Collusion as Public Monitoring Becomes Noisy: Experimental Evidence*

Masaki Aoyagi† Guillaume R. Fréchette‡
Osaka University New York University

October 13, 2005

Abstract

This paper studies collusion in an infinitely repeated game when the opponent's past actions are observed only through a noisy public signal. Attention is focused on a *threshold strategy,* which switches between cooperation and punishment phases based on the comparison between the realized public signal and a threshold. The paper develops a simple theory on when such a strategy supports the most efficient symmetric (perfect public) equilibrium, and characterizes its payoff as a function of noise in monitoring. The theoretical predictions are then tested in laboratory experiments. It is found that subjects' payoffs (i) decrease as noise increases, and (ii) are lower than the theoretical maximum for small noise, but exceed it for large noise. It is also estimated that the subjects' strategies are best described by a simple threshold strategy that looks only at the most recent public signal.

Keywords: Repeated games, collusion, cooperation.

*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, New York University, Université de Montréal, Carnegie Mellon University, RIETI, Rutgers, and the Stockholm School of Economics, as well as participants of the 2003 International Meeting of the Economic Science Association, and the 14th Summer Festival on Game Theory at Stony Brook. Alex Brown provided valuable research assistance. This research was supported in part by the Harvard Business School, by the Center for Experimental Social Science and the C.V. Starr Center. We are responsible for all remaining errors.

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

‡Department of Economics, 269 Mercer St. 7th Floor, New York, NY, 10003-6687.

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 any payoff in an important class of symmetric equilibria can be achieved through the use of 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 theoretical characterization of the maximal equilibrium payoff as a function of noise in monitoring, and then uses laboratory experiments to test the theoretical 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.

As in Green and Porter (1984), the present model has a one-dimensional public signal (price) whose distribution is monotonically related to the players' actions (quantities). While the set of repeated game strategies is large, we suppose that players coordinate their actions based on a simple rule. In particular, we suppose that players use a *threshold strategy,* which uses thresholds on the public signal 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.

Specifically, we consider a symmetric stage game with two actions for each player: As in the PD, the first action gives rise to an efficient symmetric profile, while the second action gives rise to an inefficient symmetric (one-shot) Nash equilibrium. We characterize the set of equilibrium payoffs when the two players' repeated game strategies are symmetric, sequentially rational, and public in the sense that today's actions are determined only by past public signals. In this case, it is known from the bang-bang property (Abreu *et al.* (1990)) that the highest equilibrium payoff is sustained by a grim-trigger strategy. Based on this fact, we identify sufficient conditions for the optimal grim-trigger strategy to trigger the punishment based on a threshold condition. We then proceed to 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 a simple two-person two-action PD as the stage game, and assume a simple distribution for the noise component of the public signal. The distribution allows for an explicit derivation of the maximum payoff as a function of noise in monitoring. We have five treatments depending on the level of noise from zero to infinity. 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:

1) The level of cooperation is positive for any noise level.

2) The level of cooperation is lower than the theoretical maximum for small noise

[1] This should be contrasted with the qualitative conclusion of Kandori (1992) that the set of (symmetric and asymmetric) perfect public equilibrium payoffs strictly expands as monitoring becomes more accurate. We are unaware of any quantitative characterization of equilibrium payoffs as a function of noise in monitoring.

but higher than that for large noise.

3) The level of cooperation decreases as noise increases.

4) For large noise for which theory predicts no cooperation, the level of cooperation is no higher than that in a one-shot game.

We also find a distinctive pattern in the evolution of play. In the low noise treatments for which cooperation is possible in theory, 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 their subjects cooperate even with no histories, these experiments may be interpreted as suggesting that past histories would play an insignificant role in their decisions in repeated games. 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 identify the strategy that best describes the data from a class of strategies that use thresholds. Specifically, we consider a class of strategies which start out in the cooperation phase, switch to the punishment phase when the public signal falls below a certain threshold, and return to the cooperation phase after some number of periods provided that the signal then is above another threshold. It can be seen that this class encompasses those strategies that are most frequently discussed in the literature: Included in this class are the trigger strategies as in Green and Porter

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

(1984) as well as the tit-for-tat strategy. In all but one noise treatment, we find that the subjects use only the most recent public signal in determining today's action: In other words, they choose the cooperative action today if the public signal yesterday is above a certain threshold, and choose the non-cooperative action otherwise.

It should be noted that our laboratory experiments are designed in strict accordance with the underlying assumptions of the tested theory. 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. 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.[3] 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.

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 game with imperfect public monitoring and Section 4 proves the optimality of a threshold strategy and characterizes its payoff. The experimental design is described in Section 5. Section 6 reports the results of our experiments: We first test the theoretical prediction on the players' payoffs and then estimate their strategies. We conclude in Section 7 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 cooperation and punishment phases in the repeated game.[4] 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

[3] See Section 6.

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

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.[5] 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.[6] 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.[7]

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 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

[5] See Holt (1995) for a literature review.

[6] See, for example, Sell and Wilson (1991), Feinberg and Husted (1993), and Engle-Warnick and Slonim (2002).

[7] See 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 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.[8]

In contrast to the above models, formulation of imperfect monitoring in this paper strictly matches that of the standard theory. 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 (2003) 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.[9]

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 (2001) 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 (2002) 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

[8] Cason 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.

[9] A 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.

strategy.[10]

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 Model

The set I of n players play a symmetric game infinitely often.[11] Player i's action a_i in the stage-game is chosen from the set $A_i = \{a_i^0, \hat{a}_i\}$. Let $A = A_1 \times \cdots \times A_n$ denote the set of action profiles $a = (a_1, \ldots, a_n)$ of all players. After each period, players do not observe each other's actions but observe a random public signal $z \in \boldsymbol{R}$ whose probability distribution is determined by the action profile $a \in A$ played in that period. Denote by $h(z \mid a)$ the density of the public signal z under the action profile a, and by $H(z \mid a)$ the corresponding cumulative distribution. In the Cournot model with stochastic demand, for example, a_i and z correspond to firm i's output quantity and the realized price level, respectively.

Player i's stage-payoff when the signal realization is z is given by $w_i(a_i, z)$. It should be noted that i's payoff depends on other players' actions only through the public signal z, and hence does not provide more information than z itself.

Given the action profile a, let $g_i(a)$ be player i's expected stage-payoff:

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

We assume that the stage-game $(A, (g_1, \ldots, g_n))$ is a PD:

$$g_i(a_i^0, \hat{a}_{-i}) > g_i(\hat{a}) > g_i(a^0) \geq g_i(\hat{a}_i, a_{-i}^0).$$

In other words, $\hat{a} = (\hat{a}_1, \ldots, \hat{a}_n)$ is the efficient symmetric profile, and $a^0 = (a_1^0, \ldots, a_n^0)$ is a one-shot Nash equilibrium as well as a mutual minmax profile. Furthermore, a_i^0 is a profitable one-shot deviation from the efficient profile $\hat{a}$. Denote the one-shot Nash equilibrium payoff by $g^0 = g_i(a^0)$ and the efficient payoff by $\hat{g} = g_i(\hat{a})$.

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 history of signals z and i's actions in periods

[10] It should be noted that this estimation technique is not practical for games with a long expected horizon. In Engle-Warnick and Slonim (2002), the continuation probability is such that the expected length of the repeated game is five periods. It is ten in our case.

[11] While our experiment is on two-person games, the theory is presented for n-person games as this entails little extra cost.

1 through t. The set of t-length public histories is given by $\boldsymbol{R}^t$, while the set of t-length private histories of player i is given by $\boldsymbol{R}^t \times A_i^t$. Player i's (pure) *strategy* is a function $\sigma_i : \bigcup_{t=0}^{\infty} (\boldsymbol{R}^t \times A_i^t) \to A_i$. It is a *public strategy* if it is a function of the public history alone.

Let $\delta < 1$ denote the players' common discount factor. Given a strategy profile σ, the expected payoff to player i 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, $\pi_i(\sigma) \geq \pi_i(\sigma_i', \sigma_{-i})$ for any strategy σ_i'. An equilibrium $\sigma = (\sigma_1, \ldots, \sigma_n)$ 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_{-i} for each i after every public history, and is (strongly) *symmetric* if $\sigma_1 = \cdots = \sigma_n$.

4 Symmetric Equilibrium Payoffs

Our objective is to identify the set of symmetric perfect public equilibrium payoffs. As shown by Abreu *et al.* (1990), this can be accomplished by examining the following class of "grim-trigger" equilibria. For this, suppose that player i plays a public strategy $\hat{\sigma}_i$ that 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 $\boldsymbol{R}$ but reverts to the minimax action a_i^0 forever otherwise. Let $\hat{\sigma} = (\hat{\sigma}_1, \ldots, \hat{\sigma}_n)$ be the symmetric profile of such a strategy, and v be the expected payoff associated with $\hat{\sigma}$. By the stationarity of play, v must satisfy

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

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 in set Q under the action profile $\hat{a}$. Solving (1) for v, we get

$$v = \frac{(1-\delta)\hat{g} + \delta g^0\, P(z \notin Q \mid \hat{a})}{1 - \delta P(z \in Q \mid \hat{a})}. \tag{2}$$

It is clear from (2) that the larger the value of $P(z \in Q \mid \hat{a})$, the closer to $\hat{g}$ is the payoff v. 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

$$v \geq (1-\delta) g_i(a_i, \hat{a}_{-i}) + \delta \left\{ v\, P(z \in Q \mid a_i, \hat{a}_{-i}) + g^0\, P(z \notin Q \mid a_i, \hat{a}_{-i}) \right\}. \tag{3}$$

Solving (3) for v, we get

$$v \geq \frac{(1-\delta) g_i(a_i, \hat{a}_{-i}) + \delta g^0\, P(z \notin Q \mid a_i, \hat{a}_{-i})}{1 - \delta P(z \in Q \mid a_i, \hat{a}_{-i})}. \tag{4}$$

Eliminating v from (2) and (4), we obtain

$$\frac{1 - \delta P(z \in Q \mid a_i, \hat{a}_{-i})}{1 - \delta P(z \in Q \mid \hat{a})} \geq \frac{g_i(a_i, \hat{a}_{-i}) - g^0}{\hat{g} - g^0}. \tag{5}$$

We make the following assumptions about the distribution of the public signal:

Assumption 1 *The unilateral deviation a_i^0 from the symmetric action profile $\hat{a}$ shifts down the distribution of z as measured in the likelihood ratio:*

$$\frac{h(z' \mid \hat{a})}{h(z \mid \hat{a})} \geq \frac{h(z' \mid a_i, \hat{a}_{-i})}{h(z \mid a_i, \hat{a}_{-i})} \quad \textit{for any } z' \geq z.$$

In a Cournot model, for example, Assumption 1 corresponds to assuming that a profitable increase in production lowers the distribution of the price. As seen, Assumption 1 requires that $H(\cdot \mid \hat{a})$ first-order stochastically dominate $H(\cdot \mid a_i^0, \hat{a}_{-i})$. The following theorem shows that under Assumption 1, the most efficient symmetric perfect public equilibrium payoff is replicated by a *threshold grim-trigger strategy,* which reverts to the punishment if and only if the public signal falls below a certain threshold.

Proposition 1 *Suppose that Assumption 1 holds and let v be the maximal symmetric perfect public equilibrium payoff of the repeated game. If $v > g^0$, then there exists a (pure) grim-trigger strategy profile $\hat{\sigma}$ such that $v = \pi_i(\hat{\sigma})$. Furthermore, $\hat{\sigma}_i$ can be taken to be a stationary threshold grim-trigger strategy that begins with $\hat{a}_i$ and continues with $\hat{a}_i$ as long as the realized public signal z is above some threshold k (i.e., $Q = (k, \infty)$), but reverts to a_i^0 otherwise.*

Proof. See the Appendix. ■

Let $s : A \to \boldsymbol{R}$ be a (deterministic) symmetric function of the action profile a, and x be a real-valued random variable whose distribution is independent of a. For the rest of this paper, we assume that the public signal takes the following additive form:

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

In line with the assumption that z has full support, assume that x has a strictly positive density f over $\boldsymbol{R}$. Denote the corresponding cumulative distribution by F. It follows that h and f are related through

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

It is not difficult to verify that Assumption 1 is implied by the following set of assumptions on s and f:

Assumption 2 $s(a_i^0, \hat{a}_{-i}) \leq s(\hat{a})$.

Assumption 3 $f(x-y)\, f(x'-y') \geq f(x-y')\, f(x'-y)$ *for any* $x \geq x'$ *and* $y \geq y'$.[12]

It should be noted that Assumption 3 holds for a wide class of distributions including normal and gamma distributions. We now turn to the characterization of the set of equilibrium payoffs in this setup. Proposition 1 allows us to focus on threshold grim-trigger strategies. Define

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

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

Let $\hat{\sigma} = (\hat{\sigma}_1, \ldots, \hat{\sigma}_n)$ denote the threshold grim-trigger strategy profile that begins with $\hat{a}$ and reverts to a^0 when z falls below the threshold k as specified in Proposition 1. We can rewrite (2) and (5) respectively as:

$$v = \frac{(1-\delta)\hat{g} + \delta g^0\, F(k - s(\hat{a}))}{1 - \delta\,\{1 - F(k - s(\hat{a}))\}}, \tag{6}$$

and

$$\frac{1 - \delta\,\{1 - F(k - s(\hat{a}) + d)\}}{1 - \delta\,\{1 - F(k - s(\hat{a}))\}} \geq l. \tag{7}$$

[12] The function f satisfying this condition is known as a *Polya function of degree 2* (PF_2) (Karlin, 1968). Note that f is PF_2 if and only if the function $\hat{f} : \boldsymbol{R}^2 \to \boldsymbol{R}_+$ defined by $\hat{f}(x,y) = \log\, f(x-y)$ is supermodular.

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

$$K(\delta) = \{k \in \boldsymbol{R} : \ k \text{ satisfies } (7)\}.$$

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 $K(\delta) \neq \emptyset$.[13] By (6), the optimal threshold $k = k^*(\delta)$ that maximizes v is the smallest element of $K(\delta)$:

$$k^*(\delta) = \min\ K(\delta). \tag{8}$$

It is also clear from (7) that $K(\delta) \neq \emptyset$ requires $d > 0$. The following proposition summarizes our observation.

Proposition 2 *Suppose that Assumptions 2 and 3 hold. There exists a grim-trigger equilibrium $\hat{\sigma}$ that plays the symmetric action profile $\hat{a}$ in the cooperation phase if and only if $K(\delta) \neq \emptyset$. In this case, the threshold grim-trigger equilibrium that plays $\hat{a}$ in the cooperation phase and uses the optimal threshold $k^*(\delta)$ yields the payoff*

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

Otherwise, $v^*(\delta) = g^0$.

As an index that does not depend on specific payoff numbers, we introduce the following normalization of $v^*(\delta)$:

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

It can be verified that $y^*(\delta) \in [0, 1]$ is unaffected by a positive affine transformation of the payoff numbers.[14] By Proposition 2, $y^*(\delta)$ can be explicitly written as

$$y^*(\delta) = \frac{1-\delta}{1 - \delta + \delta F(k^*(\delta) - s(\hat{a}))} \tag{9}$$

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

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

[14] That is, $y^*(\delta)$ stays the same when we add, subtract, or multiply a positive constant to the players' stage-payoffs.

5 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.[15] 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 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:

$$f(x) = \frac{1}{2\beta} e^{-\frac{|x|}{\beta}}, \text{ and } F(x) = \begin{cases} 1 - \frac{1}{2} e^{-\frac{x}{\beta}} & \text{if } x \geq 0 \\ \frac{1}{2} e^{\frac{x}{\beta}} & \text{if } x < 0, \end{cases} \tag{10}$$

where $\beta > 0$.[16] As can be readily verified, f satisfies Assumption 3. Moreover, the parameter β represents the level of noise in the public signal since

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

From (9), 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.

[15] Our 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.

[16] This distribution is simple enough to allow for the analytical derivation of the optimal threshold $k^*(\delta)$. See the Appendix.

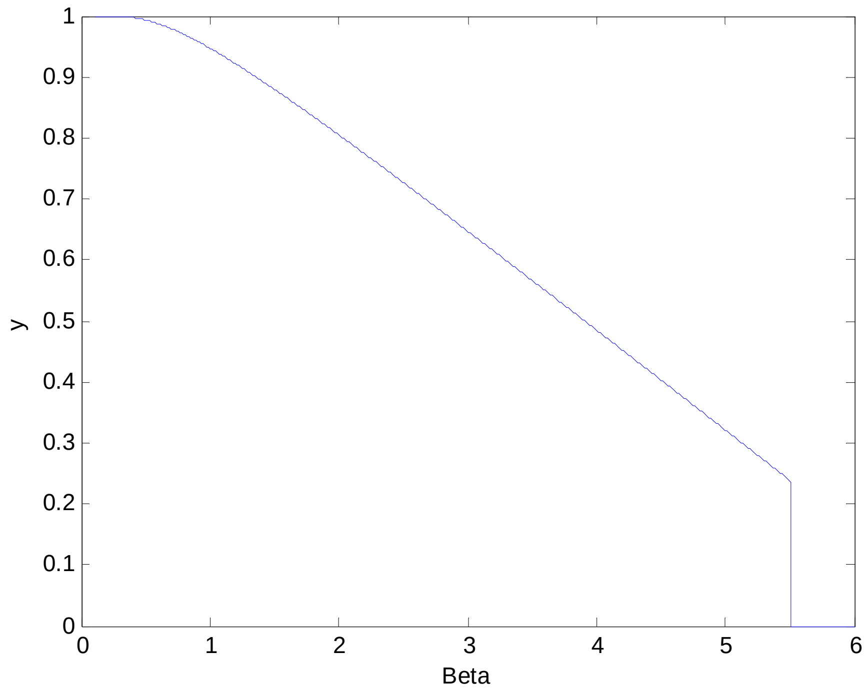

Figure 1: y^* as a function of β ($\delta = 0.9$)

In the actual treatment, we need to (i) bound the subjects' payoffs while allowing z to have full support as assumed by the theory, and (ii) end each session in finite time. For (i), 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.[17] For (ii), we interpret the discount factor as a continuation probability and terminate the game after each period with a fixed probability. We set the termination probability equal to 0.1 so that the effective discount factor $\delta = 0.9$. Note that the theory in Section 4 holds precisely under these alternative specifications.

The experiments proceed in the following steps. First, subjects are provided

[17] In 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 (2004). 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 involved presenting a complex function in the instructions to the subjects.

with the basic information about the game they will play. They are then matched 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.[18] 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 (10).[19] 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.[20] 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 a control treatment where subjects are anonymously and randomly matched after every period of play. This control hence eliminates the repeated game effect (*i.e.,* $\delta = 0$), and is designed to provide a benchmark for the first four treatments with respect to the level of cooperation. From the theoretical perspective, the one-shot environment in this treatment is equivalent to having infinitely large noise.[21] For this reason, we refer to this last treatment as $\beta = \infty$. The length of those sessions is set equal to 75 periods, which is approximately the average number of periods in the other treatments. For the control treatment, the term *cycle* refers to the block of initial 15 periods and every successive block of 10 periods.

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

[18] All games played in the same session terminate simultaneously.

[19] The instructions given to the subjects can be found at http://homepages.nyu.edu/ gf35/print/afinstructions.pdf.

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

[21] In actual implementation of the control treatment, we let the subjects observe their actions. With random and anonymous rematching after every period, however, the level of monitoring should be irrelevant.

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

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 \$20.35 to \$36.55. Details about the number of subjects and periods in each treatment are provided in Table 1.[22]

6 Results

6.1 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.[23]

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 two treatments ($\beta = 10$, ∞) 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

[22] The $\beta = 0$ session with 14 subjects had a crash after the end of cycle 8, and was re-started for two additional cycles. The $\beta = \infty$ session with 16 subjects had 1 crash.

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

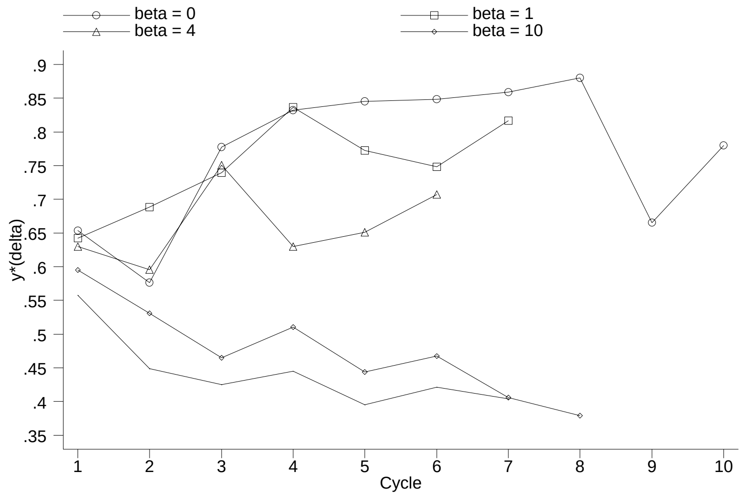

Figure 2: Evolution of $\bar{y}$ by treatments

and by cycle. (The unlabeled line corresponds to the $\beta = \infty$ treatment.) This figure shows that $\bar{y}$ has an upward trend over time for $\beta = 0$, 1, and 4, and an opposite, downward trend for $\beta = 10$ and ∞. 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:

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

indicating that $\bar{y}$ has the same relative ordering as y^*. 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.[24]

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

[24] Note 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
		(0.021)
$\beta = 1$	0.948	0.774
		(0.025)
$\beta = 4$	0.486	0.695
		(0.033)
$\beta = 10$	0	0.467
		(0.027)
$\beta = \infty$	0	0.418
		(0.017)

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 $\bar{y}(\infty)$ appear comparable. 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.[25] Every test is statistically significant at the 5% level. On the other hand, the hypothesis that $\bar{y}(10) = \bar{y}(\infty)$ cannot be rejected ($p = 0.108$, two-tailed Mann-Whitney test).[26] The results of these tests support

[25] For 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 accross cycles for a given individual.

[26] Both ANOVA and Kruscal-Wallis tests reject at the 1% level the null hypothesis that the

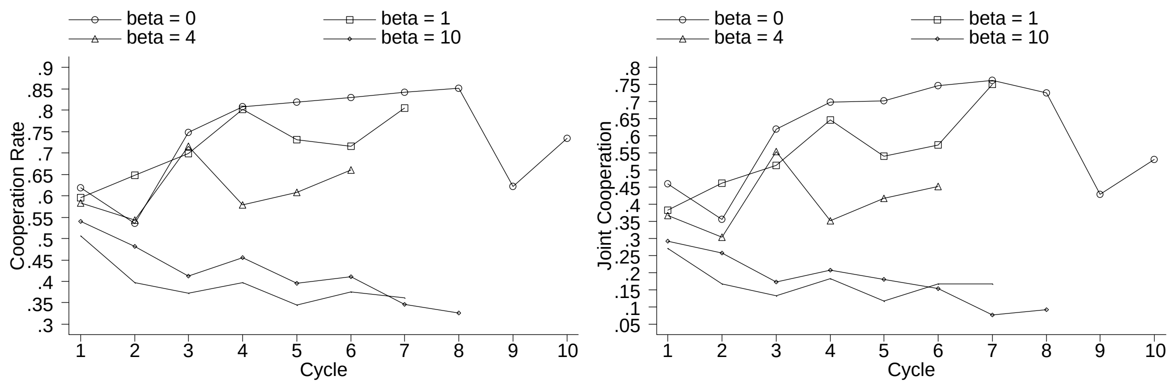

Figure 3: Evolution of cooperation: Rates of L and (L, L) by cycles

the general theoretical predictions that cooperation is easier to sustain when noise is small, and that cooperation under sufficiently large noise is as difficult as in the one-shot environment.[27] On the other hand, t-tests reject the hypothesis that $\bar{y}$ is equal to y^* at the 1% level for each treatment. In other words, the subjects' play does not conform to the most efficient symmetric equilibrium.

In contradiction to the theory, both $\bar{y}(10)$ and $\bar{y}(\infty)$ are significantly positive. This is similar to 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 (2003) 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

treatment has no effect on y for $\beta = 0$, 1, 4, and 10.

[27]This 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 $\beta = 10$ and $\beta = \infty$. The one-sided null hypothesis is rejected with a p-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.

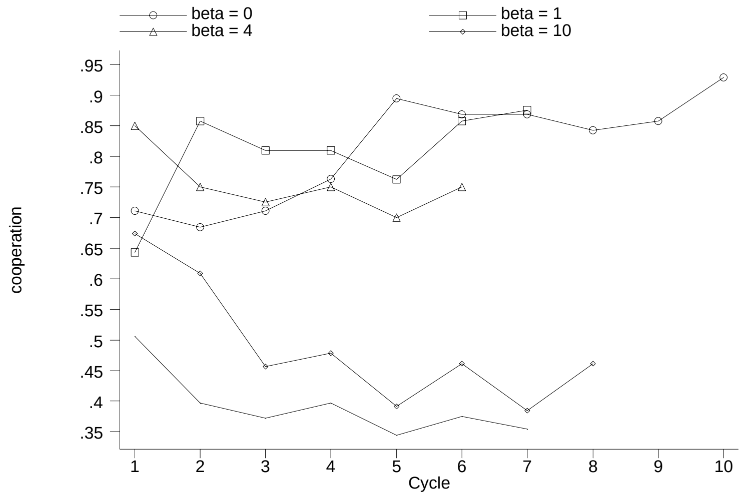

Figure 4: Rates of the cooperative choice L in period 1 by cycles

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

treatments as well.[28]

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 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

[28]According 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 (2003) and Dal Bo (2003).

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.[29] 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, t-tests reject at any conventional level the hypothesis that the rates of period 1 cooperation equal unity.

6.2 Strategies

We next turn to the analysis of the strategies. As mentioned in the Introduction, we focus on "threshold strategies" that switch between cooperation and punishment phases based on thresholds on the public signal. This class includes variations of trigger and tit-for-tat strategies that are most often discussed in the experimental analysis of repeated game strategies, as well as some form of "private" strategies that choose actions based on one's own actions in the past. 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 best fitting strategy is one that uses only the most recent public signal for transition between the phases. The analysis in this subsection excludes data from the control treatment.

Note first that any strategy that supports cooperation above the one-shot NE level 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. Figures 5, 6, and 7 shows 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. In the $\beta = 1$ and $\beta = 4$ treatments, the rate of cooperation is

[29] Neither ANOVA nor Kruscal-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).

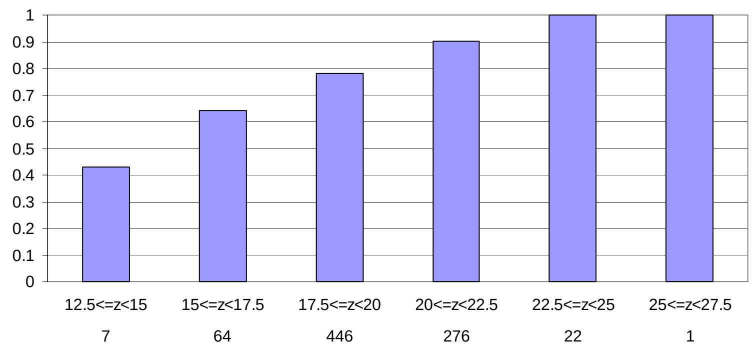

Figure 5: Rates of L as a function of the most recent public signal ($\beta = 1$), [No. of Obs. Under the Category]

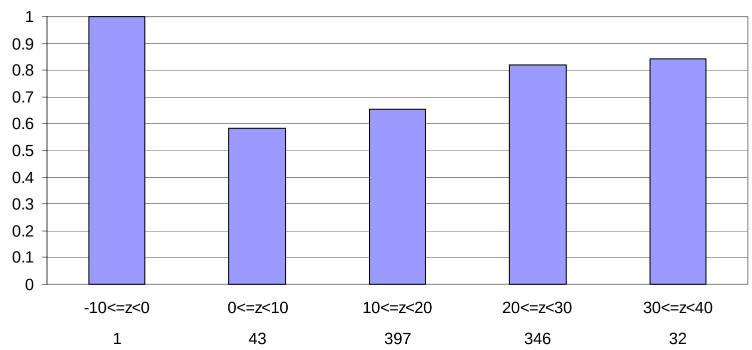

Figure 6: Rates of L as a function of the most recent public signal ($\beta = 4$), [No. of Obs. Under the Category]

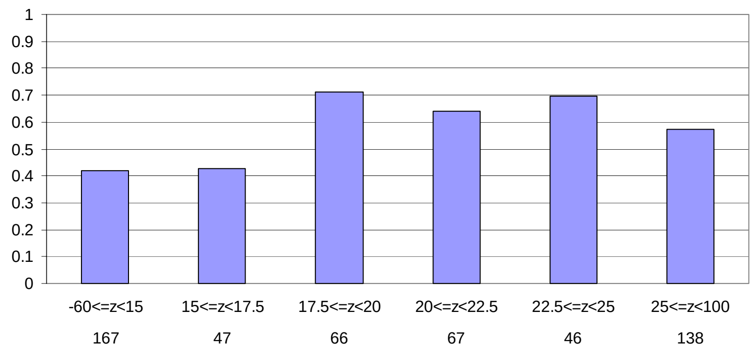

Figure 7: Rates of L as a function of the most recent public signal ($\beta = 10$), [No. of Obs. Under the Category]

clearly an increasing function of z.[30] Even in the $\beta = 10$ treatment where the theory predicts no cooperation, we see the cooperation level increase around $z = 17.5$.

6.2.1 Threshold strategies

While the above relationship between z and the rate of cooperation can be a consequence of many different strategies, we suppose that the subjects' behavior follows some simple rules. Specifically, we suppose that it is expressed by a *finite automaton,* which consists of a finite number of states as well as behavior and transition rules as follows: In the finite automaton representation of a strategy, a player is in one of the states in every period. The behavior rule determines the action he should choose today as a function of the current state, and the transition rule determines tomorrow's state as a function of the current state as well as the public signal and his own action choice today.[31] A *threshold strategy* with $T+1$ states ($T = 1, 2, \ldots$) is a finite automaton with $T+1$ states with the behavior and transition rules specified as follows. The behavior rule specifies action L (cooperation) in state 0, which corresponds to the cooperation phase, and action H (non-cooperation) in states $1, \ldots, T$, which correspond to the punishment phase. The transition rule is described as fol-

[30] In Figure 6, this observation holds with the exception of the first column, which has only one observation.

[31] In a usual specification of a finite automaton, transition is assumed independent of one's own action. We have chosen a more general specification here in order to explicitly consider a possible adjustment a player may make after his own mistake.

lows: The initial state is state 0, and the state in the next period is either state 0 or state 1 depending on the public signal z *and* the player's own action a_i in the current period.[32] When in state t ($t = 1, \ldots, T-1$), unconditional transition to state $t+1$ takes place in the next period. When in state T, the state in the next period is either state T again or state 0 depending on the public signal z and the player's own action a_i. The transition at states 0 and T can be described in more detail as follows:

State 0: Stay in state 0 if the public signal z is above threshold a when the own action is L, or if z is above another threshold $a + b_1$ when the own action is H. Move from state 0 to state 1 otherwise.

State T: Move from state T to state 0 if the public signal $z > a + b_2 + b_3$ when the own action is L, or if $z > a + b_2$ when the own action is H. Stay in state T otherwise.

In Figure 8, for example, the number above each arrow indicates the threshold condition to be satisfied for the corresponding transition when the present action is L, and the number below each arrow indicates the same when the present action is H. We allow each threshold to be $+\infty$ or $-\infty$. Note that the threshold grim-trigger strategy discussed in Section 4 is a threshold strategy with $T = 1$, $b_2 = \infty$ and $b_1 = b_3 = 0$.[33] A trigger strategy with a fixed punishment length T as discussed by Green and Porter (1984) is also a threshold strategy with $b_2 = -\infty$ and $b_1 = b_3 = 0$. Furthermore, the tit-for-tat strategy in the perfect monitoring environment is obtained by setting $T = 1$, $a \in (s(L,H), s(L,L)] = (18, 20]$, $b_1 = b_3 = -2$, and $b_2 = 0$. Note that a threshold strategy with $b_2 > -\infty$ has the punishment phase possibly longer than T periods. In general, a threshold strategy is public in the sense defined in Section 3 if $b_1 = b_3 = 0$. Otherwise, a threshold strategy is a private strategy which conditions its action choice on one's own private history. 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 with perfection after every public history, then the associated payoff is still bounded from above by v^* (or its normalization y^*) described in Section 4.

[32] Of course, when the player follows the behavior rule, then his action in state 0 is L. We let the transition depend on the own action choice in order to explicitly account for the possiblity that a player chooses an action different from that specified by the behavior rule by mistake.

[33] When the game is of finite length, an equivalent strategy is obtained by letting $b_1 = 0$ and T larger than the length of the game.

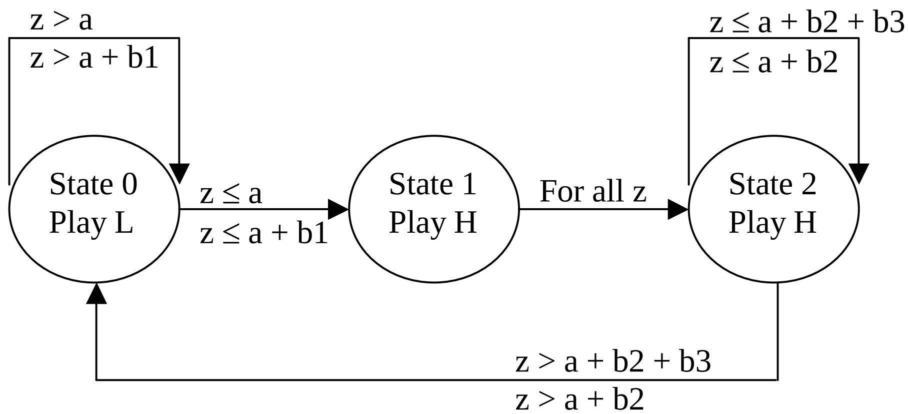

Figure 8: General threshold strategy with $T = 2$

For our discussion, we find it convenient to express the action choice by a threshold strategy without explicit reference to the states. Identify 1 with the cooperative action L and 0 with the non-cooperative action H, and for any sequences of public signals $z_1, z_2, \ldots$ and own actions $c_{i1}, c_{i2}, \ldots$, let

$$s_{it} = 1 \text{ for } t \leq 1,$$

and define $s_{i,t+1}$ and k_{it} $(t \geq 1)$ recursively by

$$k_{it} = \begin{cases} a & \text{if } s_{it} = 1 \text{ and } c_{it} = 1 \\ a + b_1 & \text{if } s_{it} = 1 \text{ and } c_{it} = 0 \\ a + b_2 & \text{if } s_{i,t-T+1} = \cdots = s_{it} = 0 \text{ and } c_{it} = 1 \\ a + b_2 + b_3 & \text{if } s_{i,t-T+1} = \cdots = s_{it} = 0 \text{ and } c_{it} = 0, \\ \infty & \text{otherwise.} \end{cases} \tag{11}$$

and

$$s_{i,t+1} = \mathbf{1}_{\{z_t > k_{it}\}}, \tag{12}$$

where for any condition A, $\mathbf{1}_{\{A\}} = 1$ if A holds and $= 0$ otherwise. It can be seen that $k_{i,t-1}$ and s_{it} $(t = 1, 2, \ldots)$ equal the threshold and action choice, respectively, prescribed by the T-state threshold strategy in state t when the public signals are $z_1, z_2, \ldots, z_{t-1}$ and the own action choices are $c_{i1}, c_{i2}, \ldots, c_{i,t-1}$ in periods $1, \ldots, t-1$.

In actual estimation, we use a *stochastic version* of the threshold strategy instead of the deterministic definition given above. This is necessary in order to explicitly incorporate the possibility of mistakes by subjects, and to capture possible asymmetry across them. Specifically, we make the following modifications to (11) and (12): First, we introduce a parameter γ_0 and reformulate the threshold condition as

$$s_{i,t+1} = \mathbf{1}_{\{\gamma_0 z_t > \kappa_{it}\}} \quad (t \geq 1). \tag{13}$$

Clearly, (13) is equivalent to (12) if $\gamma_0 > 0$ and $\kappa_{it} = \gamma_0 k_{it}$.[34] We suppose that κ_{it} is given by ($s_{it} = 1$ for $t \leq 1$):

$$\kappa_{it} = \begin{cases} \alpha + \nu_i & \text{if } s_{it} = 1 \text{ and } c_{it} = 1 \\ \alpha + \nu_i + \gamma_1 & \text{if } s_{it} = 1 \text{ and } c_{it} = 0 \\ \alpha + \nu_i + \gamma_2 & \text{if } s_{i,t-T+1} = \cdots = s_{it} = 0 \text{ and } c_{it} = 1 \\ \alpha + \nu_i + \gamma_2 + \gamma_3 & \text{if } s_{i,t-T+1} = \cdots = s_{it} = 0 \text{ and } c_{it} = 0 \\ \alpha + \nu_i + \gamma_4 & \text{otherwise.} \end{cases} \tag{14}$$

In (14), note that the parameters α, γ_1, γ_2, γ_3, γ_4 are all assumed common across subjects, and that only ν_i is indexed by i. The relationship $\kappa_{it} = \gamma_0 k_{it}$ will be restored if $\gamma_4 = \infty$ and

$$a = \frac{\alpha + \nu_i}{\gamma_0}, b_1 = \frac{\gamma_1}{\gamma_0}, b_2 = \frac{\gamma_2}{\gamma_0}, b_3 = \frac{\gamma_3}{\gamma_0}. \tag{15}$$

The term ν_i, which is the only term indexed by i, is supposed to capture possible asymmetry across the subjects: The larger is ν_i, the higher the threshold and hence more likely is subject i to play the non-cooperative action H. Technically, ν_i is treated as *correlated random effects*, and is assumed to be a random variable with the 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 common across subjects, and are to be estimated from the data. On the other hand, 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.[35] We assume that ν_i's are independent across subjects, and use them to test the symmetry assumption of the theory.

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

[35] We 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 $\nu_i \sim N\left(0, \sigma_\nu^2\right)$, the consistency of our estimate would require the (unlikely) inde-

	$\beta = 0$	$\beta = 1$	$\beta = 4$	$\beta = 10$
α	-1.437***	-0.819***	-0.569***	-0.430***
	(0.127)	(0.090)	(0.102)	(0.150)
ψ	1.698***	1.313***	1.278***	1.393***
	(0.210)	(0.379)	(0.250)	(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

We introduce another element of randomness in the form of a random shock term which represents errors made by subjects. More specifically, for a random variable u_{it}, we suppose that 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.[36] In other words, $c_{i,t+1}$ is determined by

$$c_{i,t+1} = \mathbf{1}_{\{\gamma_0 z_t > \kappa_{it} + u_{it}\}}. \tag{16}$$

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 (16) is as follows: When the realized public signal z_t is close to his threshold $\frac{\kappa_{it}}{\gamma_0}$, subject i makes errors more often than when z_t is far from it.[37] Together, (13) and (16) define a limited dependent variable model with lagged dependent variables.

6.2.2 Testing of the random choice strategy

As a benchmark, we first estimate a *random choice strategy*, for which the expected probability of each action (H and L) is the same throughout the game. Specifically,

pendence of c_{i1}'s and ν_i's. (See Wooldridge (2002) for a clear expostion 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).

[36] By modeling deviations as random shocks, we ignore any underlying motive for those deviations.

[37] An alternative specification is one where the probability of mistakes is independent of the realized value of the public signal. Such specification, however, is not compatible with the standard estimation techniques.

the estimated equation is given by

$$c_{i,t+1} = \mathbf{1}_{\{\alpha+\nu_i+u_{it}<0\}},$$

where ν_i and u_{it} are as defined in the previous section. Note that this is obtained from the general model (16) by setting $\gamma_0 = \gamma_1 = \gamma_2 = \gamma_3 = \gamma_4 = 0$. Table 5 reports the estimates of this benchmark model, where

$$\rho = \frac{\sigma_\nu^2}{\sigma_\nu^2 + 1}$$

is used as a substitute to σ_ν^2 as is customary for random effects estimates. As one would expect, the positive estimate of the coefficient ψ on ζ_i indicates 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, indicating that tendencies to cooperate vary across the subjects.[38]

6.2.3 Estimation of general threshold strategies

We now turn to the estimation of general threshold strategies. A few comments are in order: First, as noted earlier, automatic transition from state 1 to state 2, ..., from state $T-1$ to state T in the original specification of a threshold strategy is captured by setting $\gamma_4 = \infty$ in (14). However, when $\gamma_4 = \infty$, u_{it} fails to give the desired random effect through (16). For this reason, we set γ_4 equal to a finite value in our estimation.[39] Although the choice of any particular value for γ_4 is arbitrary, we set $\gamma_4 = \gamma_0\overline{z} - \alpha - \psi$, where $\overline{z}$ is the highest *realized* value of the public signal z in each treatment. The action choice implied by (16) for this γ_4 is the same as that for $\gamma_4 = \infty$ for the observed signal value z_t and the mean values of ν_i and u_{it}. However, the two action choices differ from each other when $\nu_i + u_{it}$ takes a large

[38] 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 Gutierres, Carter, and Drukker (2001).

[39] Alternatively, γ_4 could also be estimated from the data. In this case, however, the interpretation of the estimated strategy would become difficult since T would no longer be the minimum punishment length. Furthermore, it will also be seen shortly that the best estimate of T is always 1. In this sense, it does not appear important for our qualitative conclusions whether the value of γ_4 is fixed this way or is estimated.

	$\beta = 0$	$\beta = 1$	$\beta = 4$	$\beta = 10$
α	9.543	4.410***	-0.245***	-0.318*
	(42.103)	(1.157)	(0.084)	(0.168)
γ_0	0.564	0.283***	0.026***	0.008*
	(2.335)	(0.988)	(0.007)	(0.005)
γ_1	-0.237	0.010	0.804***	0.079
	(4.957)	(0.011)	(0.103)	(0.065)
γ_2	-0.167	-0.211	0.049*	0.010
	(0.580)	(0.245)	(0.026)	(0.009)
γ_3	-0.147	-0.010	0.011*	-0.003
	(0.329)	(0.024)	(0.006)	(0.002)
ψ	0.908	1.192*	0.010*	1.460***
	(1.898)	(0.126)	(0.005)	(0.498)
T	1	1	1	1
ρ	$0.116^{\S}$	$0.229^{\S}$	$0.313^{\S}$	$0.218^{\S}$
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

negative value.[40] Second, we do not place the restriction that $\gamma_0 > 0$. As will be seen, however, the estimate of γ_0 turns out to be positive in every treatment. Third, our estimation is done separately for each value of T given that every sequence of play in the data is finite. The T with the highest log likelihood is selected. It should be noted that a threshold strategy with the number of states $T+1$ greater than the longest cycle in each session is identified with a grim-trigger strategy.

Table 6 reports the estimates of the coefficients in (16). The standard errors are bootstrapped using a fixed T.[41] The estimated T is 1 in every treatment.

[40] The smaller the value we choose for γ_4, the more likely do the two action choices differ from each other. In essence, we have chosen the smallest value for which the action choices are the same for the mean values of ν_i and u_{it}. We do not expect that other choices would yield significantly different conclusions. For the strategy with $T = 1$, of course, the choice of any value for γ_4 is irrelevant.

[41] Fifty replications using the full sample size are computed. Computation of standard errors for T (a discrete variable) is omitted. It would add little to the interpretation of the results.

Once again, the random-effects specification is not rejected for any treatment. It is also worth noting that γ_0 is statistically significant in the $\beta = 10$ treatment in contradiction to the prediction of the theory: 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. We suspect that the lack of statistical significance results from the inclusion of too many parameters. This point is examined in more detail in the next section.

In general, the model has explicative power. For each $\beta \leq 4$, a likelihood ratio test rejects any of the random choice model described in Table 5 at the 1% significance level. On the other hand, for $\beta = 10$, the random choice model cannot be rejected even at the 10% level. We observe the following concerning the base threshold level a in (11): 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 (15) that the subject-independent component of the base threshold a 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.

6.2.4 Restricted threshold strategies

We now restrict some of the parameters of the general threshold specification and examine the performance of the resulting models. Specifically, the tested specifications are as follows:

Sa: $b_1 = b_2 = b_3 = 0$

Sb: $b_1 = b_2 = 0$, $b_3 = -2$

Sc: $b_1 = -2$, $b_2 = -\infty$, $b_3 = 0$

Sd: $b_1 = b_3 = -2$, $b_2 = 0$

Se: $b_1 = b_3 = 0$, $b_2 = -\infty$

The specific values are chosen based on the following considerations: The restriction $b_1 = 0$ implies that the threshold in state 0 is independent of the player's own

	γ_1	γ_2	γ_3
Sa	0	0	0
Sb	0	0	$-2\gamma_0$
Sc	$-2\gamma_0$	$\gamma_0\underline{z} - \alpha - \psi$	0
Sd	$-2\gamma_0$	0	$-2\gamma_0$
Se	0	$\gamma_0\underline{z} - \alpha - \psi$	0

$\underline{z}$ denotes the lowest realization of z.

Table 7: Parameter restrictions under each specification

action. In other words, even if the player has deviated to H in state 0 and is himself responsible for the downward shift of the distribution of the public signal, he ignores its effect. When $b_1 = -2$, on the other hand, the threshold in state 0 is lowered after an own deviation to H in state 0 to properly discount its effect on z.[42] Similarly, when $b_3 = 0$, the threshold in state T is independent of the own action, and when $b_3 = -2$, it is lowered when the player has chosen action H (as intended in state T). The restriction on b_2 concerns the length of the punishment phase: When $b_2 = -\infty$, the punishment phase is of fixed length T and transition from state T to state 0 is automatic, and when $b_2 = 0$, the threshold in state T is the same as that in state 0 and the punishment phase is of variable length. When $T = 1$, threshold strategies with $b_2 = 0$ resemble the tit-for-tat strategy in that their (on-the-path) reaction is determined only by the most recent public signal and that it does not depend on the state.

The above restrictions on b_1, b_2 and b_3 translate to those on γ_1, γ_2 and γ_3 through (15). Just as in the case of γ_4 discussed earlier, we would like to generate randomness by u_{it} even when $b_2 = -\infty$. For this reason, we replace $\gamma_2 = -\infty$ by $\gamma_2 = \gamma_0\underline{z} - \alpha - \psi$, where $\underline{z}$ denotes the lowest realization of z: 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. In sum, the restrictions for each model are those given in Table 7 as well as $\gamma_0 > 0$ and $\gamma_4 = \gamma_0\overline{z} - \alpha - \psi$ (as mentioned earlier).

Our findings can be summarized as follows: The minimal punishment length T is estimated to be 1 in each model. In terms of their fit levels, the five models are

[42]Note that the probability that i's opponent played L given the public signal z and the own action $a_i = L$ is the same as that given $z - 2$ and $a_i = H$.

	$\beta = 0$	$\beta = 1$	$\beta = 4$	$\beta = 10$
α	8.391***	4.220***	0.309*	-0.330**
	(0.599)	(0.511)	(0.168)	(0.159)
γ_0	0.504***	0.274***	0.048***	0.005*
	(0.031)	(0.027)	(0.007)	(0.003)
ψ	0.819***	1.253***	1.219***	1.389***
	(0.257)	(0.375)	(0.242)	(0.250)
ρ	$0.124^{\S}$	$0.227^{\S}$	$0.286^{\S}$	$0.184^{\S}$
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 (Sa) specification

ordered almost identically for all the treatments: (Sa) has the best fit and is followed by (Sb), (Sd), (Se) and (Sc) in this order.[43] The remaining results are obtained using likelihood ratio tests. In every treatment except for $\beta = 4$, we cannot reject at the 10% level the hypothesis that the (Sa) specification fits the data as well as the general threshold specification does. In the $\beta = 4$ treatment, however, (Sa) is rejected at the 1% level. With a few exceptions, the corresponding hypothesis for each one of the other specifications is rejected at the 1% level in every treatment. We depict in Figure 9 the (Sa) specification. It can be checked that given our parameter values, there exists a threshold k for which (Sa) is a symmetric Nash equilibrium strategy if and only if the noise level $\beta \leq 2\left(\log \frac{5\delta}{5\delta-2}\right)^{-1} \approx 3.4026$.[44] This may in part explain the rejection of this strategy in the $\beta = 4$ treatment.

The fact that (Sa) 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

[43] (Sd) and (Se) are reversed for $\beta = 1$ and (Se) and (Sc) are reversed for $\beta = 10$.

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

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

It can also be verified that (Sa) cannot be a symmetric perfect equilibrium strategy for any β.

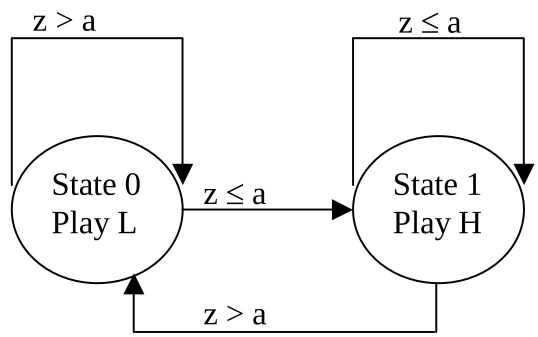

Figure 9: The (Sa) specification

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 the (Sa) specification in the $\beta = 4$ treatment indicates that at least one of the parameters b_1, b_2, and b_3 is non-zero, we test the following additional specifications:

Sa-i: $b_2 = b_3 = 0$

Sa-ii: $b_1 = b_3 = 0$

Sa-iii: $b_1 = b_2 = 0$

Both (Sa-ii) and (Sa-iii) are rejected at the 1% level, whereas (Sa-i) is not rejected at the 10% level.[45] The estimates for γ_0 and γ_1 both turn out positive, implying that $b_1 = \gamma_1/\gamma_0 > 0$, or 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.

As a further examination of the (Sa) specification, we report below the results of two alternative tests that do not rely on the parametric assumptions used for the estimation of the threshold strategies.

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

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

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. Assuming $T = 1$, 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. They in particular imply that the specifications (Sb)-(Sd), which all have $(b_1, b_3) \neq (0,0)$, are unlikely.[46]

To give a better idea on how well the (Sa) strategy fits, we check how well the deterministic specification of the (Sa) strategy (*i.e.,* $b_1 = b_2 = b_3 = 0$ in (11) and (12)) 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.[47] 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.

[46] While (Se) has $(b_1, b_3) = (0,0)$, it performs extremely poorly in the likelihood ratio test.

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

The results of these alternative tests provide a stronger support for the use of the (Sa) strategy by the subjects.

7 Discussions

As discussed in the Introduction, this paper analyzes cooperation in infinitely repeated games with imperfect public monitoring under the exact 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 other 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 comments 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.[48] 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 treatment where the subjects' payoffs are strictly positive. In theory, however, the unique (symmetric or asymmetric) equilibrium in this treatment is the repetition of the one-shot Nash equilibrium, which yields zero.

[48] In 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).

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.

In comparison with the real industrial setting, the subjects in our experiments play in an extremely simple environment with 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. By the bang-bang property of a perfect public equilibrium (Abreu *et al.* (1990, Theorem 3)), the maximal symmetric perfect public equilibrium payoff v can be generated by a stationary grim-trigger equilibrium σ which plays the symmetric action profile $\hat{a}$ throughout the cooperation phase and reverts to a^0 if and only if $z \notin Q$ for some $Q \subset \boldsymbol{R}$.[49] Consider an alternative grim-trigger strategy profile $\hat{\sigma}$ which begins with $\hat{a}$ and reverts to a^0 if and only if $z \leq k$, where k is such that

$$P(z \in Q \mid \hat{a}) = \int_Q h(z \mid \hat{a})\, dz = \int_k^\infty h(z \mid \hat{a})\, dz = P(z > k \mid \hat{a}). \tag{17}$$

It then follows from (2) that $\hat{\sigma}$ and σ yield the same payoff. On the other hand, the incentive constraint (5) for this strategy can be written as

$$\frac{1 - \delta\{1 - H(k \mid a_i^0, \hat{a}_{-i})\}}{1 - \delta\{1 - H(k \mid \hat{a})\}} \geq \frac{g_i(a_i^0, \hat{a}_{-i}) - g^0}{\hat{g} - g^0}. \tag{18}$$

In what follows, we show that $\hat{\sigma}$ is also an equilibrium by verifying (18). Denote $K = (k, \infty)$ and write for $a_i \in A_i$,

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

[49] Generation of a perfect public equilibrium payoff below v requires the use of a different Q in period 1.

Note that $M(\hat{a}_i) = N(\hat{a}_i)$ by (17), and that

$$P(z > k \mid a_i, \hat{a}_{-i}) = P(z \in Q \mid a_i, \hat{a}_{-i}) + M(a_i) - N(a_i).$$

Note that (18) follows from $P(z > k \mid a_i^0, \hat{a}_{-i}) \leq P(z \in Q \mid a_i^0, \hat{a}_{-i})$, or equivalently, $M(a_i^0) \leq N(a_i^0)$. Assumption 1 implies that

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

and that

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

Therefore,

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

This completes the proof. //

Characterization of $K(\delta)$ and $k^*(\delta)$. Proposition 3 below provides a characterization of the set $K(\delta)$ of thresholds for which player i finds playing $\hat{a}_i$ incentive compatible.

Assumption 4 *f is continuous, and $f(0) = \max_{x \in \boldsymbol{R}} f(x)$.*

Assumption 4 holds for many standard distributions as well as the one (10) 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 3 *Suppose that Assumptions 2 and 3 hold. Then $K(\delta)$ is a (possibly empty) closed interval. If, in addition, Assumption 4 holds, then the optimal threshold $k^*(\delta) < s(\hat{a})$ when $K(\delta) \neq \emptyset$.*

Proof. Define

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

After some algebra, we see that (7) is equivalent to

$$W(k, a_i) \geq \frac{l-1}{\delta} \tag{19}$$

Since f satisfies Assumption 3 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 $W(k, a_i)$, $W(k', a_i) \geq \frac{l-1}{\delta}$. Then for any k'' between k and k',

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

and

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

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 $W(k'', a_i) \geq \frac{l-1}{\delta}$. This implies that the set of k's which satisfy (8) is convex. That $K(\delta)$ is closed follows from the continuity of W in k.

Suppose now that Assumption 4 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 $W(k, a_i) \geq \frac{l-1}{\delta}$ for some $k \geq s(\hat{a})$, then $W(s(\hat{a}) - \gamma, a_i) > W(s(\hat{a}), a_i) \geq W(k, a_i) \geq \frac{l-1}{\delta}$ 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 (10) 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 $\log l < \lambda < \mu < \nu$. The following three cases arise depending on the value of the discount factor δ:

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

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

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

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

where

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

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

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

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) $k_4 > s(\hat{a}) \Leftrightarrow \frac{d}{\beta} > \nu$.

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., Pearce, D. and Stacchetti, E.: 1990, Toward a theory of discounted repeated games with imperfect monitoring, *Econometrica* **58**, 1041–1063.

Bereby-Meyer, Y. and Roth, A. E.: 2004, The fragility of cooperation. mimeo.

Casari, M., Ham, J. C. and Kagel, J. H.: 2004, Selection bias, demographic effects and ability effects in common value auction experiments. mimeo.

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.

Dal-Bo, P.: 2003, Cooperation under the shadow of the future: experimental evidence from infinitely repeated games. mimeo.

Duffy, J. and Ochs, J.: 2003, 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. and Slonim, R. L.: 2002, Inferring repeated-game strategies from actions: Evidence from trust game experiments. mimeo.

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 Techical 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.: 1992, 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.: 2001, In support of trigger strategies: Experimental evidence from two-person non-cooperative games. mimeo.

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.

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.