Flexibility and Collusion with Imperfect Monitoring

Maria Bigoni and Jan Potters and Giancarlo Spagnolo*

March 2, 2012

Abstract

Flexibility - the ability to react swiftly to others' choices - facilitates collusion by reducing gains from defection before opponents react. Under imperfect monitoring, however, flexibility may also hinder collusion by inducing punishment after too few noisy signals. The combination of these forces predicts a non-monotonic relationship between flexibility and collusion. To test this subtle prediction we implement in the laboratory an indefinitely repeated Cournot game with noisy price information and vary how long players have to wait before changing output. We find that (i) the facilitating role of flexibility is lost under imperfect monitoring, and (ii) with learning, collusion unravels with low or high flexibility, but not with intermediate flexibility.

JEL: C73, C92, D43, L13, L14.

Keywords: Collusion, Cooperation, Flexibility, Imperfect monitoring, Oligopoly, Repeated games.

*Bigoni: University of Bologna, piazza Scaravilli 2, Bologna, Italy, maria.bigoni@unibo.it. Potters: Tilburg University and CentER, P.O. Box 90153 5000 LE Tilburg, The Netherlands, j.j.m.potters@uvt.nl. Spagnolo: SITE-Stockholm School of Economics, Tor Vergata, EIEF and CEPR, Sveavagen 65, 113 83 Stockholm, Sweden, giancarlo.spagnolo@hhs.se. We thank Andrzej Skrzypacz for his precious advice, and Charles Angelucci, Jeff Butler, Gabriele Camera, Martin Dufwenberg, Tobias Klein, participants at the 2009 European ESA Meeting in Innsbruck, the M-BEES 2011 Workshop in Maastricht, the 2011 Industrial Organization Workshop in Otranto, the CRESSE conference in Rhodes 2011, and seminars at Stockholm School of Economics, and Tilburg University for valuable comments. Giancarlo Spagnolo also gratefully acknowledges research funding from Konkurrensverket (the Swedish Competition Authority).

Flexibility, the ability to react swiftly to others' choices, is commonly seen as a factor that facilitates cooperation. The intuitive logic behind this belief is that flexibility reduces gains from unilateral defections by drawing punishment nearer. Axelrod, in his *Evolution of Cooperation* (1984, p. 129), puts it as follows: "[One] way to enlarge the shadow of the future is to make the interactions more frequent. In such a case the next interaction occurs sooner, and hence the next move looms larger than it otherwise would". A recent experiment by Friedman and Oprea (forth.) offers strong support for a positive effect of flexibility on cooperation. They implement a 60-seconds finite horizon repeated Prisoner's Dilemma under perfect monitoring and find a strong positive monotonic relationship between the speed at which subjects could adjust their actions and the rate of cooperation.[1]

Though appealing and intuitive, this established role of flexibility is theoretically robust only for games in which players can perfectly observe each others' actions. As first shown by Abreu, Milgrom and Pearce (1991), with imperfect monitoring flexibility may actually harm cooperation. The reason is that when imperfect information arrives frequently, high flexibility forces players to react to 'bad news' early, when it is still very noisy. This generates many costly mistakes which erode the value of cooperation. This negative effect of flexibility counteracts the positive effect which is due to the reduced gains from defection. Sannikov and Skrzypacz (2007) show in a variety of oligopolistic environments that for high levels of flexibility the negative effect dominates the positive one and renders collusion impossible altogether! For low levels of flexibility, on the other hand, the positive effect tends to dominate the negative one and an increase in flexibility will make collusion easier. The remarkable consequence is that the impact of flexibility on the sustainability of collusion is non-monotonic.

The channel through which flexibility hinders cooperation under imperfect monitoring is subtle though.[2] Certainly the intuition for the negative effect is less straightforward than the positive one that defection can be punished sooner. One may therefore question the behavioral relevance of the negative effect of flexibility. Does cooperation really unravel

[1] In the extreme case in which subjects could adjust their actions almost continuously the median rate of cooperation was as high as 90%. At the other extreme, in which subjects could adjust their actions only once, cooperation rates were close to zero.

[2] Vives (2009) describes this result as 'counterintuitive'. One might speculate, for example, that it is possible to delay punishments until more convincing information becomes available that the other player is really defecting. Such a strategy unravels though since such a delay will strengthen the incentives of the other player to defect.

with sufficiently high flexibility when monitoring is imperfect, or will the positive effect still dominate? Can one really observe a non-monotonic effect of flexibility on collusion in this case?

This paper presents an experimental study designed to start answering these questions. We implement in the laboratory an indefinitely repeated quantity-setting duopoly game in discrete time with imperfect monitoring, analogous to that studied by Sannikov and Skrzypacz (2007). Players do not observe each other's quantity choices; they only observe price which is a noisy signal of total quantity. Across treatments we vary inflexibility (Δ), that is, the number of periods players have to wait before they can change quantity. The game is set up such that collusion can be supported as an equilibrium of the repeated game when $\Delta = 2$, but not when $\Delta = 1$ or $\Delta = 3$. This allows us to examine the empirical support for a non-monotonic relationship between (in)flexibility and collusion in the laboratory. In order to give it a fair chance, we allow for learning by having subjects play seven repetitions of the indefinitely repeated game.[3]

The results show no support for the non-monotonicity of flexibility in the first few plays of the repeated game. Average cooperation rates are highest when $\Delta = 1$, intermediate when $\Delta = 2$, and lowest when $\Delta = 3$, which is in line with the effect that inflexibility increases the period over which defection can go unpunished. With repetition, however, cooperation rates show a significant downward trend when $\Delta = 1$ but not when $\Delta = 2$. In the last two repetitions of the repeated game, cooperation rates are higher for $\Delta = 2$ than for $\Delta = 1$ and for $\Delta = 3$, in line with the predicted non-monotonicity (although the first difference is not significant). This suggests that it takes some time for the disruptive effect of flexibility to reveal its force, so that the theoretical prediction may indeed bite if there is sufficient scope for learning. Additional evidence in this direction comes from our results on the profitability of cooperation, which increases significantly in time with $\Delta = 2$ (i.e., when cooperation is an equilibrium) while it does not when $\Delta = 1$ or $\Delta = 3$ (i.e., when cooperation is not an equilibrium).

Our results reiterate the need to reconsider the 'common wisdom' that in general cooperation is stabilized by the ability to react quickly. Both theory and experiment suggest that flexibility is a mixed blessing with imperfect monitoring. This has important implications for our understanding of collusion and for antitrust policy, as the relevant textbooks

[3]Several experimental papers document strong learning effects in repeated game settings (Selten and Stoecker, 1986; Camerer and Weigelt, 1988; Dal Bò and Fréchette, 2011).

present frequent interaction as a facilitating factor independently of the information agents have access to.[4] It also has implications for other settings where imperfect observability is the rule rather than the exception, such as team production, principal-agent relationships and international agreements. If effort can only be observed with noise, frequent monitoring and very short reaction lags are not generally conducive to cooperation, nor in fact will be very lax monitoring and unresponsiveness to new information. The optimal design may well involve some intermediate level of flexibility.

A number of recent experimental studies examines how cooperation is affected by imperfect monitoring *per se*. Bereby-Meyer and Roth (2006) study how imperfect monitoring interferes with learning, showing that it considerably reduces subjects' ability to learn to cooperate in repeated PD games (and to defect in one-shot ones). Aoyagi and Fréchette (2009) ask whether subjects' ability to cooperate falls when information becomes more noisy in a repeated game, finding significant support for this theoretical prediction. Fudenberg et al. (forth.) look at the prevailing strategies in repeated games where subjects' choices are implemented with mistakes, highlighting the success of strategies that are "lenient" (do not punish the first deviation) and "forgiving" (return to cooperation after a short punishment phase).[5] These studies indicate that strategies and outcomes are significantly affected by the presence of imperfect monitoring. We take imperfect monitoring as given, and examine how it affects the comparative statics of a key structural variable.

1 Experimental Design

The design of our experiment aims at replicating in the simplest possible setting Sannikov and Skrzypacz (2007) [S&S]'s intriguing theoretical results. To do this, we adapt and further simplify their analysis of collusion in a Cournot supergame (they also obtain the result in more complex set ups).

[4]See Tirole (1988, p.240); Church and Ware (2000, p.343); Martin (2001, p.192); Ivaldi et al. (2003); Motta (2004, p.145), Belleflamme and Peitz (2010, p.254).

[5]See also the earlier work by Cason and Khan (1999), who compare perfect monitoring with perfect but delayed monitoring; and by Feinberg and Snyder (2002) and Holcomb and Nelson (1997), who study the effects of different types of imperfect but private monitoring on cooperation.

The Game

Two players interact repeatedly, in discrete time, in a Cournot market with homogeneous products. Players simultaneously set quantities (q_{1t}, q_{2t}) and the resulting price depends on total quantity ($Q_t = q_{1t} + q_{2t}$) and a random shock (ϵ_t). Specifically, price P_t in period t is given by the following demand function:

$$P_t(Q_t) = a - Q_t + \epsilon_t, \;\; \epsilon_t \sim N(0, \sigma^2)$$

Monitoring is imperfect because players receive information about each period's price P_t, but not about total quantity Q_t or the random shock (ϵ_t). We restrict the action set to $q_{it} \in \{3, 4\}$. We set $a = 12$ in the demand function, marginal cost equal to 0, and per period fixed cost equal to 16. Table 1 presents expected price and profits of the stage game that result from this parameterization. Note that the expected profits are those of a Prisoner's Dilemma.[6] Finally, we set the standard deviation of random price equal to $\sigma = 1.3$.[7]

		q_2	
	Price	3 units	4 units
q_1	3 units	6	5
	4 units	5	4

		q_2	
	Profits	3 units	4 units
q_1	3 units	2, 2	-1, 4
	4 units	4, -1	0, 0

Table 1: Expected price and expected profits of the stage game

Prices and profits materialize in every period, but quantities can be adjusted only every Δ periods. So, the quantities chosen in period t, together with the random shocks, determine prices and profits for the following Δ periods. A larger Δ implies that it takes longer before players can react to a (bad) price signal, but also that when they react they will have observed more signals about their opponent's quantity choice.

In the model of S&S the game is infinitely repeated and players discount future profits using a common interest rate r. The per-period discount factor is then equal to $\delta = e^{-r\Delta}$. In the experiment we do not implement discounting, but we implement a repeated game of indeterminate length with a continuation probability equal to δ. This is a practical

[6]A convenient feature is that the Nash and minmax payoff are the same. This makes it easier to derive the best possible collusive outcomes.

[7]This implies for example that there is a probability of 44% that the realized price will deviate by at least 1 unit from the expected price.

way to implement infinitely repeated games in the lab, and theoretically innocuous under risk neutrality (which we need to assume). We choose $r = 0.10$ so that we have $\delta = e^{-r}$ $= 0.90$ for Δ=1, $\delta = 0.82$ for Δ=2, and $\delta = 0.74$ for Δ=3. The continuation probability decreases with Δ, but conditional on the game being continued the number of additional periods increases with Δ. The expected number of periods is $\frac{\Delta}{1-\delta}$, which is equal to 10, 11.1, and 11.5 for $\Delta = 1$, 2, and 3, respectively.

In the experiment we implement three different treatments in which Δ takes the values 1, 2, and 3, respectively. A smaller value of Δ (higher flexibility) has two contrasting effects. One is that the discount factor is higher. This implies that defection can be punished more effectively, which generates the usual positive effect on collusion. A smaller value of Δ also implies that the players attain fewer noisy (price) signals about the other player's previous action before making the next choice. This has a negative impact on the scope for collusion, since it generates a high rate of "false positives". The analysis of S&S implies that the interplay between these two effects generates a non-monotonic effect of Δ. Collusion can be supported as an equilibrium in the repeated game only for intermediate values of Δ. When Δ is large the gains from defection are too attractive; when Δ is small the stochastic variation in prices erodes the gains from collusion by triggering too frequent punishments. In Appendix A we outline how this result can be derived. Applied to our game, assuming risk neutral payoff maximization, collusion is sustainable when $\Delta = 2$, but not when $\Delta = 1$ or $\Delta = 3$. It is this theoretical prediction that we explore in our experiment.

Procedure

The experiment was run in the CentERlab at Tilburg University in March 2009. There were six sessions, two for each treatment, with 16 subjects in each. Within each session, there were two matching groups of 8 subjects and subjects interacted only with other subjects in their matching group. This gives us four independent observations per treatment. Subjects were recruited through an e-mail list of students interested to participate in experiments. The experiment was computerized and programmed with zTree (Fischbacher 2007). Interaction between subjects in the experiment was anonymous.

Upon entering the lab, subjects were randomly seated at tables separated by partitions. Written instructions were distributed and read aloud. See Appendix B for a copy of the instructions. Subjects were given ample time to study the instructions at their own pace and to privately ask questions. A short quiz was conducted to check their understanding.

During the experiment profits were denoted in points; after the experiment points were converted into cash at a rate of 8 points = 1 Euro. To accommodate for potential losses, subjects were given a starting endowment of 80 points.[8] Sessions lasted on average one hour and 45 minutes, including instructions and payment, and subjects received an average payment of 17 Euro.[9]

Subjects were randomly matched to one other subject to play the repeated game. Below we will refer to one play of the repeated game as a *match*. In the first period of each match, subjects had to determine the quantity ($q_{it} \in \{3, 4\}$) they wanted to produce. Depending on the treatment, quantities were fixed for the next Δ periods. At the end of each block of Δ periods, subjects received information about the realized prices and their own profits in the last Δ periods. The random price shock which was drawn independently for each period from a normal distribution with zero mean and standard deviation $\sigma = 1.3$.[10] After each block of Δ periods, there was a probability δ that the game continued, and a probability $1 - \delta$ that the game ended. When the game continued, subjects had to choose the quantity for the next Δ periods.

When a game ended, a subject was rematched to a new subject to play the repeated game anew. To facilitate this rematching, the realization of the continuation probability was common across all pairs of subjects in the same session. Rematching took place 6 times. So, each subject played the indefinitely repeated game exactly 7 times, and this was common knowledge. To exclude reputation building across matches, we adopted a matching protocol that ensured that two subjects never interacted together in more than one match.

We carefully explained the details of the game and the procedure to the subjects. In

[8]In keeping with the model by S&S, we used an unbounded support for the price shocks. This implied that a subject could in theory attain negative cumulative profits. The probability of this happening when subjects always defect is 3%, but it is much lower when subjects cooperate. In 3 out of 672 matches, it occurred that a subjects had a negative total payoff at some point, but it never happened that a subject ended the experiment with a negative balance. The opposite possibility, that a subject would gain an unreasonably high profit of above, say, 100 Euro, is negligible and comparable to the probability of a power failure.

[9]If subjects had played non-cooperatively throughout, they would have had zero profits on average and earned only the starting endowment, amounting to 10 Euro. Playing cooperatively throughout would earn them 29.4 Euro in total in expectation. Playing the best cooperative equilibrium in $\Delta = 2$ would generate expected total earnings of 23.8 Euro. This illustrates that the incentives to cooperate are substantial.

[10]The random shocks on prices were generated by the software at the beginning of each period. They were different for each couple of subjects, and changed across sessions.

particular, we took great care to explain the role of the random price shocks, the random determination of the number of periods, and the (re)matching procedure.

2 Results

Our main interest is in how the rate of cooperation (collusion) varies across the three treatments. A look at aggregate cooperation frequencies across treatments, reported in Table 2, suggests at first hand that the theoretical prediction is not borne out by our experimental data: there are no significant differences in aggregate cooperation frequencies across the three treatments.[11]

Table 2: average cooperation frequencies, by treatment.

Treatment	Cooperation rate
$\Delta = 1$	0.247
$\Delta = 2$	0.254
$\Delta = 3$	0.205
Total	0.235

However, we know that aggregate comparisons may not tell the right story because learning effects are often very important in non-trivial experimental settings (Seleten and Stoeker, 1986; Camerer and Weigelt, 1988; Roth and Erev, 1995; and Dal Bò and Frechétte, 2011). We also know that the presence of imperfect monitoring may make learning even more difficult in set ups similar to ours (Bereby Meyer and Roth, 2006).

The positive effect of flexibility on the sustainability of collusion is more intuitive than the negative one and less related to the information structure. We therefore expect that subjects need to gain more experience with the game before the negative effect displays its force than for the positive one to act. We will see that this expectation is largely confirmed by the experimental data.

Figure 1 presents the development of the rate of cooperation ($q_t = 3$) over the matches for each of the three treatments. The top panel displays the average rates of cooperation in the first period of each match; the bottom panel gives average rates of cooperation across

[11] We are counting one independent observation per matching group, 12 observations in total. According to a Mann-Whitney Wilcoxon test, differences across treatments are not significant at any standard significance level (p-value > 0.3 for all the three comparisons).

all periods of a match. In both figures we see that in the early matches cooperation is most frequent in treatment $\Delta = 1$. The positive effect of flexibility on cooperation seems to dominate in the beginning of the experiment. The decline in cooperation over the matches, however, is more pronounced for $\Delta = 1$ and $\Delta = 3$ than for $\Delta = 2$. As a consequence, in later matches cooperation rates are higher in treatment $\Delta = 2$ than in treatments $\Delta = 1$ and $\Delta = 3$.[12]

The decline in cooperation rates over matches in treatments $\Delta = 1$ and $\Delta = 3$ is confirmed by the regressions in Table 3. The dependent variable is the binary decision between cooperation ($q_t = 3$) and defection ($q_t = 4$), while the independent variables are the treatment dummies (with $\Delta = 2$ being the reference treatment), the Match (which takes values 1 to 7), and interactions between the two.[13]

Table 3: Panel regression for rates of cooperation

	Cooperation first period		Cooperation all periods	
Match	0.001	(0.015)	−0.008	(0.012)
$\Delta = 1$	0.295**	(0.135)	0.134	(0.084)
$\Delta = 3$	0.076	(0.110)	0.035	(0.071)
Match x $\Delta = 1$	−0.051**	(0.025)	−0.033**	(0.015)
Match x $\Delta = 3$	−0.046***	(0.017)	−0.022*	(0.013)
Period			−0.002*	(0.001)
Constant	0.286***	(0.085)	0.306***	(0.078)
Observations	672		6272	
R2-overall	0.061		0.024	
R2-within	0.052		0.025	
R2-between	0.073		0.018	

Note: Standard errors in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.
Linear probability model with random effects at the subject level. Robust standard errors for data clustered on matching groups.

[12] The average cooperation rate over matches 4 to 7 is significantly higher in $\Delta = 2$ than in $\Delta = 3$ if we take the four matching groups in each treatment as independent observations ($p < 0.05$ with a Mann-Whitney rank-sum test). The difference between $\Delta = 1$ and $\Delta = 2$ is not significant. This holds both for cooperation rates in the first period and for cooperation rates across all periods in a match. The same results are obtained if we look at cooperation rates in only the last two matches.

[13] In this and in the following regressions, standard errors are computed clustering on matching groups. Clustering on sessions would produce qualitatively similar results (see Table C.1 in Appendix C). Here, as well as below, we use a linear probability model. This makes it easier to calculate and report the coefficients and significance levels of interaction effect (see Ai and Norton, 2003). Results for a logit model are reported in Table C.2 in Appendix C.

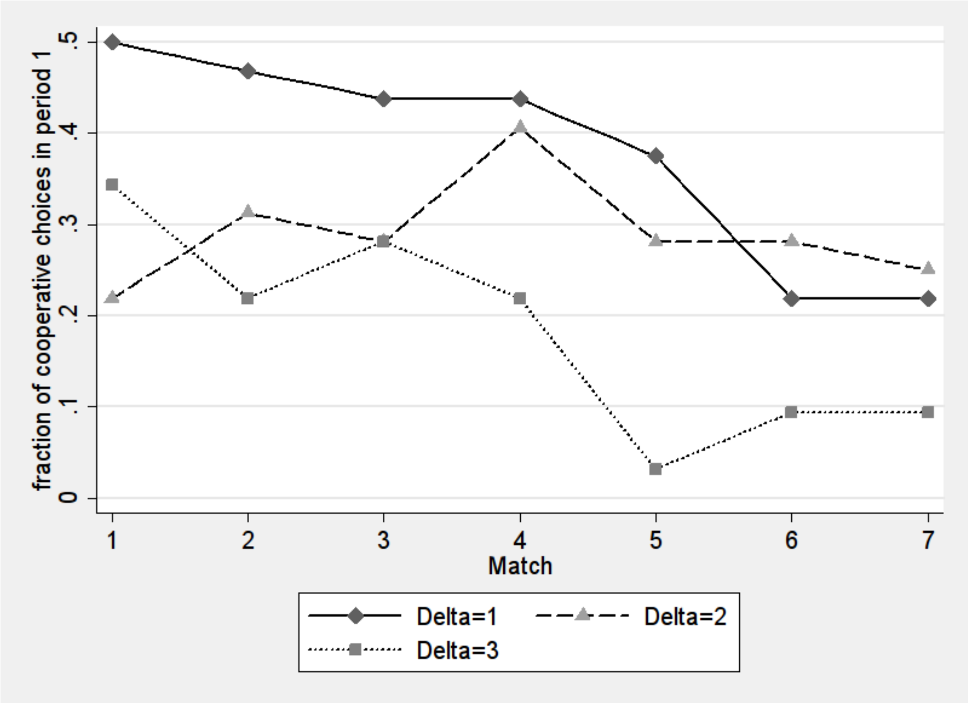

(a) First period.

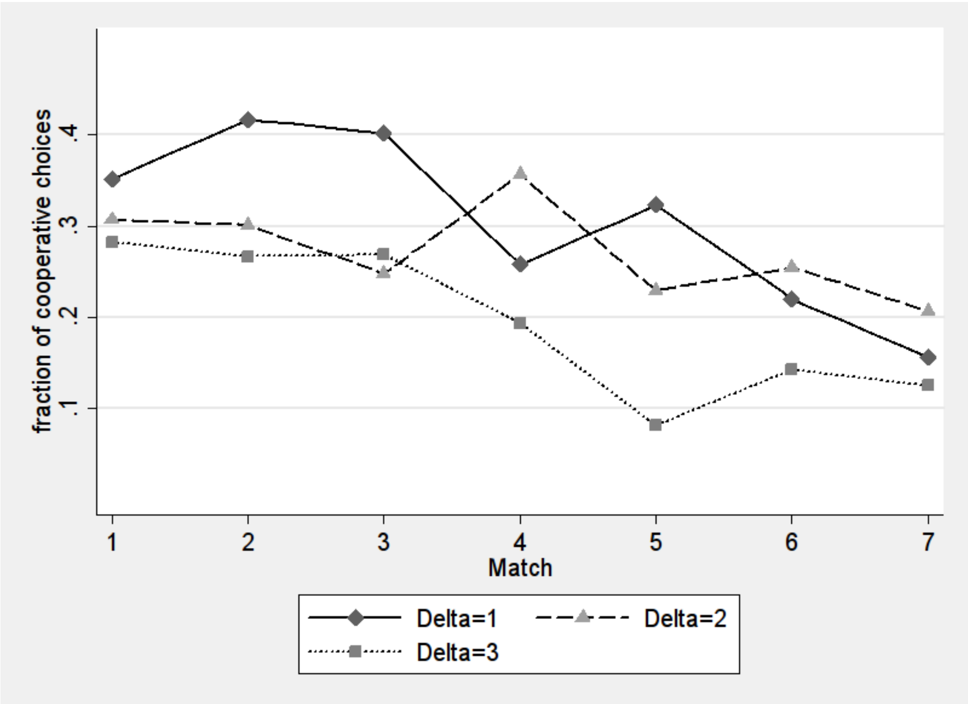

(b) All periods.

Figure 1: Frequency of cooperation by match.

The first regression considers only the first period decision, whereas the second regression considers the decisions of all periods in each match. The latter regression also includes the period in a match as an independent variable.[14] While no significant trend in cooperation across matches is found for the reference treatment ($\Delta = 2$), for the other two treatments the rate of cooperation decreases significantly with the match number.

Result 1. *The rate of cooperation decreases significantly with experience in treatments* $\Delta = 1$ *and* $\Delta = 3$*, but not in treatment* $\Delta = 2$*.*

Next we examine how profitable cooperation is in the different treatments. We hypothesize that cooperation will be more remunerative in an environment in which it can be sustained as an equilibrium outcome ($\Delta = 2$) than in environments in which it cannot ($\Delta = 1$, $\Delta = 3$). Figure 2 presents the average per period profit for subjects who cooperated in the first period of a match and those who did not. It shows that on average cooperation is more profitable than defection in $\Delta = 2$ but not in the other two treatments.

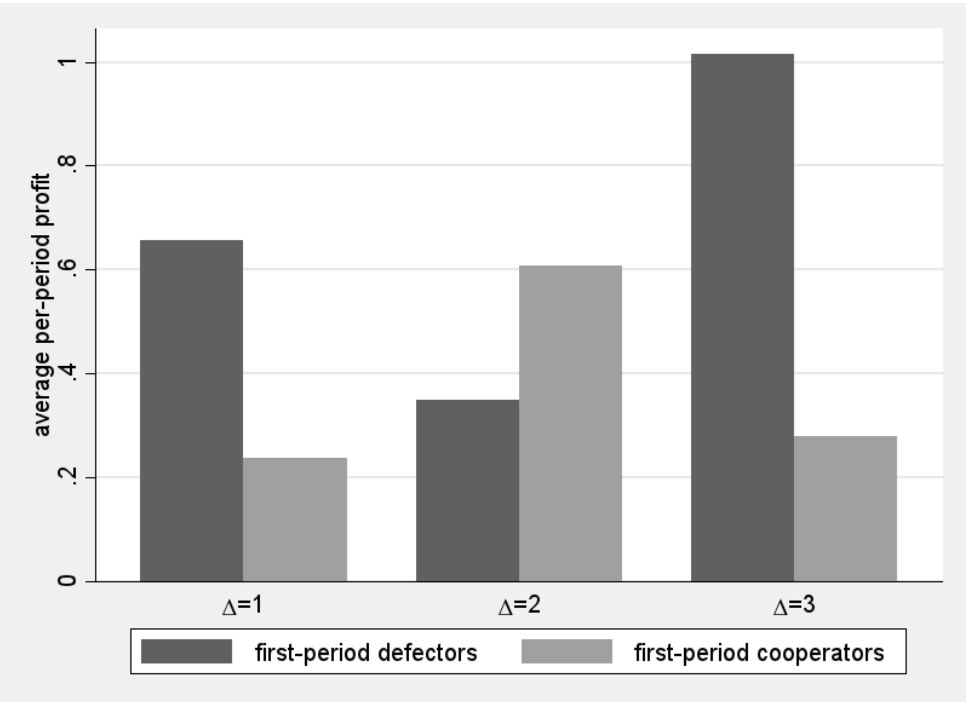

Figure 2: Difference in average profit for first-period cooperators and defectors.

Table 4 presents the results of a linear panel regression. The dependent variable is the average profit of each subject in each match. Thus we have 7 observations per subject.

[14]Table C.3 in Appendix C reports results for two alternative specifications of the model: the first one does not control for the trend across periods within a match, the second one allows for different trends across treatments.

The regressors consist of the subject's decision whether or not to cooperate in the first period of the match (Coop(t=1)), dummies for the treatment ($\Delta = 1$ and $\Delta = 3$), the Match, and their interactions. The regression analyzes whether it pays off to play cooperatively in the first period of a match, and how this depends on the treatment and the number of the match.

Table 4: Panel regression of average profits

	Average period profit	
Constant	0.765^{*}	(0.390)
Coop(t=1)	-0.465	(0.283)
$\Delta = 1$	0.364	(0.942)
$\Delta = 3$	0.395	(0.780)
Coop(t=1) x $\Delta = 1$	0.867	(1.260)
Coop(t=1) x $\Delta = 3$	0.473	(0.723)
Match	-0.104	(0.101)
Match x		
Coop(t=1)	0.181^{***}	(0.064)
$\Delta = 1$	-0.006	(0.189)
$\Delta = 3$	0.070	(0.152)
Coop(t=1) x $\Delta = 1$	-0.444^{**}	(0.222)
Coop(t=1) x $\Delta = 3$	-0.440^{***}	(0.134)
Observations	672	
R2-overall	0.019	
R2-within	0.016	
R2-between	0.043	

Note: Standard errors in parentheses * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.
Panel regression with random effects. Robust standard errors for data clustered on matching groups.

Results indicate that for treatment $\Delta = 2$ there is a strong significant increase in the profitability of cooperation over the matches (the positive coefficient of Match x Coop(t=1) in the bottom panel). The other two treatments do not show such a trend. The profitability of cooperation decreases over the matches in $\Delta = 1$ and $\Delta = 3$, and significantly so in $\Delta = 3$.[15]

[15] The sum of coefficients for Match x Coop(t=1) and for Match x Coop(t=1) x $\Delta = 3$ is significantly different from 0 at the 5% level. The sum of coefficients for Match x Coop(t=1) and for Match x

Result 2. *The profitability of cooperation increases significantly as subjects gain experience in treatment* $\Delta = 2$*, while no such trend is visible for* $\Delta = 1$ *and* $\Delta = 3$*.*

When collusion is an equilibrium – as in treatment $\Delta = 2$ of our experiment – it is supported by trigger strategies prescribing to cooperate as long as observed price is above a cut-off level and to switch to defection otherwise (see Appendix A).

To examine how subjects' choices depended on realized prices we run a regression where the dependent variable is the binary decision between cooperation ($q_t = 3$) and defection ($q_t = 4$), while the independent variables are the treatment dummies (with $\Delta = 2$ being the reference treatment), the Match number, the action chosen by the subject in the previous Δ periods (Lcoop and Ldefect), the average price observed in the last Δ periods (Lprice), and interactions between these variables. Results are presented in Table 5, while Figure 3 displays the marginal effect of the observed price (Lprice) on cooperation, depending on the treatment and the action taken by the subject in the previous period.

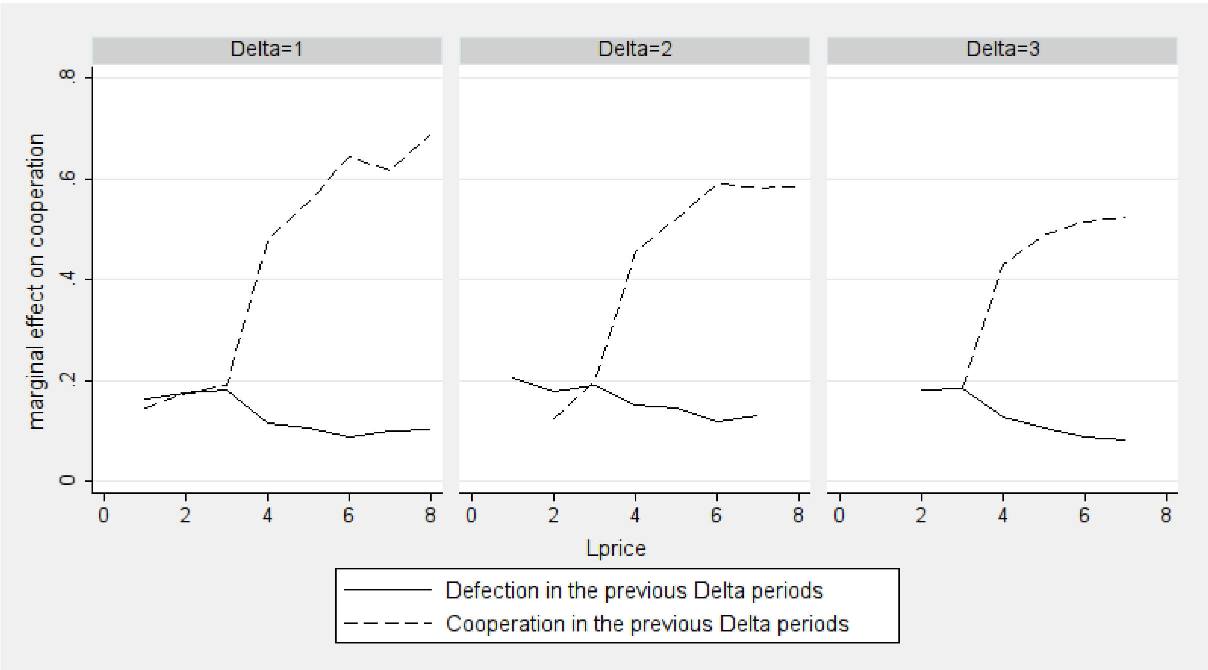

Figure 3: Cooperation rates and prices observed over the previous Δ periods.

Figure 3 reveals that in all treatments subjects who cooperated in the previous Δ periods are more likely to cooperate for the next Δ periods the higher is the observed price in the last Δ periods. Subjects who defected previously exhibit a much weaker

Coop(t=1) x $\Delta = 1$ is not significantly different from 0 (p-value> 0.1).

(negative) reaction to the observed price level. This is in line with the prediction that the cooperative regime requires decisions to be dependent on the price signals, while the defective regime does not.

Results from the regression in Table 5 confirm that there is a positive correlation between a subject's decision to keep cooperating and the price observed in the previous Δ periods, but they also point out an important learning effect. The tendency to react positively to prices when in a cooperative mode increases with experience (Match x Lprice x Lcoop). However, this tendency is significantly weaker when $\Delta = 1$ as shown by the negative coefficient of Match x Lprice) x Lcoop x $\Delta = 1$. This may partly explain why cooperation becomes less profitable over time in this treatment (see Result 2).

Result 3. *When a subject is in the "cooperative mode" (i.e. when he cooperated in the previous Δ period), continued cooperation depends positively on observed prices. This dependency becomes stronger with experience when $\Delta = 2$ and $\Delta = 3$, while it weakens when $\Delta = 1$.*

Table 5: Probit regression with random effects table

	Cooperation in all periods > 1	
Constant	0.273***	(0.089)
Lprice x Lcoop	0.046***	(0.005)
Lprice x Ldefect	−0.031*	(0.018)
$\Delta = 1$	−0.021	(0.099)
Lprice x Lcoop x $\Delta = 1$	0.040***	(0.012)
Lprice x Ldefect x $\Delta = 1$	0.020	(0.022)
$\Delta = 3$	0.073	(0.095)
Lprice x Lcoop x $\Delta = 3$	−0.027	(0.020)
Lprice x Ldefect x $\Delta = 3$	−0.013	(0.019)
Match	−0.026**	(0.012)
Match x		
Lprice x Lcoop	0.006***	(0.002)
Lprice x Ldefect	0.007*	(0.004)
$\Delta = 1$	0.011	(0.014)
Lprice x Lcoop x $\Delta = 1$	−0.010***	(0.003)
Lprice x Ldefect x $\Delta = 1$	−0.008**	(0.004)
$\Delta = 3$	−0.011	(0.015)
Lprice x Lcoop x $\Delta = 3$	0.004	(0.005)
Lprice x Ldefect x $\Delta = 3$	−0.000	(0.004)
Observations	3024	
R2-overall	0.162	
R2-within	0.078	
R2-between	0.862	

Note: Standard errors in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.
Linear probability model with random effects at the subject level. Robust standard errors for data clustered on matching groups.

3 Conclusions

We illustrate the behavioral relevance of the argument, first pointed out by Abreu et al. (1991), that the incentives to cooperate may be eroded by the ability to respond quickly to noisy information about other players' actions, when there is sufficient scope for learning. As the subjects in our experiment gain experience, the rate of cooperation decreases in the treatment with high flexibility. Cooperation rates also fall with experience at very low levels of flexibility, illustrating that the gains from defection should not go unpunished for too long either. So we indeed end up observing a non-linear relationship between flexibility and cooperation as predicted by theory under imperfect monitoring (Sannikov and Skrzypacz, 2007). Consistent with these results, we find that in the intermediate flexibility treatment where collusion is an equilibrium, collusion becomes more remunerative over time (while this is not the case when collusion is not an equilibrium), and that subjects appear to react to noisy price information much in line with the trigger strategies suggested by theory.

Several avenues for further experimental research suggest themselves. Our design varies decision flexibility, while keeping the rate of information feedback constant. An alternative design would be to keep action flexibility fixed, while varying the frequency of information arrival. It would be interesting to see whether such an alternative design would produce similar results, in the sense that cooperation is hindered both when information arrives very frequently and when the information lag is large (Abreu et al., 1991). Another variation would be to implement different types of noisy information signals. Whereas Sannikov and Skrzypacz (2007) focus on information which arrives continuously without shocks, another relevant environment is where signals arrive discontinuously (at a Poisson rate). A testable prediction is that the impact of flexibility on cooperation will depend on the type of signal (Abreu et al., 1991). High flexibility is more harmful if the arrival rate of a signal is increasing in the rate of cooperation (the 'good news' case) than in case the arrival rate of a signal is decreasing in the rate of cooperation (the 'bad news' case). It would be important to examine the behavioral support for this prediction as it will further our understanding of the properties of optimal monitoring and enforcement schemes.

Appendices

A Theoretical analysis

In this Appendix we outline how the non-monotonic effect of flexibility (Δ) on the sustainability of collusion is derived. The sustainability of collusion depends on the punishment strategy adopted by the players. Nonetheless, the authors prove that it is possible to compute a robust lower bound by finding the best symmetric equilibrium with Nash reversion as a punishment, and a robust upper bound by finding the best symmetric equilibrium with the minimax payoff of 0 as a punishment. Both bounds are valid for both symmetric and asymmetric equilibria. With our parameters, the upper and the lower bound coincide, as the Nash equilibrium profit coincides with the minimax payoff in our game (both are equal to zero).

We can follow Sannikov and Skrzypacz (2007) to show that in our set up collusion ($q_i = 3$) can be sustained when $\Delta = 2$, but not when $\Delta = 1$ or $\Delta = 3$. From Abreu *et al.* (1986) we know that the best strongly-symmetric equilibrium payoff of this game can be achieved by the following strategy profile:

- Players start in the collusive state and choose quantities q_C, q_C (for us it will be $(3, 3)$).
- As long as the realized price is in region P_+, players remain in the collusive state. If the price is outside this region, they move to the punishment state forever after.
- Because in our game mini-max has the same payoffs as the static Nash equilibrium, in an optimal equilibrium once the players reach the punishment state they play $(4, 4)$ forever.

We now characterize the region P_+ and it's complement P_-. Let $G(Q)$ be the probability that the price will be in P_+, and V the expected profit of the collusive equilibrium. Each player's IC constraint is:

$$\pi(q_D, q_C)(1-\delta) + \delta\{V * G(q_D + q_C) + 0 * [1 - G(q_D + q_C)]\} \leq \\ \pi(q_C, q_C)(1-\delta) + \delta\{V * G(2q_C) + 0 * [1 - G(2q_C)]\}, \quad (1)$$

which can be re-written as:

$$\delta V\left[G\left(2q_C\right)-G\left(q_D+q_C\right)\right]-(1-\delta)\left[\pi\left(q_D,q_C\right)-\pi\left(q_C,q_C\right)\right]\geq 0. \tag{2}$$

If the IC constraints are satisfied, then the expected profit in this equilibrium is:

$$V=(1-\delta)\,\pi\left(q_C,q_C\right)+\delta\left[VG\left(2q_C\right)+0*\left(1-G\left(2q_C\right)\right)\right],$$

which yields:

$$V=\pi\left(q_C,q_C\right)\frac{1-\delta}{1-\delta G\left(2q_C\right)}.$$

Note that V is decreasing in δ and increasing in $G(2q_C)$.

Sannikov and Skrzypacz (2007) show that the optimal P_+ region (that maximizes V) corresponds is a tail test. There is a cutoff $\hat{p}$ such that above $\hat{p}$ are in P_+ and prices below are in P_-.

If a tail test is adopted, then

$$G(Q)=\int_{\hat{p}}^{\infty}\phi\left[p(Q),\frac{\sigma^2}{\Delta},p\right]dp,$$

where $\phi(\mu,\sigma^2,x)$ is the probability density function of a normal distribution with mean μ and variance σ^2, evaluated at x. Using the parametrization in our experiment, with $p(Q)=12-q_1-q_2$ and $\sigma=1.3$, we can rewrite the IC-constraint as a function of $\hat{p}$ and calculate when it can be satisfied at different levels of Δ.

Numerical calculations show that the left-hand side of the IC-constraint (2) is convex, and that when $\Delta=2$ it is positive for cutoff prices $\hat{p}\in[4.758,5.060]$, while it takes negative values for any $\hat{p}\geq 0$ when $\Delta=1$ or $\Delta=3$.

This implies that in the infinite horizon Cournot duopoly game with imperfect public monitoring collusion is sustainable in equilibrium when $\Delta=2$, while no collusive equilibrium is sustainable when $\Delta=1$ or $\Delta=3$.

B Instructions

(for on-line publication)

In this Appendix, we report the experimental instructions for the treatment with $\Delta = 3$. Instructions for the other two treatments only change where strictly necessary, and are available from the authors upon request.

Welcome to our experiment. Please follow the instructions carefully. During the experiment your earnings are denoted in points. At the beginning of the experiment you will receive an initial endowment of 80 points. In addition, you will make decisions that can make you earn or lose points. The number of points you earn depends on your decisions, the decisions of other participants, and chance. At the end, we will exchange your points into Euro according a conversion rate of 1 point = 12.5 Eurocent, which means that 8 points = 1 Euro. You will receive your payment privately at the end of the experiment. We guarantee anonymity with respect to other participants and we do not record any information connecting your name to your decisions or earnings.

Please be quiet during the entire experiment and do not talk to your neighbors. If you have a question please raise your hand and you will be answered privately.

Your task

Production:

You will make decisions for a firm in this experiment. For a number of periods you have to determine the quantity that your firm will produce. You can decide to produce a low level of 3 units or a high level of 4 units. Your firm operates in a market with one other firm. In each period, your profits (in points) will depend on the number of units you produce and the number of units produced by other firm. The decisions for this firm will be made by another participant. You cannot know who this participant is, nor can this participant know who you are. We will refer to this other participant as “the other firm”. We will now explain how your profits depend on the number of units you produce and the number of units the other firm produces.

Costs:

Production involves costs. Every period, you have to pay a fixed cost of 16 points. These costs are independent of whether you produce 3 units or 4 units.

Price:

The market price in a period is the same for your firm and the other firm. The market price depends on the total production in a period. The total production is the sum of the number of units you produce, and the number of units produced by the other firm. The larger total production, the lower the market price. The expected market price is as follows:

Expected price = 12−(number of units you produce)−(number of units other firm produces)

For convenience the following table summarizes how the expected market price depends on the number of units produced by your firm and the other firm.

Expected price		Production of other firm	
		3 units	4 units
Your production	3 units	6	5
	4 units	5	4

Profit:

Each period, your profits are equal to your revenue minus your cost, where your revenue is equal to the number of units you produce multiplied by the market price. Hence, your profit is:

*Expected profit = Expected price * (number of units you produce) − 16*

Recall that the price depends on your production and the production of the other firm. For convenience, the table below calculates how your expected profit and the expected profit of the other firm depend on your production and the production of the other firm. The first entry in each cell represents your profit, while the second entry (in gray) represents the profit of the other firm.

Expected profits		Production of other firm	
		3 units	4 units
Your production	3 units	2, 2	-1, 4
	4 units	4, -1	0, 0

For example, you can read in the table that if in a period you produce 3 units and the other firm produces 3 units, your expected profit will be equal to 2. You can check this as follows:

- Expected price $= 12 - 3 - 3 = 6$
- Expected profit $= 6 * 3 - 16 = 2$

You can also read in the table that if you produce 4 units and the other firm produces 4 units, your expected profit will be equal to 0. You can check this as follows:

- Expected price $= 12 - 4 - 4 = 4$
- Expected profit $= 4 * 4 - 16 = 0$

Note that profit can be negative. In the unlikely event, that the total amount of points you earn in the experiment is lower than 0, you will not receive any money, but you will not have to pay any money either.

Price shocks: You may have noted that until now, we have talked about the expected price and expected profits. Due to unobservable variations in demand, the market price in a period is affected by a random shock. Specifically, the market price is the expected price plus the shock:

Price = 12−(number of units you produce)−(number of units other firm produces)+ shock

The price shock in one period is independent of the price shock in another period. The shock in each period is normally distributed with a mean of zero and a standard deviation of 1.3. This means that the shock is equally likely to be positive or negative. The probability that the shock attains a value in a certain range is summarized in the following table.

Range of shock values	below -1	-1 to 0	0 to 1	above 1
Probability	22%	28%	28%	22%

Since the mean value of the shock is zero, the expected price and the expected profit depend on the number of units produced by you and the other firm, as indicated in the tables above. The actual price and the actual profit, however, will differ as a result of the shock. For example, if your firm produces 3 units, and the other firm produces 3 units, the price will be equal to 12-3-3+shock = 6+shock, which means that the price will be

- below 5 with probability 22%
- between 5 and 6 with probability 28%
- between 6 and 7 with probability 28%
- above 7 with probability 22%.

Now suppose the actual price shock is -0.5. Then the actual price will be 6 -0.5= 5.5 and your actual profit will be 5.5 * 3 $-$ 16 = 0.5, while your expected profit was 2. Therefore, the price depends on the number of units produced by you, the number of units produced by the other firm, and the shock as follows:

Price

		Production of other firm	
		3 units	4 units
Your production	3 units	6 + shock	5 + shock
	4 units	5 + shock	4 + shock

By reducing the price a negative price shock also reduces your revenue and your profits. Conversely, a positive price shock increases your revenue and your profit. Your profits will then be:

Profit = *(expected price + shock) * (number of units you produce)* $-$ *16*

It is important to realize that you have no influence whatsoever on the price shock. It is truly random. The number of units produced by you and the other firm affect the expected price, which will be higher the lower is the total production. But the actual price is also affected by the random price shock.

Periods and markets

- You will be randomly paired to another participant for a sequence of periods, referred to as a market. This other participant will make the decisions for the other firm.
- During the whole experiment you will participate in a total of 7 markets.
- In these 7 markets you will be paired to another participant at most once.
- Every 3 periods you will have to decide how many units your firm produces in each period. This means that you will not be able to change the number of units you produce every period, but only once every 3 periods. The same holds for the other producer.
- How many periods a market will last is randomly determined. Each time three periods have been completed, the computer will randomly draw a number between 1 and 100. If the number is below or equal to 74, the market will continue for another three periods. Hence, the probability that the market continues with the same participant for at least three more periods is 74%. If the number is above 74, a new market will start in which you will be randomly paired to another participant; unless you have already participated in 7 markets in which case the experiment will end.

Information

At the end of each period you will be informed about the number of units you produced, the price and your profits. For the periods in which you do not make a decision, this information is shown only shortly. After every block of three periods, you will also receive information on the average price and your average profits for the last three periods. Information from all previous periods is presented in the so-called History Table in lower part of your screen.

It is important to note that you do not receive information on the number of units produced by the other firm. You do get information on the market price, but because of the random price shock you cannot infer exactly how many units the other firm produced, nor how much profit the other firm made. Still, the price does give you some imperfect indication about the number of units produced by the other firm.

On the top left of the screen you can see how many points you have earned until now in the current market, and in the top right you can see how many points you have earned during the whole experiment, including the initial endowment of 80 points.

Summary

1. You decide how many units you wish to produce in the next three periods.
2. The number of units you produce, the number of units the other firm produces, and the price shock determine your profit in a period.
3. You are paired to one other participant for a sequence of periods, called a market.
4. After each block of three periods, there is a probability of 74% that you remain paired to the same participant for another three periods and a probability of 26% that the present market ends.
5. If a market ends you will be randomly paired to another participant and new market will start, until you have participated in 7 markets in total.
6. The total profits you accumulate over all markets, together with the starting endowment of 80 points determine your earnings for the experiment. 8 points will be converted into 1 Euro.

Procedure and questions

You are now given some time to study the instructions on your own and to ask clarifying questions (if any). After that, you will be asked to answer a few control questions to check your understanding. The first market will start as soon as all the participants have correctly answered the control questions.

Please be reminded that you are not allowed to talk or communicate to other participants during the experiment. If you have a question, please raise your hand and I will come to your table.

C Additional regressions

(not for publication)

Table C.1: Panel regression for rates of cooperation

	Cooperation first period		Cooperation all periods	
Match	0.001	(0.011)	-0.008	(0.014)
$\Delta = 1$	0.295^{*}	(0.152)	0.134	(0.092)
$\Delta = 3$	0.076	(0.109)	0.035	(0.071)
Match x $\Delta = 1$	-0.051^{*}	(0.031)	-0.033^{*}	(0.019)
Match x $\Delta = 3$	-0.046^{***}	(0.017)	-0.022	(0.015)
Period			-0.002	(0.002)
Constant	0.286^{***}	(0.042)	0.306^{***}	(0.081)
Observations	672		6272	
R2-overall	0.061		0.024	
R2-within	0.052		0.025	
R2-between	0.073		0.018	

Standard errors robust for clustering at the session level are reported in parentheses.
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table C.2: Logit regression for rates of cooperation

	Cooperation first period		Cooperation all periods	
coop				
Match	0.008	(0.104)	−0.056	(0.086)
$\Delta = 1$	2.019**	(1.020)	0.956	(0.648)
$\Delta = 3$	0.777	(0.825)	0.352	(0.559)
Match x $\Delta = 1$	−0.368**	(0.183)	−0.211*	(0.109)
Match x $\Delta = 3$	−0.492***	(0.132)	−0.201*	(0.106)
Period			−0.018*	(0.010)
Constant	−1.399**	(0.690)	−1.268**	(0.619)
Subj1				
Constant	1.932***	(0.314)	1.662***	(0.224)
Observations	672		6272	
ll	−322.557		−2700.272	

Standard errors in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Logit regression with random effects.

Robust standard errors for data clustered on matching groups.

Table C.3: Panel regression for rates of cooperation

	without period trend		different trends	
Match	-0.005	(0.011)	-0.007	(0.012)
$\Delta = 1$	0.160^{*}	(0.083)	0.186^{*}	(0.102)
$\Delta = 3$	0.051	(0.072)	0.009	(0.085)
Match x $\Delta = 1$	-0.036^{**}	(0.015)	-0.030^{*}	(0.016)
Match x $\Delta = 3$	-0.024^{**}	(0.012)	-0.022^{*}	(0.013)
Period			-0.002^{*}	(0.001)
Period x $\Delta = 1$			-0.012^{***}	(0.002)
Period x $\Delta = 3$			0.003^{*}	(0.002)
Constant	0.271^{***}	(0.071)	0.302^{***}	(0.081)
Observations	6272		6272	
R2-overall	0.021		0.027	
R2-within	0.023		0.031	
R2-between	0.011		0.012	

Standard errors in parentheses. Linear probability model with random effects at the subject level. Robust standard errors for data clustered on matching groups.
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

References

Abreu, D., P. Milgrom, and D. Pearce (1991), "Information and Timing in Repeated Partnerships", *Econometrica*, 59, 1713-1733.

Abreu, D., D. Pearce, and E. Stacchetti (1986), "Optimal Cartel Equilibria with Imperfect Monitoring", *Journal of Economic Theory*, 39, 251-269.

Ai, C., and E. Norton (2003), "Interaction terms in logit and probit models", *Economics Letters*, 80, 123-129.

Aoyagi, M., and G. Fréchette (2009), "Collusion as Public Monitoring Becomes Noisy: Experimental Evidence," *Journal of Economic Theory*, 144, 1135-1165.

Axelrod, R. (1984), *The Evolution of Cooperation*, New York: Basic.

Belleflamme, P., and M. Peitz, (2010), *Industrial Organization: Markets and Strategies*, Cambridge: Cambridge University Press.

Bereby-Meyer, Y., and A. Roth, (2006), "The Speed of Learning in Noisy Games: Partial Reinforcement and the Sustainability of Cooperation", *The American Economic Review*, 96, 1029-1042.

Camerer, C., and K. Weigelt (1988), "Experimental Tests of a Sequential Equilibrium Reputation Model," *Econometrica*, 56, 1-36

Cason, T., and F. Khan (1999), "A Laboratory Study of Voluntary Public Goods Provision with Imperfect Monitoring and Communication," *Journal of Development Economics*, 58, 533-552.

Church, J., and R. Ware (2000), *Industrial Organization. A Strategic Approach*, McGrawHill.

Dal Bò, P., and G. Fréchette (2011), "The Evolution of Cooperation in Infinitely Repeated Games: Experimental Evidence", *American Economic Review*, 101, 411-429.

Feinberg, R., and C. Snyder (2002), "Collusion with secret price cuts: An experimental investigation," *Economics Bulletin*, 3, 1-11.

Fischbacher, U. (2007), “z-Tree: Zurich Toolbox for Ready-made Economic Experiments”, *Experimental Economics*, 10, 171-178.

Friedman, D., and R. Oprea (forthcoming), “A Continuous Dilemma”, *American Economic Review.*

Fudenberg, D., D. Rand, and A. Dreber (forthcoming), “Slow to Anger and Fast to Forgive: Cooperation in an Uncertain World”, *American Economic Review.*

Holcomb, J.H., and P.S. Nelson (1997), “The Role of Monitoring in Duopoly Market Outcomes”, *Journal of Socio-Economics*, 26(1), 79-93.

Ivaldi, M., B. Jullien, P. Rey, P. Seabright, and J. Tirole (2003), “The Economics of Tacit Collusion,” *Final Report to DG Competition, European Commission*, av. from *http://ec.europa.eu/competition/mergers/studies_reports/ the_economics_of_tacit_collusion_en.pdf*

Martin, S. (2001), *Industrial Organization in Context*, Oxford: Oxford University Press.

Motta, M. (2004), *Competition Policy: Theory and Practice,* Cambridge: Cambridge University Press.

Roth, A. & Erev, I. (1995), “Learning in extensive-form games: Experimental data and simple dynamic models in the intermediate term”, *Games and Economic Behavior* **8**(1), 164--212.

Sannikov, Y., and A. Skrzypacz (2007), “Impossibility of Collusion under Imperfect Monitoring with Flexible Production,” *American Economic Review*, 97, 1794-1823.

Selten, R., and R. Stoecker (1986), “End Behavior in Sequences of Finite Prisoner’s Dilemma Supergames: A Learning Theory Approach,” *Journal of Economic Behavior and Organization*, 7, 47-70.

Tirole, J. (1988), *The Theory of Industrial Organization.* Cambridge, Massachusetts: MIT Press.

Vives, X. (1999) *Oligopoly Pricing. Old ideas and new tools.* Cambridge, Massachusetts: MIT Press.