). Comparison of overall cooperation rates under public and perfect monitoring in the literature is not conclusive, and suggests the influence of the specific monitoring technology used. A finding similar to the one above is reported by Fudenberg et al. (2012), who observe a statistically significant increase in overall cooperation rates when small noise is introduced into monitoring. On the other hand, movement in the opposite direction with the introduction of small noise into monitoring is reported by Aoyagi & Fréchette (2009).³²

The key finding in Observation 1 is cooperation under private monitoring. We see again in the left panel of Figure 2 that in the last four supergames, there is 61% cooperation in round 1, which is statistically different from 0 at the 1% level. The cross comparison across the treatments reveals that the only difference between round 1 cooperation rates is between public and private.³³ In particular, the rates

³²In fact, Aoyagi & Fréchette (2009) report a monotonically decreasing relationship between noise and cooperation rates in all rounds. They observe no statistical difference in round 1 cooperation rates in the treatments where cooperation is theoretically feasible.

³³In addition to the description in footnote 30, statistical tests involving comparisons across treatments control for the random sequence of supergames. This is done to take into account the potential correlations due to the fact that the realized lengths of supergames has been shown to affect choices

are *not* statistically different between perfect and private. The cooperation rates over all rounds in the last four supergames are 46% under private monitoring as depicted in the right panel of Figure 2. It is again statistically different from zero ($p < 0.01$). There is no statistical difference between perfect (46% cooperation) and private, or between public (58% cooperation) and private. These observations are summarized below:

Observation 2 Whether in round one or all rounds, the cooperation rates under private monitoring are not different from those under perfect monitoring. The rates in all rounds under private monitoring are not different from those under public monitoring.

We should emphasize that Observation 2 is remarkable considering the theoretical difficulty in supporting cooperation under private monitoring. Again, a useful comparison is with the level of cooperation in one-shot PD. Dal Bó & Fréchette (2014) assemble a data set of 122,233 choices in eleven infinitely repeated and one-shot prisoner’s dilemma experiments. They run a probit regression of round one cooperation rates on the payoff parameters g and ℓ , discount factor δ , and indicators for subgame perfection and risk dominance.³⁴ Using this regression, we can predict the level of cooperation in one-shot PD under our parametrization. We find that the observed round one cooperation rate in the private monitoring treatment is higher by 26% (significant) than the prediction. A similar regression using only data from four one-shot PD experiments with 11,038 choices also shows that the round one cooperation rate in our private monitoring treatment is higher by 35% (significant) than the predicted value. This shows that Observation 2 cannot be simply explained by the behavioral hypotheses often used to explain cooperation in one-shot PD.

7.2 Coordination

As mentioned previously, the critical feature of private monitoring is the lack of a coordination device. In theory, players can perfectly coordinate their actions under both perfect and public monitoring, but not under private monitoring. In this sense, it is interesting to see if the subjects indeed have difficulty coordinating their actions under private monitoring. Figure 3 extracts the first five rounds of each supergame in the three treatments and presents the values of $\Pr(a^t = (C, C))$ and $\Pr(a^t = (D, D))$ as well as their sum. It also depicts the values of $\Pr(a_i^t = C)^2$ and $\Pr(a_i^t = D)^2$, which would be the coordination rates should the subjects choose their

(see for instance Dal Bó & Fréchette (2011)).

³⁴This is estimated supergame by supergame.

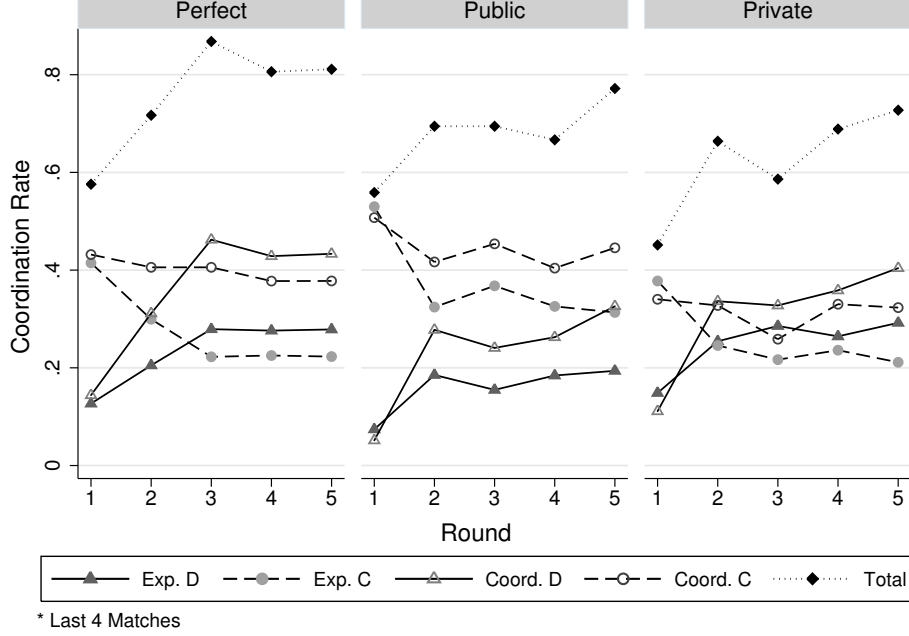


Figure 3: Coordination rates: Implied by independent action choice and realized

actions independently. Coordination rates $\Pr(a^t = (C, C) \text{ or } (D, D))$ are 0.722 and 0.712 under perfect and public monitoring, respectively, and no statistical difference exists between them. In comparison, the coordination rates are 0.660 under private monitoring, and there is a statistical difference between public and private. However, the difference is relatively small and it is surprising to see how much coordination is achieved under private monitoring after the initial round despite the difficulty implied by the theory. When we compare $\Pr(a^t = (C, C))$ and $\Pr(a^t = (D, D))$ with $\Pr(a_i^t = C)^2$ and $\Pr(a_i^t = D)^2$, respectively, we find that the former is always higher except in the first round. In fact, the coordination rates $\Pr(a^t = (C, C) \text{ or } (D, D))$ are higher by 13 percentage points than $\Pr(a_i^t = C)^2 + \Pr(a_i^t = D)^2$ for all rounds after round 1, corresponding to a 41% difference for perfect and 84% for public. For private, the rates are higher by 12 percentage points, corresponding to a 40% difference. This suggests that, to a certain extent, subjects have the correct expectation about the other player's action.

Observation 3 Coordination rates under private monitoring are (insignificantly)

lower than those under perfect and public monitoring, but are positive and significantly higher than the level implied by independent action choices.

7.3 Conditional Cooperation

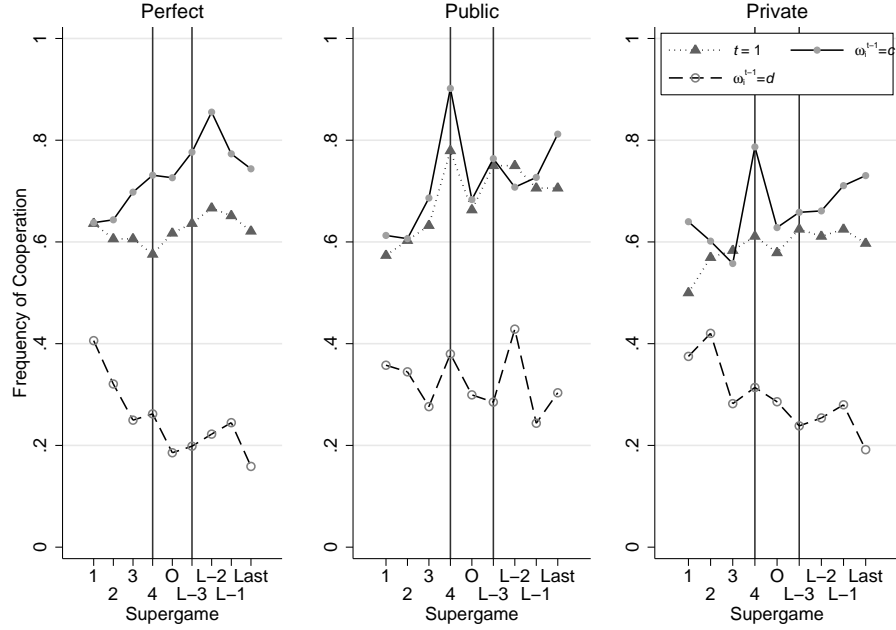


Figure 4: Cooperation conditional on the previous signal

In this subsection we explore further how cooperation rates vary with the action-signal pair in the preceding round in each monitoring environment. We begin by focussing on cooperation rates conditional only on the signal in the preceding round. Figure 4 shows the rates with which player i chooses $a_i^t = C$ in round $t \geq 2$ when his signal in round $t - 1$ is c (labeled $\omega_i^{t-1} = c$), when it is d (labeled $\omega_i^{t-1} = d$), and when $t = 1$ (labeled $t = 1$). Clearly, cooperation rates following a good signal are much higher than following a bad one ($p < 0.01$ in all cases). Another striking point is that this difference increases as the subjects accumulate experience. For instance, in the first supergame, the difference in cooperation rates following the two signals is between 23 and 26 percentage points in any treatment, whereas in the last supergame, the corresponding difference is 59 percentage points under per-

fect monitoring, 51 percentage points under public monitoring, and 54 percentage points under private monitoring.

Observation 4 In every monitoring treatment, the rate of the cooperative action C is higher after a good signal c about the opponent's action than after a bad signal d about it.

Figure 4 also shows that (1) responsiveness (= difference in cooperation rates following a c signal and a d signal defined by (6)) varies across treatments, and that (2) round one cooperation rates are about the same as cooperation rates following a good signal under perfect and public monitoring, whereas they are different under private monitoring. If we suppose that the subjects play the memory-one belief-free equilibrium described in Section 4, then responsiveness should in theory be lowest under perfect monitoring at approximately 0.292, and about 0.365 under either public or private monitoring. Our data show, however, that responsiveness under perfect monitoring is higher than that under public or private monitoring: The numbers are 0.354, 0.249, and 0.295 for the perfect, public, and private treatments, respectively.³⁵ A joint test reveals that responsiveness in the perfect monitoring treatment is statistically different ($p < 0.05$) from that in the other two. Compared with the theoretical prediction in each case, the observed responsiveness is significantly different ($p < 0.01$) in the private monitoring treatment, different but not as significantly ($p=0.094$) in the public monitoring treatment, and not different in the perfect monitoring treatment.³⁶ We also note that responsiveness is rather sensitive to the choice of a specific sample. For instance, although the same predictions should apply to all rounds after the first, if we compute responsiveness in round 2 only, it is 0.181, 0.236, and 0.305 for perfect, public, and private monitoring, respectively. Notice that responsiveness is now lower under perfect monitoring than under public and private monitoring (not statistically different however). The observed discrepancy from the theoretical prediction based on the memory-one belief-free equilibrium may come from a number of different sources. One important consideration is that strategies condition on events beyond the most recent signal. This point is examined in more detail later.

While the above analysis only considers the action choice conditional on the most recent signal, it may as well depend on one's own action in the previous round. The relationship between the action choice in the present round and one's

³⁵These numbers are obtained by first computing responsiveness for each subject and then performing t-tests on these subject averages with clustering at the session level.

³⁶Responsiveness is significantly higher under perfect monitoring than under public monitoring ($p < 0.01$), and (insignificantly) higher under perfect monitoring than under private monitoring ($p = 0.104$). Furthermore, the levels are statistically different between public and private monitoring ($p < 0.05$).

	perfect	public	private
$\Pr(a_i^t = C \mid Cc)$	0.946	0.922	0.921
$\Pr(a_i^t = C \mid Cd)$	0.414	0.553	0.469
$\Pr(a_i^t = C \mid Dc)$	0.223	0.470	0.351
$\Pr(a_i^t = C \mid Dd)$	0.114	0.135	0.105

Table 2: Signal and cooperation rate conditional on previous choice and signal combination

own action choice and signal in the previous round is summarized in the Table 2, where $(a_i^{t-1}, \omega_i^{t-1}) = (C, c)$ is abbreviated as Cc , *etc.*

First note that cooperation rates across treatments are about the same (jointly not statistically different) when the subjects previously cooperated and received a good signal (the first row) and when they previously defected and received a bad signal (the fourth row). The main differences are in the cooperation rates after Cd and Dc . Note that the higher cooperation rate after Cd implies more leniency, whereas the higher rate after Dc corresponds to more forgiveness. With this interpretation, strategies under public monitoring exhibit more leniency and forgiveness than those under perfect monitoring (both are statistically different), and strategies under private monitoring come somewhere in between in both dimensions. We should note, however, that this is only a rough measure of leniency and forgiveness as the strategies may condition on events beyond the previous round. More direct analysis of leniency and forgiveness is performed in conjunction with the estimation of strategies. Table 2 also shows that subjects under private monitoring react to a negative signal more strongly when they cooperated: the difference between the cooperation rates after Cc and after Cd is the largest in this treatment.

Observation 5 In every treatment, the rates of cooperative action C substantially vary with the signal in the previous round as well as the action-signal pair in the previous round. The cooperation rates after Cd and Dc are the highest under public monitoring and lowest under perfect monitoring.

7.4 Leniency and Forgiveness

To further investigate the question of leniency and forgiveness, we study behavior after some key histories that are possibly longer than one round.³⁷

³⁷See Fudenberg et al. (2012) for a similar exercise.

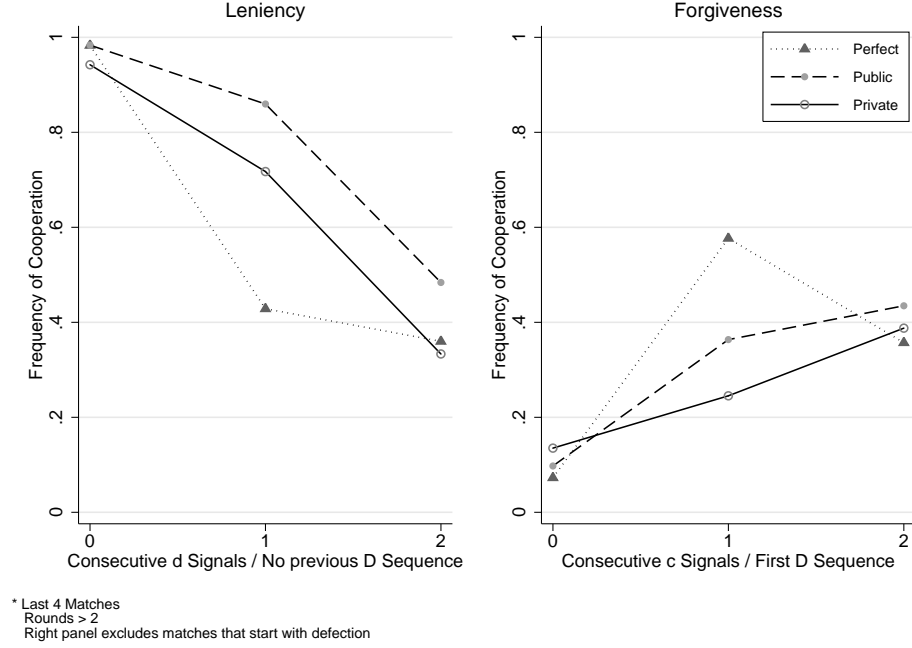


Figure 5: Cooperation rates after d signals when cooperating (Left), and after c signals when punishing (Right)

Figure 5 presents the cooperation rates after histories along which a subject has consistently chosen $a_i^t = C$ but has observed either one bad signal in the previous round or two consecutive bad signals in the two previous rounds. In other words, the height of the three points on the graph in the left panel corresponds to the values of³⁸

$$\begin{aligned} &\Pr(a_i^t = C \mid Cc, \dots, Cc), \\ &\Pr(a_i^t = C \mid Cc, \dots, Cc, Cd), \quad \text{and} \\ &\Pr(a_i^t = C \mid Cc, \dots, Cc, Cd, Cd), \end{aligned}$$

where the history h_i^{t-1} such that $(a_i^1, \omega_i^1) = \dots = (a_i^{t-1}, \omega_i^{t-1}) = (C, c)$ is abbreviated as Cc, \dots, Cc , etc. As can be seen, the drop in the cooperation rates following a single d signal is most conspicuous under perfect monitoring, suggesting the use

³⁸The figure only considers action choices in rounds three and above to allow for the observation of at least two signals.

of non-lenient strategies by the subjects. The rates under private monitoring are similar to those under public monitoring but slightly lower. There is a statistical difference between perfect and either public or private, but there is no statistical difference between the last two.³⁹

Although identification of histories that are relevant to forgiveness is less straightforward, we select the following histories as relevant. For subjects who started by cooperating: look at their first sequence of defection and compute the probability that they cooperate when they observe one c signal or two consecutive c signals. The height of the three points on the graph in the right panel of Figure 5 corresponds to the values of

$$\begin{aligned} \Pr(a_i^t = C \mid C^*, \dots, C^*, Dd, \dots, Dd), \\ \Pr(a_i^t = C \mid C^*, \dots, C^*, Dd, \dots, Dd, Dc), \quad \text{and} \\ \Pr(a_i^t = C \mid C^*, \dots, C^*, Dd, \dots, Dd, Dc, Dc), \end{aligned}$$

where C^* implies either Cc or Cd . As can be seen, there is more forgiveness under perfect than public following one cooperate signal, but less following two (these are jointly statistically different from each other). Directionally, the comparison between perfect and private is similar, but the rates following two c signals are much closer than perfect and public are. On the other hand, the difference between public and private is not statistically significant. However, ranking the treatment in terms of forgiveness on the basis of this exercise is difficult since the ranks vary following one versus two cooperate signals.

For both leniency and forgiveness, the above analysis neglects most of the choices in a supergame and hence is limited. We will return to the issue of leniency and forgiveness in the next section upon estimating strategies based on the subjects' action choices in all rounds of a supergame.

Observation 6 Conditional on some key histories, the levels of leniency are such that

$$\text{public} > \text{private} \gg \text{perfect}.$$

The ordering in terms of the forgiveness levels is not as clear except that after a single good signal in a punishment phase, a return to cooperation is most likely under perfect monitoring.

³⁹This is established by regressing cooperation on a dummy if there was one d signal and interacted with a dummy for the type of monitoring (as well as other controls).

7.5 Estimation of Strategies

We now turn to the direct estimation of the subjects' strategies. Our analysis is based on the Strategy Frequency Estimation Method (SFEM) developed in Dal Bó & Fréchette (2011). SFEM has now been used in multiple papers to estimate the strategies in repeated games, and its use is supported by Fudenberg et al. (2012) who conduct Monte-Carlo simulations to evaluate its performance, and Dal Bó & Fréchette (2013) who find that the strategies identified as most popular by SFEM are also the most popular strategies elicited from the subjects using an alternative method.⁴⁰ In essence, SFEM uses the maximum likelihood to estimate a mixture over a given set of strategies.⁴¹ The parameters that are recovered represent the estimated fraction ϕ^k of strategy k in the set, and the variance in the distribution of the error term. Instead of reporting the parameter capturing the variance in the error term, γ , we report the implied probability $\beta \equiv \frac{1}{1 + \exp(\frac{-1}{\gamma})}$ that a cooperative action would be taken when it is prescribed by a strategy. This gives an idea of how well the model fits the data since $\beta \rightarrow 1$ as $\gamma \rightarrow 0$, and $\beta \rightarrow \frac{1}{2}$ (a coin toss) as $\gamma \rightarrow \infty$.

We consult previous studies that use SFEM in PD games to determine which strategies to include. Specifically, we include all strategies that were found in a statistically significant proportion in any of the following papers: Dal Bó & Fréchette (2011); Fudenberg et al. (2012) and their re-analysis of Dal Bó & Fréchette (2011) and Dreber et al. (2008); Dal Bó & Fréchette (2013); Fréchette & Yüksel (2013); and Embrey et al. (2013). The strategies included in our analysis are listed in Table 3 and are formally described by finite automata in Appendix A.2.

The three strategies in the top panel of Table 3 do not condition on the history: “always cooperate” (AllC), “always defect” (AllD), and a strategy that cooperates in the first round and defects in all other rounds (CDDD). The second panel includes the well-known strategies of “grim trigger” (Grim) and “tit-for-tat” (TFT) as well as their variants that either do not trigger a punishment after a single d or do not immediately return to cooperation following a single c : Grim2, Grim3, TF2T, TF3T, 2TFT, and 2TF2T.⁴² Also in the second panel is the Sum2 strategy that counts the numbers of good and bad signals: It has an internal counter that is

⁴⁰Other papers using SFEM in different contexts include Vespa & Wilson (2015) on dynamic games and Bigoni et al. (2015) on continuous-time games.

⁴¹Intuitively, the method can be described as looking for the strategy from some given set that best explains the observed choices of a subject in multiple supergames. It then looks for the frequency of each strategies in the entire sample. See Dal Bó & Fréchette (2011) for the details.

⁴²TFT in the imperfect monitoring environment starts by cooperating and then chooses C if and only if $\omega_i = c$. In short, Grim- k is a variant of Grim that reverts to D after k consecutive d signals, and m-TF- n -T is a variant of TFT that plays D in at least m consecutive rounds after n consecutive d signals.

	Lenient	Forgiving	Complexity
AllC			1
AllD			1
CDDD			2
WSLS		✓	2
Sum2	✓	✓	4
Grim			2
Grim2	✓		3
Grim3	✓		4
TFT		✓	2
2TFT		✓	3
TF2T	✓	✓	3
2TF2T	✓	✓	4
TF3T	✓	✓	4
STFT		✓	2
SSum2	✓	✓	4

Table 3: Properties of strategies

initially set equal to zero. The counter is increased by 1 every time a good signal is observed and the current value is below 2, and is decreased by 1 every time a bad signal is observed and the current value is above -2 . The counter is unchanged in other cases. Sum2 plays C if the counter is ≥ 0 and D otherwise. Sum2 was first explored in Embrey et al. (2013). The second panel also lists the “win-stay, lose-shift” strategy (WSLS, sometimes referred to as “Pavlov” or “perfect tit-for-tat”) that is known to have some desirable properties in environments with noise (Imhof et al. (2007)) but has not been found in statistically significant proportions in any experiment we are aware of. Every strategy in the second panel yields a sequence of (C, C) ’s when matched against itself. On the other hand, the third panel lists the *suspicious* versions of TFT and Sum2 (STFT and SSum2) that start by defecting, and yield a sequence of (D, D) ’s when matched against itself. The check marks in Table 3 show whether these strategies are lenient and/or forgiving.⁴³ The last column of the table also shows the level of complexity of each strategy by the scale 1-4: It equals the number of states required when they are expressed as a finite automaton.⁴⁴

The results for perfect and public monitoring reproduce some results documented in the literature. First, under perfect monitoring, the majority of the data can be accounted for by the three strategies: AllD, Grim, and TFT (Dal Bó & Fréchette (2014)). Second, lenient and forgiving strategies are more popular under imperfect monitoring than under perfect monitoring (Fudenberg et al. (2012)). Third, despite its theoretical appeal, WSLS is not observed in any significant proportion in any treatment. Fourth, as in Embrey et al. (2013), Sum2 is observed in a statistically significant proportion. In other words, our results show that the findings in the literature are robust with respect to the specifications of perfect and public monitoring such as randomly generated payoffs for perfect monitoring, introduction of noise into observation rather than into action choice under public monitoring, and the cardinality and dimension of the signal space.

With respect to private monitoring, we first notice the prevalence of Sum2, which was first documented in Embrey et al. (2013) in the public monitoring environment with a linearly ordered binary signal. Second, Grim and TFT are much less popular than under perfect monitoring. In fact, TFT and all of its variants are not very popular, and none of them is statistically significant individually (nor are they jointly significant). Third, the lenient versions of Grim (Grim2 and Grim3) are more popular than in any other treatment. Although Grim2 is not statistically significant on its own, its frequency is relatively high at 9% and Grim2 and Grim3 are

⁴³Classification of leniency and forgiveness is applied only to those strategies that have transitions that depend on signals. Note that Sum2 and SSum2 are not lenient in all situations. For instance, they will play D after a single d signal in the first round.

⁴⁴See Rubinstein (1986).

	Perfect	Public	Private
AlIC	0.024 (0.032)	0.196** (0.067)	0.032 (0.038)
AlID	0.314** (0.062)	0.191** (0.057)	0.271** (0.077)
CDDD	0.000 (0.004)	0.000 (0.025)	0.000 (0.014)
WSLS	0.022 (0.045)	0.029 (0.028)	0.000 (0.040)
Sum2	0.000 (0.007)	0.114* (0.066)	0.195** (0.056)
Grim	0.117** (0.050)	0.035 (0.028)	0.0138 (0.040)
Grim2	0.046 (0.034)	0.000 (0.000)	0.090 (0.056)
Grim3	0.023 (0.038)	0.025 (0.043)	0.097** (0.035)
TFT	0.176** (0.042)	0.000 (0.034)	0.075 (0.046)
2TFT	0.000 (0.000)	0.039 (0.034)	0.056 (0.066)
TF2T	0.108** (0.042)	0.129** (0.059)	0.062 (0.049)
2TF2T	0.079* (0.045)	0.157** (0.068)	0.000 (0.010)
TF3T	0.076* (0.042)	0.059 (0.074)	0.061 (0.049)
STFT	0.015 (0.029)	0.000 (0.045)	0.032 (0.030)
SSum2	0.000	0.027*	0.014
Alpha	0.471** (0.040)	0.474** (0.038)	0.569** (0.042)
β	0.893	0.892	0.853
Cooperative	0.671	0.783	0.683
Noncooperative	0.329	0.217	0.317
Lenient	0.332	0.484	0.505
Forgiving	0.461	0.526	0.450
Complexity = 1	0.338	0.387	0.304
= 2	0.330	0.064	0.121
= 3	0.155	0.168	0.208
= 4	0.178	0.381	0.367

*** Stat. sig. at the 1%, ** 5%, and * 10% levels.

Top panel are unconditional strategies, 2nd panel are conditional and cooperative, and 3rd panel are conditional and defective.

Bottom panel are total frequencies by feature.

Table 4: Estimates of proportion of each strategy

jointly statistically significant ($p < 0.01$). This is in sharp contrast with perfect and public monitoring, where neither of these two strategies is ever statistically significant. Fourth, going from perfect to private, strategies become more lenient, but not much more forgiving: 51% of strategies are lenient under private monitoring while 33% are lenient under perfect monitoring ($p = 0.13$). On the other hand, 15% are non-lenient under private monitoring whereas 32% are non-lenient under perfect monitoring ($p < 0.1$). In fact, the proportion of lenient strategies under private monitoring (51%) is similar to that under public monitoring (48%) (p -value = 0.86 for the equality between public and private). As for forgiveness, the proportion of forgiving strategies is lower under private monitoring (45%) than under public monitoring (53%) (insignificant at $p = 0.52$), and the proportion of unforgiving strategies is higher under private monitoring (20%) than under public monitoring (6%) (again insignificant at $p = 0.12$). On the other hand, the forgiveness level under private monitoring is similar to that under perfect monitoring.

This finding on forgiveness hence is at odds with the result of the reduced form approach in Section 7.3 that the highest forgiveness level after a single c signal was under perfect monitoring. The discrepancy may come from the fact that the analysis in Section 7.3 is restricted to behavior after particular histories. The conflicting findings are however partly reconciled by the observation that the estimated fraction of the TFT variants that return to cooperation immediately after a single c signal (TFT, TF2T, and TF3T) is higher (0.36 in total) under perfect monitoring than under public monitoring (0.19) or private monitoring (0.20).

In relation to leniency and forgiveness, we find that the strategies become more complex when monitoring becomes imperfect whether it is public or private. Specifically, the estimated proportion of strategies which has just two states in the automaton representation is higher under perfect monitoring (33%) than under either public (6%, $p < 0.01$) or private (12%, $p < 0.05$) monitoring. Likewise, the estimated proportion of strategies with three or four states is one third under perfect monitoring but slightly more than a half in both public and private monitoring ($p < 0.05$ for both comparisons). The average number of states equals 2.172, 2.544, and 2.638 for perfect, public, and private, respectively.

The most important differences across treatments can be gleaned by focusing on the top three strategies in each treatment as listed in Table 5.

Notice that top three strategies represent more than 50% in proportion in all three treatments. While the non-cooperative strategy AllD is always very popular, characteristics of cooperative strategies are markedly different in the three treatments. Under perfect monitoring, both Grim and TFT are non-lenient. Under perfect monitoring, 2TF2T is both lenient and forgiving. Under private monitoring, both Sum2 and Grim3 are lenient, but the latter is non-forgiving. It is only under private monitoring that the intuitive strategy of Sum2 that counts the numbers of c 's and d 's is

Popularity	Monitoring		
	Perfect	Public	Private
1 st	AllD	AllC	AllD
2 nd	TFT	AllD	Sum2
3 rd	Grim	2TF2T	Grim3

Table 5: Top strategies by treatment

in the top three.

Observation 7 Strategies under public and private monitoring are more complex than those under perfect monitoring. The increased complexity comes mainly from the lenient and forgiving variants of TFT under public monitoring, and from the lenient (but not forgiving) variants of Grim under private monitoring. Under both public and private monitoring, Sum2, which counts the numbers of good and bad signals and is sometimes lenient and forgiving, is important.

8 Conclusion

While theory suggests the importance of a monitoring structure on the play of a repeated game, experimental work on the subject is still limited. This paper presents one approach to the problem by comparing three major monitoring structures using the same PD as a stage game.

Our findings from the perfect and public monitoring treatments serve as robustness checks of earlier results in the experimental literature as follows. First, we find that cooperation is sustained even when payoffs are randomly determined if the actions are perfectly monitored. This observation is in line with the finding of Rand et al. (2015) on an infinitely repeated PD that subjects for the most part ignore random outcomes if they know their opponents' intentions. Second, we observe cooperation under imperfect public monitoring which is specified differently from that in prior experiments.⁴⁵ Third, we confirm the key finding from Fudenberg et al. (2012) that strategies become more lenient and forgiving under public monitoring than under perfect monitoring. We find this true under an alternative specification of each monitoring structure as well as under an additional control on the expected stage payoffs across the two monitoring treatments.

The primary focus of our analysis is on the comparison of private monitoring with perfect and public monitoring. While theory suggests the difficulty of co-

⁴⁵See Dal Bó & Fréchette (2014) for a review.

operation under private monitoring, we observe that the subjects maintain almost the same level of cooperation under private monitoring as under perfect and public monitoring. Even more surprisingly, the rates of coordination on either (C, C) or (D, D) is significantly higher than the hypothetical rates that would be obtained when those actions are chosen independently at the observed rates.

We also find substantial difference in behavior under the three monitoring structures. Specifically, a reduced form approach based on choices after key histories, and strategy estimation both reveal that the behavior is more lenient under private monitoring than under perfect monitoring. In both cases, the leniency level is similar under private and public monitoring. On the other hand, the comparison of forgiveness levels is less clear. While strategy estimation suggests that similar forgiveness levels under perfect and private monitoring, and a higher level under public monitoring, the reduced form approach suggests the highest forgiveness level under perfect monitoring.

In relation to leniency and forgiveness, we find that the strategies become more complex when monitoring becomes imperfect whether it is public or private. In particular, the average number of states in the finite automaton representation increases as we move from perfect to public, and from public to private. Under private monitoring, Sum2, which is complex and uses four states, is found to be one of the top three strategies. It is interesting to note that as in the public monitoring treatment of Embrey et al. (2013), where Sum2 is first observed, our private monitoring treatment has the feature that the signal is binary and can be interpreted as either good or bad.

As in Matsushima & Toyama (2013), our finding regarding responsiveness is not necessarily consistent with the play of a memory-one belief-free equilibrium, which is widely used in the repeated game literature. Specifically, if subjects played the memory-one belief-free equilibrium in all treatments, then the responsiveness level should be lowest under perfect monitoring and the same under public and private monitoring. We observe instead that the level under perfect monitoring is the highest. This implies at least that such an equilibrium is not played in all treatments. Along with the results of our strategy estimation, we suspect that the restriction to memory one is among other reasons for the observed deviation.

A full account of the behavior reported in this experiment would require the development of a new theory based on the combination of such elements as the complexity cost of strategies, preference for efficiency, and the importance of intentions: The fact that the estimated strategies are more complex in a more complex environment suggests that the complexity of a strategy is perceived as a cost by the subjects. In other words, simple strategies are preferred so long as they entail no efficiency loss. The Sum2 strategy is popular in the private monitoring treatment perhaps because it is considered the simplest rule of thumb that works in the en-

vironment. Substantial rates of cooperation and coordination in every treatment and leniency of strategies in the imperfect monitoring treatments both imply preference for efficiency. The higher responsiveness under perfect monitoring suggests that the subjects have a stronger incentive to react to the opponent's action when his intention is clearer. The reaction can also be used to discipline the opponent when there is little or no noise. We view these as interesting insights to guide future theory work on repeated games.

References

- Aoyagi, M. & Fréchette, G. R. (2009), 'Collusion as public monitoring becomes noisy: Experimental evidence', *Journal of Economic Theory* **144**(3), 1135–1165.
- Battaglini, M., Nunnari, S. & Palfrey, T. R. (2015), The dynamic free rider problem: A laboratory study. Working Paper.
- Bernard, M., Fanning, J. & Yüksel, S. (2014), Finding cooperators: sorting through repeated interaction. Working Paper.
- Bigoni, M., Casari, M., Skrzypacz, A. & Spagnolo, G. (2015), 'Time horizon and cooperation in continuous time', *Econometrica* **83**(2), 587–616.
- Cabral, L. M. B., Ozbay, E. & Schotter, A. (2014), 'Intrinsic and instrumental reciprocity: An experimental study', *Games and Economic Behavior* **87**, 100–121.
- Camera, G. & Casari, M. (2009), 'Cooperation among strangers under the shadow of the future', *American Economic Review* **99**(3), 979–1005.
- Camera, G. & Casari, M. (2011), 'The coordination value of monetary exchange: experimental evidence', *American Economic Journal: Microeconomics* **6**(1), 290–314.
- Camera, G., Casari, M. & Bigoni, M. (2012), 'Cooperative strategies in anonymous economies: an experiment', *Games and Economic Behavior* **75**(2), 570–586.
- Cason, T. N. & Khan, F. U. (1999), 'A laboratory study of voluntary public goods provision with imperfect monitoring and communication', *Journal of Development Economics* **58**, 533–552.

- Cason, T. N. & Mui, V.-L. (2015), 'Coordinating collective resistance through communication and repeated interaction', *Journal of Experimental Political Science* **2**(1), 94–106.
- Cooper, D. J. & Kühn, K.-U. (2014), 'Communication, renegotiation, and the scope for collusion', *American Economic Journal: Microeconomics* **6**(2), 247–278.
- Dal Bó, P. (2005), 'Cooperation under the shadow of the future: experimental evidence from infinitely repeated games', *American Economics Review* pp. 1591–1604.
- Dal Bó, P. & Fréchette, G. (2011), 'The evolution of cooperation in infinitely repeated games: Experimental evidence', *American Economic Review* **101**(1), 411–429.
- Dal Bó, P. & Fréchette, G. (2013), Strategy choice in the infinitely repeated prisoners dilemma. Working Paper.
- Dal Bó, P. & Fréchette, G. (2014), On the determinants of cooperation in infinitely repeated games: A survey. Working Paper.
- Dreber, A., Rand, D. G., Fudenberg, D. & Nowak, M. A. (2008), 'Winners don't punish', *Nature* **452**, 348–351.
- Duffy, J. & Ochs, J. (2009), 'Cooperative behavior and the frequency of social interaction', *Games and Economic Behavior* **66**(2), 785–812.
- Ely, J. & Välimäki, J. (2002), 'A robust folk theorem for the prisoner's dilemma', *Journal of Economic Theory* **102**(1), 84–105.
- Embrey, M. S., Fréchette, G. R. & Stacchetti, E. (2013), An experimental study of imperfect public monitoring: Renegotiation proofness vs efficiency. Working Paper.
- Engle-Warnick, J. & Slonim, R. L. (2004), 'The evolution of strategies in a repeated trust game', *Journal of Economic Behavior & Organization* **55**(4), 553–573.
- Engle-Warnick, J. & Slonim, R. L. (2006a), 'Inferring repeated-game strategies from actions: Evidence from repeated trust game experiments', *Economic Theory* **28**(3), 603–632.
- Engle-Warnick, J. & Slonim, R. L. (2006b), 'Learning to trust in indefinitely repeated games', *Games and Economic Behavior* **54**(1), 95–114.

- Feinberg, R. M. & Husted, T. A. (1993), 'An experimental test of discount-rate effects on collusive behaviour in duopoly markets', *Journal of Industrial Economics* **41**(2), 153–160.
- Feinberg, R. M. & Snyder, C. (2002), 'Collusion with secret price cuts: An experimental investigation', *Economics Bulletin* **3**(6), 1–11.
- Fréchette, G. R. (2012), 'Session-effects in the laboratory', *Experimental Economics* **15**(3), 485–498.
- Fréchette, G. R. & Yüksel, S. (2013), Infinitely repeated games in the laboratory: Four perspectives on discounting and random termination. Working Paper.
- Fudenberg, D., Rand, D. G. & Dreber, A. (2012), 'Slow to anger and fast to forget: Leniency and forgiveness in an uncertain world', *American Economic Review* **102**(2), 720–749.
- Fudenberg, D., Rand, D. G. & Dreber, A. (2014), 'Who cooperates in repeated games: The role of altruism, inequity aversion, and demographics', *Journal of Economic Behavior & Organization* **98**, 41–55.
- Harrington, J. & Skrzypacz, A. (2007), 'Collusion under monitoring of sales', *The RAND Journal of Economics* **38**(2), 314–331.
- Harrington, J. & Skrzypacz, A. (2011), 'Private monitoring and communication in cartels: Explaining recent collusive practices', *American Economic Review* **101**(6), 2425–2449.
- Holcomb, J. H. & Nelson, P. S. (1997), 'The role of monitoring in duopoly market outcomes', *Journal of Socio-Economics* **26**(1), 79–93.
- Holt, C. A. (1985), 'An experimental test of the consistent-conjecture hypothesis', *American Economic Review* **75**(3), 314–325.
- Imhof, L. A., Fudenberg, D. & Nowak, M. A. (2007), 'Tit-for-tat or win-stay, lose-shift?', *Journal of Theoretical Biology* **247**(3), 574–580.
- Kandori, M. (1992), 'The use of information in repeated games with imperfect monitoring', *Review of Economic Studies* **59**, 581–593.
- Levenstein, M. & Suslow, V. (2006), 'What determines cartel success?', *Journal of Economic Literature* **44**(1), 43–95.

- Matsushima, H. & Toyama, T. (2013), Monitoring accuracy and retaliation in infinitely repeated games with imperfect private monitoring: Theory and experiments. Working Paper.
- Murnighan, J. K. & Roth, A. E. (1983), 'Expecting continued play in prisoner's dilemma games', *Journal of Conflict Resolution* **27**(2), 279–300.
- Palfrey, T. R. & Rosenthal, H. (1994), 'Repeated play, cooperation and coordination: An experimental study', *The Review of Economic Studies* **61**(3), 545–565.
- Piccione, M. (2002), 'The repeated prisoner's dilemma with imperfect private monitoring', *Journal of Economic Theory* **102**(1), 70–83.
- Rand, D. G., Fudenberg, D. & Dreber, A. (2015), 'It's the thought that counts: The role of intentions in noisy repeated games', *Journal of Economic Behavior and Organization* **116**, 481–499.
- Roth, A. E. (1995). "Bargaining Experiments" in *Handbook of Experimental Economics*, ed. by John H. Kagel and Alvin E. Roth. Princeton: Princeton University Press, 253–348.
- Roth, A. E. & Murnighan, J. K. (1978), 'Equilibrium behavior and repeated play of the prisoner's dilemma', *Journal of Mathematical Psychology* **17**, 189–198.
- Rubinstein, A. (1986), 'Finite automata play the repeated prisoner's dilemma', *Journal of Economic Theory* **39**, 83–96.
- Schwartz, S. T., Young, R. A. & Zvinakis, K. (2000), 'Reputation without repeated interaction: A role for public disclosures', *Review of Accounting Studies* **5**(4), 351–375.
- Sekiguchi, T. (1997), 'Efficiency in repeated prisoner's dilemma with private monitoring', *Journal of Economic Theory* **76**(2), 345–361.
- Sherstyuk, K., Tarui, N., Ravago, M.-L. V., & Saijo, T. (2013), 'Payment schemes in random termination experimental games', *Experimental Economics* **16**, 125–153.
- Stigler, G. (1964), 'A theory of oligopoly', *The Journal of Political Economy* **72**(1), 44–61.
- Vespa, E. (2015), An experimental investigation of strategies in the dynamic common pool game. Working Paper.
- Vespa, E. & Wilson, A. J. (2015), Dynamic games and markov perfection: Putting the "conditional" in conditional cooperation. Working Paper.

A Appendix

A.1 Cooperation under perfect monitoring in Dal Bó & Fréchet (2013).

In Dal Bó & Fréchet (2013), each session in the $\delta = 0.5$ treatment has at least 19 supergames, while the three sessions in the $\delta = 0.9$ treatment have 12, 18 and 19 supergames. Given that there are at most 19 supergames in the current experiments, Figure 6 includes at most 19 supergames to make comparison easier. Dal Bó & Fréchet (2013) specify the stage-payoffs as $u_i(D, C) = 50$, $u_i(C, C) = 32$, $u_i(D, D) = 25$, and $u_i(C, D) = 12$, making the stage-game strategically equivalent to (2) for $g = \frac{25}{7} - 1 \approx 2.57$ and $\ell = \frac{13}{7} \approx 1.86$.

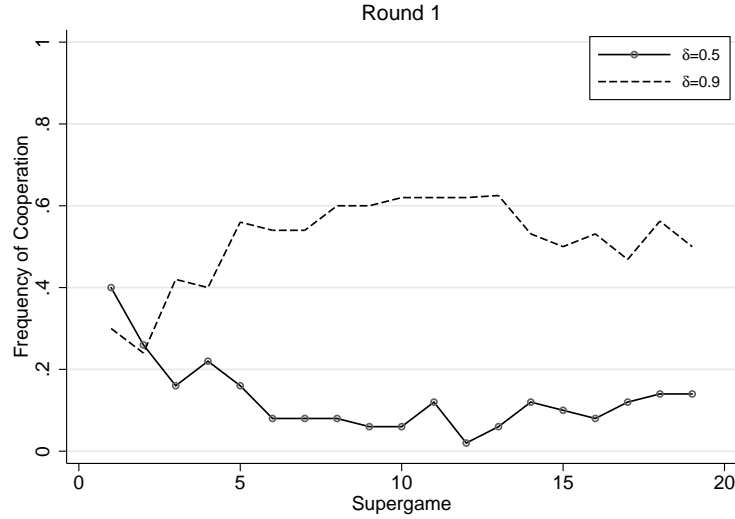


Figure 6: Cooperation rates in Dal Bó & Fréchet (2013) by supergame

A.2 Strategies included in the estimation



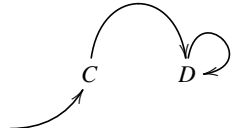
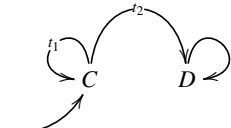
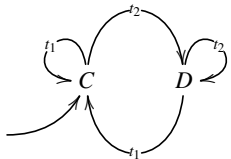
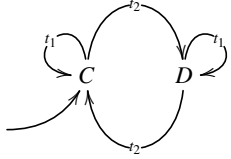
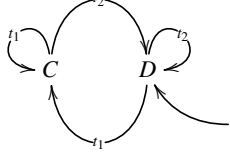
Automaton name in text	Diagram	Perfect and Public	Private
AllC			
AllD			
CDDD			
Grim		$t_1 = \{a_i = C, \omega = (c, c)\}$ $t_2 = \neg t_1$	$t_1 = \{a_i = C, \omega_i = c\}$ $t_2 = \neg t_1$
TFT		$t_1 = \{\omega_i = c\}$ $t_2 = \neg t_1$	
WSLS		$t_1 = \{\omega_i = c\}$ $t_2 = \neg t_1$	
STFT		$t_1 = \{\omega_i = c\}$ $t_2 = \neg t_1$	

Table 6: Unconditional and Two-States Automata

Automaton name in text	Diagram	Perfect and Public	Private
Grim2		$t_1 = \{a_i = C, \omega = (c, c)\}$ $t_2 = \neg t_1$	$t_1 = \{a_i = C, \omega_i = c\}$ $t_2 = \neg t_1$
Grim3		$t_1 = \{a_i = C, \omega = (c, c)\}$ $t_2 = \neg t_1$	$t_1 = \{a_i = C, \omega_i = c\}$ $t_2 = \neg t_1$
TF2T			$t_1 = \{\omega_i = c\}$ $t_2 = \neg t_1$
TF3T			$t_1 = \{\omega_i = c\}$ $t_2 = \neg t_1$
2TFT			$t_1 = \{\omega_i = c\}$ $t_2 = \neg t_1$
2TF2T			$t_1 = \{\omega_i = c\}$ $t_2 = \neg t_1$
Sum2			$t_1 = \{\omega_i = c\}$ $t_2 = \neg t_1$
SSum2			$t_1 = \{\omega_i = c\}$ $t_2 = \neg t_1$

Table 7: Automata with More Than Two States