

Collusion as Public Monitoring Becomes Noisy: Experimental Evidence*

Masaki Aoyagi[†]
Osaka University

Guillaume Fréchette[‡]
New York University

June 20, 2008

Abstract

This paper uses laboratory experiments to test the implications of the theory of repeated games on equilibrium payoffs and estimate strategies in an infinitely repeated prisoners' dilemma game with imperfect public monitoring. We find that subjects' payoffs (i) decrease as noise increases, and (ii) are lower than the theoretical maximum for low noise, but exceed it for high noise. Under the assumption that the subjects' strategy uses thresholds on the public signal for transition between cooperation and punishment states, we find that the best fitting strategy simply compares the most recent public signal against a single threshold.

Keywords: Repeated games, imperfect public monitoring, collusion, cooperation.

*This is a substantially revised version of Aoyagi and Fréchette (2005). Special thanks are given to John Kagel, who helped with comments and guidance throughout the development of the paper. We have benefited from discussions with Al Roth and Marco Casari and comments by Pedro Dal Bo, Georg Weizsacker and seminar participants at Université Laval, Harvard University, Kyoto University, New York University, Université de Montréal, Carnegie Mellon University, RIETI, Rutgers, and the Stockholm School of Economics, participants of the 2003 International Meeting of the Economic Science Association, participants at the 14th Summer Festival on Game Theory at Stony Brook, two referees and the associate editor. Alex Brown provided valuable research assistance. We thank Jo Ducey, Matthew Embrey, Sotiris Georganas, John Lightle, Karen Manukyan, John-David Slaughter, and Yong Yu for help with the additional control sessions. This research was supported in part by the NSF via grants SES-0519045 and SES-0721111, by the Harvard Business School, by the Center for Experimental Social Science, by the C.V. Starr Center, and by the Zengin Foundation for Studies on Economics and Finance.

[†]ISER, 6-1 Mihogaoka, Ibaraki, Osaka 567-0047, Japan.

[‡]Department of Economics, 19 West 4th Street, 6th Floor, New York, NY, 10012.

1 Introduction

Imperfection in monitoring is an integral part of many competitive situations in reality. It is a particularly important topic in the field of industrial organization, where the signals are subject to various external shocks. Firms' ability to cooperate in such environments is of clear interest to researchers as well as regulation authorities. A key prediction of the theory of repeated games is that when players are patient, a simple strategy can sustain collusion even under *imperfect public monitoring*, where opponents' actions are observed only through a noisy public signal. In the repeated prisoners' dilemma (PD), for example, it is known that a wide range of symmetric equilibrium payoffs can be achieved by a grim-trigger strategy, which reverts to the one-shot Nash equilibrium in the event of particular signal realizations.

This paper is aimed at testing the theory of imperfect public monitoring in infinitely repeated games: It provides a characterization of the maximal equilibrium payoff as a function of noise in monitoring, and then uses laboratory experiments to test these predictions. The main focus of the paper is on the comparative statics of the effect of noise on the players' payoffs, and on the strategies they use to achieve collusion.

While imperfect monitoring has attracted much attention in economic theory, empirical work on the subject has been limited because of some fundamental difficulties. For example, it is not easy to identify the exact public signal the firms use to coordinate their actions: it could be price, shares in a nationwide or regional market, industry output, or the combination of any of these and other indicators. There are also difficulties with identifying the firms' strategic variables and their payoff structure. Free from these problems, a laboratory experiment in a controlled environment is considered an ideal alternative for the study of the subject.

Green and Porter (1984) are the first to provide a theoretical analysis of repeated games with imperfect public monitoring: In their model of quantity-setting oligopoly, the market price serves as a noisy public signal of firms' output quantities because of demand fluctuations. They put forth an equilibrium based on the *trigger strategy* as follows: The firms produce at the jointly monopolistic level as long as the realized price is above a certain threshold, but revert to the one-shot Cournot quantity for a fixed number of periods when it falls below the threshold. Because of the random component in demand, periodic price wars occur on the equilibrium path.

In this paper, the stage game is given by a standard symmetric PD with two

players: The action set of each player consists of an efficient cooperative action and an inefficient action which gives rise to a (one-shot) Nash equilibrium. We characterize the set of symmetric equilibrium payoffs when the two players' repeated game strategies are public in the sense that today's actions are determined only by past public signals, and perfect in the sense that they form a Nash equilibrium after every public history. In each period, the players' actions determine the distribution of a public signal in that period. In line with the price-quantity interpretation of Green and Porter (1984), we suppose that the public signal is a one-dimensional continuous variable, and that the higher the public signal, the more likely the opponent has chosen the cooperative action. More precisely, we assume that the likelihood ratio of the cooperative action versus the non-cooperative action monotonically increases with the signal. Under this assumption, it is natural to suppose that players' action choice is determined by whether the observed signal is high or low. In other words, a natural class of strategies is given by those that use thresholds as a coordination device: The players agree to take one action profile if the realized public signal is above a certain value but take another action profile if it falls below it. We refer to such a strategy as a *threshold strategy* in this paper. Using the bang-bang property (Abreu *et al.* (1990)), which states in this case that the highest equilibrium payoff is sustained by a grim-trigger strategy, we identify sufficient conditions for the optimal grim-trigger strategy to be a threshold strategy. We then derive an explicit and testable link between the maximal symmetric equilibrium payoff and the level of noise in monitoring.¹

In our experiments, we use an exponential signal distribution for analytical tractability and parameterize it by the level of noise in monitoring. We have five treatments varying the level of noise from zero to infinity in addition to a one-shot game treatment with perfect monitoring. In theory, positive cooperation is possible for the three low noise treatments, while the best equilibrium entails no cooperation for the two high noise treatments. We have two major objectives in analyzing data from our experiments. First, we compare the players' actual payoffs against the theoretical maximum derived as above, and examine how they change with the noise in monitoring. Our findings are as follows:

¹Kandori (1992b) shows that the set of (symmetric and asymmetric) perfect public equilibrium payoffs expands as monitoring becomes more accurate. Note, however, that his conclusion is qualitative and not directly testable. When the public signal has a Poisson distribution, Abreu *et al.* (1991) describe the relationship between the equilibrium payoff and monitoring accuracy. Sannikov (2007) also analyzes the effect of noise on the equilibrium payoff set in continuous time.

- 1) The level of cooperation is positive for any noise level.
- 2) The level of cooperation is lower than the theoretical maximum for low noise but higher than that for high noise.
- 3) The level of cooperation decreases as noise increases.
- 4) For high but finite noise, the level of cooperation is no higher than that in the one-shot game.
- 5) For infinite noise (*i.e.*, no public signal), the level of cooperation is lower than that in the one-shot game.

We also find a distinctive pattern in the evolution of play. In the low noise treatments, we observe that subjects increase the level of cooperation as they accumulate experience. In the high noise treatments, on the other hand, the subjects begin to behave more non-cooperatively as they become more experienced.

As stated above, we find that reducing noise in monitoring increases the level of cooperation. It should be emphasized that this is the first experiment to identify such a relationship between the noise in monitoring and the level of cooperation in the standard imperfect public monitoring setting. This result is far from obvious for the following reasons: First, as is well known, experiments on the one-shot PD often generate positive levels of cooperation. Given that the subjects of those experiments cooperate even with no histories, these results may be interpreted as suggesting that decision making in repeated games is independent of histories and that cooperation is generated through some other mechanism. The present experiment rejects this view. Second, our finding is at odds with those of some experiments on imperfect monitoring. These experiments, which model imperfect monitoring in ways significantly different from the present one, find no increase in cooperation when monitoring becomes “more accurate.”² Third, sustaining cooperation in a repeated game requires the ability to perform non-trivial reasoning even if it only involves a simple strategy. Our finding suggests that the subjects do in fact possess such capabilities.

Our second objective is to estimate the subjects’ repeated game strategies. For this, we suppose that they play *threshold strategies with regime shifts* between the cooperation and punishment states. These strategies start out in the cooperation

²As discussed in the next section, these experiments deviate in some important ways from the standard models of imperfect public monitoring.

state, switch to the punishment state when the public signal falls below a certain threshold, and return to the cooperation state when the signal exceeds another threshold. Specifically, the threshold may depend on the current state and the own action choice. In all but one noise treatment, we find that the data is best described by the simplest threshold strategies which only have a single threshold. In other words, those strategies simply check the most recent public signal and plays the cooperative action if it is above some threshold, and the non-cooperative action otherwise.

It should be noted that our laboratory experiments replicate the framework of the tested theory as closely as possible. First, by the standard identification of the continuation probability with the discount factor, an infinitely repeated game is replaced by a repeated game with a random termination point.³ The noise is taken to be independent and identically distributed across periods and has full support regardless of actions. Payments to subjects are designed so that they are bounded and reveal no more information than the public signal during the course of play.⁴ Each pair of subjects understand that they observe the same stochastic signal after every period, and how its probability distribution is related to their actions. In particular, while the exponential signal distribution may not necessarily be something familiar to the undergraduate subjects, the instruction includes a simple chart illustrating how the likelihood of the opponent's action is related to the observed signal and their own action.

The organization of the paper is as follows. In the next section, we present a brief review of the related literature. Section 3 presents a model of a repeated PD with imperfect public monitoring and characterizes the highest symmetric perfect public equilibrium payoff. Section 4 describes the design of our experiments and Section 5 reports their results: We first test the theoretical prediction on the players' payoffs and then estimate their strategies. We conclude in Section 6 with some discussions.

2 Related Literature

Most empirical work on repeated games with imperfect monitoring analyzes the data from the 1880's Joint Executive Committee (JEC) railroad cartel, with a special emphasis on the specification of the timing of regime shifts, *i.e.*, switches between

³As is true with all repeated game experiments, however, the length of the game must be constrained by the practical time limit for each session.

⁴See Section 6.

cooperation and punishment phases in the repeated game.⁵ Early work assumes that regime shifts follow a Bernoulli distribution (Porter (1983), Lee and Porter (1984)), while some later work uses the Markovian chain in the estimation (Cosslett and Lee (1985)). Porter (1985) takes a detailed look into what triggers the regime shifts, the effect of market structure, and the determinants of price war duration. Ellison (1994), again using the JEC data, tests the Green and Porter (1984) model. In contrast to the prior estimates that were closer to the Cournot level, he finds collusive behavior to be much closer to the monopoly level. He identifies several factors as statistically significant determinants (*i.e.* triggers) of regime shifts. However, the estimated mechanism is not strong enough to deter cheating. Furthermore, he finds evidence of secret price cutting, which is not predicted by the model.

Experimental economics has focused much attention on the possibility of cooperation in various models of oligopolies including the PD and public goods games.⁶ For repeated games with *perfect* monitoring, the results of laboratory experiments generally indicate that repeated play generates cooperation strictly above the one-shot Nash equilibrium level and below the first-best level. However, there is no definitive conclusion on the strategies that players use to achieve cooperation. For example, there exists conflicting evidence as to the use of trigger strategies.⁷ It should also be noted that most of the early results need to be interpreted with caution as they pertain to repeated games with an “unknown horizon,” where subjects are not informed of how long the game will last.⁸

Experiments on imperfect monitoring include Feinberg and Snyder (2002), Cason and Khan (1999), and Holcomb and Nelson (1997). These papers introduce monitoring imperfection in rather specific ways. Cason and Khan (1999) study a repeated public good experiment and compare standard perfect monitoring with perfect, but delayed monitoring of past actions. They do not find any significant difference in the levels of contributions between the two treatments. Feinberg and Snyder (2002) consider a version of the repeated PD where each subject observes his own payoff

⁵Green and Porter (1984) suggest that such regime shifts follow a Markov process of order equal to the length of punishment periods.

⁶See Holt (1995) for a literature review.

⁷See, for example, Sell and Wilson (1991), Feinberg and Husted (1993), and Engle-Warnick and Slonim (2006).

⁸See Roth and Murnighan (1978), who point out that such a game yields significantly different results from an infinitely repeated game. They propose, to properly replicate an infinitely repeated game with discounting in an experimental setting, to terminate the game after each period with a fixed probability.

in each period. They introduce imperfection by occasionally manipulating those payoff numbers, and compare the treatments with and without the ex post revelation of such manipulation. Less collusive behavior is found in the latter treatment. Holcomb and Nelson (1997), on the other hand, study a repeated duopoly model in which information about the opponent's quantity choice is randomly changed 50% of the time. They conclude that such manipulation "does significantly affect market outcomes" (p. 79). It should be noted that the formulation of imperfect monitoring in these papers is not in line with the assumptions of the standard theory. For example, in Feinberg and Snyder (2002) and Holcomb and Nelson (1997), monitoring is imperfect but private since players in these models do not necessarily observe the same signal.⁹

As mentioned earlier, formulation of imperfect monitoring in this paper follows the standard theory much more closely than that in the above models. The distinguishing feature of our model is the assumption that the public signal is a one-dimensional real variable. This assumption has the following advantages: First, it closely replicates the oligopoly models where price serves as the public signal. Second, this specification naturally incorporates a monotone relationship between the signal and action: the higher the public signal, the more likely the other player has cooperated. This relationship is easy for the subjects to understand, and also justifies the use of the threshold strategy.

Besides imperfection in monitoring, some recent experiments look at factors that affect players' cooperative behavior. Duffy and Ochs (2007) study the effects of fixed versus random pairing in a repeated game. For parameter values that can sustain cooperation even with random matching, they find cooperation emerge only in the fixed pairing case. Dal Bo (2003) compares a repeated game with random termination against that with a fixed and known length. He finds that cooperation in the former treatment is at a higher level.¹⁰ Dal Bo and Fréchette (2007) investigate the conditions that lead subjects to coordinate on a cooperative or non-cooperative equilibrium.¹¹ They find that if coordination cannot be supported in equilibrium,

⁹Cason and Khan (1999) and Holcomb and Nelson (1997) use finite horizon games but do not specify what information was given to their subjects about the duration of the game.

¹⁰A much earlier experiment by Roth and Murnighan (1978) also studies the effect of the continuation probability on the level of cooperation. However, their experiment matches subjects to computerized opponents.

¹¹See the references therein for a more complete list of experiments on infinitely repeated games with perfect monitoring.

subjects learn to defect, while if coordination is an equilibrium they may or may not coordinate on an equilibrium that supports cooperation. They also find that coordination may not arise even if it can be supported as part of a risk-dominant equilibrium.

One of the main focuses of the present paper is the analysis of the players' strategies in repeated games. This is the subject of some recent experiments as follows. Mason and Phillips (2002) study the use of trigger strategies in a repeated Cournot duopoly game with perfect monitoring. They estimate the duration and severity of punishment by changing the stage payoffs corresponding to deviations. They conclude that evidence is consistent with the use of trigger strategies and that behavior is more consistent with the use of a strategy with long and mild punishment phases. Engle-Warnick and Slonim (2006) study the strategies played by subjects in repeated trust games with perfect monitoring. They look for the strategy that best fits the observed play from the set of pure strategies that can be expressed as deterministic finite automata, and conclude that some subjects use a grim-trigger strategy.^{12,13}

In comparison with the above papers, the threshold specification in the present paper allows for a direct estimation of the subjects' strategies based on a standard econometric technique.

3 Prisoners' Dilemma with Imperfect Public Monitoring

Two players play a symmetric PD game infinitely often. Player i 's action a_i in the stage-game is chosen from the set $A_i = \{a_i^0, \hat{a}_i\}$, where \hat{a}_i is the cooperative action and a_i^0 is the non-cooperative action. Let $A = A_1 \times A_2$ denote the set of action profiles $a = (a_1, a_2)$. After each period, players observe only a random public signal $z \in \mathbf{R}$ whose probability distribution is determined by the action profile $a \in A$ in that period. In a Cournot model with stochastic demand, for example, a_i and z correspond to firm i 's output quantity and the realized price level, respectively. Denote by $h(z \mid a)$ the density of the public signal z under the action profile a .

¹²It should be noted that this estimation technique is not practical for games with a long expected horizon. In Engle-Warnick and Slonim (2006), the continuation probability is such that the expected length of the repeated game is five periods. It is ten in our case.

¹³Engle-Warnick, McCausland, and Miller (2004) propose a Bayesian method for the estimation of repeated game strategies in a perfect monitoring environment.

Player i 's stage-payoff equals $w_i(a_i, z)$ when his own action is a_i and the signal realization is z . Assume that w_i is continuous in z , and note that the payoff does not provide more information about the other player's action than the public signal z itself. Player i 's expected stage-payoff from any action profile a , $g_i(a)$, can be computed as

$$g_i(a) = \int_{\mathbf{R}} w_i(a_i, z) h(z | a) dz.$$

The assumption that the stage-game is a PD translates to

$$g_i(a_i^0, \hat{a}_j) > g_i(\hat{a}) > g_i(a^0) \geq g_i(\hat{a}_i, a_j^0).$$

As noted above, $\hat{a} = (\hat{a}_1, \hat{a}_2)$ is the symmetric efficient profile, and $a^0 = (a_1^0, a_2^0)$ is the one-shot Nash equilibrium of the PD. Note also that a_i^0 is a minmax action as well as a profitable one-shot deviation from the efficient profile \hat{a} . Denote $g^0 = g_i(a^0)$ and $\hat{g} = g_i(\hat{a})$. A t -length *public history* is the history of signals z in periods 1 through t . A t -length *private history* of player i is the sequence of i 's actions in periods 1 through t . The set of t -length public histories is given by \mathbf{R}^t , while the set of t -length private histories of player i is given by A_i^t . Player i 's (pure) *strategy* is a function $\sigma_i : \bigcup_{t=0}^{\infty} (\mathbf{R}^t \times A_i^t) \rightarrow A_i$. It is a *public strategy* if it is a function of the public history alone. Let $\delta < 1$ denote the players' common discount factor. Player i 's expected payoff in the repeated game from a strategy profile σ is given by

$$\pi_i(\sigma) = (1 - \delta) \sum_{t=1}^{\infty} \delta^{t-1} g_i^t,$$

where g_i^t is the expected stage-payoff in period t under the probability distribution induced by σ . The strategy profile σ is a (pure) *equilibrium* if for every i , $\pi_i(\sigma) \geq \pi_i(\sigma'_i, \sigma_j)$ for any strategy σ'_i . An equilibrium $\sigma = (\sigma_1, \sigma_2)$ is *public* if σ_i is a public strategy for each i . A public equilibrium is *perfect* if σ_i is a best response to σ_j for each i after every public history, and is (strongly) *symmetric* if $\sigma_1 = \sigma_2$.

Throughout the paper, we assume that the public signal z is related to the action profile a in the following additive way:

$$z = s(a) + x,$$

where $s : A \rightarrow \mathbf{R}$ is a deterministic function of a , and x is a real-valued random variable whose distribution is independent of a . We assume that s is symmetric ($s(a) = s(a')$ if $(a'_1, a'_2) = (a_2, a_1)$), and that x has a strictly positive density f over

\mathbf{R} . Denote the corresponding cumulative distribution by F . The density $h(z \mid a)$ of z under a can be expressed in terms of f and s as

$$h(z \mid a) = f(z - s(a)) \quad \text{for each } z \in \mathbf{R} \text{ and } a \in A.$$

We make the following assumptions on s and f :

Assumption 1 $s(a_i^0, \hat{a}_j) \leq s(\hat{a})$.

Assumption 2 $f(x - y) f(x' - y') \geq f(x - y') f(x' - y)$ for any $x \geq x'$ and $y \geq y'$.¹⁴

It can be readily verified that these assumptions together imply that the deviation a_i^0 from \hat{a} shifts down the distribution of the public signal as measured in the likelihood ratio. In other words,

$$\frac{h(z' \mid \hat{a})}{h(z \mid \hat{a})} \geq \frac{h(z' \mid a_i^0, \hat{a}_j)}{h(z \mid a_i^0, \hat{a}_j)} \quad \text{for any } z' \geq z.$$

In a Cournot model, for example, these assumptions imply that an increase in the production level lowers the distribution of the price. It should be also noted that Assumption 2 holds for a wide class of distributions including normal and gamma distributions. Define

$$l = \frac{g_i(a_i^0, \hat{a}_j) - g^0}{\hat{g} - g^0} > 1 \quad \text{and} \quad d = s(\hat{a}) - s(a_i^0, \hat{a}_j) \geq 0.$$

Namely, l is the (normalized) one-shot gain from a deviation from \hat{a} to a_i^0 , while d is the sensitivity of the public signal measured by the change in its expected value following such a deviation.

We now turn to the characterization of the set of equilibrium payoffs in this setup. As shown by Abreu *et al.* (1990), this can be accomplished by examining the following class of “grim-trigger” equilibria: player i starts with \hat{a}_i , and keeps playing \hat{a}_i as long as the realized public signal z falls in a certain (Borel) subset Q of \mathbf{R} but reverts to the minimax action a_i^0 forever otherwise. As verified in the proof of Proposition 1 below, under Assumptions 1 and 2, we can take the set Q of admissible signals to be of the form $[k, \infty)$ for some k . That is, the highest

¹⁴The function f satisfying this condition is known as a *Polya function of degree 2* (PF₂) (Karlin, 1968). Note that f is PF₂ if and only if the function $\hat{f}: \mathbf{R}^2 \rightarrow \mathbf{R}_+$ defined by $\hat{f}(x, y) = \log f(x - y)$ is supermodular.

symmetric perfect public equilibrium payoff is generated by a *threshold grim-trigger strategy*, which reverts to the punishment if and only if the public signal falls below a certain threshold k . Suppose then that $\hat{\sigma} = (\hat{\sigma}_1, \hat{\sigma}_2)$ is the threshold grim-trigger strategy profile that begins with \hat{a} and reverts to a^0 when z falls below k . As seen in the proof of Proposition 1 in the Appendix, the payoff v associated with $\hat{\sigma}$ can be obtained as a solution to the standard recursive equation:

$$v = \frac{(1 - \delta)\hat{g} + \delta g^0 F(k - s(\hat{a}))}{1 - \delta \{1 - F(k - s(\hat{a}))\}}. \quad (1)$$

It is also a standard exercise to verify that the incentive condition for player i to choose \hat{a}_i on the equilibrium path can be written as:

$$\frac{1 - \delta \{1 - F(k - s(\hat{a}) + d)\}}{1 - \delta \{1 - F(k - s(\hat{a}))\}} \geq l. \quad (2)$$

Let $K(\delta)$ denote the set of thresholds k for which the above incentive compatibility condition holds:

$$K(\delta) = \{k \in \mathbf{R} : k \text{ satisfies (2)}\}.$$

By construction, $K(\delta)$ is a closed set. In the Appendix, it is also shown that $K(\delta)$ is an interval. There exists a threshold grim-trigger equilibrium that supports the action profile \hat{a} if and only if $K(\delta) \neq \emptyset$.¹⁵ By (1), the optimal threshold $k = k^*(\delta)$ that maximizes v is the smallest element of $K(\delta)$:

$$k^*(\delta) = \min K(\delta). \quad (3)$$

It is also clear from (2) that $K(\delta) \neq \emptyset$ requires $d > 0$. The following proposition summarizes our observation.

Proposition 1 *Suppose that Assumptions 1 and 2 hold and let $v^*(\delta)$ denote the maximal symmetric perfect public equilibrium payoff of the repeated PD. If $K(\delta) = \emptyset$, $v^*(\delta) = g^0$. On the other hand, if $K(\delta) \neq \emptyset$, then*

$$v^*(\delta) = \frac{(1 - \delta)\hat{g} + \delta g^0 F(k^*(\delta) - s(\hat{a}))}{1 - \delta + \delta F(k^*(\delta) - s(\hat{a}))}.$$

In the latter case, $v^(\delta)$ is achieved by a (pure) stationary grim-trigger strategy profile $\hat{\sigma}$ which begins with \hat{a}_i and continues with \hat{a}_i as long as the realized public signal z is at or above the threshold $k^*(\delta)$ (i.e., $Q = [k^*(\delta), \infty)$), but reverts to a_i^0 otherwise.*

¹⁵By Proposition 1, if there exists no threshold grim-trigger equilibrium, then there exists no grim-trigger equilibrium which supports $v > g^0$.

Proof. See the Appendix. ■

In interpreting the results of our experiments described in Section 5, the following points should be noted. First, while the grim-trigger equilibrium yields a clear description of the bound, the bound applies to all types of perfect public equilibria whether grim-trigger or not. Conversely, the theorem does not preclude the (near) efficiency of other types of perfect public equilibria, which may return to cooperation after some contingency.

For comparison across different payoff numbers, the following normalization of $v^*(\delta)$ is useful:

$$y^*(\delta) = \frac{v^*(\delta) - g^0}{\hat{g} - g^0}.$$

It can be verified that $y^*(\delta) \in [0, 1]$ is unaffected by a positive affine transformation of the payoff numbers.¹⁶ By Proposition 2, $y^*(\delta)$ can be explicitly written as

$$y^*(\delta) = \frac{1 - \delta}{1 - \delta + \delta F(k^*(\delta) - s(\hat{a}))} \quad (4)$$

if $K(\delta) \neq \emptyset$ and $y^*(\delta) = 0$ otherwise.

4 Experimental Design

The experiment tests the theory developed in the previous sections in the following environment. The expected stage payoffs are specified as follows:

$1 \setminus 2$	L	H
L	25, 25	15, 28
H	28, 15	16, 16

As seen, L (“low output”) corresponds to the cooperative action and H (“high output”) represents the opportunity for a profitable one-shot deviation.¹⁷ Note that $\hat{a} = (L, L)$ and $a^0 = (H, H)$ in the notation of the previous section.

The public signal z is generated through $z = s(a) + x$ with the deterministic function s of the action profile and a random variable x specified as follows: The

¹⁶That is, $y^*(\delta)$ stays the same when a positive constant is added or multiplied to, or subtracted from, the stage-payoffs.

¹⁷Our choice of this particular stage-game is based on the fact that it scores high on the indexes proposed by Rapoport and Chammah (1965) and Roth and Murnighan (1978) that correlate with the level of cooperation in the infinitely repeated PD with perfect monitoring.

function s is given by

$1 \setminus 2$	L	H
L	20	18
H	18	16

and the random variable x has the following distribution:

$$f(x) = \frac{1}{2\beta} e^{-\frac{|x|}{\beta}}, \text{ and } F(x) = \begin{cases} 1 - \frac{1}{2} e^{-\frac{x}{\beta}} & \text{if } x \geq 0 \\ \frac{1}{2} e^{\frac{x}{\beta}} & \text{if } x < 0, \end{cases} \quad (5)$$

where $\beta > 0$. As can be readily verified, s satisfies Assumption 1 and f satisfies Assumption 2. Moreover, the parameter β represents the level of noise in the public signal since

$$E[x] = 0 \quad \text{and} \quad \text{Var}(x) = 2\beta^2.$$

From (4), we can write the normalized maximum symmetric equilibrium payoff as

$$y^*(\delta) = \begin{cases} \frac{1-\delta}{1-\delta+\frac{\delta}{2} e^{\frac{k^*(\delta)-20}{\beta}}} & \text{if } K(\delta) \neq \emptyset \\ 0 & \text{otherwise.} \end{cases}$$

The explicit description of the set $K(\delta)$ of admissible thresholds is found in the Appendix.

In the actual treatment, we need to bound the subjects' payoffs while allowing z to have full support as assumed by the theory. For this, we have the subjects receive the expected stage payoffs $g_i(a)$ at the conclusion of the experiment instead of having them receive the (stochastic) payoffs $w_i(a_i, z)$ after each period.¹⁸ We simulate the infinitely repeated game with the discount factor $\delta = 0.9$ by terminating the game after each period with a fixed probability of 0.1.¹⁹

The experiments proceed in the following steps. First, subjects are provided with the basic information about the game they will play. They are then matched

¹⁸In this specification, hence, the public signal z indicates the opponent's action choice but does *not* directly affect the players' payoffs. This design hence abstracts from the psychological impact of the payoffs as analyzed by Bereby-Meyer and Roth (2006). Alternatively, we could have paid the subjects $w_i(a_i, z)$ after each period of play for some (bounded) function w_i . This payment method, however, would have added complexity to the instructions to the subjects.

¹⁹One possible concern with a fixed time design is that subjects hurry their decisions so as to increase the payment. To address this concern, we forced a pause between periods. The pause began only after every subject in the session had finished their play, and lasted for 25 seconds during which subjects were asked to write down their decision and the public signal z on a record sheet.

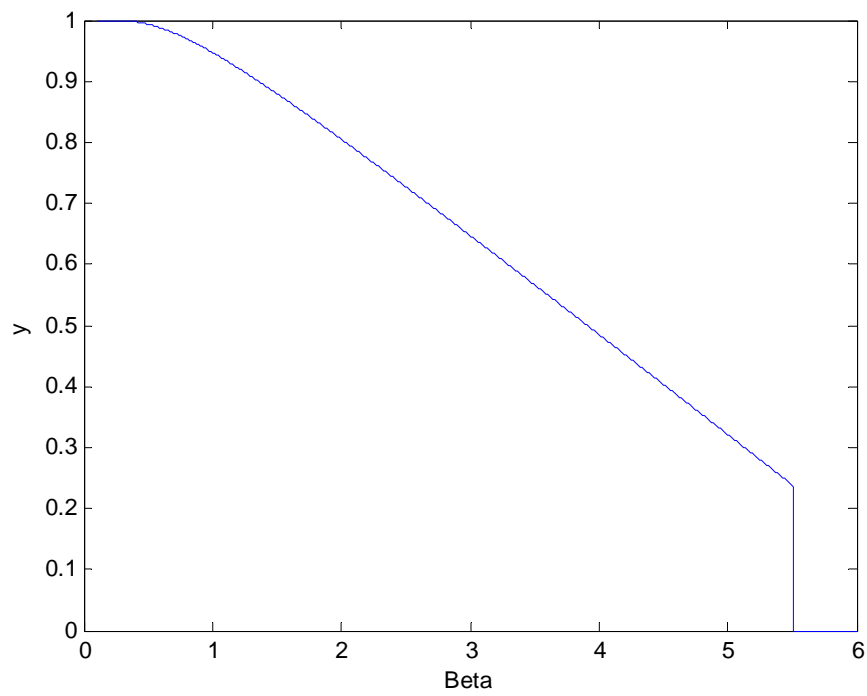


Figure 1: y^* as a function of β ($\delta = 0.9$)

in pairs to play a repeated game with imperfect public monitoring. As mentioned earlier, this repeated game has stochastic length and terminates in finite time almost surely.²⁰ The sequence of play between any pair of subjects is referred to as a *cycle*. At the conclusion of every cycle, the subjects see on the screen their own payoffs in that cycle. They are then randomly rematched to play a new cycle. The information provided to the subjects at the outset includes: (i) The length of a cycle is randomly determined by the termination probability 0.1. (ii) They play against a randomly chosen subject in the session. (iii) The distribution of the random shocks to the public signal is given by (5).²¹ The random matching for each new cycle is done in a round robin manner: A subject is matched with someone new as long as it is possible, and matched with someone they have played with previously thereafter.²² The first cycle to end after one hour of play marks the end of the session. Therefore, different sessions have different numbers of cycles.

In the experiments, we use four different values of β : They are $\beta = 0$ (no noise = perfect monitoring), $\beta = 1$ (low noise), $\beta = 4$ (medium noise), and $\beta = 10$ (high noise). Figure 1 plots $y^* \equiv y^*(0.9)$ as a function of β . In addition to the above four treatments, we conduct two control treatments. The first control admits perfect monitoring of the opponent's action but matches subjects anonymously and randomly after every period of play. Since this control eliminates the repeated game effect (*i.e.*, $\delta = 0$), it is referred to as "one-shot" and symbolized as $\beta = \text{OS}$.²³ The second control uses the same continuation probability as in the first four treatments but eliminates the public signal. Since this treatment is equivalent to having subjects observe infinitely noisy public signal, it is referred to as $\beta = \infty$.²⁴ Given that both controls support no cooperation in theory, they should provide a benchmark for the first four treatments with respect to the level of cooperation.

Subjects were recruited through announcements in undergraduate classes, advertisements in the student newspaper, flyers posted on campus, and e-mail advertise-

²⁰ All games played in the same session terminate simultaneously.

²¹ The instructions given to the subjects can be found at http://homepages.nyu.edu/~gf35/print/Aoyagi_2008a_inst.pdf.

²² As it happened, the subjects were matched with someone they played with previously in only 14% of the cases.

²³ For the one-shot treatment, the length of each session equals 75 periods, approximately the average number of periods in the other treatments. For this control, the term *cycle* refers to the block of initial 15 periods and every successive block of 10 periods.

²⁴ The second control was suggested by one of the referees. The two controls differ from each other also by the existence of the feedback of outcomes.

	Treatments											
	$\beta = 0$		$\beta = 1$		$\beta = 4$		$\beta = 10$		$\beta = \text{OS}$		$\beta = \infty$	
Sessions	2		2		2		2		2		2	
Subjects	24	14	26	16	20	20	20	26	16	20	20	18
Cycles	8	10	6	7	4	5	5	8	—	—	10	10
Periods	74	91	73	78	66	75	73	69	74	75	69	83

Table 1: Subjects, cycles and periods per session

ment at the Ohio State University. This resulted in recruiting a broad cross section of undergraduate students. At the end of each experimental session, subjects were paid \$0.017 for every point they accumulated in the experiment. Earnings ranged from \$20.20 to \$36.55. Details about the number of subjects and periods in each treatment are provided in Table 1.²⁵

5 Results

5.1 Payoffs

We first examine the subjects' payoffs. Note that in the continuation probability formulation, the sum of stage payoffs for the duration of the game corresponds to the average discounted payoff of the infinitely repeated game. For each given value of β , let $\bar{v}(\beta)$ be the sum of stage payoffs averaged over all cycles and sessions, and let

$$\bar{y}(\beta) = \frac{\bar{v}(\beta) - g^0}{\hat{g} - g^0}$$

be the normalization of $\bar{v}(\beta)$. When the subjects play a symmetric equilibrium of the repeated game, then $\bar{y}(\beta)$ should lie between the one-shot NE level 0 and the maximal symmetric equilibrium level $y^*(\beta)$ for each β .²⁶

We refer to the three treatments ($\beta = 0, 1, 4$) for which cooperation is possible according to the theory (*i.e.*, $y^*(\beta) > 0$) as the *cooperation treatments*, and the three treatments ($\beta = 10, \infty$, and OS) for which it is not (*i.e.*, $y^*(\beta) = 0$) as the *non-cooperation treatments*. The general evolution of play between the cooperation and non-cooperation treatments is strikingly different: Figure 2 plots $\bar{y}(\beta)$ by treatment

²⁵The $\beta = 0$ session with 14 subjects had a crash after the end of cycle 8, and was re-started for two additional cycles. The $\beta = \text{OS}$ session with 16 subjects had 1 crash.

²⁶In the experiment, β is changed while $\delta = 0.9$ is fixed. Hence, y^* is indexed by β instead of δ .

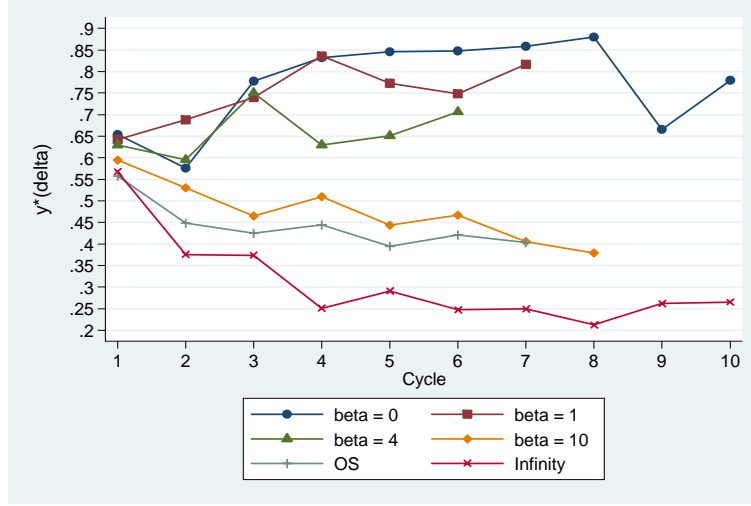


Figure 2: Evolution of \bar{y} by treatments

and by cycle. This figure shows that \bar{y} has an upward trend over time for $\beta = 0, 1$, and 4 , and an opposite, downward trend for $\beta = 10, \infty$, and OS. In other words, subjects appear to improve their ability to cooperate over time when cooperation is theoretically possible, but learn not to cooperate otherwise. As seen, $\bar{y}(\beta)$ is much higher in the first set of treatments, and seems to increase as noise decreases. As for the second set, both treatments have relatively close $\bar{y}(\beta)$. Diagrammatically, we have:

$$\bar{y}(0) > \bar{y}(1) > \bar{y}(4) \gg \bar{y}(10) \approx \bar{y}(\text{OS}) > \bar{y}(\infty) > 0,$$

indicating that \bar{y} has the same relative ordering as y^* except for OS and $\beta = \infty$, which should create no difference in theory. A few more aspects of this figure are worth noting: For all treatments, the average payoffs start out almost identically, yet begin to differ substantially by cycle 3. Furthermore, by cycle 3 they almost reach the level they will eventually keep in the end.

In order to concentrate on stable behavior, the analysis in what follows excludes data from the first two cycles or the cycles that occur in the first 20 periods of play.²⁷

Table 2 lists the values of $y^*(\beta)$ and $\bar{y}(\beta)$ for each treatment. $\bar{y}(\beta)$ is positive

²⁷Note that 20 is the expected number of periods for 2 cycles. In all but one session, the first 2 cycles last at least 20 periods. The exception is one session of the $\beta = 0$ treatment where the first 2 cycles only had a total of 3 periods of play.

Treatment	$y^*(\beta)$	$\bar{y}(\beta)$
$\beta = 0$	1	0.845 (0.129)
$\beta = 1$	0.948	0.774 (0.163)
$\beta = 4$	0.486	0.695 (0.208)
$\beta = 10$	0	0.467 (0.182)
OS	0	0.418 (0.101)
$\beta = \infty$	0	0.270 (0.149)

Standard deviations in parentheses.

Table 2: $y^*(\beta)$ and $\bar{y}(\beta)$ by treatments

Treatments	$\beta = 1$	$\beta = 4$	$\beta = 10$
$\beta = 0$	0.017	0.000	0.000
$\beta = 1$		0.031	0.000
$\beta = 4$			0.000

Table 3: p-values of the one-sided Mann-Whitney test that \bar{y} decreases with noise

for both the cooperation and non-cooperation treatments at the 1% significance level. $\bar{y}(\beta)$ lies in the predicted interval $[0, y^*(\beta)]$ for two of the three cooperation treatments, but not for the non-cooperation treatments. In either case, it is not close to $y^*(\beta)$: it is too low when $\beta = 0$ and 1, and too high for all other treatments. It should also be noted that $\bar{y}(10)$ and $\bar{y}(\text{OS})$ appear comparable, but that $\bar{y}(\infty)$ is lower. Each of these observations will be analyzed in turn.

That $\bar{y}(\beta)$ decreases with β is formally established in Table 3, which gives the p-values of a Mann-Whitney test of the hypothesis that the $\bar{y}(\beta)$ for β on the left (row) is equal to the $\bar{y}(\beta)$ for β in the top (column) against the one sided hypothesis that the former is greater than the latter.²⁸ Every test is statistically significant

²⁸For all such tests, per subject averages are used instead of all the subject-cycle averages because it is likely that \bar{y}_i are correlated across cycles for a given individual.

at the 5% level. On the other hand, the hypothesis that $\bar{y}(10) = \bar{y}(\text{OS})$ cannot be rejected ($p = 0.108$, two-tailed Mann-Whitney test) while the hypothesis that $\bar{y}(10) = \bar{y}(\infty)$ is rejected ($p\text{-value} < 0.01$).²⁹ The results of these tests support the general theoretical predictions that cooperation is easier to sustain when noise is low, and that cooperation under high but finite noise is as difficult as in the one-shot case.³⁰ It does not immediately follow from the theory however that cooperation under infinite noise (*i.e.*, no public signal) is more difficult than in the one-shot case. On the other hand, a sign test rejects the hypothesis that \bar{y} is equal to y^* at the 1% level for each treatment, implying that the subjects' play does not conform to the most efficient symmetric equilibrium.

In contradiction to the theory, both $\bar{y}(10)$, $\bar{y}(\infty)$, and $\bar{y}(\text{OS})$ are significantly positive. This result is in line with the positive levels of cooperation observed in experiments on the one-shot PD. To be precise, Figure 3 describes the evolution of the rate of the cooperative action $a_i = L$ as well as that of the action profile (L, L) . The observed level of cooperation is relatively high when compared to those in related experiments on the PD. For example, it is significantly higher than that reported by Duffy and Ochs (2007) or Dal Bo (2003): In the random matching treatment of Duffy and Ochs, the rate of cooperation drops to almost 0% by the end. In the one-shot treatment of Dal Bo, the rate of cooperation is a little more than 5% by the end. We believe that this difference is attributed to the selection of the payoff matrix: Our payoff numbers, which are designed to generate high levels of cooperation under the perfect monitoring treatment, raised the level of cooperation in the non-cooperation treatments as well.³¹

When the subjects play the most efficient symmetric equilibrium, their period 1 action should equal $a_i = L$ when $\beta = 0, 1$, or 4 . Under this hypothesis, hence, the data must always exhibit action L in period 1 of any cycle in any cooperation

²⁹Both ANOVA and Kruskal-Wallis tests reject at the 1% level the null hypothesis that the treatment has no effect on y for $\beta = 0, 1, 4$, and 10 .

³⁰This conclusion does not change even if each session is treated as the unit of observation: If observations are correlated within a session, one could argue that each session should be treated as a single data point. To address this concern we can average \bar{y} by session and use a Mann-Whitney test to show that \bar{y} is higher for $\beta = 1$ and $\beta = 4$ than for $\beta = 10$, $\beta = \infty$, and OS. The one-sided null hypothesis is rejected with a $p\text{-value}$ of 0.01 . In fact, \bar{y} in any session of the cooperation treatments is higher than that in any session of the non-cooperation treatments.

³¹According to all four indices proposed by Rapoport and Chammah (1965) and Roth and Murnighan (1978), our PD is expected to generate more cooperation than any of the three PD matrices used by Duffy and Ochs (2007) and Dal Bo (2003).

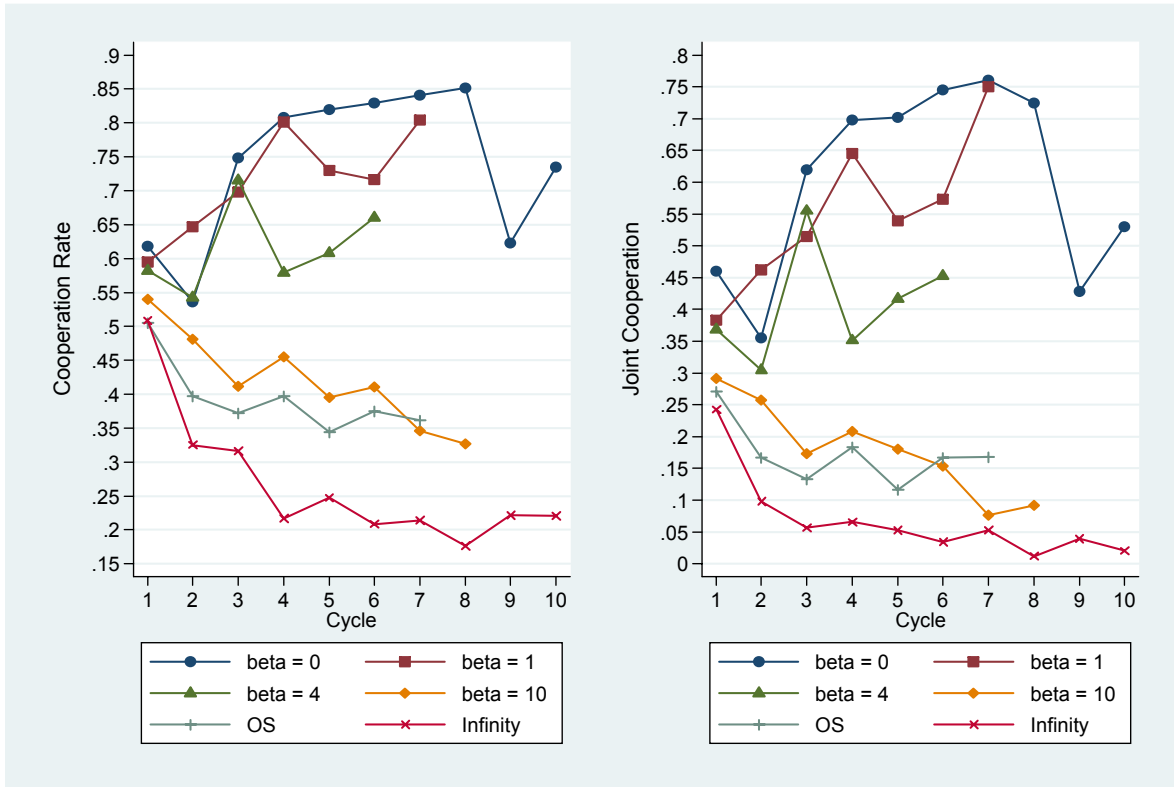


Figure 3: Evolution of cooperation: Rates of L and $(L^c L)$ by cycles

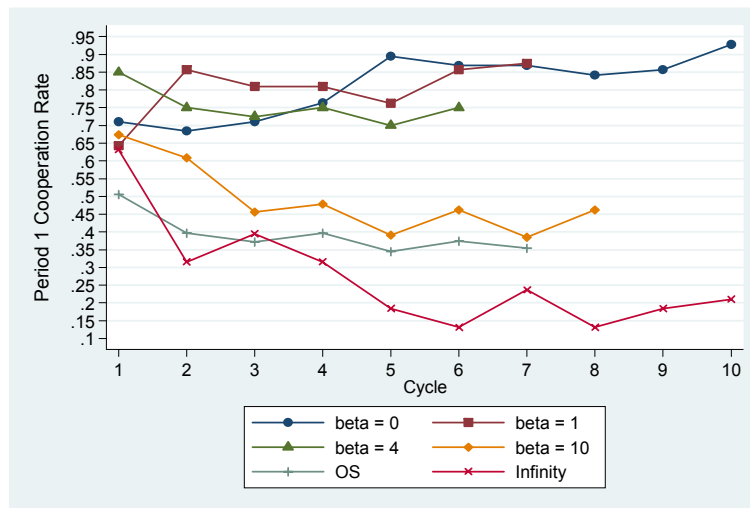


Figure 4: Rates of the cooperative choice L in period 1 by cycles

Treatments	$\beta = 1$	$\beta = 4$	$\beta = 10$	$\beta = \infty$
$\beta = 0$	0.630	0.287	0.000	0.000
$\beta = 1$		0.441	0.000	0.000
$\beta = 4$			0.000	0.000
$\beta = 10$				0.230

Table 4: p-values of the two-sided Mann-Whitney test that the rates of the cooperative choice L in period 1 are equal across treatments

treatment. Having found that the rate of cooperative action in period 1 is only 39.1% in their repeated public goods experiments, Sell and Wilson (1991) reject the hypothesis that their subjects use trigger strategies. In comparison, the rates of period 1 cooperation in our cooperation treatments are higher. Furthermore, the level of period 1 cooperation in those treatments increases over the course of each session (Figure 4). By the last cycle, the rate of period 1 cooperation is 85% in the cooperation treatments. On the other hand, the rate of period 1 cooperation in the non-cooperation treatments is much lower at 41% in the last cycle. Table 4 reports the p-values for the test of the hypothesis that period 1 cooperation is the same across different treatments. The hypothesis that they are the same across all cooperation treatments cannot be rejected.³² Neither can be the hypothesis that they are the same across the non-cooperation treatments. On the other hand, we can reject the hypothesis that they are the same between the cooperation and non-cooperation treatments. However, for all treatments, sign tests reject at any conventional level the hypothesis that the rates of period 1 cooperation equal unity.

5.2 Strategies

We next turn to the analysis of the strategies. A particular focus is on threshold strategies with regime shifts between cooperation and punishment states. This class includes trigger and tit-for-tat strategies, which are most often discussed in the experimental analysis of repeated game strategies. We use standard likelihood ratio tests to examine if any particular specification of a threshold strategy describes the observed pattern of play. In all but one treatment, we find that the data is best described by the strategy which uses the same threshold in both the cooperation

³²Neither ANOVA nor Kruskal-Wallis rejects the null hypothesis that the treatment has no impact on the rate of period 1 cooperation for $\beta = 0, 1$, and 4 (p-value > 0.1).

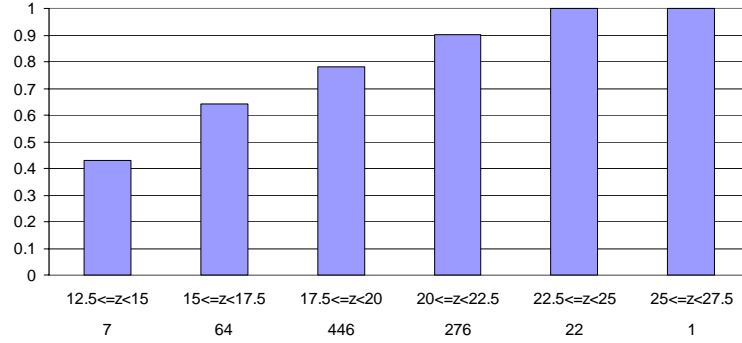


Figure 5: Rates of L as a function of the most recent public signal ($\beta = 1$), [No. of Obs. Under the Category]

and punishment states.³³ For numerical analysis in this subsection, we identify 1 with the cooperative action L and 0 with the non-cooperative action H . We also denote by $c_{it} \in \{0, 1\}$ subject i 's actual action choice in period t .

Note first that any strategy that supports cooperation must condition the current choice on the past public signals. In fact, this is what we observe in this experiment. In the perfect monitoring ($\beta = 0$) treatment, for example, if both players cooperated in the last period, each player cooperates in the current period 94% of the time. On the other hand, if one player cooperated and the other defected in the last period, then the rate of cooperation by the former in the current period decreases to 43%. The same trend can be found in the imperfect monitoring treatments. Figure 5 shows for the $\beta = 1$ treatment how a player who cooperated in the previous period chooses his action in the current period as a function of the most recent public signal z . The increasing rate of cooperation with z is found in the $\beta = 1$ and $\beta = 4$ treatments, and also to some extent in the $\beta = 10$ treatment where the theory predicts no cooperation.

We next examine the impact of the public signal on future actions through a probit model which regresses subject i 's action choice $c_{it} \in \{0, 1\}$ in period t on the following variables:³⁴

$$z_{t-1} \mathbf{1}_{\{c_{i,t-1}=1\}}, z_{t-1} \mathbf{1}_{\{c_{i,t-1}=0\}}, z_{t-2} \mathbf{1}_{\{c_{i,t-2}=1\}}, z_{t-2} \mathbf{1}_{\{c_{i,t-2}=0\}}. \quad (6)$$

³³The analysis in this subsection excludes data from the control treatments.

³⁴For any condition A , $\mathbf{1}_A$ denotes the indicator function which takes value 1 if the condition holds and value 0 otherwise.

That is, the estimated equation measures the impact of the pair (public signal, action choice) in each one of the last two periods on the current action choice.³⁵ In all but one treatment, the first two regressors in (6) are statistically significant while the other two are not.³⁶ In other words, the impact of the public signal on the subjects' action choice is immediate but not persisting. The sole exception to this is the $\beta = 0$ treatment where the other two regressors are also significantly positive. However, their impacts are much smaller with the marginal effects of the last two variables about one third of those of the first two. Note that the finding here in particular suggests that the subjects' strategies are not described by the grim-trigger type.

5.2.1 Threshold strategies with regime shifts

Our analysis of strategies in this section focuses on a *threshold strategy with regime shifts*, or simply a *threshold strategy*, which has two states 0 and 1, chooses L (cooperation) in state 0 and H (non-cooperation) in state 1, begins in state 0, and has the following rules of transition between the two states:

State 0: Stay in state 0 if the public signal $z > a$ and the own action $c_i = 1$ ($\Leftrightarrow L$), or if $z > a + b_1$ and $c_i = 0$ ($\Leftrightarrow H$). Move to state 1 otherwise.

State 1: Move to state 0 if $z > a + b_2$ and $c_i = 1$, or if $z > a + b_2 + b_3$ and $c_i = 0$. Stay in state 1 otherwise.

The second transition rule in each state accounts for a possible mistake where a player fails to choose an intended action. The threshold strategy is illustrated in Figure 6, where the condition above each arrow applies when the own action $c_i = 1$, and the condition below applies when $c_i = 0$. We allow each parameter to be $+\infty$ or $-\infty$. The threshold grim-trigger strategy discussed in Section 4 is a threshold strategy with $b_2 = \infty$ and $b_1 = b_3 = 0$. In general, a threshold strategy is public if $b_1 = b_3 = 0$. On the other hand, the tit-for-tat strategy in the perfect monitoring environment is a leading example of a private threshold strategy. Note, however, that the on-the-path action is publicly determined even if $b_1 \neq 0$ and/or $b_3 \neq 0$. Therefore, if a threshold strategy gives rise to a symmetric equilibrium

³⁵Three or more lags are omitted since they lead to the loss of too many data. See the Appendix for the precise specification of the estimated equation as well as for the table of estimation results (Table 9).

³⁶In all treatments with $\beta \geq 1$, a Wald test that both regressors for $t - 2$ have no impact cannot be rejected at the 10% level.

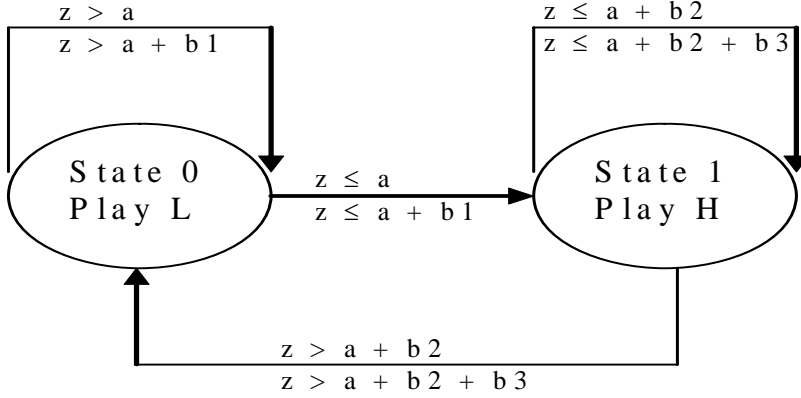


Figure 6: Threshold strategy with regime shifts

with perfection after every public history, then the associated payoff is still bounded from above by v^* (or its normalization y^*) defined in Section 3. Proposition 3 in the Appendix identifies when a pair of threshold strategies constitute a symmetric perfect public equilibrium. With the identification of each state with the action intended in it, we obtain an alternative representation of the threshold strategy used in our analysis: Let z_t be the public signal, and c_{it} be i 's actual action choice in period t . Then the action s_{it} intended by the threshold strategy in period t is given by $s_{i1} = 1$ and for $t \geq 1$,

$$s_{i,t+1} = \mathbf{1}_{\{z_t > k_{it}\}}, \quad (7)$$

where

$$k_{it} = \begin{cases} a & \text{if } s_{it} = 1 \text{ and } c_{it} = 1, \\ a + b_1 & \text{if } s_{it} = 1 \text{ and } c_{it} = 0, \\ a + b_2 & \text{if } s_{it} = 0 \text{ and } c_{it} = 1, \\ a + b_2 + b_3 & \text{if } s_{it} = 0 \text{ and } c_{it} = 0. \end{cases} \quad (8)$$

Based on (7) and (8), we estimate a limited dependent variable model with lagged dependent variables which incorporates the possibilities of mistakes and asymmetry

among subjects as follows ($t \geq 0$):³⁷

$$s_{i,t+1} = \mathbf{1}_{\{\gamma_0 z_t > \kappa_{it}\}}, \quad (9)$$

$$c_{i,t+1} = \mathbf{1}_{\{\gamma_0 z_t > \kappa_{it} + u_{it}\}}, \quad (10)$$

where

$$\kappa_{it} = \begin{cases} \alpha + \nu_i & \text{if } s_{it} = 1 \text{ and } c_{it} = 1 \\ \alpha + \nu_i + \gamma_1 & \text{if } s_{it} = 1 \text{ and } c_{it} = 0 \\ \alpha + \nu_i + \gamma_2 & \text{if } s_{it} = 0 \text{ and } c_{it} = 1 \\ \alpha + \nu_i + \gamma_2 + \gamma_3 & \text{if } s_{it} = 0 \text{ and } c_{it} = 0. \end{cases} \quad (11)$$

In (11), note that the parameters α , γ_1 , γ_2 , and γ_3 are all assumed common across subjects, and that only ν_i is indexed by i . (9) is equivalent to (7) if $\gamma_0 > 0$ and $\kappa_{it} = \gamma_0 k_{it}$, which will be restored if³⁸

$$a = \frac{\alpha + \nu_i}{\gamma_0}, b_1 = \frac{\gamma_1}{\gamma_0}, b_2 = \frac{\gamma_2}{\gamma_0}, b_3 = \frac{\gamma_3}{\gamma_0}. \quad (12)$$

The term ν_i represents the *correlated random effects* and captures possible asymmetry across the subjects: The larger is ν_i , the higher the threshold and hence the more likely is subject i to play the non-cooperative action H . Each ν_i is assumed to have an independent normal distribution $N(\psi\zeta_i, \sigma_\nu)$ for some constants ψ , σ_ν and ζ_i . The variance σ_ν and the factor of proportion ψ are assumed common across subjects, and estimated from the data, while ζ_i is set equal to the fraction of times that subject i chooses H in period 1 of each cycle under estimation: ζ_i serves as a proxy for i 's tendency to play the non-cooperative action given that any threshold strategy would play action L in period 1.³⁹ On the other hand, the term u_{it} in

³⁷We have also tested a more general version of the threshold strategy in which the regime shifts depend on the public signal up to T periods ago ($T > 1$). They did not have better fits than the present specification with $T = 1$. The interested reader is referred to our working paper (Aoyagi and Fréchette (2005)).

³⁸ γ_0 is required for a technical reason: It will allow the error term u_{it} introduced later to have the unit variance.

³⁹We assume that the mean of ν_i is proportional to ζ_i in order to deal with the initial conditions problem. (See Heckman (1981) or Chamberlain (1980) for the static case.) Under an alternative assumption that $\nu_i \sim N(0, \sigma_\nu^2)$, the consistency of our estimate would require the (unlikely) independence of c_{i1} 's and ν_i 's. (See Wooldridge (2002) for a clear exposition of the initial conditions problem and solutions to it.) The log likelihood is estimated using quadrature techniques with a twelve points Gauss Hermit quadrature. Weights and abscisae can be found in Abramowitz and Stegun (1972).

	$\beta = 0$	$\beta = 1$	$\beta = 4$	$\beta = 10$
α	-1.437*** (0.127)	-0.819*** (0.090)	-0.569*** (0.102)	-0.430*** (0.150)
ψ	1.698*** (0.210)	1.313*** (0.379)	1.278*** (0.250)	1.393*** (0.250)
ρ	0.265 [§]	0.363 [§]	0.295 [§]	0.183 [§]
LL	-781.234	-577.886	-809.018	-774.420
Obs.	1908	1112	1360	1258

*, **, *** indicate statistical significance at 10%, 5%, and 1% respectively.

[§] indicates statistical significance at 1% using a likelihood ratio test.

Table 5: Estimates of the random choice strategy

(10) represents a random error made by subjects: i 's action choice $c_{i,t+1}$ in period $t + 1$ is determined in reference to $\kappa_{it} + u_{it}$ rather than κ_{it} itself. We assume that u_{it} is independent across subjects and across periods, and has the standard normal distribution $N(0, 1)$. One intended effect of specifying the shock term u_{it} as in (10) is the stronger tendency to make errors when the realized public signal z_t is closer to the threshold $k_{it} = \frac{\kappa_{it}}{\gamma_0}$.

5.2.2 Estimation of threshold strategies

We first estimate a benchmark model of a *random choice strategy*, which plays each action (H and L) with a constant probability throughout the game. Formally, this strategy is obtained by setting $\gamma_0 = \gamma_1 = \gamma_2 = \gamma_3 = 0$ in (9) and (10). Table 5 reports the estimates of this benchmark model.⁴⁰

Table 6 reports the estimation results for the general threshold specification

⁴⁰In Table 5, $\rho = \frac{\sigma_\nu^2}{\sigma_\nu^2 + 1}$ is used as a substitute for σ_ν^2 as is customary for random effects estimates. The estimate of the coefficient ψ on ζ_i is positive, indicating that someone who is less likely to cooperate in period 1 is less likely to do so in any other period. The constant term α increases with noise, which implies that increasing noise tends to decrease cooperation. The random-effects specification is not rejected in any treatment. As for the base threshold level a in (8), the coefficient estimates of α and γ_0 both decrease with noise, and the ratio of those estimates ($\frac{\alpha}{\gamma_0}$) decreases with noise. This indicates from (12) that the component of the base threshold a which is common across subjects decreases with noise. On the other hand, the ratio of the estimates of ψ and γ_0 (i.e., $\frac{\psi}{\gamma_0}$) increases with noise, indicating that the weight on the subject-specific component of a increases with noise. In other words, as noise increases, the gap in the base threshold levels widens across individuals.

(9) and (10).⁴¹ In comparison with the benchmark model, the general model has explicative power. For each $\beta \leq 4$, a likelihood ratio test rejects any of the random choice model with p-value=0.01. On the other hand, for $\beta = 10$, the random choice model cannot be rejected even at p-value=0.10. It should be noted that the random-effects specification is not rejected for any treatment, suggesting that tendencies to cooperate vary across the subjects.⁴² It is also worth noting that γ_0 is statistically significant in the $\beta = 10$ treatment in contradiction to the theoretical prediction: This suggests that the subjects use the public signal even when they should not. Few coefficient estimates are statistically significant for the $\beta = 0$, $\beta = 1$ and $\beta = 10$ treatments, while all regressors are statistically significant for the $\beta = 4$ treatment.

Given the observations above, our next objective is test some parameter restrictions on the general threshold specification. First, we are interested in testing whether the subjects' strategies are public in the sense that their actions depend only on the past public signals. Recall that a threshold strategy is public if and only if $b_1 = b_3 = 0$. Alternatively, when the strategies are private and choose actions based also on the own action choice in the past, b_1 and/or b_3 are non-zero. In this case, we suppose that b_1 and/or b_3 are chosen so that they will cancel out the effect of a player's own deviation on the public signal. For example, suppose that a player has deviated to H in state 0. This deviation causes the signal distribution to shift down by 2 given our specification of $s(a)$. In other words, the probability that $z \leq a$ when the own action is L (as intended in state 0) exactly equals the probability that $z \leq a - 2$ when he has deviated to H . Hence, we use $b_1 = -2$ as an alternative to $b_1 = 0$. Likewise, we use $b_3 = -2$ as an alternative to $b_3 = 0$. In summary, we test four alternative specifications (Sa)-(Sd) as listed in part (1) of Table 7. As a further motivation of these alternatives, consider the tit-for-tat strategy in the perfect monitoring environment ($\beta = 0$). Since the signals are 20, 18 and 16 for sure after the action profiles (L, L) , (H, L) , and (H, H) , respectively, tit-for-tat is expressed as the threshold strategy with $a \in [18, 20)$, $b_1 = -2$, $b_2 = 0$ and $b_3 = -2$. If a subject plays tit-for-tat, hence, (Sd) should best fit the data. We find that in all treatments with or without random effects, (Sa) fits the data best of all the four alternatives.

⁴¹ Although we do not place the restriction that $\gamma_0 > 0$, it will be seen that the estimate of γ_0 turns out to be positive in every treatment. The standard errors are bootstrapped: Fifty replications using the full sample size are computed.

⁴² When testing for the significance of the random effects specification, the fact that the null hypothesis is at the boundary of the parameter space is properly dealt with. See Gutierrez, Carter, and Drukker (2001).

	$\beta = 0$	$\beta = 1$	$\beta = 4$	$\beta = 10$
α	9.543 (42.103)	4.410*** (1.157)	-0.245*** (0.084)	-0.318* (0.168)
γ_0	0.564 (2.335)	0.283*** (0.988)	0.026*** (0.007)	0.008* (0.005)
γ_1	-0.237 (4.957)	0.010 (0.011)	0.804*** (0.103)	0.079 (0.065)
γ_2	-0.167 (0.580)	-0.211 (0.245)	0.049* (0.026)	0.010 (0.009)
γ_3	-0.147 (0.329)	-0.010 (0.024)	0.011* (0.006)	-0.003 (0.002)
ψ	0.908 (1.898)	1.192* (0.126)	0.010* (0.005)	1.460*** (0.498)
ρ	0.116 [§]	0.229 [§]	0.313 [§]	0.218 [§]
LL	-636.384	-520.464	-775.364	-771.442
Obs.	1908	1112	1360	1258

*, **, *** indicate statistical significance at 10%, 5%, and 1% respectively.

[§] indicates statistical significance at 1% using a likelihood ratio test.

Table 6: Parameter estimates of the general threshold strategy

	Names	b_1	b_2	b_3
(1)	Sa	0	—	0
	Sb	0	—	−2
	Sc	−2	—	0
	Sd	−2	—	−2
(2)	S ₀	0	0	0
	S ₁	0	−∞	0
(3)	Se	—	0	0
	Sf	0	0	—

The restrictions on b_1 and b_3 translate to those on γ_1 and γ_3 through (12):

$$b_1 = 0 \Leftrightarrow \gamma_1 = 0, \quad b_1 = -2 \Leftrightarrow \gamma_1 = -2\gamma_0,$$

$$b_3 = 0 \Leftrightarrow \gamma_3 = 0, \quad b_3 = -2 \Leftrightarrow \gamma_3 = -2\gamma_0.$$

We also impose $\gamma_0 > 0$.

Table 7: Parameter restrictions on threshold strategies

This, combined with the estimation results on b_2 presented below, provides general support for the use of public strategies except when $\beta = 4$.⁴³ Among the other three alternatives, (Sc) has the second best fit in all but one case ($\beta = 0$ with random effects) where it has the third best fit. (Sd) has the worst fit in all treatments with random effects, and in $\beta = 0$ without them. In the other cases it has the third best fit and (Sb) has the worst fit.

Given the above results, we now assume that $b_1 = b_3 = 0$ and examine the hypotheses on the value of b_2 , which has to do with the intended duration of the punishment state. Specifically, we test (S₀) and (S₁) in part (2) of Table 7: (S₀) uses the same threshold in both states 0 and 1, implying that it chooses L if and only if the most recent signal exceeds a . It is hence the simplest and most intuitive of all the threshold strategies. Just like tit-for-tat or grim-trigger, a player can play (S₀) without explicit recognition of its formal definition. It can be seen from (21) (in the Appendix) that when $a = 20$, a pair of (S₀) is a symmetric perfect public equilibrium if $\beta \leq 2/\log 6 \approx 1.12$. Furthermore, there exists a for which a pair of

⁴³For $\beta = 4$, see the discussion later in this section.

(S₀) is a symmetric Nash equilibrium if and only if $\beta \leq 2\left(\log \frac{5\delta}{5\delta-2}\right)^{-1} \approx 3.4026$.⁴⁴ On the other hand, (S₁) stays in state 1 just for one period and returns to state 0 afterward regardless of the realized public signal.⁴⁵ (S₁) can also be thought of as a special case of the trigger strategy discussed by Green and Porter (1983).

In every treatment except for $\beta = 4$, we cannot reject at the 10% level the hypothesis that (S₀) fits the data as well as the general threshold specification does. The results of two non-parametric tests reported in the Appendix also provide support for the use of the (S₀) strategy for $\beta = 0$. On the other hand, (S₁) fits the data poorly: The hypothesis that it fits the data as well as the general specification does is rejected at the 1% level for most treatments. The $\beta = 4$ treatment presents the sole exception to the general rule with (S₀) rejected at the 1% level. The rejection may in part be explained by the fact that $\beta = 4$ exceeds the maximal noise level admitting a Nash equilibrium as mentioned above.

The fact that (S₀) is not rejected for $\beta = 10$ indicates that some subjects may condition their behavior on the public signal even though the theory suggests that they should not. Again the random-effects and ψ both turn out to be statistically significant, implying the existence of asymmetry across subjects. As with the general threshold specification, the coefficient estimates of α and γ_0 , as well as the ratio of those estimates ($\frac{\alpha}{\gamma_0}$) all decrease with noise, while the estimate of ψ and the ratio of the estimates of ψ and γ_0 both increase with noise.

Since the rejection of (S₀) in the $\beta = 4$ treatment indicates that at least one of the parameters b_1 , b_2 , and b_3 is non-zero, we choose specifications (Se) and (Sf) in part (c) of Table 7 as possible alternatives. As it turns out, both (Sa) and (Sf) are rejected at the 1% level, whereas (Se) is not rejected at the 10% level.⁴⁶ The estimates for γ_0 and γ_1 both turn out positive, implying that $b_1 = \gamma_1/\gamma_0 > 0$, or

⁴⁴For $\delta > 2/5$ and $\beta \leq 2\left(\log \frac{5\delta}{5\delta-2}\right)^{-1}$, the optimal threshold is given by

$$k^*(\beta, \delta) = 18 + \beta \log \frac{\frac{1}{\delta} - 1 + \sqrt{(\frac{1}{\delta} - 1)^2 + e^{-\frac{2}{\beta}}(3 - 4e^{-\frac{2}{\beta}})}}{3 - 4e^{-\frac{2}{\beta}}} \leq 18.$$

⁴⁵Note that $b_2 = -\infty$ would translate to $\gamma_2 = -\infty$ through (12). However, in order to generate randomness by u_{it} even when $b_2 = -\infty$, we replace $\gamma_2 = -\infty$ by $\gamma_2 = \gamma_0 \underline{z} - \alpha - \psi$, where \underline{z} denotes the lowest realization of z in the data: When $\gamma_2 = \gamma_0 \underline{z} - \alpha - \psi$, the inequality between $\gamma_0 z_t$ and $\kappa_{it} + u_{it}$ stays the same as when $\gamma_2 = -\infty$ for each realization of z_t if ν_i and u_{it} are at or below their mean values, but is reversed when $\nu_i + u_{it}$ takes a large positive value.

⁴⁶The coefficient estimates are $\alpha = 0.554$, $\gamma_0 = 0.059$, $\gamma_1 = 0.196$, $\psi = 1.048$, and $\rho = 0.339$.

	$\beta = 0$	$\beta = 1$	$\beta = 4$	$\beta = 10$
α	8.391*** (0.599)	4.220*** (0.511)	0.309* (0.168)	-0.330** (0.159)
γ_0	0.504*** (0.031)	0.274*** (0.027)	0.048*** (0.007)	0.005* (0.003)
ψ	0.819*** (0.257)	1.253*** (0.375)	1.219*** (0.242)	1.389*** (0.250)
ρ	0.124 [§]	0.227 [§]	0.286 [§]	0.184 [§]
LL	-639.219	-520.528	-784.588	-772.695
Obs.	1908	1112	1360	1258

*, **, *** indicate statistical significance at 10%, 5%, and 1% respectively.

[§] indicates statistical significance at 1% using a likelihood ratio test.

Table 8: Parameter estimates of the (S_0) strategy

that subjects have a higher threshold following an own deviation to H in state 0. It is not clear why the estimated strategy only in this treatment conditions on private actions.

6 Discussions

As discussed in the Introduction, this paper analyzes cooperation in infinitely repeated PD with imperfect public monitoring under the conditions of the standard theory. It studies the effect of noise levels in a standard oligopolistic setting where deviations monotonically shift the distribution of the one-dimensional public signal. Our findings suggest that subjects do cooperate in such an environment, and their payoffs are a decreasing function of the level of noise as predicted. The paper also analyzes the subjects' strategies by focusing attention on threshold strategies, which encompass trigger and tit-for-tat strategies that have been frequently discussed in the literature. Our estimates suggest that the subjects' strategies in most treatments have a remarkably simple representation: In every period, this strategy chooses the cooperative action when the public signal in the last period is above a certain threshold, and chooses the punitive action otherwise.

While the present paper limits its theory to symmetric equilibrium payoffs, the data suggests a certain degree of asymmetry in the subjects' strategies. A few com-

ments are in order regarding this point. First, a theory of asymmetric equilibrium payoffs as a function of noise would be enormously complex. We think that the simplifying assumption of symmetry provides a good approximation to our qualitative findings on noise and cooperation.⁴⁷ Second, while the consideration of asymmetric strategy profiles raises efficiency, it appears to add little to our analysis of payoffs: In the low noise treatments, the observed payoffs are lower than the maximal symmetric equilibrium payoff. Hence, they are also lower than the maximal asymmetric equilibrium payoff. In the high noise treatments, on the other hand, the observed payoffs do exceed the maximal symmetric equilibrium payoff. However, it is not easy to explain this through asymmetric equilibria either. The case in point is the control treatments where the subjects' payoffs are strictly positive. In theory, however, the unique (symmetric or asymmetric) equilibrium in those treatments is the repetition of the one-shot Nash equilibrium, which yields zero.

One possible explanation for the observed deviations of the subjects' payoffs can be provided by trembling in choosing actions. Suppose, for example, that the noise is low so that the efficient equilibrium strategy entails the cooperative action L most of the time on the equilibrium path. Then, when a player trembles, he switches from L to H more frequently, triggering a punishment and lowering his payoff from the level without trembling. On the other hand, if the noise is high, the efficient equilibrium strategy entails H most of the time, and trembling causes switching from H to L more often. This raises the player's payoff from the level without trembling.

It comes to our surprise that the infinite noise treatment generated a significantly lower level of cooperation than the one-shot treatment. One possible reason is that the subjects make their decision contingent on the public signal in the one-shot control as in the $\beta = 10$ treatment, but cannot do so in the infinite noise control. Another possibility, perhaps less likely, is that the effectively smaller group size associated with more frequent rematching in the one-shot control induces the subjects to consider the consequence of their own action on the group.⁴⁸

⁴⁷In fact, symmetry is the working assumption of much of the experimental literature, which often finds asymmetries across subjects. See, for example, Roth (1995) and Casari, Ham and Kagel (2004).

⁴⁸Such behavior, first explored by Kandori (1992a), however runs counter to the experimental finding of Duffy and Ochs (2007) that cooperation cannot be sustained in a randomly matching population even though the group size is small. A yet different explanation has to do with the fact that in all treatments but the one-shot control, the subjects were reminded before each new cycle

Our observation that the subjects in most treatments do not discount the threshold after an own deviation indicates that they use the signal more as a switching device rather than the source of information about the opponent's play. This is consistent with our interpretation of the public signal as a coordination device. On the other hand, if only a private signal of the opponent's play is available, we suspect that the subjects would give more weight to its informational content and possibly discount the effect of an own action on the signal.

The paper's discussion is based on a one-dimensional continuous signal distribution. This formulation, coupled with the monotone likelihood ratio (Assumption 2), justifies the use of thresholds in constructing strategies. While the subjects of the experiments may find the distribution function of a continuous signal with full support less accessible than that of, say, a finite signal, we believe that this problem is outweighed by the advantage of being able to focus on intuitive thresholds.

In comparison with the real industrial setting, the subjects in our experiments play in an extremely simple environment with only two stage actions and a single public signal. It remains to be seen whether the paper's observation continues to hold in a more complex environment that mimics the reality. In this sense, more analysis is required for the discussion of its implications on social welfare.

Appendix

Proof of Proposition 1. Since $h(\cdot \mid a)$ has full support for any a and w_i is continuous, the bang-bang property holds for a perfect public equilibrium (Abreu *et al.* (1990, Theorem 3)), and the maximal symmetric perfect public equilibrium payoff v can be generated by a stationary grim-trigger strategy profile σ , which plays \hat{a} throughout the cooperation phase and reverts to a^0 if and only if $z \notin Q$ for some $Q \subset \mathbf{R}$.⁴⁹

By the stationarity of play, v must satisfy

$$v = (1 - \delta)\hat{g} + \delta \{v P(z \in Q \mid \hat{a}) + g^0 P(z \notin Q \mid \hat{a})\}, \quad (13)$$

where $P(z \in Q \mid \hat{a}) = \int_{z \in Q} h(z \mid \hat{a}) dz$ is the probability that the public signal falls

that they would be matched with somebody randomly.

⁴⁹To achieve a payoff below the maximal level v , we would need to consider a (semi) non-stationary strategy for which the set of admissible signals in period 1 is different from that in subsequent periods.

in set Q under the action profile \hat{a} . Solving (13) for v , we get

$$v = \frac{(1 - \delta)\hat{g} + \delta g^0 P(z \notin Q \mid \hat{a})}{1 - \delta P(z \in Q \mid \hat{a})}. \quad (14)$$

For $\hat{\sigma}$ to be an equilibrium, playing \hat{a}_i in the cooperation phase must be incentive compatible for player i : For any alternative action $a_i \neq \hat{a}_i$, v and Q must satisfy

$$v \geq (1 - \delta)g_i(a_i^0, \hat{a}_j) + \delta \{v P(z \in Q \mid a_i^0, \hat{a}_j) + g^0 P(z \notin Q \mid a_i^0, \hat{a}_j)\}. \quad (15)$$

Solving (15) for v , we get

$$v \geq \frac{(1 - \delta)g_i(a_i^0, \hat{a}_j) + \delta g^0 P(z \notin Q \mid a_i^0, \hat{a}_j)}{1 - \delta P(z \in Q \mid a_i^0, \hat{a}_j)}. \quad (16)$$

Eliminating v from (14) and (16), we obtain

$$\frac{1 - \delta P(z \in Q \mid a_i^0, \hat{a}_j)}{1 - \delta P(z \in Q \mid \hat{a})} \geq \frac{g_i(a_i^0, \hat{a}_j) - g^0}{\hat{g} - g^0} = l. \quad (17)$$

Consider an alternative grim-trigger strategy profile $\hat{\sigma}$ which begins with \hat{a} and reverts to a^0 if and only if $z < k$, where k is such that

$$P(z \in Q \mid \hat{a}) = \int_Q h(z \mid \hat{a}) dz = \int_k^\infty h(z \mid \hat{a}) dz = P(z \geq k \mid \hat{a}). \quad (18)$$

It then follows from (14) that $\hat{\sigma}$ and σ yield the same payoff. On the other hand, the incentive constraint (17) for this strategy can be written as

$$\frac{1 - \delta\{1 - H(k \mid a_i^0, \hat{a}_j)\}}{1 - \delta\{1 - H(k \mid \hat{a})\}} \geq l, \quad (19)$$

where $H(\cdot \mid a)$ is the cumulative distribution corresponding to $h(\cdot \mid a)$. In what follows, we show that $\hat{\sigma}$ is also an equilibrium by verifying (19). Denote $K = [k, \infty)$ and write for $a_i \in A_i$,

$$M(a_i) = \int_{K \setminus Q} h(z \mid a_i^0, \hat{a}_j) dz \quad \text{and} \quad N(a_i) = \int_{Q \setminus K} h(z \mid a_i^0, \hat{a}_j) dz.$$

Note that $M(\hat{a}_i) = N(\hat{a}_i)$ by (18), and that

$$P(z \geq k \mid a_i^0, \hat{a}_j) = P(z \in Q \mid a_i^0, \hat{a}_j) + M(a_i^0) - N(a_i^0).$$

Note that (19) follows from $P(z \geq k \mid a_i^0, \hat{a}_j) \leq P(z \in Q \mid a_i^0, \hat{a}_j)$, or equivalently, $M(a_i^0) \leq N(a_i^0)$. Assumption 1 implies that

$$\begin{aligned}
M(a_i^0) &= \int_{K \setminus Q} h(z \mid a_i^0, \hat{a}_j) dz \\
&= \int_{K \setminus Q} h(k \mid a_i^0, \hat{a}_j) \frac{h(z \mid a_i^0, \hat{a}_j)}{h(k \mid a_i^0, \hat{a}_j)} dz \\
&\leq \int_{K \setminus Q} h(k \mid a_i^0, \hat{a}_j) \frac{h(z \mid \hat{a})}{h(k \mid \hat{a})} dz \\
&= \frac{h(k \mid a_i^0, \hat{a}_j)}{h(k \mid \hat{a})} M(\hat{a}_i),
\end{aligned}$$

and that

$$\begin{aligned}
N(a_i^0) &= \int_{Q \setminus K} h(z \mid a_i^0, \hat{a}_j) dz \\
&= \int_{Q \setminus K} h(k \mid a_i^0, \hat{a}_j) \frac{h(z \mid a_i^0, \hat{a}_j)}{h(k \mid a_i^0, \hat{a}_j)} dz \\
&\geq \int_{Q \setminus K} h(k \mid a_i^0, \hat{a}_j) \frac{h(z \mid \hat{a})}{h(k \mid \hat{a})} dz \\
&= \frac{h(k \mid a_i^0, \hat{a}_j)}{h(k \mid \hat{a})} N(\hat{a}_i).
\end{aligned}$$

Therefore,

$$M(a_i^0) - N(a_i^0) \leq \frac{h(k \mid a_i^0, \hat{a}_j)}{h(k \mid \hat{a})} \{M(\hat{a}_i) - N(\hat{a}_i)\} = 0.$$

This completes the proof. ■

Characterization of $K(\delta)$ and $k^*(\delta)$. Proposition 2 below provides a characterization of the set $K(\delta)$ of thresholds for which player i finds playing \hat{a}_i incentive compatible.

Assumption 3 f is continuous, and $f(0) = \max_{x \in \mathbf{R}} f(x)$.

Assumption 3 holds for many standard distributions as well as the one (5) used in our experiment. The following proposition shows that the optimal threshold under such a distribution is always below the expected value of the public signal.

Proposition 2 Suppose that Assumptions 1 and 2 hold. Then $K(\delta)$ is a (possibly empty) closed interval. If, in addition, Assumption 3 holds, then the optimal threshold $k^*(\delta) < s(\hat{a})$ when $K(\delta) \neq \emptyset$.

Proof. Define

$$W(k) = \int_{k-s(\hat{a})}^{\infty} \left\{ l - \frac{f(x+d)}{f(x)} \right\} f(x) dx.$$

After some algebra, we see that (2) is equivalent to

$$W(k) \geq \frac{l-1}{\delta} \quad (20)$$

Since f satisfies Assumption 2 and $d \geq 0$, it can be verified that $\frac{f(x+d)}{f(x)}$ is weakly decreasing in x . Take any k and k' such that $k < k'$ and $W(k), W(k') \geq \frac{l-1}{\delta}$. Then for any k'' between k and k' ,

$$W(k'') = W(k) - \int_{k-s(\hat{a})}^{k''-s(\hat{a})} \left\{ l - \frac{f(x+d)}{f(x)} \right\} f(x) dx,$$

and

$$W(k'') = W(k') + \int_{k''-s(\hat{a})}^{k'-s(\hat{a})} \left\{ l - \frac{f(x+d)}{f(x)} \right\} f(x) dx.$$

Since the quantity inside the brackets in each integrand is weakly increasing in x , if the first integral is positive, so is the second, and equivalently, if the second integral is negative, so is the first. In either case, we have $W(k'') \geq \frac{l-1}{\delta}$. This implies that the set of k 's which satisfy (8) is convex. That $K(\delta)$ is closed follows from the continuity of W .

Suppose now that Assumption 3 holds. We then have

$$l - \frac{f(d)}{f(0)} > 0.$$

Since $l - \frac{f(x+d)}{f(x)}$ is weakly increasing and continuous in x , if $W(k) \geq \frac{l-1}{\delta}$ for some $k \geq s(\hat{a})$, then $W(s(\hat{a}) - \gamma, a_i) > W(s(\hat{a})) \geq W(k) \geq \frac{l-1}{\delta}$ for a sufficiently small $\gamma > 0$ as well. This shows that $k^*(\delta) < s(\hat{a})$. ■

Description of $K(\delta)$. When the distribution of the random variable x is as specified in (5) with $\beta > 0$, the set $K(\delta)$ of effective thresholds is explicitly given as follows. Let

$$\begin{aligned} \lambda &= \log \frac{\delta^2 l}{(\delta l + 1 - l)^2}, \\ \mu &= \log \frac{\delta l}{\delta(2l - 1) - 2(l - 1)}, \\ \nu &= \log \frac{\delta}{\delta l - 2(l - 1)}. \end{aligned}$$

Note that μ is well-defined when $\delta > \frac{2(l-1)}{2l-1}$, and ν is well-defined when $\frac{2(l-1)}{l} < \delta < 1$. Furthermore, whenever these quantities are well-defined, we have $\log l < \lambda < \mu < \nu$. There are the following three cases depending on the value of the discount factor δ :

$$1) \quad \delta \in \left(0, \frac{2(l-1)}{2l-1}\right].$$

$$K(\delta) = \emptyset.$$

$$2) \quad \delta \in \left(\frac{2(l-1)}{2l-1}, \min \left\{\frac{2(l-1)}{l}, 1\right\}\right).$$

$$K(\delta) = \begin{cases} [k_3, k_2] & \text{if } \frac{d}{\beta} \in [\mu, \infty), \\ [k_1, k_2] & \text{if } \frac{d}{\beta} \in [\lambda, \mu), \\ \emptyset & \text{if } \frac{d}{\beta} \in (0, \lambda), \end{cases}$$

where

$$\begin{aligned} k_1 &= s(\hat{a}) + \beta \log l^{-1/2} \left\{ e^{-\lambda/2} - (e^{-\lambda} - e^{-\frac{d}{\beta}})^{1/2} \right\}, \\ k_2 &= s(\hat{a}) + \beta \log l^{-1/2} \left\{ e^{-\lambda/2} + (e^{-\lambda} - e^{-\frac{d}{\beta}})^{1/2} \right\}, \\ k_3 &= s(\hat{a}) + \beta \log \frac{2(1-\delta)(l-1)}{\delta \left\{ e^{\frac{d}{\beta}} - l \right\}}. \end{aligned}$$

In this case, we have (a) $k_1 > s(a_i^0, \hat{a}_{-i}) \Leftrightarrow \frac{d}{\beta} < \mu$, (b) $k_3 < s(a_i^0, \hat{a}_{-i}) \Leftrightarrow \frac{d}{\beta} > \mu$, and (c) $k_2 < s(\hat{a})$.

$$3) \quad \delta \in \left[\min \left\{\frac{2(l-1)}{l}, 1\right\}, 1\right).$$

$$K(\delta) = \begin{cases} [k_3, k_4] & \text{if } \frac{d}{\beta} \in [\nu, \infty), \\ [k_3, k_2] & \text{if } \frac{d}{\beta} \in [\mu, \nu), \\ [k_1, k_2] & \text{if } \frac{d}{\beta} \in [\lambda, \mu), \\ \emptyset & \text{if } \frac{d}{\beta} \in (0, \lambda), \end{cases}$$

where k_1 , k_2 and k_3 are defined as above and

$$k_4 = s(a_i^0, \hat{a}_{-i}) + \beta \log \frac{\delta \left\{ l e^{\frac{d}{\beta}} - 1 \right\}}{2(l-1)}.$$

In this case, we have (a) and (b) above, and (c') $k_2 < s(\hat{a}) \Leftrightarrow \frac{d}{\beta} < \nu$, and (d) $k_4 > s(\hat{a}) \Leftrightarrow \frac{d}{\beta} > \nu$.

	$\beta = 0$	$\beta = 1$	$\beta = 4$	$\beta = 10$
$z_{t-1}\mathbf{1}_{\{c_{t-1}=1\}}$	0.520*** (0.050)	0.253*** (0.033)	0.053*** (0.007)	0.012*** (0.004)
$z_{t-1}\mathbf{1}_{\{c_{t-1}=0\}}$	0.537*** (0.055)	0.245*** (0.036)	0.031*** (0.008)	-0.001 (0.004)
$z_{t-2}\mathbf{1}_{\{c_{t-2}=1\}}$	0.145*** (0.051)	0.014 (0.031)	-0.006 (0.007)	0.003 (0.004)
$z_{t-2}\mathbf{1}_{\{c_{t-2}=0\}}$	0.161*** (0.057)	0.014 (0.035)	-0.006 (0.008)	0.003 (0.004)
Constant	-12.259*** (1.232)	-5.169*** (0.874)	-1.171*** (0.301)	-0.999*** (0.176)
ψ	0.676** (0.309)	1.064*** (0.358)	0.982*** (0.280)	1.163*** (0.269)
ρ	0.153 [§] (0.050)	0.198 [§] (0.064)	0.225 [§] (0.059)	0.169 [§] (0.050)
Observations	1680	928	1240	1042
Number of subjects	38	42	40	46

, * indicate statistical significance at 5%, and 1% respectively.

[§] indicates statistical significance at 1% using a likelihood ratio test.

Table 9: Effects of Past Public Signals and Choices on the Current Choice

Estimation results for (6) Besides those given in (6), we include in the explanatory variables the correlated random effects term ν_i and the constant term. ν_i is assumed to have the normal distribution $N(\psi\zeta_i, \sigma_a)$, where ζ_i is set equal to the fraction of times subject i chooses L in period 1 among all the cycles he plays. In other words, the mean of the random-effects component is allowed to vary linearly with the ratio of choice L in period 1. The coefficient ψ on ζ_i and the variance σ_a are both estimated from data. The parameter estimates are given in Table 9, where $\rho = \frac{\sigma_a^2}{\sigma_a^2 + 1}$ measures the degree of asymmetry across subjects.⁵⁰

Proposition 3 *Let σ be a public (i.e., $b_1 = b_3 = 0$) threshold strategy with regime*

⁵⁰ $\rho = 0$ if there is no individual subject effect and $\rho = 1$ if all the variance is explained by individual subject effects.

shifts. Then (σ, σ) is a perfect public equilibrium if and only if a and b_2 satisfy:

$$\begin{aligned}
& \{F(a - s(L, H)) - F(a - s(L, L))\}(\hat{g} - g^0) \\
& + \left\{F(a + b_2 - s(H, H)) - F(a - s(L, L)) - \frac{1}{\delta}\right\} \{g_i(H, L) - \hat{g}\} \geq 0, \\
& \{F(a + b_2 - s(H, H)) - F(a + b_2 - s(L, H))\}(\hat{g} - g^0) \\
& + \left\{F(a + b_2 - s(H, H)) - F(a - s(L, L)) - \frac{1}{\delta}\right\} \{g^0 - g_i(L, H)\} \leq 0,
\end{aligned} \tag{21}$$

where $\hat{g} = g_i(L, L)$ and $g^0 = g_i(H, H)$.

For our parametrization, a pair of threshold strategies is a perfect public equilibrium if

$$\begin{aligned}
& 3\{F(a - 18) - F(a - 20)\} + F(a + b_2 - 16) - F(a - 20) \geq 0.9^{-1} \\
& 9\{F(a + b_2 - 16) - F(a + b_2 - 18)\} + F(a + b_2 - 16) - F(a - 20) \leq 0.9^{-1}.
\end{aligned}$$

For example, when $a \in [18, 20)$ and $a + b_2 \in \mathbf{R} \setminus [12, 18]$ or when $a = 20$ and $a + b_2 > 18$, there exists $\bar{\beta}(a, b_2) > 0$ such that a pair of public threshold strategies σ with (a, b_2) is a perfect public equilibrium under the noise level $\beta < \bar{\beta}(a, b_2)$.

Proof. For simplicity, write

$$\begin{aligned}
p &= P(z \leq a \mid L, L), & p' &= P(z \leq a + b_2 \mid L, L), \\
q &= P(z \leq a \mid L, H), & q' &= P(z \leq a + b_2 \mid L, H), \\
r &= P(z \leq a \mid H, H), & r' &= P(z \leq a + b_2 \mid H, H).
\end{aligned}$$

Let v_0 and v_1 be the (average) repeated game payoff beginning from states 0 and 1, respectively, when both players play the threshold strategy σ . Then we have the following recursive equations:

$$\begin{aligned}
v_0 &= (1 - \delta) \hat{g} + \delta\{(1 - p)v_0 + pv_1\}, \\
v_1 &= (1 - \delta) g^0 + \delta\{(1 - r')v_0 + r'v_1\}.
\end{aligned} \tag{22}$$

On the other hand, the incentive compatibility conditions can be written as:

$$\begin{aligned}
v_0 &\geq (1 - \delta) g_i(H, L) + \delta\{(1 - q)v_0 + qv_1\}, \\
v_1 &\geq (1 - \delta) g_i(L, H) + \delta\{(1 - q')v_0 + q'v_1\}.
\end{aligned} \tag{23}$$

Solving (22) for v_0 and v_1 , we obtain

$$\begin{bmatrix} v_0 \\ v_1 \end{bmatrix} = (1 - \delta) \begin{bmatrix} 1 - \delta(1 - p) & -\delta p \\ -\delta(1 - r') & 1 - \delta r' \end{bmatrix}^{-1} \begin{bmatrix} \hat{g} \\ g^0 \end{bmatrix}. \tag{24}$$

Upon substituting the right-hand side of (22) into the left-hand side of (23), we also obtain the following inequalities in matrix form:

$$\delta \begin{bmatrix} q - p & -q + p \\ -r' + q' & r' - q' \end{bmatrix} \begin{bmatrix} v_0 \\ v_1 \end{bmatrix} \geq (1 - \delta) \begin{bmatrix} g_i(H, L) - \hat{g} \\ g_i(L, H) - g^0 \end{bmatrix}. \quad (25)$$

Substituting (24) into (25) and simplifying, we obtain (21). ■

Non-parametric tests of the (S₀) strategy. First, we test the hypothesis that $b_1 = b_3 = 0$ in the perfect monitoring treatment ($\beta = 0$) using the sign test (Snedecor and Cochran 1980), which requires no parametric assumption. Suppose that a player uses a threshold strategy with $b_1 = 0$ in the perfect monitoring game. Then the condition of transition from state 0 to state 1 should be neutral with respect to the identity of the deviator. In other words, if the opponent's deviation in state 0 moves the player to state 1, then so does his own deviation. This can be tested as follows. Take any subject i and consider the following two sequences of play: In the first sequence, both i and his opponent j play L in period $t - 2$, and i plays L and j plays H in period $t - 1$. In the second sequence, both play L in $t - 2$ and i plays H and j plays L in $t - 1$. When $b_1 = 0$, i 's action in period t should be the same conditional on either sequence. We compare the rate that subject i plays action H after the first sequence with that after the second sequence using a sign test. The null hypothesis that they are the same cannot be rejected (p-value = 0.51): 6 subjects play H more often when they played L in period $t - 1$, 3 less often and 1 exactly the same number of times. We can also test if $b_3 = 0$ by comparing the rate that subject i plays L after the sequence $((H, H), (H, L))$ with that after the sequence $((H, H), (L, H))$. Again, the null hypothesis that they are the same cannot be rejected (p-value = 0.45): 2 subjects played L more often when they played L in $t - 1$, 5 less often, and 1 the same number of times. These results support the findings from the likelihood ratio test. In particular they imply that the specifications (Sb)-(Sd), which all have $(b_1, b_3) \neq (0, 0)$, are unlikely.

We next check how well the deterministic specification of (S₀) (*i.e.*, $b_1 = b_2 = b_3 = 0$ in (8) and (7)) describes the data. Specifically, pick any subject, fix the base threshold level a , and consider the sequence of actions this strategy would generate given the sequence of the public signals and his own actions. For each value of a , we compare the actions thus generated against the actions actually chosen by the subject, and count the number of periods in which the former matches the latter. We then choose a so as to maximize the hit rate, *i.e.*, the ratio of periods for which

the two action choices coincide. It is 93% for the median player and 88% on average in the $\beta = 0$ treatment.⁵¹ Likewise, the median and average hit rates are 80% and 82% in the $\beta = 1$ treatment, 77% and 77% in the $\beta = 4$ treatment, and 67% and 71% in the $\beta = 10$ treatment. All these numbers are statistically different from 50% (a coin toss) at the 1% level.

⁵¹In other words, when the subjects are ranked by their hit rates, the hit rate of the median subject is 93%.

References

- Abramowitz, M. and Stegun, I. A. (eds): 1972, *Handbook of mathematical functions with formulas, graphs, and mathematical tables*, U.S. Government Printing Office, Washington, D.C.
- Abreu, D., Milgrom, P. and Pearce, D.: 1991, Information and timing in repeated partnerships, *Econometrica* **59**, 1713–1733.
- Abreu, D., Pearce, D. and Stacchetti, E.: 1990, Toward a theory of discounted repeated games with imperfect monitoring, *Econometrica* **58**, 1041–1063.
- Aoyagi, M. and Fréchette, G. R.: 2005, Collusion as public monitoring becomes noisy: Experimental evidence. http://homepages.nyu.edu/~gf35/print/Aoyagi_2005a.pdf.
- Bó, P. D.: 2003, Cooperation under the shadow of the future: experimental evidence from infinitely repeated games. mimeo.
- Bó, P. D. and Fréchette, G. R.: 2007, The evolution of cooperation in infinitely repeated games: Experimental evidence. mimeo.
- Bereby-Meyer, Y. and Roth, A. E.: 2006, Learning in noisy games: Partial reinforcement and the sustainability of cooperation, *American Economic Review* **96**(4), 1029–1042.
- Cason, T. N. and Khan, F. U.: 1999, A laboratory study of voluntary public goods provision with imperfect monitoring and communication, *Journal of Development Economics* **58**, 533–552.
- Chamberlain, G.: 1980, Analysis of covariance with qualitative data, *Review of Economic Studies* **47**, 225–238.
- Cosslett, S. R. and Lee, L.-F.: 1985, Serial correlation in latent discrete variable models, *Journal of Econometrics* **27**, 79–97.
- Duffy, J. and Ochs, J.: 2007, Cooperative behavior and the frequency of social interaction. mimeo.
- Ellison, G.: 1994, Theories of cartel stability and the joint executive committee, *Rand Journal of Economics* **25**, 37–57.

- Engle-Warnick, J., McCausland, W. J. and Miller, J. H.: 2004, The ghost in the machine: Inferring machine-based strategies from observed behavior. mimeo.
- Engle-Warnick, J. and Slonim, R. L.: 2006, Inferring repeated-game strategies from actions: Evidence from repeated trust game experiments, *Economic Theory* **28**(3), 603–632.
- Feinberg, R. M. and Husted, T. A.: 1993, An experimental test of discount-rate effects on collusive behaviour in duopoly markets, *Journal of Industrial Economics* **41**(2), 153–160.
- Feinberg, R. M. and Snyder, C.: 2002, Collusion with secret price cuts: An experimental investigation, *Economics Bulletin* **3**(6), 1–11.
- Green, E. J. and Porter, R. H.: 1984, Noncooperative collusion under imperfect price information, *Econometrica* **52**, 87–100.
- Gutierrez, R. G., Carter, S. and Drukker, D. M.: 2001, sg160: On boundary-value likelihood-ratio tests, *Stata Technical Bulletin Reprints* **10**, 269–273.
- Heckman, J. J.: 1981, *Structural Analysis of Discrete Data with Econometric Applications*, MIT Press, Cambridge, MA, chapter The Incidental Parameters Problem and the Problem of Initial Conditions in Estimating a Discrete Time-Discrete Data Stochastic Process, pp. 179–195.
- Holcomb, J. H. and Nelson, P. S.: 1997, The role of monitoring in duopoly market outcomes, *Journal of Socio-Economics* **26**(1), 79–93.
- Holt, C. A.: 1995, *Handbook of Experimental Economics*, Princeton University Press, Princeton, chapter Industrial Organization: A Survey of Laboratory Research, pp. 349–444.
- Kandori, M.: 1992a, Social norms and community enforcement, *Review of Economic Studies* **59**, 63–80.
- Kandori, M.: 1992b, The use of information in repeated games with imperfect monitoring, *Review of Economic Studies* **59**, 581–593.
- Karlin, S.: 1968, *Total Positivity*, Stanford University Press, Stanford.

- Lee, L.-F. and Porter, R. H.: 1984, Switching regression models with imperfect sample separation information – with an application to cartel stability, *Econometrica* **52**, 391–418.
- Mason, C. F. and Phillips, O. R.: 2002, In support of trigger strategies: Experimental evidence from two-person noncooperative games, *Journal of Economics and Management Strategy* **11**(4), 685–716.
- Porter, R. H.: 1983, Optimal cartel trigger price strategies, *Journal of Economic Theory* **29**, 313–338.
- Porter, R. H.: 1985, On the incidence and duration of price wars, *Journal of Industrial Economics* **33**, 415–426.
- Rapoport, A. and Chammah, A. M.: 1965, *Prisoner's Dilemma: A Study in Conflict and Cooperation*, The University of Michigan Press, Ann Arbor.
- Roth, A. E.: 1995. “Bargaining Experiments” in *Handbook of Experimental Economics*, ed. by John H. Kagel and Alvin E. Roth. Princeton: Princeton University Press, 253–348.
- Roth, A. E. and Murnighan, J. K.: 1978, Equilibrium behavior and repeated play of the prisoner's dilemma, *Journal of Mathematical Psychology* **17**, 189–198.
- Sannikov, Y.: 2007, Games with imperfectly observed actions in continuous time, *Econometrica* **75**, 1285–1329.
- Sell, J. and Wilson, R. K.: 1991, Trigger strategies in repeated-play public goods games: Forgiving, non-forgiving or non-existent. Working Paper, Rice University.
- Snedecor, G. W. and Cochran, W. G.: 1980, *Statistical Methods*, Iowa State University Press, Ames, Iowa.
- Wooldridge, J. M.: 2002, *Econometric Analysis of Cross Section and Panel Data*, The MIT Press, Cambridge.