

The Impact of Monitoring in Infinitely Repeated Games: Perfect, Public, and Private[†]

By MASAKI AOYAGI, V. BHASKAR, AND GUILLAUME R. FRÉCHETTE*

This paper uses a laboratory experiment to study the effect of the monitoring structure on the play of the infinitely repeated prisoner's dilemma. Keeping the strategic form of 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 across the three treatments. Specifically, the strategies under imperfect monitoring are both more complex and more lenient than those under perfect monitoring. The results show how the changes in strategies across monitoring structures mitigate the effect of noise in monitoring on efficiency. (JEL C72, C73, C92, D82, D83)

Many economic situations involve repeated interactions among players. While the players can perfectly monitor other players' past play in some of those interactions, they often observe only noisy information about it (*imperfect monitoring*). Imperfect monitoring is further classified into (*imperfect*) *public monitoring*, where the noisy information is publicly observable, and (*imperfect*) *private monitoring*, where it is only privately observed by each player. Under private monitoring, hence, the players do not know what signals other players have observed about their own play. For example, the classic model of collusion in a quantity-setting oligopoly by Green and Porter (1984) is an instance of public monitoring, where the firms publicly observe the market price that imperfectly signals their quantity choices, whereas the model of collusion in a price-setting oligopoly by Stigler (1964) is an instance of private monitoring, where the firms privately observe the demand for their own good that imperfectly signals the other

*Aoyagi: Osaka University, ISER, Osaka University, 6-1 Mihogaoka, Ibaraki, Osaka 567-0047, Japan (email: aoyagi@iser.osaka-u.ac.jp); Bhaskar: University of Texas at Austin, 2225 Speedway Stop C3100, Austin, TX 78731 (email: v.bhaskar@austin.utexas.edu); Fréchette: New York University, 19 West 4th St., New York, NY 10012 (email: frechette@nyu.edu). We wish to thank the referees, Matthew Embrey, Ryan Stevens, 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. Finally, Bhaskar would like to thank the NSF for its support via grant SES-201503942; Fréchette would like to thank the NSF for its support via grant SES-1225779, the CV Starr Center, and CESS; and Aoyagi would like to thank the JSPS (grant numbers: 21653016, 24653048, 15K13006, 22330061, and 15H03328), and the Joint Usage/Research Center at ISER for research support. We are responsible for all errors.

[†]Go to <https://doi.org/10.1257/mic.20160304> to visit the article page for additional materials and author disclosure statement(s) or to comment in the online discussion forum.

firms' prices. While there is now extensive theory of repeated games in the three monitoring environments, there is little empirical work on what effects they have on the players' ability to cooperate/collude and their choice of repeated game strategies. The scarcity of empirical work is in part attributed to the difficulty of obtaining data in the field: publicly available datasets are unlikely to include the information the firms use to collude in the repeated interactions, or 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 explore subjects' behavior in a laboratory in the three distinct monitoring environments.

In our experiments, the subjects play the repeated prisoners' dilemma (PD) game in the three monitoring environments described as follows: in the *perfect* monitoring treatment, at the end of each round, player i observes player j 's action choice a_j in that round. In the two imperfect monitoring treatments, player j 's action a_j generates a signal for player i , ω_i . This signal is correct and equals player j 's action a_j with probability $1 - \varepsilon$, but is incorrect with probability ε , where $\varepsilon = 0.1$. At the end of each round, in the public monitoring treatment, player i observes the pair of noisy signals so generated, (ω_i, ω_j) . In the *private* monitoring treatment, he only observes ω_i .¹ The signal profile in the public monitoring treatment is hence common knowledge since it is observed by both players, whereas no event that is informative about a player's action choice is common knowledge in the private monitoring treatment. In order to focus on the effect of the difference in the monitoring structures, we keep fixed other elements of the game as much as possible: the three treatments have the same expected stage payoffs and the same continuation probability $\delta = 0.9$, which is interpreted as the discount factor. Furthermore, the payoffs in the perfect monitoring treatment are determined randomly by the same probability distribution as in the imperfect monitoring treatments to control for the effect of uncertainty.

Our first question concerns the ability of experimental subjects to sustain cooperation in the three monitoring environments. Since previous laboratory experiments find a positive level of cooperation under perfect and public monitoring, our central focus is on the subjects' ability to cooperate under private monitoring. As Kandori (2002) notes, and by now is well recognized, private monitoring implies the absence of crucial common knowledge events: the players cannot simultaneously begin the punishment or return to cooperation based on their signals precisely because those signals are private. This is in contrast to the perfect monitoring environment where the history of past actions is common knowledge, and the public monitoring environment where the history of public signals is common knowledge. In both of these environments, play can be coordinated based on these common knowledge events. From a theoretical perspective, cooperation in the absence of common knowledge of signals is difficult. Unlike in the perfect and public monitoring environments where cooperation can be sustained by simple strategies, existing constructions of cooperative equilibria under private

¹ Accordingly, the information observed by a player takes one of two values in the private monitoring treatment, whereas it takes one of four values in the public monitoring treatment.

monitoring entail the play of intricate strategies.² Hence, it would not be surprising if laboratory subjects fail to cooperate under private monitoring. We attempt to carefully isolate the effect of the absence of common knowledge of histories by using the private monitoring treatment that differs from the other treatments only in the monitoring structure.

We next examine if and how the subjects' behavior in the three monitoring environments differ. We approach this problem in two different ways: we first examine if their behavior after certain histories is different under different monitoring structures, and then 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 the punishment after a single bad signal, or *forgiving* in the sense that it returns to cooperation after punishing the opponent? Our analysis also provides an indirect test of the theory of private monitoring through the examination of *responsiveness*, which is the difference in the likelihood of the cooperative action after two different signals about the opponent's action choice.

Our findings can be summarized as follows. First, the cooperation rate under private monitoring is comparable to those under perfect and public monitoring, and significantly higher than that predicted in the one-shot environment. 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, the subjects play 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. Although comparisons in terms of forgiveness are less conclusive, there is some suggestive evidence that strategies used under private monitoring are not as forgiving as those under public monitoring. Furthermore, when the complexity of each strategy is measured by the number of states in their finite automaton representation, the strategies under private monitoring are more complex than those under perfect monitoring. These findings suggest that subjects find ways to cooperate and coordinate using a different mechanism under each monitoring structure.

To examine the effects of increasing noise in monitoring on the subjects' ability to cooperate, we conduct additional treatments in which the noise level is doubled: $\varepsilon = 0.2$.³ In these high noise treatments, we observe a substantial drop in the level of cooperation in all three monitoring environments. The levels of cooperation in

²In particular, randomization is used either to generate correlation and coordination of continuation play (Sekiguchi 1997, Bhaskar and Obara 2002, and Bhaskar and van Damme 2002), or to make such coordination unnecessary, as in "belief-free" equilibria (Piccione 2002; Ely and Välimäki 2002; and Ely, Hörner, and Olszewski 2005).

³Recall that our perfect monitoring treatments entail random payoffs whose distribution depends on ε . A larger ε hence implies larger uncertainty in payoff realizations.

these high noise treatments are again comparable to one another, but are statistically indistinguishable from what can be expected in a one-shot PD. In other words, despite the theoretical possibility for cooperation, the high noise treatments with any monitoring structure fail to create dynamic incentives required to sustain cooperation. This finding from the additional treatments indicates that the ability to cooperate is sensitive to the environment, and can sometimes be more significantly affected by the degree of randomness in outcomes than by the monitoring structure. A deeper investigation of this merits further work.

The organization of the paper is as follows: in the next section, we give a brief review of the literature. Section II formulates a model of repeated PD, Section III provides a theoretical background, and Section IV describes the experimental design. The questions our analysis attempts to answer are listed in Section V, and the results are presented in Section VI. The findings from the high noise treatments are discussed in Section VII. Section VIII concludes with a discussion.

I. 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 and Suslow (2006) identifies joint sales agencies and industry associations as mechanisms that help cartels through the collection and dissemination of information. Harrington and Skrzypacz (2007, 2011) finds 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 interfirm sales as a way to transfer profits to sustain collusion. Their work suggests that when firms collude in private monitoring environments, they make arrangements for making information public, and also make side payments, thus attesting to the difficulty of sustaining collusion under private monitoring.

Our primary objective in 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.⁴

Early experimental studies find some cooperation when subjects engage in repeated interactions under perfect monitoring.⁵ Further evidence of cooperation in repeated games was provided by Engle-Warnick and Slonim (2004, 2006a, b), Dal Bó (2005), Aoyagi and Fréchette (2009), and Duffy and Ochs (2009) in various settings (these subsequent studies differ from the earlier ones in that they

⁴Experiments on infinitely repeated games address a number of different questions. They include, to mention a few, Schwartz, Young, and Zvinakis (2000); Dreber et al. (2008) on modified PD; Cason and Mui (2014) on a collective resistance game; Cooper and Kühn (2014) on the role of communication and renegotiation; Dreber, Fudenberg, and Rand (2014) on the relationship between behavior in the dictator game and that in an infinitely repeated game; Cabral, Ozbay, and Schotter (2014) on reciprocity; and Bernard, Fanning, and Yuksel (2017) on a gift exchange game. Other forms of dynamic games are studied by Battaglini, Nunnari, and Palfrey (2015) and Vespa (2016).

⁵See Roth (1995). Early studies include Roth and Murnighan (1978), Murnighan and Roth (1983), Feinberg and Husted (1993), Holt (1985), and Palfrey and Rosenthal (1994).

allow subjects to play multiple repeated games). Dal Bó and Fréchette (2011) finds in perfect monitoring games that cooperation rates by experienced subjects are very low when cooperation is theoretically infeasible, and higher when it is theoretically feasible, and very high for certain parameter values. Furthermore, Dal Bó and Fréchette (2018b) finds 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 always D (defect); grim-trigger, which begins with C (cooperate) but switches to D forever following a defection; and tit-for-tat (TFT), which begins with C and thereafter mimics the other player's action in the previous round.

On games with imperfect public monitoring, Aoyagi and Fréchette (2009) finds that subjects cooperate in an environment with a continuous public signal, and that their payoff decreases with the level of noise in the public signal in line with the theoretical prediction on the maximal symmetric *perfect public equilibrium* payoff.⁶ Fudenberg, Rand, and Dreber (2012) studies a model of repeated PD under imperfect public monitoring that is close to our public monitoring treatment, and examines the effects of stage payoffs and noise levels on subjects' behavior. Fudenberg, Rand, and Dreber (2012) finds no systematic difference in the levels of cooperation (initial and overall) as the level of noise in public monitoring is increased (starting from perfect monitoring), but that the subjects' strategies under public monitoring are more *lenient* and more *forgiving* than under perfect monitoring. They support this finding both by analyzing specific histories and using the strategy frequency estimation method proposed in Dal Bó and Fréchette (2011). These results prompt us to study the leniency and forgiveness properties of strategies in the private monitoring environment.⁷

The present formulation of public monitoring differs from those in Aoyagi and Fréchette (2009) and Fudenberg, Rand, and Dreber (2012) in a few important ways. As mentioned above, the model of Aoyagi and Fréchette (2009) has a *continuous* public signal whose distribution depends on the *sum* of the two players' actions. This, in particular, implies that statistical identification of a deviator is not possible in Aoyagi and Fréchette (2009) unlike in the present model where it is possible since the public signal consists of two components that correspond to each player's action choice. This distinction is known to be very important theoretically. Fudenberg, Rand, and Dreber (2012) formulates the public signal in the same way as in the present paper, but suppose that the public signal not only signals a player's action choice, but also determines his payoff. We instead suppose that a player's

⁶ Some, including Cason and 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 and Murnighan (1978). See Fréchette and Yuksel (2017) for some alternative termination methods and Sherstyuk, Tarui, and Saijo (2013) for alternative payment methods.

⁷ Other related studies of infinitely repeated PD games with imperfect public monitoring include Rojas (2012); Embrey, Fréchette, and Stacchetti (2013); and Rand, Fudenberg, and Dreber (2015). Rojas (2012) is interested in comparing behavior in an environment à la Green and Porter (1984) versus one à la Rotemberg and Saloner (1986) and thus varies monitoring accordingly. Embrey, Fréchette, and Stacchetti (2013) explores the explanatory power of renegotiation proofness within an imperfect public monitoring environment with communication. Starting with the public monitoring technology used in Fudenberg, Rand, and Dreber (2012), Rand, Fudenberg, and Dreber (2015) studies the impact of revealing the intended play of each players, thus making the treatment one with perfect monitoring.

payoff is determined by his own action and the component of the public signal that corresponds to the other player's action.

To the best of our knowledge, Kayaba, Matsushima, and Toyama (2016) offers the only other study of private monitoring games in a standard setting.⁸ Kayaba, Matsushima, and Toyama (2016) conducts experiments with two treatments that vary the accuracy of players' private signal, and compare behavior with predictions by focusing on a family of strategies they refer to as *generous TFT* strategies, which choose a mixed action as a function only of the most recent signal (i.e., have *memory-one*), and play *C* with a higher probability after a *c* signal than after a *d* signal.^{9,10} There are a few important differences between the present paper and Kayaba, Matsushima, and Toyama (2016): first, our primary objective is to examine the possible difference in behavior across different monitoring environments including private monitoring. For this purpose, our experiments are designed so as to yield the clearest evidence. In particular, we consult the previous studies on the topic and choose parameter values that allow us to statistically distinguish the cooperation levels in repeated games (under perfect and public monitoring) from that in the one-shot games.¹¹ Second, it is recognized in this literature that there are often important changes in behavior for inexperienced subjects. In view of this, our analysis of subjects' behavior focuses on what happens after they accumulate experience.¹² Third, in terms of the analysis, we do find that our subjects condition behavior on histories longer than one round. Unlike Kayaba, Matsushima, and Toyama (2016), we take candidate strategies from a more general class than the memory-one class. Our results indicate that the complexity of strategies, such as memory length, is in fact a key difference in the subjects' behavior across different monitoring environments.

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 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 that cooperates as long as all past interactions have resulted in (C, C) , but defects otherwise. Thus, a single defection results in the

⁸ Earlier experimental studies that featured private monitoring in repeated games did so in environments that do not lend themselves to exploring standard concepts from the theoretical work on the topic. Holcomb and Nelson (1997) observes 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." Feinberg and Snyder (2002) studies the effect of occasional manipulation of payoff numbers in a modified repeated PD (a third choice is added), and finds less collusive behavior when such manipulation is *ex post* not revealed than when it is.

⁹ As seen in Section III, these strategies form the building block of the belief-free equilibrium. Fudenberg, Rand, and Dreber (2012) (online Appendix) finds none of the generous TFT strategies in significant proportions in their strategy estimation for public monitoring games.

¹⁰ Kayaba, Matsushima, and Toyama (2016) has two treatments: a player's private signal equals the other player's action with probability 0.9 in the high-accuracy treatment, and with probability 0.6 in the low-accuracy treatments. In both treatments, $g = \ell = 2/9$ according to our notation.

¹¹ This is important since for some payoff combinations and level of experience, positive cooperation rates are observed in experiments in contradiction to the theoretical prediction. See Dal Bó and Fréchet (2018b).

¹² See also Dal Bó and Fréchet (2018b). This issue is not addressed in Kayaba, Matsushima, and Toyama (2016), who let the subjects play only three supergames under each parametrization.

breakdown of cooperation through the contagion process. Duffy and Ochs (2009) finds that even when the group size is fairly small—as low as six—their subjects cannot sustain cooperation in the random matching environment.¹³ Subsequent experiments by Camera and Casari (2009); Camera, Casari, and Bigoni (2012); and Camera and Casari (2014) also confirm that cooperation in the random re-matching environment is fragile and is possible only for very small groups, of size four. The difficulty of supporting cooperation with random re-matching in small groups naturally poses a question as to whether players can sustain cooperation in bilateral interactions with private monitoring.

II. 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 $\Pr(\omega_j = c | a_i = C) = \Pr(\omega_j = d | a_i = D) = 1 - \varepsilon$. The two signals ω_1 and ω_2 are independent conditional on the action profile $a = (a_1, a_2)$.

The payoff of player i , in the imperfect monitoring treatments, depends on 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

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

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

$$(2) \quad \begin{array}{c} \\ C \\ D \end{array} \begin{array}{cc} C & D \\ \hline 1 & 1 & -\ell & 1 + g \\ 1 + g & -\ell & 0 & 0 \end{array}.$$

In the perfect monitoring treatment, the expected stage game payoffs are also given by the above table. Furthermore, in order to ensure commonality across treatments, the realized payoffs under perfect monitoring given any action profile are random, and have the same distribution as under imperfect monitoring. In our experiments, the parameters g and $\ell > 0$ are chosen to satisfy $g = \ell$.¹⁴

Under *perfect monitoring*, player i observes j 's action a_j at the end of each round. 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 .

¹³Their results are replicated in Figure A1 in Appendix A.

¹⁴The equality $g = \ell$ implies that the expected payoff table has a benefit-cost form à la Fudenberg, Rand, and Dreber (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 Kayaba, Matsushima, and Toyama (2016). The benefit-cost (b/c) ratio is given by $(1 + g)/g$.

Let x_i^t denote player i 's information at the end of round t regarding events in the round as described above.¹⁵ 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 $\pi_i(\sigma|h_i)$ be i 's expected continuation payoff under σ following history $h_i \in H_i$.¹⁶ A strategy profile $\sigma = (\sigma_1, \sigma_2)$ is a *perfect Bayesian equilibrium* (PBE, or simply an *equilibrium*) of the repeated game if for $i = 1, 2$,

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

for any alternative strategy σ'_i and any private history $h_i \in H_i$.¹⁷ 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 the public history $(\omega^1, \dots, \omega^t)$ 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. A PPE is *strongly symmetric* if σ_1 and σ_2 entail the same action after every public history.

III. 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.

One essential observation concerns the relation between efficiency and the severity of punishments. Under perfect monitoring, cooperation in every round 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 monitoring, however, 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. However, the situation is significantly more complex in the specific case

¹⁵ As mentioned earlier, the payoffs in the perfect monitoring treatment are randomly generated. Specifically, player i 's action a_i generates a random signal ω_j with the same distribution as under imperfect monitoring, and j 's payoff is determined by his action a_j and ω_j . Player i observes a_j but not ω_j or j 's payoff. In the case of perfect monitoring, x_i^t hence includes the realizations of his random payoff. See Section IV for details of the actual implementation.

¹⁶ Throughout, we consider the average discounted payoff, which equals the sum of discounted stage payoffs multiplied by $1 - \delta$.

¹⁷ Under public and private monitoring, every history of signals occurs with strictly positive probability under any history of actions. For this reason, we omit reference to beliefs when discussing a PBE.

of 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. However, 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. The 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.¹⁸

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

Under public monitoring, a pair (σ_G, σ_G) of grim-trigger strategies that revert to the punishment when the history $h_i^t \neq (Ccc, \dots, Ccc)$ is also a PPE for δ sufficiently large and ε sufficiently small.¹⁹ However, such an equilibrium entails a significant efficiency loss since permanent defection is triggered with probability $1 - (1 - \varepsilon)^2$ in every round.²⁰ We can verify that among strongly symmetric PPE, the highest payoff is achieved by a lenient strategy profile that triggers a punishment with probability $\frac{1}{\varepsilon(1 - 2\varepsilon - \varepsilon g)} \frac{1 - \delta}{\delta} g$ after a pair of bad signals $\omega = dd$, and is given by

$$(3) \quad V^{public} = 1 - \frac{\varepsilon}{1 - 2\varepsilon} g,$$

provided that $\varepsilon < \frac{1}{2+g}$ and $\delta > \frac{g}{\varepsilon(1 - 2\varepsilon) + (1 - \varepsilon)(1 + \varepsilon)g}$.

In the case of private monitoring, the lack of common knowledge of histories becomes a major obstacle for cooperation. As mentioned in the introduction, two approaches to the problem have been developed in the literature as detailed below.

The belief-based approach (Sekiguchi 1997, Bhaskar and Obara 2002) 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

¹⁸Under private monitoring with conditional independence as assumed in this paper, Matsushima (1991) shows that the only equilibrium of the repeated game is the repetition of a one-shot Nash equilibrium (NE) if the players' strategies are restricted to the following type: i plays the same action after h_i and h_i' if his belief about j 's private histories h_j conditional on h_i is the same as that conditional on h_i' .

¹⁹As before, $(a_i, \omega_j, \omega_j) = (C, c, c)$ is abbreviated as Ccc . (σ_G, σ_G) is a PPE if $(1 - 2\varepsilon)(1 - \varepsilon) - \varepsilon(2 - \varepsilon)g \geq 1 - \delta/\delta g$.

²⁰The 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 $\varepsilon = 0.1$).

Grim and the strategy AllD of choosing D always. Note that the continuation strategy of such a mixed strategy after each history is either AllD or again a mixture of Grim and AllD. The initial probability weights on Grim and AllD are chosen so that after every history, it is incentive compatible to revert to AllD if and only if a player observes a d signal.²¹ 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.

The belief-free approach (Ely and Välimäki 2002, Piccione 2002) supposes that players play a behavioral 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. This approach is useful since it yields the only class of equilibria under private monitoring for which explicit characterization of behavior and payoffs is possible for a fixed discount factor. Based on Ely and Välimäki (2002) and Piccione (2002), Appendix B illustrates a *memory-one* belief-free equilibrium in which the choice of a mixed action in every round depends only on the signal realization of the previous round. One behavioral prediction when the subjects play such an equilibrium is as follows. Let the *responsiveness* of a strategy be defined by the difference between the probability that a subject cooperates ($a_i = C$) after a good signal ($s_i = c$) and that after a bad signal ($s_i = d$). As seen in equation (B3) in Appendix B, it is expressed in terms of the underlying parameters as

$$(4) \quad \Pr(a_i^{t+1} = C | \omega_i^t = c) - \Pr(a_i^{t+1} = C | \omega_i^t = d) = \frac{g}{\delta(1 - 2\varepsilon)(1 + g)}.$$

In this equilibrium, hence, responsiveness increases with the noise level ε . This 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. We can also verify that the highest payoff achieved by the class of memory-one belief-free equilibria is given by

$$V^{private} = 1 - \frac{\varepsilon}{1 - 2\varepsilon} g,$$

²¹ It is typically the case that with high discount factors, the players do not have an incentive to switch to AllD 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 supgame into several segments so that each segment is played only once in several rounds. See, for example, Sekiguchi (1997).

provided that $\varepsilon < \frac{1}{2(1+g)}$ and $\delta > \frac{g}{(1-2\varepsilon)(1+g)}$.²²

IV. Experimental Design

The original experiment has three treatments corresponding to the three monitoring structures described above. The public and private monitoring treatments use the payoff function $g_i(a_i, \omega_i)$ given by

$$(5) \quad \begin{array}{c|cc} a_i \backslash \omega_i & c & d \\ \hline C & 46 & 8 \\ \hline D & 54 & 16 \end{array} .$$

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 $g_i(C, c)$ with probability $1 - \varepsilon$ and $g_i(C, d)$ 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

$$(6) \quad \begin{array}{c|cc} a_1 \backslash a_2 & C & D \\ \hline C & 42.2, 42.2 & 11.8, 50.2 \\ \hline D & 50.2, 11.8 & 19.8, 19.8 \end{array} .$$

Note that the payoff matrix (6) is strategically equivalent to (2) for

$$(7) \quad g = \ell = \frac{5}{14} \approx 0.357.²³$$

²² It happens to be the case that $V^{private} = V^{public}$ in (3), which is derived under the assumption of strong symmetry. While the belief-free equilibrium is not strongly symmetric, V^{public} can also be interpreted as the highest belief-free equilibrium payoff under public monitoring. All the conditions on ε and δ stated in this section hold in our experiments including the high noise treatments discussed in Section VII.

²³ Simply apply the affine transformation $22.4 u_i(a) + 19.8$. The benefit-cost ratio mentioned in footnote 14 hence equals $(1 + g)/g = 3.8$.

In each of the three treatments, these parameter values ensure that there exist equilibria in which the players cooperate with strictly positive probability at least initially.

The experiments use a between-subject design so that each subject participates in one and only one treatment. Sessions were conducted at the Center for Experimental Social Science lab at New York University.²⁴ In each session, after the instructions are read aloud, subjects are randomly and anonymously paired via computer with another subject to play a supergame.²⁵ All supergames in a session are simultaneously terminated after every round with probability 0.1, and subjects are randomly rematched to play another supergame.²⁶ After each round t , subject i sees on his screen his own action choice a_i^t and other information that varies across different monitoring structures: under perfect monitoring, i 's screen shows the opponent's action choice a_j^t as well as the realization of the random draw that determines i 's own payoff. Under public monitoring, i 's screen shows the pair of signals (ω_i^t, ω_j^t) , and under private monitoring, it just shows his signal ω_i^t about the opponent's action choice. In each case, subject i also sees the possible realizations of the payoff pair highlighted in the payoff matrix.²⁷ 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).²⁸ The sessions were approximately 1 hour and 40 minutes, and subjects earned between \$19.78 and \$48.12 with an average earning of \$33.21.²⁹ These and other summary statistics are provided in Table 1.³⁰

²⁴ Subjects who had participated in previous experiments with randomly terminated games or a PD as a stage-game were excluded.

²⁵ Instructions and screenshots can be found at http://cess.nyu.edu/frechette/print/Aoyagi_2016a_inst.pdf. In the experimental instructions, the term "match" is used in place of "supergame." The experimental interface was programmed using z-tree (Fischbacher 2007).

²⁶ The length of a supergame in each session of the perfect monitoring treatment was determined by a random number generator whose seed is tied to the computer's internal clock. 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 and Fréchette (2011) and Engle-Warnick and Slonim (2006b) both document the impact of the length of supergames on behavior.

²⁷ As mentioned in footnote 15, the realization of a random payoff under perfect monitoring is privately observed as in the other cases. The screen also displays a table describing subject i 's history $h_i^{t-1} = (x_i^1, \dots, x_i^{t-1})$ up to round $t - 1$ within the supergame.

²⁸ The difference in the average number of rounds results from the variation in number of supergames between sessions.

²⁹ Points are converted to dollars at a preannounced 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.

³⁰ Given 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. As stated in footnote 26, the length of the k th supergame in session n is the same regardless of the treatment.

TABLE 1—SUMMARY STATISTICS

Treatments	Subjects	Sessions	Subjects per session	Supergames per session	Rounds per supergame		Subject earnings (US dollars)	
Perfect	66	4	16, 18, 16, 16	11, 12, 19, 11	Min	1	Min	19.78
					Avg.	10.3	Avg.	32.91
					Max	37	Max	43.67
Public	68	4	18, 20, 14, 16	11, 11, 22, 11	Min	1	Min	23.40
					Avg.	10.1	Avg.	34.87
					Max	37	Max	48.12
Private	72	4	20, 18, 14, 20	12, 9, 19, 11	Min	1	Min	22.34
					Avg.	10.4	Avg.	31.92
					Max	37	Max	44.12

V. Directions of Analysis

As discussed in the introduction, we are primarily interested in answering the following questions: (1) Do subjects cooperate under private monitoring? If so, how does the level of cooperation compare with those under public and perfect monitoring? (2) Are there any differences in the subjects' behavior in the three treatments? We stress that presenting testable hypotheses is difficult given the lack of comprehensive theory about the equilibrium of private monitoring games. Although our investigations are more of an exploratory nature for this reason, we will relate the results to insights and predictions provided by the 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 III. 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 III, efficiency considerations imply that strategies under imperfect monitoring should be more lenient and forgiving than under perfect monitoring. Indeed, previous work confirms this view under public monitoring.

In the perfect monitoring environment, Dal Bó and Fréchette (2018b) finds the grim-trigger, which is not lenient, among one of the three most frequently observed strategies. To the contrary, both Fudenberg, Rand, and Dreber (2012) and Embrey, Fréchette, and Stacchetti (2013) find that the subjects' strategies are more lenient and more forgiving in imperfect public monitoring environments. A similar observation can be made on the strategies that best describe the subjects' behavior in Aoyagi and Fréchette (2009): as noise in public information increases, the range of a "bad" signal that causes transition from the cooperation phase to the punishment phase shrinks, and the range of a "good" signal that causes transition from the punishment phase to the cooperation phase widens.³¹

Similar considerations with respect to mitigating efficiency losses associated with a strategy such as grim-trigger apply to private monitoring, but additional considerations also come into play. Under private monitoring, a strategy may be more lenient than under perfect monitoring if a player believes that his opponent plays *C* with high probability: in such a case a player will be reluctant to punish a bad signal since it increases the chance of a punishment from the opponent who himself does not know that his choice was transmitted as a bad signal. 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.³²

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

A memory-one belief-free strategy profile as described in Section III is an equilibrium in every monitoring treatment we consider. If the subjects indeed play such an equilibrium, their responsiveness as defined in equation (4) should be lower when the monitoring is more accurate.

VI. Results

We present our results on the original set of treatments in two parts. The first part is a direct analysis of cooperation and coordination rates as well as action choices

³¹ The 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 and Fréchette (2009) show that this threshold decreases as noise increases.

³² As previously mentioned, leniency 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.

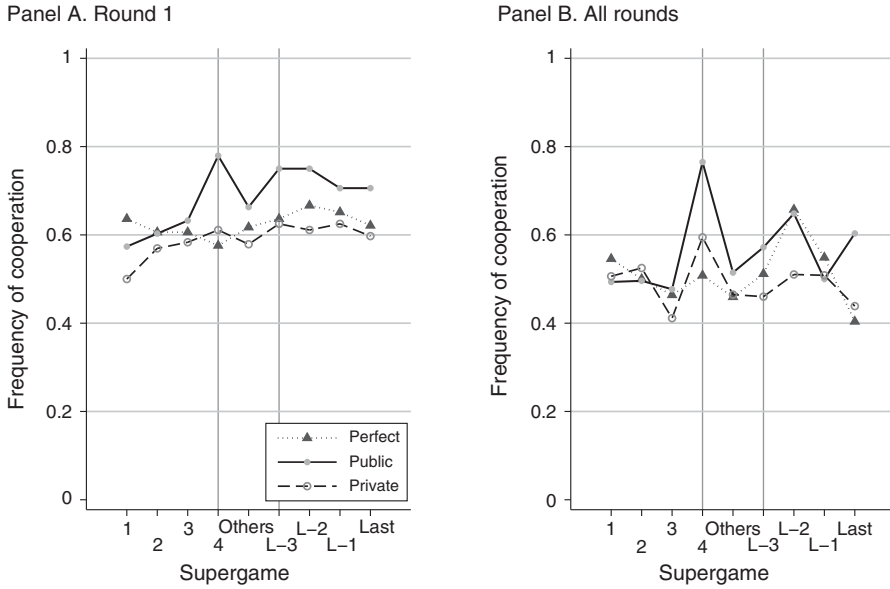


FIGURE 1. COOPERATION RATES BY SUPERGAME

conditional on some histories. The second part is an analysis based on the estimation of strategies.

A. Cooperation Rates

Cooperation rates in the three treatments can be assessed visually in the two panels of Figure 1. 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 1 (with the exception of the middle point) represents the average of four supergames, one from each session.

Observation 1: Subjects cooperate under perfect, public, and private monitoring.

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 percent, which is statistically different from 0 at the 1 percent level.³³

Using no cooperation as a benchmark can be misleading since a positive degree of cooperation is typically observed even in one-shot PD experiments. However, the

³³ Throughout 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 Fréchet (2012). When results are referred to as not statistically significant, it implies a *p*-value greater than 10 percent.

observed level of cooperation in our perfect monitoring treatment is substantially higher than those in repeated PD experiments with similar stage-game payoffs but a lower discount factor that does not support cooperation in equilibrium. In fact, the level of cooperation here is similar to that observed by Dal Bó and Fréchet (2018b) in perfect monitoring repeated PD for a sufficiently high discount factor, and far above that for a low discount factor.^{34,35} Our finding is also in line with that of Rand, Fudenberg, and Dreber (2015), who let subjects in a repeated PD experiment observe the intended action choice of the other player, but introduce payoff randomness through errors in the implementation of action choices; they find that cooperation rates in this treatment are similar to those in the standard perfect monitoring treatment without payoff randomness.

Turning now to public monitoring, we see in panel A of Figure 1 that the round 1 cooperation rate in the last four supergames is 73 percent, which is again statistically different from 0 at the 1 percent level. Positive cooperation in our experiments is in line with the findings in the literature on various forms of public monitoring (Fudenberg, Rand, and Dreber 2012, implementation errors in the subjects' action choices; Aoyagi and Fréchet 2009, one-dimensional continuous signal with infinite support that does not statistically identify the deviator; Embrey, Fréchet, and Stacchetti 2013, outcomes depend probabilistically on the subjects' action choices and the public signal is binary).

There is no statistical difference in the round 1 cooperation rates between perfect and public. However, a difference exists in cooperation rates over all rounds. Panel B of Figure 1 shows that the cooperation rate over all rounds in the last four supergames is 46 percent under perfect monitoring and 58 percent under public monitoring. The rates are both statistically different from 0 ($p < 0.01$), and statistically different from each other ($p < 0.01$). The existing literature provides no clear-cut conclusions on the comparison of overall cooperation rates between public and perfect, and indicates the importance of the particular specification of public monitoring. Fudenberg, Rand, and Dreber (2012) reports a small and non-significant increase in overall cooperation rates when small noise is introduced into monitoring. Rojas (2012) examines perfect and public monitoring under different pairs of the continuation probability δ and the payoffs, and finds a mixed effect of monitoring: Moving from perfect to public, cooperation rates increase in two (δ , payoff) treatments but decrease in the other two treatments. Aoyagi and Fréchet (2009) reports a monotonically decreasing relationship between noise and all-round cooperation rates.³⁶

The key finding in Observation 1 is cooperation under private monitoring. We see again in panel A of Figure 1 that in the last four supergames, there is 61 percent cooperation in round 1, which is statistically different from 0 at

³⁴ See Appendix C. In Dal Bó and Fréchet (2018b), stage payoffs are similar to those used here; there is no randomness in outcomes, and cooperation is an SPE outcome if and only if $\delta \geq 0.72$. As seen in Figure C1, cooperation rates for $\delta = 0.9$ and $\delta = 0.5$ diverge as the subjects accumulate experience.

³⁵ See also the discussions after Observation 2 and in Section VII for alternative ways to evaluate the cooperation rates observed here in comparison with those in a one-shot PD.

³⁶ Aoyagi and Fréchet (2009) observes no statistical difference in round 1 cooperation rates in the treatments where cooperation is theoretically feasible.

the 1 percent level. The cross comparison across treatments in terms of round 1 cooperation rates reveals that the only difference is between public and private ($p < 0.01$), and that there is no statistical difference between perfect and private.³⁷ Panel B of Figure 1 shows that the cooperation rates over all rounds in the last four supergames are 46 percent under private monitoring, which is again statistically different from 0 ($p < 0.01$). In terms of cooperation rates over all rounds, there is no statistical difference between perfect (46 percent cooperation) and private, or between public (58 percent cooperation) and perfect, but public is higher than private weakly significantly ($p < 0.1$). These observations, which answer Question 1, are summarized below.

Observation 2: All three monitoring structures yield strictly positive cooperation rates. Whether in round 1 or in all rounds, cooperation rates under private monitoring are lower than under public monitoring, but are not different from those under perfect monitoring.

We should emphasize that Observation 2 is remarkable considering the theoretical difficulties in sustaining cooperation under private monitoring. A useful comparison is with the level of cooperation in one-shot PD. Dal Bó and Fréchette (2018a) assembles a dataset of 157,170 choices in 15 infinitely repeated (perfect monitoring) and one-shot prisoner's dilemma experiments. Using their dataset, we estimate a probit regression of round 1 cooperation rates on the payoff parameters g and ℓ , discount factor δ , and indicators for subgame perfection and risk dominance (as defined in Dal Bó and Fréchette 2011).³⁸ Using this regression, we can predict the level of cooperation in one-shot PD under our parametrization. We find that the observed round 1 cooperation rate in the private monitoring treatment is higher by 25 percent (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 1 cooperation rate in our private monitoring treatment is higher by 43 percent (significant) than the predicted value.³⁹ This shows that Observation 2 cannot be simply explained by social preferences or other behavioral hypotheses that are often used to explain cooperation in one-shot PD.

B. Coordination

As mentioned previously, the critical feature of private monitoring is the lack of a common knowledge of histories that hinders coordination. In theory, players can perfectly coordinate their actions under both perfect and public monitoring, but not under private monitoring. Thus, it is interesting to see if the subjects indeed have

³⁷ In addition to the description in footnote 33, statistical tests involving comparisons across treatments control for the random sequence of supergames (the variable of interest is regressed on a treatment dummy and on indicator variables for each random sequence). This is done to take into account the potential correlations due to the fact that the realized lengths of supergames has been shown to have an effect (albeit small) on choices (see, for instance, Dal Bó and Fréchette 2011).

³⁸ This is estimated supergame by supergame.

³⁹ The test is a two sample t -test that allows to account for both the variance in the data of our experiment, but also for the prediction error from the estimation on the meta data.

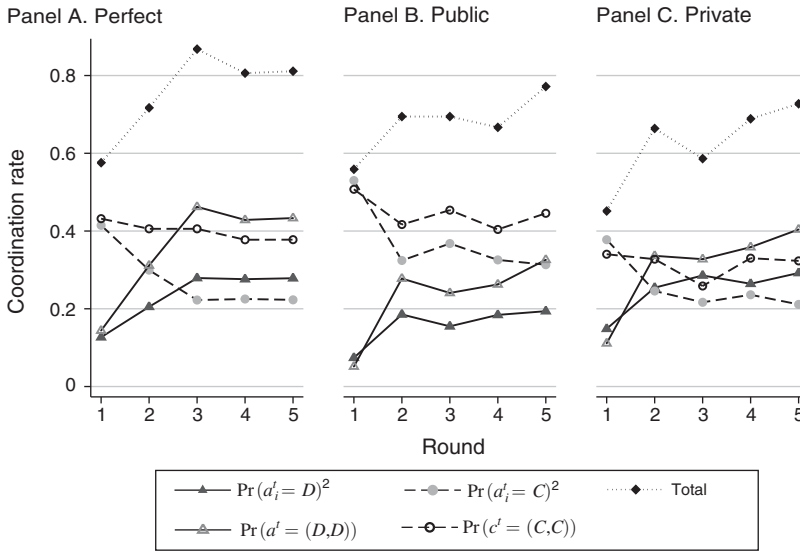


FIGURE 2. COORDINATION RATES IMPLIED BY INDEPENDENT ACTION CHOICE AND REALIZED

Note: All panels A, B, and C are based on the action choice in the last four matches.

difficulty coordinating their actions under private monitoring. Figure 2 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 actions independently. Coordination rates $\Pr(a^t = (C, C))$ or $\Pr(a^t = (D, D))$ are 0.722, 0.712, and 0.660 under perfect, public, and private monitoring, respectively, and no statistical difference exists among them according to the joint test of the three monitoring treatments. In terms of pairwise comparison, a statistical difference exists not between perfect and public, but between public and private ($p < 0.1$). However, the difference between public and private 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 $\Pr(a^t = (C, C))$ or $\Pr(a^t = (D, D))$ 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 percent difference for perfect and 84 percent for public. For private, the rates are higher by 12 percentage points, corresponding to a 40 percent difference.⁴⁰ This suggests that, to a certain extent, subjects have the correct expectation about the other player's action even under private monitoring.

⁴⁰ The difference between actual and expected is statistically significant at the 1 percent level for perfect and at the 10 percent level for public and private.

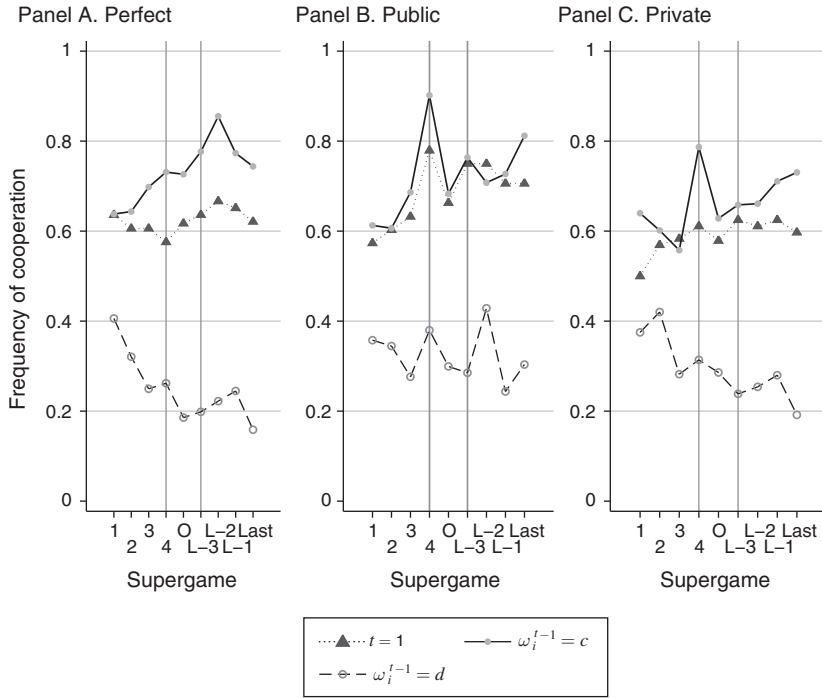


FIGURE 3. COOPERATION CONDITIONAL ON THE PREVIOUS SIGNAL

Observation 3: Coordination rates under private monitoring are close to those under perfect and public monitoring. In particular, they are positive and significantly higher than the level implied by independent action choices.

C. Conditional Cooperation

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 focusing on cooperation rates conditional only on the signal in the preceding round. Figure 3 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 perfect monitoring, 51 percentage points under public monitoring, and 54 percentage points under private monitoring.

Observation 4: In every 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 3 also shows that responsiveness (the difference in cooperation rates following a c signal and a d signal defined by (4)) varies across treatments, and that round 1 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 III, 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.⁴¹ A joint test reveals that responsiveness in the perfect monitoring treatment is statistically different ($p < 0.05$) from that in the other two (answering Question 4). 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.⁴² 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 own action choice and signal in the previous round is summarized in 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 at similar levels 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).⁴³ 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

⁴¹ These numbers are obtained by first computing responsiveness for each subject. t -tests are performed on these subject averages with clustering at the session level.

⁴² Responsiveness 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 ($p < 0.05$).

⁴³ However, they are nonetheless jointly statistically different following Cc ($p < 0.01$), but not following Dd .

TABLE 2—SIGNAL AND COOPERATION RATE CONDITIONAL ON PREVIOUS CHOICE AND SIGNAL COMBINATION

	Perfect	Public	Private
$\Pr(a_i^j = C Cc)$	0.946	0.922	0.921
$\Pr(a_i^j = C Cd)$	0.414	0.553	0.469
$\Pr(a_i^j = C Dc)$	0.223	0.470	0.351
$\Pr(a_i^j = C Dd)$	0.114	0.135	0.105

public monitoring exhibit more leniency and forgiveness than those under perfect monitoring ($p < 0.1$ and $p < 0.01$, respectively), and strategies under private monitoring come somewhere in between in both dimensions (the rates under private monitoring following Dc are statistically different from those under perfect, $p < 0.01$, and public, $p < 0.05$; but not following Cd). 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 vary substantially with the signal in the previous round as well as the action-signal pair in the previous round.

The analysis in this section, and Table 2 more specifically, ignores the possibility that a subject in the public monitoring treatment may condition his behavior on the public signal of his own action. This possibility under public monitoring is examined below.

Momentarily restricting attention to public monitoring, we note that common knowledge of public histories in this environment enables players to coordinate their behavior and simultaneously revert to the punishment or return to cooperation. Our question in this subsection is if the subjects in our public monitoring treatment do indeed use public signals as a coordination device. We examine this question by considering pairs of player i 's private histories that differ from each other only in the realization of ω_j , i.e., the public signal corresponding to i 's own action.⁴⁴ For example, consider a pair of player i 's private histories h_i and \hat{h}_i at the end of round 1 such that $h_i = (a_i, \omega_i, \omega_j) = Ccc$ and $\hat{h}_i = (a_i, \omega_i, \hat{\omega}_j) = Ccd$. In other words, the public history is cc in one case, and cd in the other case, and i 's private action C is the same across these private histories. If the players cooperate by coordinating behavior using the public signal as assumed in the standard construction of a perfect public equilibrium, then continuation play will differ across these histories. However, we find that this is *not* the case. In fact, for *no* pair of such histories in round 1 do we find significant difference in player i 's action choice in round 2. In order to examine i 's action choice in a later round, we need to account for the possibility that it conditions on what happened in two or more

⁴⁴ There are 863 such cases out of 8,804 choices. However, the sample is much smaller when we focus on the last four supergames.

TABLE 3—IMPACT OF THE SIGNAL ABOUT ONE'S OWN ACTION ON COOPERATION RATES

$a_i \omega_j \omega_i$		Signal (ω_j)		
$t - 2$	$t - 1$	c		d
<i>Panel A</i>				
<i>Ddd</i>	$D\omega_j d$	10.64	>*	5.42
<i>Ddc</i>	$D\omega_j d$	9.52	≈	14.17
<i>Ccc</i>	$C\omega_j d$	80.58	<**	92.00
<i>Ccc</i>	$C\omega_j c$	96.25	≈	95.06
<i>Panel B</i>				
$D\omega_j d$	<i>Ddd</i>	4.76	≈	5.42
$C\omega_j d$	<i>Ccd</i>	46.00	≈	60.00
$C\omega_j c$	<i>Ccd</i>	80.58	≈	87.10
$C\omega_j c$	<i>Ccc</i>	96.25	≈	97.18

rounds ago. Because of data restriction, we constrain our analysis to the relationship between i 's action choice in round t and the public signals in rounds $t - 1$ and $t - 2$. More specifically, we consider a pair of histories h_i and \hat{h}_i such that $h_i = ((a_i^{t-2}, \omega_i^{t-2}, \omega_j^{t-2}), (a_i^{t-1}, \omega_i^{t-1}, \omega_j^{t-1}))$, and $\hat{h}_i = ((a_i^{t-2}, \omega_i^{t-2}, \hat{\omega}_j^{t-2}), (a_i^{t-1}, \omega_i^{t-1}, \hat{\omega}_j^{t-1}))$ with either $\omega_j^{t-2} \neq \hat{\omega}_j^{t-2}$ and $\omega_j^{t-1} = \hat{\omega}_j^{t-1}$, or $\omega_j^{t-2} = \hat{\omega}_j^{t-2}$ and $\omega_j^{t-1} \neq \hat{\omega}_j^{t-1}$. Namely, i 's action gives rise to two different public signals in round $t - 1$ or $t - 2$, but not both. Table 3 reports the average rate of $a_i^t = C$ following the most common histories.⁴⁵ As seen, the signal about one's own action has no impact on the player's action choice in all but two comparisons.⁴⁶ With the caution that the results are based on a rather small sample, we summarize our findings as follows.

Observation 6: Subjects' choices under public monitoring appear to ignore the public signal of their own action in many instances.

If the subjects ignore the signal about their own action and do not use the public signal as a coordination device, then the resulting play can be replicated under private monitoring where they do not observe the signal about their own action.⁴⁷ Observation 6 hence indicates the presence of a cooperation mechanism not discussed in the public monitoring literature that identifies the public signal as a coordination device, and also has a strong implication for a similar degree of cooperation under private monitoring. Observation 6, furthermore, is largely

⁴⁵ We count the number of observations for each combination of choices and signals in two consecutive rounds and take the four most common for each one of ω_j^{t-1} and ω_j^t . This results in at least 20 observations for each case.

⁴⁶ Of the two exceptions, a lower cooperation rate after $h_i = (Ccc, Ccc)$ than after $\hat{h}_i = (Ccc, Cdc)$ is rather surprising. One way to make sense of it is that some subjects try to take advantage of the leniency of their opponents by defecting once in a while after a c signal, but is careful not to do so after a d signal.

⁴⁷ In other words, such strategies under public monitoring are measurable with respect to the signal about the opponent's action, and can be implemented under private monitoring.

consistent with our strategy estimation results where the most popular strategies under public monitoring are those that condition only on the signal about the *opponent's* action choice.⁴⁸

D. Leniency and Forgiveness

Leniency and forgiveness (as they are documented here) require looking beyond what happened in the previous round. Hence, we first examine whether the subjects' action choice depends on the event from two or more rounds ago. Table D1 in Appendix D presents the regression of a subject's action choice in round 3 on the outcomes of the first two rounds. The analysis separates the data between the two cases where a subject chooses *C* or *D* in round 1, and includes as regressors the subject's own action choice in round 2, and his signal about the other player's action choice in the first two rounds.⁴⁹ As can be seen, for both public and private monitoring, when a subject cooperates in round 1, the opponent's round 1 choice has a statistically significant impact on the choice made in round 3 (controlling for the outcome in round 2).

To further investigate the question of leniency and forgiveness, we study behavior after some key histories that are possibly longer than one round.⁵⁰

Figure 4 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 panel A corresponds to the values of:

$$\Pr(a_i^t = C | Cc, \dots, Cc),$$

$$\Pr(a_i^t = C | Cc, \dots, Cc, Cd), \quad \text{and}$$

$$\Pr(a_i^t = C | Cc, \dots, Cc, Cd, Cd),$$

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.⁵¹ As can be seen, the drop in the cooperation rates following a single *d* signal is most conspicuous under perfect monitoring, suggesting the use 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

⁴⁸ However, the strategy estimation shows that the popular strategies under public and private are different.

⁴⁹ Analysis of the action choice in round 3 provides the cleanest evidence for the memory length. To reduce the complexity of the regression, analysis in the public monitoring treatment is restricted to the case where the public signal about the own action is correct in the first two rounds.

⁵⁰ See Fudenberg, Rand, and Dreber (2012) for a similar exercise.

⁵¹ The figure only considers action choices in rounds three and above to allow for the observation of at least two signals.

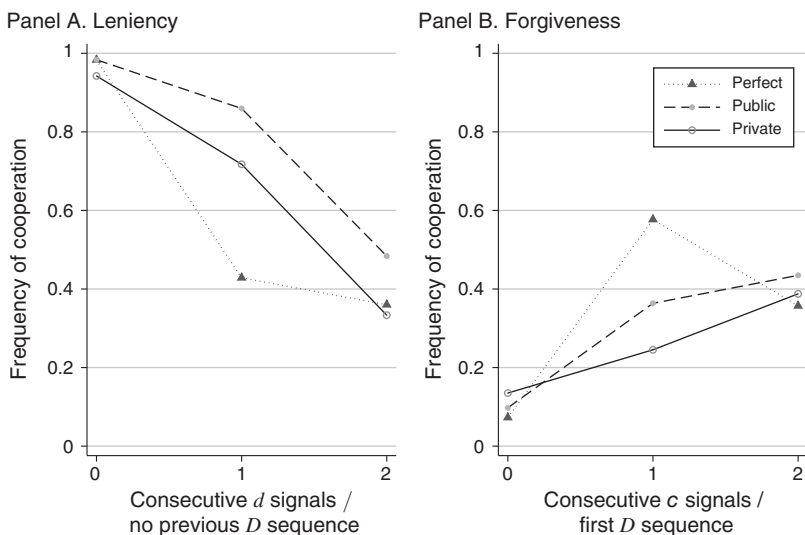


FIGURE 4. COOPERATION RATES AFTER d SIGNALS WHEN COOPERATING (left), AND AFTER c SIGNALS WHEN PUNISHING (right)

Notes: Both panels A and B are based on the action choice in and after round 3 in the last four matches. Panel B excludes matches that start with defection.

($p < 0.01$ and $p < 0.05$, respectively), but there is no statistical difference between the last two.⁵²

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 panel B of Figure 4 corresponds to the values of

$$\Pr(a_i^t = C | C^*, \dots, C^*, Dd, \dots, Dd),$$

$$\Pr(a_i^t = C | C^*, \dots, C^*, Dd, \dots, Dd, Dc), \text{ and}$$

$$\Pr(a_i^t = C | C^*, \dots, C^*, Dd, \dots, Dd, Dc, Dc),$$

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 $p < 0.05$). Directionally, the comparison between perfect and private is similar, but they are much closer to each other in terms of the rates following two c signals than perfect and public are.

⁵²This is established by regressing cooperation on dummies if there was one or two d signals and interacted with a dummy for the type of monitoring (as well as indicators for the sequence of random terminations).

Nonetheless, 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.

We should note, however, that the analysis in this section provides only a rough measure of leniency and forgiveness as the strategies may condition on events beyond the previous two rounds. We perform a more direct analysis of leniency and forgiveness in the next section by estimating strategies based on the subjects' action choices in all rounds of a supergame. Nonetheless, we can summarize the evidence, which is a partial answer to Question 2, as follows.

Observation 7: Conditional on some key histories, the levels of leniency are such that

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

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.

E. 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 and Fréchette (2011). SFEM has now been used in multiple papers to estimate the strategies in repeated games, and its use is supported by Fudenberg, Rand, and Dreber (2012) and Dal Bó and Fréchette (2018b) who conduct Monte-Carlo simulations to evaluate its performance, and in addition, Dal Bó and Fréchette (2018b) finds that the strategies identified as most popular by the SFEM are also the most popular strategies elicited from the subjects using an alternative method.⁵³ In essence, SFEM uses maximum likelihood to estimate a mixture over a given set of strategies.⁵⁴ 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 $\beta \equiv \frac{1}{1 + \exp(-1/\hat{\gamma})}$ 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 $\beta \rightarrow 1/2$ (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 and Fréchette (2011); Fudenberg, Rand, and Dreber (2012) and their reanalysis

⁵³ Other papers using the SFEM in different contexts include Vespa and Wilson (2016) on dynamic games and Bigoni et al. (2015) on continuous-time games.

⁵⁴ Intuitively, 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 strategy in the entire sample. See Dal Bó and Fréchette (2011) for details.

TABLE 4—PROPERTIES OF STRATEGIES

	Cooperative	Responsive	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

of Dal Bó and Fréchette (2011) and Dreber et al. (2008); Dal Bó and Fréchette (2018b); Fréchette and Yuksel (2017); and Embrey, Fréchette, and Stacchetti (2013). The strategies included in our analysis are listed in Table 4 and their finite automaton representations are given in Appendix E.

The three strategies in the top panel of Table 4 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). Such strategies will be termed *nonresponsive*. All other strategies we consider are *responsive* in that transitions depend on signals. 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.⁵⁵ 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 initially set equal to 0. 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, Fréchette, and Stacchetti (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, Fudenberg, and Nowak 2007) but has almost never been found in statistically significant proportions in previous experiments. Every strategy in the second panel yields a sequence of (*C, C*)s when matched against itself. However, 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

⁵⁵ TFT 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.

TABLE 5—ESTIMATES OF PROPORTION OF EACH STRATEGY

	Perfect	Public	Private
<i>Panel A</i>			
AllC	0.024 (0.032)	0.196 (0.067)	0.032 (0.038)
AllD	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)
<i>Panel B</i>			
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)
<i>Panel C</i>			
STFT	0.015 (0.029)	0.000 (0.045)	0.032 (0.030)
SSum2	0.000	0.027	0.014
<i>Panel D</i>			
γ	0.471 (0.040)	0.474 (0.038)	0.569 (0.042)
β	0.893	0.892	0.853
<i>Panel E</i>			
Cooperative	0.671	0.783	0.683
Noncooperative	0.329	0.217	0.317
<i>Panel F</i>			
Leniency ratio	0.514	0.824	0.776
Forgiving ratio	0.712	0.897	0.692
<i>Panel G</i>			
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

Notes: Listed in panel A are unconditional strategies, panel B are responsive and cooperative, and panel C are responsive and noncooperative. Panels E, F, and G show total frequencies by feature. The leniency ratio is the ratio, amongst cooperative and responsive strategies, of lenient to lenient and non-lenient. The forgiving ratio is the ratio, amongst cooperative and responsive strategies, of forgiving to forgiving and non-forgiving.

in Table 4 show whether these strategies are lenient and/or forgiving.⁵⁶ 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.⁵⁷ A *cooperative* strategy is a strategy that starts with cooperation and produces cooperation with positive probability in subsequent rounds when matched against itself. When discussing leniency and forgiveness, we focus on cooperative strategies that are responsive.

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 and Fréchette 2018a). Second, lenient and forgiving strategies are more popular under public monitoring than under perfect monitoring (Fudenberg, Rand, and Dreber 2012). Third, despite its theoretical appeal, WSLs is not observed in any significant proportion in any treatment. Fourth, as in Embrey, Fréchette, and Stacchetti (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, Fréchette, and Stacchetti (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 percent and Grim2 and Grim3 are 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 more forgiving: the ratio of lenient strategies $\left(= \frac{\text{lenient}}{\text{cooperative and responsive}} \right)$ is 78 percent under private monitoring, while it is 51 percent under perfect monitoring ($p = 0.00$). In fact, the ratio of lenient strategies under private monitoring (78 percent) is similar to that under public monitoring (82 percent) ($p = 0.83$ for the equality between public and private). As for forgiveness, the ratio of forgiving strategies $\left(= \frac{\text{forgiving}}{\text{cooperative and responsive}} \right)$ is lower under private monitoring (69 percent) than under public monitoring (90 percent) (insignificant at $p = 0.29$). However, the forgiveness level under private monitoring (69 percent) is similar to that under perfect monitoring (71 percent) (insignificant at $p = 0.70$).

⁵⁶ Classification of leniency and forgiveness is applied only to responsive strategies. 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.

⁵⁷ See Rubinstein (1986).

This finding on forgiveness hence is at odds with the result of the reduced-form approach in Section VIC, which shows 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 VIC 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 that has just two states in the automaton representation is higher under perfect monitoring (33 percent) than under either public (6 percent, $p < 0.01$) or private (12 percent, $p < 0.05$) 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 6. Notice that the top three strategies represent more than 50 percent in proportion in all three treatments. While the noncooperative 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 in the top three.

The evidence with respect to the strategy used answers to Question 2 and Question 3 and is summarized in the following.

Observation 8: 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.

With these estimation results at hand, we now return to the most surprising finding from our experiment: that the level of cooperation under private monitoring is comparable to that under perfect and public monitoring. Although a complete analysis of the mechanism behind this finding is beyond the scope of this paper, we consider how the changes in strategies mitigate the effect of the change in the monitoring structure. First, the fraction of subjects who mostly defect varies across treatments: 19 percent under public monitoring, 27 percent under private monitoring, and 31 percent under perfect monitoring. In other words, the fraction of subjects adopting some form of cooperative strategies is higher under imperfect

TABLE 6—TOP STRATEGIES BY TREATMENT

Popularity	Monitoring		
	Perfect	Public	Private
First	AllD	AllC	AllD
Second	TFT	AllD	Sum2
Third	Grim	2TF2T	Grim3

monitoring. Ignore this effect for the moment and fix the fraction of subjects that adopt the strategy AllD of always defecting at 25 percent. We will study the effect of the choice of a cooperative strategy adopted by the remaining 75 percent. As seen in Table 6, two of the most popular cooperative strategies are TFT and Grim under perfect monitoring. If we assume that the 75 percent of subjects who cooperate use either of these strategies, then the average cooperation rate (given $\delta = 0.9$) is 58 percent under perfect monitoring. If the cooperative strategy under public and private monitoring were also given by either TFT or Grim, then the average cooperation rate would be 42 percent in the case of TFT (under both public and private, a drop of 16 percentage points), and it would be 23 percent (under public, a drop of 35 points) and 24 percent (under private, a drop of 34 points) in the case of Grim. However, if we suppose that the cooperative strategy is 2TF2T under public monitoring and Sum2 under private monitoring (as indicated in Table 6), then the cooperation rate is 57 percent under public monitoring and 56 percent under private monitoring.⁵⁸ That is, the observed difference in the subjects' strategies almost completely compensates for the effect of the monitoring structure on cooperation rates. We can only speculate on the reason why the subjects adopt different cooperative strategies under different monitoring structures, but as this example shows, such changes can mitigate the effects of noise on cooperation rates.

VII. High Noise Treatments

To study the impact of noise on cooperation under private monitoring, we conducted additional treatments, again one for each monitoring structure. In particular, the noise is increased to $\varepsilon = 0.2$ as opposed to $\varepsilon = 0.1$ in the original treatments. The original goal was to see whether the differences observed in the first set of experiments in terms of strategies (in particular with respect to leniency and forgiveness) would remain at this increased level of noise. As will become obvious, the results for these treatments are not informative with respect to our original goal. Hence, these are only briefly discussed here and details are provided in Appendix F.

⁵⁸ Although the most popular cooperative strategy under public monitoring is AllC, we consider 2TF2T since efficiency is clearly unaffected by changes in the monitoring technology when a nonresponsive strategy such as AllC is considered.

The round 1 cooperation rates in the three $\varepsilon = 0.2$ treatments are not statistically different from one another, and more importantly, from the prediction for a one-shot game based on previous PD experiments.⁵⁹ These results indicate that there are little to no dynamic incentives to generate cooperation when $\varepsilon = 0.2$.

In summary, increasing noise in monitoring from $\varepsilon = 0.1$ to $\varepsilon = 0.2$ poses a substantial difficulty to the subjects' ability to cooperate, and this difficulty appears to be caused by the randomness in the payoffs coupled with the larger value of $g = \ell = 0.541$.

Observation 9: Increasing noise from $\varepsilon = 0.1$ to $\varepsilon = 0.2$ lowers cooperation rates in every monitoring treatment to levels similar to what is predicted in one-shot PD. In particular, the level of cooperation under perfect monitoring with $\varepsilon = 0.2$ is significantly lower than what is predicted under perfect monitoring without random payoffs. There is no significant difference in the levels of cooperation among the three high noise treatments.

VIII. Conclusion

While theory suggests the importance of the 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. Specifically, we confirm the key findings from Fudenberg, Rand, and Dreber (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 game 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 cooperation 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) are significantly higher than the hypothetical rates that would be obtained when the two players choose their actions independently at the observed rates. The high levels of cooperative behavior in our private monitoring environment are in sharp contrast with the findings of Duffy and Ochs (2009)—that cooperation is hard to sustain under random matching, even for relatively modest group sizes (see Section I). The difference suggests that experimental subjects find the constancy of a relationship with their partner more important than the monitoring structure when it comes to the decision on whether to cooperate or not.

⁵⁹This is true either if the prediction is based only on the sample of one-shot game experiments from Dal Bó and Fréchette (2018a) (the one indicated in Figure F1), or on the fitted relation used to form the prediction for $\delta = 0.9$ (see footnote 39).

We should emphasize that our results do not imply that cooperation will emerge whenever it is theoretically feasible. Indeed, low cooperation rates in our high noise treatments are indistinguishable from what is expected in one-shot games. Hence, it is most appropriate to interpret our results as demonstrating that in *some* private monitoring environments, cooperation is possible at a level comparable to that under perfect and public monitoring. Our results also suggest that in order to identify the environments conducive to cooperation, it is particularly important to isolate the effect of monitoring from the effect of randomness in payoffs.⁶⁰

In the case where dynamic incentives emerge, we also find differences 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. However, the comparison of forgiveness levels is less clear. While strategy estimation suggests 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 4 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, Fréchet, and Stacchetti (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.

We also confirm the finding of Kayaba, Matsushima, and Toyama (2016), that responsiveness is not necessarily consistent with the play of a memory-one belief-free equilibrium across all treatments. 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 the 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. In this regard, a theory of action choice based on the mental state as analyzed by Compte and Postlewaite (2008) may provide one direction for

⁶⁰ Perfect monitoring experiments in the literature with no randomness in outcomes find a substantially higher level of cooperation. Similarly, Rand, Fudenberg, and Dreber (2015) finds that randomness in outcomes have at most a modest impact on behavior. However, Bereby-Meyer and Roth (2006) finds that randomness slows down (reinforcement) learning in one-shot PD and finitely repeated PD.

research. 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 environment. 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.

APPENDIX A: RANDOM RE-MATCHING IN GROUPS




FIGURE A1. RANDOM RE-MATCHING IN GROUPS

Source: Duffy and Ochs (2009)

APPENDIX B: RESPONSIVENESS UNDER MEMORY-ONE BELIEF-FREE EQUILIBRIUM

Consider a memory-one behavioral strategy σ_j of player j such that the probability $\Pr(a_j^t = C | h_j^{t-1})$ that it plays C in round $t \geq 2$ depends only on his signal ω_j^{t-1} in round $t - 1$. Any such strategy σ_i can be expressed by the two probabilities p and q such that

$$\Pr(a_j^t = C | a_i^{t-1}, \omega_i^{t-1}) = \begin{cases} p & \text{if } \omega_i^{t-1} = c \\ q & \text{if } \omega_i^{t-1} = d \end{cases}$$

Let now $W(c)$ denote i 's continuation payoff from round t on when j observes $\omega_j^{t-1} = c$ in round $t - 1$, and $W(d)$ denote i 's continuation payoff when j observes $\omega_j^{t-1} = d$ in round $t - 1$. Since σ_j makes player i indifferent between playing C and D , we should have

$$(B1) \quad (1 - \delta)g = \delta(1 - 2\varepsilon)[W(c) - W(d)],$$

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 $1 - 2\varepsilon$.⁶¹ Next, if player j observes $\omega_j^{t-1} = c$ in round $t - 1$ and player i plays D in round t , then i 's continuation payoff from round t on is given by

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

However, if player j observes $\omega_j^{t-1} = d$ in round $t - 1$ and player i plays D in round t , then 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

$$(B2) \quad W(c) - W(d) = (1 - \delta)(p - q)(1 + g).$$

Combining (B1) and (B2), we obtain

$$(B3) \quad p - q = \frac{g}{\delta(1 - 2\varepsilon)(1 + g)}.$$

When player j plays σ_j satisfying (B3), player i is indifferent between playing C and playing D at every history. It follows that the strategy profile (σ_i, σ_j) with

⁶¹ Note that the gain from playing D does not depend on j 's action when $g = \ell$ as assumed.

both σ_i and σ_j satisfying (B3) is an equilibrium. When σ_i and σ_j both begin with C in round 1, we can verify that a player’s expected payoff in this equilibrium is given by

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

The highest equilibrium payoff is hence achieved when $p = 1$, and is given by $1 - \frac{\varepsilon}{1 - 2\varepsilon}g$.

APPENDIX C: COOPERATION UNDER PERFECT MONITORING IN DAL BÓ AND FRÉCHETTE (2018B)

In Dal Bó and Fréchette (2018b), each session in the $\delta = 0.5$ treatment has at least 19 supergames, while the 3 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 C1 includes at most 19 supergames to make comparison easier. Dal Bó and Fréchette (2018b) specifies 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 $g = 25/7 - 1 \approx 2.57$ and $\ell = 13/7 \approx 1.86$.




FIGURE C1. COOPERATION RATES IN DAL BÓ AND FRÉCHETTE (2018B) BY SUPERGAME

APPENDIX D: CHOICES IN ROUND THREE

TABLE D1—DOES THE OPPONENT'S CHOICE TWO ROUNDS AGO AFFECT CHOICES?

Dependent variable: Cooperation in round 3 ($\mathbf{1}_{\{a_i^3=C\}}$)						
Case	Perfect		Public [†]		Private	
	Defect in round 1	Coop. in round 1	Defect in round 1	Coop. in round 1	Defect in round 1	Coop. in round 1
Action C in round 2 ($\mathbf{1}_{\{a_i^2=C\}}$)	0.219 (0.050)	0.243 (0.071)	0.155 (0.146)	0.203 (0.045)	0.081 (0.106)	0.373 (0.118)
Signal c in round 2 ($\mathbf{1}_{\{\omega_i^2=c\}}$)	-0.105 (0.010)	0.270 (0.102)	-0.028 (0.041)	0.300 (0.111)	-0.021 (0.049)	0.407 (0.095)
Pair (C, c) in round 2 ($\mathbf{1}_{\{(a_i^2, \omega_i^2)=(C, c)\}}$)	0.411 (0.087)	0.189 (0.111)	0.367 (0.200)	0.143 (0.087)	0.195 (0.110)	-0.119 (0.145)
Signal c in round 1 ($\mathbf{1}_{\{\omega_i^1=c\}}$)	0.018 (0.061)	0.048 (0.056)	0.064 (0.040)	0.166 (0.062)	-0.003 (0.081)	0.108 (0.023)
Constant	0.130 (0.074)	0.203 (0.071)	0.138 (0.033)	0.190 (0.041)	0.262 (0.049)	0.150 (0.048)
Observations	246	406	196	343	290	402

Note: Clustered standard errors are in parentheses.

[†] Restricted to cases where the public signal about i 's action ω_j equals his action a_i for $t = 1$.

APPENDIX E: STRATEGIES INCLUDED IN THE ESTIMATION

TABLE E1—UNCONDITIONAL AND TWO-STATES AUTOMATA

Automaton name in text	Diagram	Perfect and public	Private
<i>Panel A</i>			
AIC			
AID			
CDDD			
<i>Panel B</i>			
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 E2—AUTOMATA WITH MORE THAN TWO STATES

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		$t_1 = \{\omega_i = c\}$ $t_2 = \neg t_1$	

APPENDIX F: COOPERATION AND COORDINATION IN THE HIGH NOISE
TREATMENTS

With the payoff function g_i held fixed as in (5), and the noise $\varepsilon = 0.2$, the stage game in these high noise treatments have $g = \ell = 20/37 \approx 0.541$ as opposed to 0.357 in the original treatments, and are explicitly given as follows:⁶²

(F1)

$a_1 \backslash a_2$	C	D
C	38.4, 38.4	15.6, 46.4
D	46.4, 15.6	23.6, 23.6

As g and ℓ increase, the meta study of Dal Bó and Fréchet (2018b) predicts a modest decrease in cooperation rates in infinitely repeated PD under perfect monitoring without randomly generated payoffs. In line with the expectation, the cooperation rates in the high noise treatments are lower than in the original treatments for every monitoring structure (Figure F1): the drop in the round 1 cooperation rates from $\varepsilon = 0.1$ to $\varepsilon = 0.2$ is significant with $p < 0.01$ for each case. However, the size of the reduction is much larger than predicted by Dal Bó and Fréchet (2018b): the cooperation rates for $\varepsilon = 0.2$ are statistically lower than the meta-analysis based prediction in all three monitoring treatments ($p < 0.01$). This is in contrast with the $\varepsilon = 0.1$ perfect and public monitoring treatments where the cooperation rates are not statistically different from the prediction.⁶³

As noted in Section VII, round 1 cooperation rates in those treatments are not different from what is predicted for one-shot PD games. This diminished role of dynamic considerations in behavior is also visible in Figure F2 on the frequency of cooperation rate after a good versus bad signal, and in Figure F3 on the realized coordination rates. Specifically, Figure F2 shows very little separation between cooperation rates following good versus bad signals. Compared with the differences of 56 (perfect), 44 (public), and 35 (private) percentage points in the original $\varepsilon = 0.1$ treatments, the differences in the $\varepsilon = 0.2$ treatments are 23 (perfect), 11 (public), and -4 (private) percentage points with only the one for perfect being statistically significant ($p < 0.1$).

Figure F3 is the counterpart of Figure 2 and shows the realized coordination rates for $\varepsilon = 0.2$ as well as what would be expected if choices were independent within a pair. Consistent with the observations made so far, it shows that the realized coordination rates are nearly identical to what would be implied by independent action choices. In fact, comparing the total coordination rates to the sum $\Pr(a_i^t = C)^2 + \Pr(a_i^t = D)^2$, we find statistical difference only for perfect monitoring (difference = 0.06, $p < 0.01$). The figure also shows that total coordination is lower than for $\varepsilon = 0.1$, and most of it is accounted for by coordination on defection.

⁶² A total of 96 subjects participated in 6 sessions, 2 each for perfect, public, and private monitoring treatments, with the number of subjects per session equal to either 12 or 24. The subjects played at least 11 supergames in every session, and the data analysis below focuses on supergames 8 through 11. In two of the six sessions, a software bug ended one supergame earlier than it should have.

⁶³ The rate is lower than the prediction ($p < 0.05$) in the $\varepsilon = 0.1$ private monitoring treatment.




FIGURE F1. ROUND ONE COOPERATION RATES BY TREATMENT IN THE LAST FOUR SUPERGAMES

Note: Predictions generated using the metadata from Dal Bó and Fréchette (2018a).

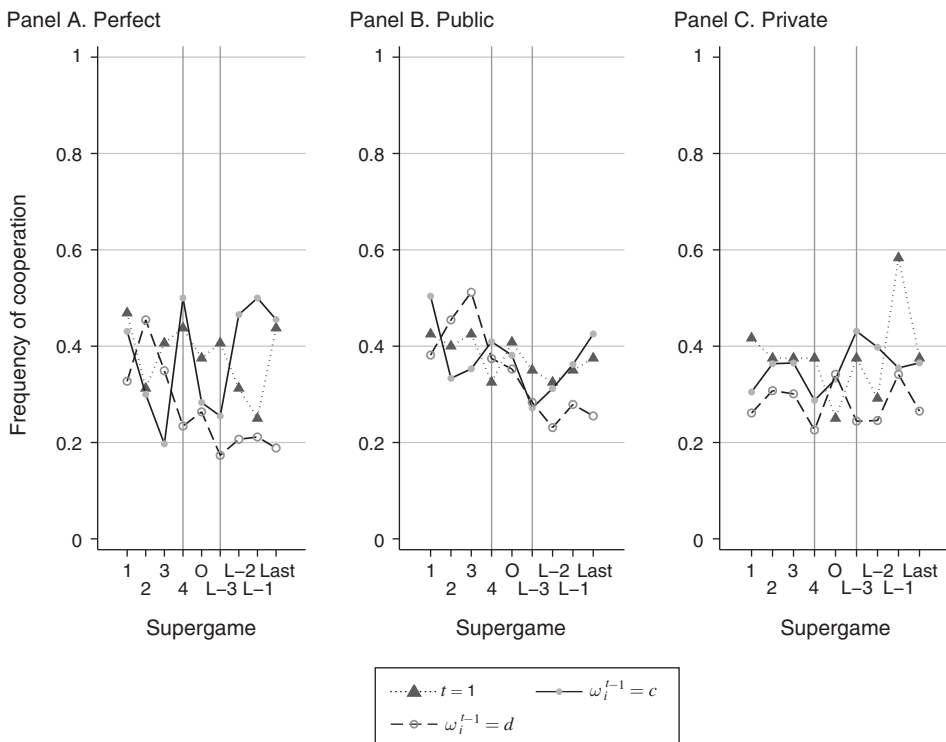


FIGURE F2. COOPERATION CONDITIONAL ON THE PREVIOUS SIGNAL WITH $\epsilon = 0.2$ (equivalent to Figure 3)

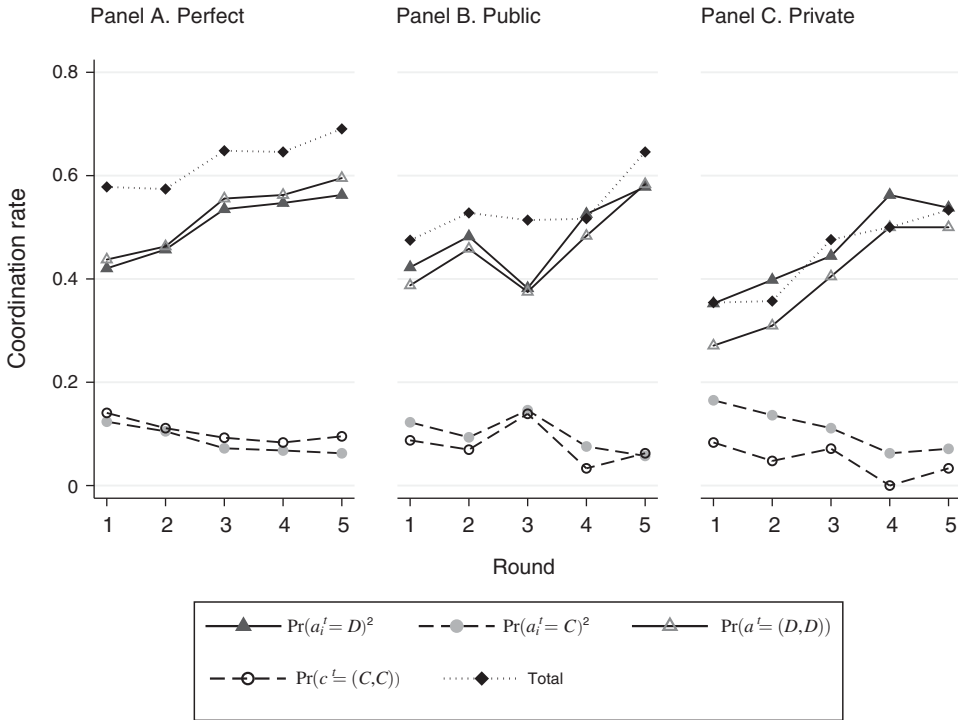


FIGURE F3. COORDINATION RATES IMPLIED BY INDEPENDENT ACTION CHOICE AND REALIZED WITH $\varepsilon = 0.2$

Notes: The figure corresponds to Figure 2. All panels A, B, and C are based on the action choice in the last four matches.

Considering all rounds together, we find cooperation rates under public and private monitoring not statistically different from each other, and both higher than those under perfect monitoring ($p < 0.1$ and $p < 0.01$, respectively). The differences, however, are not substantial, being at most 5 percentage points between perfect (27 percent) and private (32 percent).

REFERENCES

- Aoyagi, Masaki, V. Bhaskar, and Guillaume R. Fréchet. 2019. "The Impact of Monitoring in Infinitely Repeated Games: Perfect, Public, and Private: Dataset." *American Economic Journal: Microeconomics*. <https://doi.org/10.1257/mic.20160304>.
- Aoyagi, Masaki, and Guillaume Fréchet. 2009. "Collusion as public monitoring becomes noisy: Experimental evidence." *Journal of Economic Theory* 144 (3): 1135–65.
- Battaglini, Marco, Salvatore Nunnari, and Thomas R. Palfrey. 2015. "The Dynamic Free Rider Problem: A Laboratory Study." http://www.its.caltech.edu/~trp/dyngvc_16May2015.pdf.
- Bereby-Meyer, Yoella, and Alvin E. Roth. 2006. "The Speed of Learning in Noisy Games: Partial Reinforcement and the Sustainability of Cooperation." *American Economic Review* 96 (4): 1029–42.
- Bernard, Mark, Jack Fanning, and Sevgi Yuksel. 2017. "Finding cooperators: Sorting through repeated interaction." http://econ.ucsb.edu/~sevgi/BernardFanningYuksel_April2017.pdf.
- Bhaskar, V., and Ichiro Obara. 2002. "Belief-Based Equilibria in the Repeated Prisoners' Dilemma with Private Monitoring." *Journal of Economic Theory* 102 (1): 40–69.

- Bhaskar, V., and Eric van Damme.** 2002. "Moral Hazard and Private Monitoring." *Journal of Economic Theory* 102 (1): 16–39.
- Bigoni, Maria, Marco Casari, Andrzej Skrzypacz, and Giancarlo Spagnolo.** 2015. "Time Horizon and Cooperation in Continuous Time." *Econometrica* 83 (2): 587–616.
- Cabral, Luis, Erkut Y. Ozbay, and Andrew Schotter.** 2014. "Intrinsic and instrumental reciprocity: An experimental study." *Games and Economic Behavior* 87: 100–121.
- 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.** 2012. "Cooperative strategies in anonymous economies: An experiment." *Games and Economic Behavior* 75 (2): 570–86.
- Cason, Timothy N., and Feisal U. Khan.** 1999. "A laboratory study of voluntary public goods provision with imperfect monitoring and communication." *Journal of Development Economics* 58 (2): 533–52.
- Cason, Timothy N., and Vai-Lam Mui.** 2014. "Coordinating Resistance Through Communication and Repeated Interaction." *Economic Journal* 124 (574): F226–56.
- Compte, Olivier, and Andrew Postlewaite.** 2008. "Repeated Relationships with Limits on Information Processing." <https://pdfs.semanticscholar.org/74ca/7a72843c4dc96c9b1c67bd79296442920a00.pdf>.
- Cooper, David J., and Kai-Uwe Kühn.** 2014. "Communication, Renegotiation, and the Scope for Collusion." *American Economic Journal: Microeconomics* 6 (2): 247–78.
- Dal Bó, Pedro.** 2005. "Cooperation under the Shadow of the Future: Experimental Evidence from Infinitely Repeated Games." *American Economic Review* 95 (5): 1591–1604.
- Dal Bó, Pedro, and Guillaume R. Fréchet.** 2011. "The Evolution of Cooperation in Infinitely Repeated Games: Experimental Evidence." *American Economic Review* 101 (1): 411–29.
- Dal Bó, Pedro, and Guillaume R. Fréchet.** 2018a. "On the Determinants of Cooperation in Infinitely Repeated Games: A Survey." *Journal of Economic Literature* 56 (1): 60–114.
- Dal Bó, Pedro, and Guillaume R. Fréchet.** 2018b. "Strategy Choice In The Infinitely Repeated Prisoners' Dilemma." http://cess.nyu.edu/frechette/print/Dal_Bo_2018b.pdf.
- Dreber, Anna, Drew Fudenberg, and David G. Rand.** 2014. "Who cooperates in repeated games: The role of altruism, inequity aversion, and demographics." *Journal of Economic Behavior and Organization* 98: 41–55.
- Dreber, Anna, David G. Rand, Drew Fudenberg, and Martin A. Nowak.** 2008. "Winners don't punish." *Nature* 452: 348–51.
- Duffy, John, and Jack Ochs.** 2009. "Cooperative behavior and the frequency of social interaction." *Games and Economic Behavior* 66 (2): 785–812.
- Ely, Jeffrey C., Johannes Hörner, and Wojciech Olszewski.** 2005. "Belief-Free Equilibria in Repeated Games." *Econometrica* 73 (2): 377–415.
- Ely, Jeffrey C., and Juuso Välimäki.** 2002. "A Robust Folk Theorem for the Prisoner's Dilemma." *Journal of Economic Theory* 102 (1): 84–105.
- Embrey, Matthew, Guillaume R. Fréchet, and Ennio Stacchetti.** 2013. "An Experimental Study of Imperfect Public Monitoring: Efficiency versus Renegotiation-Proofness." http://cess.nyu.edu/frechette/print/Embrey_2013a.pdf.
- Engle-Warnick, Jim, and Robert L. Slonim.** 2004. "The evolution of strategies in a repeated trust game." *Journal of Economic Behavior and Organization* 55 (4): 553–73.
- Engle-Warnick, Jim, and Robert L. Slonim.** 2006a. "Inferring repeated-game strategies from actions: Evidence from trust game experiments." *Economic Theory* 28 (3): 603–32.
- Engle-Warnick, J., and Robert L. Slonim.** 2006b. "Learning to trust in indefinitely repeated games." *Games and Economic Behavior* 54 (1): 95–114.
- Feinberg, Robert M., and Thomas A. Husted.** 1993. "An Experimental Test of Discount-Rate Effects on Collusive Behaviour in Duopoly Markets." *Journal of Industrial Economics* 41 (2): 153–60.
- Feinberg, Robert M., and Christopher Snyder.** 2002. "Collusion with secret price cuts: An experimental investigation." *Economics Bulletin* 3 (6): 1–11.
- Fischbacher, Urs.** 2007. "z-Tree: Zurich toolbox for ready-made economic experiments." *Experimental Economics* 10 (2): 171–78.
- Fréchet, Guillaume R.** 2012. "Session-effects in the laboratory." *Experimental Economics* 15 (3): 485–98.
- Fréchet, Guillaume R., and Sevgi Yuksel.** 2017. "Infinitely repeated games in the laboratory: Four perspectives on discounting and random termination." *Experimental Economics* 20 (2): 279–308.

- Fudenberg, Drew, David G. Rand, and Anna Dreber.** 2012. "Slow to Anger and Fast to Forgive: Cooperation in an Uncertain World." *American Economic Review* 102 (2): 720–49.
- Green, Edward J., and Robert H. Porter.** 1984. "Noncooperative Collusion under Imperfect Price Information." *Econometrica* 52 (1): 87–100.
- Harrington, Joseph E., Jr., and Andrzej Skrzypacz.** 2007. "Collusion under monitoring of sales." *RAND Journal of Economics* 38 (2): 314–31.
- Harrington, Joseph E., and Andrzej Skrzypacz.** 2011. "Private Monitoring and Communication in Cartels: Explaining Recent Collusive Practices." *American Economic Review* 101 (6): 2425–49.
- Holcomb, James H., and Paul S. Nelson.** 1997. "The role of monitoring in duopoly market outcomes." *Journal of Socio-Economics* 26 (1): 79–93.
- Holt, Charles A.** 1985. "An Experimental Test of the Consistent-Conjectures Hypothesis." *American Economic Review* 75 (3): 314–25.
- Imhof, Lorens A., Drew Fudenberg, and Martin A. Nowak.** 2007. "Tit-for-tat or win-stay, lose-shift?" *Journal of Theoretical Biology* 247 (3): 574–80.
- Kandori, Michihiro.** 2002. "Introduction to Repeated Games with Private Monitoring." *Journal of Economic Theory* 102 (1): 1–15.
- Kayaba, Yutaka, Hitoshi Matsushima, and Tomohisa Toyama.** 2016. "Accuracy and Retaliation in Repeated Games with Imperfect Private Monitoring: Experiments and Theory." Center for International Research on the Japanese Economy Discussion Paper CIRJE-F-1004.
- Levenstein, Margaret C., and Valerie Y. Suslow.** 2006. "What Determines Cartel Success?" *Journal of Economic Literature* 44 (1): 43–95.
- Matsushima, Hitoshi.** 1991. "On the theory of repeated games with private information: Part I: anti-folk theorem without communication." *Economics Letters* 35 (3): 253–56.
- Murnighan, J. Keith, and Alvin E. Roth.** 1983. "Expecting Continued Play in Prisoner's Dilemma Games: A Test of Several Models." *Journal of Conflict Resolution* 27 (2): 279–300.
- Palfrey, Thomas R., and Howard Rosenthal.** 1994. "Repeated Play, Cooperation and Coordination: An Experimental Study." *Review of Economic Studies* 61 (3): 545–65.
- Piccione, Michele.** 2002. "The Repeated Prisoner's Dilemma with Imperfect Private Monitoring." *Journal of Economic Theory* 102 (1): 70–83.
- Rand, David G., Drew Fudenberg, and Anna Dreber.** 2015. "It's the thought that counts: The role of intentions in noisy repeated games." *Journal of Economic Behavior and Organization* 116: 481–99.
- Rojas, Christian.** 2012. "The role of demand information and monitoring in tacit collusion." *RAND Journal of Economics* 43 (1): 78–109.
- Rotemberg, Julio J., and Garth Saloner.** 1986. "A Supergame-Theoretic Model of Price Wars during Booms." *American Economic Review* 76 (3): 390–407.
- Roth, Alvin E.** 1995. "Bargaining Experiments." In *Handbook of Experimental Economics*, edited by John H. Kagel and Alvin E. Roth, 253–348. Princeton: Princeton University Press.
- Roth, Alvin E., and J. Keith Murnighan.** 1978. "Equilibrium behavior and repeated play of the prisoner's dilemma." *Journal of Mathematical Psychology* 17 (2): 189–98.
- Rubinstein, Ariel.** 1986. "Finite automata play the repeated prisoner's dilemma." *Journal of Economic Theory* 39 (1): 83–96.
- Schwartz, Steven T., Richard A. Young, and Kristina Zvinakis.** 2000. "Reputation Without Repeated Interaction: A Role for Public Disclosures." *Review of Accounting Studies* 5 (4): 351–75.
- Sekiguchi, Tadashi.** 1997. "Efficiency in Repeated Prisoner's Dilemma with Private Monitoring." *Journal of Economic Theory* 76 (2): 345–61.
- Sherstyuk, Katerina, Nori Tarui, and Tatsuyoshi Saijo.** 2013. "Payment schemes in infinite-horizon experimental games." *Experimental Economics* 16 (1): 125–53.
- Stigler, George J.** 1964. "A Theory of Oligopoly." *Journal of Political Economy* 72 (1): 44–61.
- Vespa, Emanuel.** 2016. "An Experimental Investigation of Strategies in the Dynamic Common Pool Game." <https://drive.google.com/file/d/0BwGkg6vwKfDISkc4NUxhMDF5Qik/view>.
- Vespa, Emanuel, and Alistair J. Wilson.** 2016. "Experimenting with Equilibrium Selection in Dynamic Games." <https://drive.google.com/file/d/0BwGkg6vwKfDIUjJJeGRSOFdwbkU/view>.

This article has been cited by:

1. Edoardo Gallo, Yohanes E. Riyanto, Nilanjan Roy, Tat-How Teh. 2022. Cooperation and punishment mechanisms in uncertain and dynamic social networks. *Games and Economic Behavior* **134**, 75-103. [[Crossref](#)]
2. Andreas Nicklisch, Louis Putterman, Christian Thöni. 2021. Trigger-happy or precisionist? On demand for monitoring in peer-based public goods provision. *Journal of Public Economics* **200**, 104429. [[Crossref](#)]
3. Andreas Freitag, Catherine Roux, Christian Thöni. 2021. COMMUNICATION AND MARKET SHARING: AN EXPERIMENT ON THE EXCHANGE OF SOFT AND HARD INFORMATION. *International Economic Review* **62**:1, 175-198. [[Crossref](#)]
4. Masaki Aoyagi, Guillaume R. Frechette, Sevgi Yuksel. 2021. Beliefs in Repeated Games. *SSRN Electronic Journal* **33**. . [[Crossref](#)]
5. Masaki Aoyagi, Guillaume R. Frechette, Sevgi Yuksel. 2021. Beliefs in Repeated Games. *SSRN Electronic Journal* **33**. . [[Crossref](#)]
6. Yutaka Kayaba, Hitoshi Matsushima, Tomohisa Toyama. 2020. Accuracy and retaliation in repeated games with imperfect private monitoring: Experiments. *Games and Economic Behavior* **120**, 193-208. [[Crossref](#)]
7. Emanuel Vespa. 2020. AN EXPERIMENTAL INVESTIGATION OF COOPERATION IN THE DYNAMIC COMMON POOL GAME. *International Economic Review* **61**:1, 417-440. [[Crossref](#)]
8. Fabian Dvorak. 2020. Stratest: Strategy Estimation in R. *SSRN Electronic Journal* . [[Crossref](#)]
9. Hitoshi Matsushima. 2019. Behavioral Theory of Repeated Prisoner's Dilemma: Generous Tit-For-Tat Strategy. *The B.E. Journal of Theoretical Economics* **20**:1. . [[Crossref](#)]
10. Pedro Dal Bó, Guillaume R. Fréchette. 2019. Strategy Choice in the Infinitely Repeated Prisoner's Dilemma. *American Economic Review* **109**:11, 3929-3952. [[Abstract](#)] [[View PDF article](#)] [[PDF with links](#)]
11. Shun Kurokawa, Xiudeng Zheng, Yi Tao. 2019. Cooperation evolves more when players keep the interaction with unknown players. *Applied Mathematics and Computation* **350**, 209-216. [[Crossref](#)]
12. Timothy N. Cason, Vai-Lam Mui. 2019. Individual versus group choices of repeated game strategies: A strategy method approach. *Games and Economic Behavior* **114**, 128-145. [[Crossref](#)]
13. James R. Bland. 2019. Heterogeneous Trembles and Model Selection in the Strategy Frequency Estimation Method. *SSRN Electronic Journal* . [[Crossref](#)]
14. Tetsuya Kawamura, Tsz Kwan Tse. 2019. Equilibrium Selection and Cognitive Ability: Experimental Evidence from an Infinitely Repeated Public Goods Game. *SSRN Electronic Journal* . [[Crossref](#)]
15. Julian Romero, Yaroslav Rosokha. 2018. Mixed Strategies in the Indefinitely Repeated Prisoner's Dilemma. *SSRN Electronic Journal* . [[Crossref](#)]
16. Brown Economics Working Paper Series, Thomas Markussen, Louis G. Putterman. 2017. Governing Collective Action in the Face of Observational Error. *SSRN Electronic Journal* . [[Crossref](#)]
17. Fabian Dvorak, Sebastian Fehrler. 2017. Negotiating Cooperation Under Uncertainty: Communication in Noisy, Indefinitely Repeated Interactions. *SSRN Electronic Journal* . [[Crossref](#)]

18. Julian Romero, Yaroslav Rosokha. 2016. Constructing Strategies in the Indefinitely Repeated Prisoner's Dilemma Game. *SSRN Electronic Journal* . [[Crossref](#)]
19. David G. Rand, Drew Fudenberg, Anna Dreber. 2013. It's the Thought that Counts: The Role of Intentions in Reciprocal Altruism. *SSRN Electronic Journal* . [[Crossref](#)]