Flexibility and Cooperation with Imperfect Monitoring*

Maria Bigoni Jan Potters Giancarlo Spagnolo

March 15, 2011

Preliminary version: please do not circulate.

Abstract

Flexibility - the ability to react swiftly to others' choices - facilitates cooperation by reducing the gains a deviating party can secure before opponents react. Under imperfect monitoring however, flexibility may also hinder cooperation by inducing players to react too early, after observing too few and noisy signals. The interaction between these forces predicts a non-monotonic relationship between flexibility and collusion. We implement in the laboratory an indefinitely repeated 2x2 Cournot game in discrete time where players only observe noisy price information. Across treatments we vary the number of periods firms have to wait before changing their output, to test if the predicted non-linear relationship can be observed with real world subjects.

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

JEL: C73, C92, D43, L13, L14

*We thank Andrzej Skrzypacz for his precious advice, and Gabriele Camera, Tobi Klein, and participants to the 2009 European ESA Meeting in Innsbruck and to seminar at the Stockholm School of Economics for valuable comments.

Flexibility, agents' ability to react swiftly to others' choices, is widely regarded as a factor that facilitates cooperation. This "common wisdom" derives from the intuitive argument that a faster reaction reduces the gains a unilateral defection generates anticipating its punishment.[1] Axelrod, for example, writes:

> "There are two basic ways of doing this [enlarging the shadow of the future]: by making the interactions more durable, and by making them more frequent." (The Evolution of Cooperation, 1984, p. 129).

However, a positive monotonic relationship between flexibility and cooperation is only theoretically justified under complete information. In a pioneering study Abreu, Milgrom and Pearce (1991) showed that in repeated games with imperfect monitoring an increase in the frequency of interaction is not theoretically equivalent to an increase in the intertemporal discount factor and may make cooperation harder to sustain. In the same vein, Sannikov and Skrzypacz (2007) showed that under imperfect monitoring high levels of flexibility make collusion impossible in a wide range of situations, from repeated Cournot duopolies to non-stationary competitive environments. The reason is that with imperfect monitoring high flexibility makes deviations harder to detect without mistakes and leads agents to react too early to still too noisy 'bad news'. These frequent mistakes trigger punishments too often eroding the value of cooperation or collusion. This negative effect counteracts the positive effect of reduced gains from defections discussed above, and at sufficiently high flexibility dominates it leading cooperation to unravel.[2]

Starting from sufficiently low flexibility levels, under imperfect monitoring we should then expect cooperation to first increase and then decrease when flexibility grows. The channel through which flexibility hinders cooperation under imperfect monitoring is subtle however. Certainly it is less easy to grasp than the more intuitive positive effect discussed at the beginning. One may therefore wonder how strong this negative effect of flexibility is in reality in comparison to the positive effect. Does collusion really unravel with sufficient flexibility when monitoring is imperfect? Can one really observe a non-monotonic effect of flexibility on collusion in this case?

This paper presents an experimental study designed to answer these questions. We simulated in the laboratory an indefinitely repeated oligopoly game in discrete time with imperfect monitoring analogous to that studied by Sannikov and Skrzypacz (2007) and tested whether a negative effect of increased flexibility on collusion and a non-linear relationship between the two could be observed with human subjects.

While the first effect in Axelrod's quote has been documented by many laboratory experiments, much less attention has been given to the second.[3] A recent experiment by Friedman and Oprea (2010) filled this gap, offering striking evidence on the positive effect of flexibility on cooperation. They implemented a 60-seconds finite horizon repeated Prisoner's Dilemma under perfect monitoring varying the speed at which subjects could adjust their actions. Their results show a strong positive monotonic relationship

[1]Based on this intuition academics and practitioners regularly argue that cooperation and collusion are more likely where interaction is more frequent; see for example Ivaldi *et al.* (2003) and DoJ's "a Primer on Collusion". Shapiro (1983) is the first to model this effect of flexibility on reputation for product quality.

[2]See also Fudenberg and Levine (2009), and Sannikov and Skrzypacz (forth.)

[3]Experimental studies showing that "making the interactions more durable" (i.e. increasing the continuation proability) increases cooperation rates in repeated games include Roth and Murnighan (1978, 1982), Dal Bò (2005), Blonski, Ockenfels and Spagnolo (forthcoming), and Dal Bò and Frechétte (forthcoming), among many others.

between action flexibility and the rate of cooperation.[4] On the other hand, a number of recent experimental studies looked at the effects of imperfect monitoring on cooperation, without considering flexibility. Bereby-Meyer and Roth (2006) studied 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 Frechétte (2009) asked whether subjects' ability to cooperate falls when information becomes more noisy in a repeated game, finding significant support for this theoretical prediction. Fudenberg, Rand and Dreber (2010) looked 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] Our study is at the intersection of these two streams of literature. Inspired by Abreu, Milgrom and Pearce (1991) and Sannikov and Skrzypacz (2007) we examine whether the positive effect of flexibility on cooperation observed under perfect monitoring needs to be qualified when monitoring is imperfect.

1 Experimental Design

The design of our experiment implements the stationary repeated Cournot game analyzed by Sannikov and Skrzypacz (2007) [S&S]. We chose specifications and parameters so as to allow, in the simplest possible setting, for S&S's theoretical results.

The Game

Two players interact repeatedly in a Cournot market with homogeneous products. Time is discrete, the time horizon is indefinite and players discount future profits with a common interest rate r. Players 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:

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

Quantities can be adjusted only every Δ periods. A larger Δ implies that agents cannot always react immediately to a signal, but also that when they can react they will have observed more signals. Prices and profits materialize in every period, so the quantities chosen in period t, together with the random shocks, determine prices and profits for the following Δ periods. As a consequence, a player will discount future profits with a factor $\delta = e^{-r\Delta}$.

We set $a = 12$ in the demand function and let the noise be generated with $\sigma = 1.3$. We choose marginal cost equal to 0, and per period fixed cost equal to 16. Finally, we restrict the players' actions set to $q_{it} \in \{3, 4\}$, i.e. players can only choose between the fully collusive output, and the Cournot-Nash equilibrium output. These parameters ensure that cooperation is sustainable between players maximizing expected monetary gains only for $\Delta = 2$, as shown in Appendix A.

[4] 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.

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

		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

Δ	$\delta = e^{-r\Delta}$	*exp. # of periods*
1	0.90	10
2	0.82	11.2
3	0.74	11.4

Table 2: Treatments

Table 1 presents expected price and profits of the stage game for every possible combination of actions. Note that the expected profits resemble a Prisoner's Dilemma.

Treatments and predictions

In the experiment we implement three different treatments in which Δ takes the values 1, 2, and 3, respectively. As a consequence, both the theoretical discount factor $\delta = e^{-r\Delta}$ and the number of price signals observed after each action vary across treatments. To implement the three different treatments, we need to induce the corresponding discount factors in the lab. We do this by letting subjects play a repeated game of indeterminate length, with a continuation probability equal to the theoretical discount factor δ. Table 2 reports the values of Δ in the three treatments, the corresponding values of δ (recall that $r = 0.10$), and the number of periods the repeated game is expected to last ($\frac{1}{1-\delta}$).

A smaller value of Δ implies a larger discount factor, 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 his 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. The best players can do to support collusion is to use grim trigger strategies in which they collude as long as the price signals they receive are above a certain threshold value. This generates an incentive compatibility constraint which cannot be satisfied for small values of Δ, nor for large values 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. Applied to our game, collusion is sustainable when $\Delta = 2$, but not when $\Delta = 1$ or $\Delta = 3$. It is this remarkable theoretical prediction that we aim to test in our experiment. In Appendix A we outline how this result can be derived.

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. Sessions lasted on average one hour and 45 minutes, including instructions and payment, and subjects received an average payment of 17 Euros.

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.

Subjects were randomly matched to one other subject to play the repeated game. For each period, they 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 prices in each period were determined by total quantity and a random shock which was drawn independently for each period from a normal distribution with zero mean and standard deviation $\sigma = 1.3$. So, prices were only a noisy signal of the other subject's quantity choice. 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 a session. Rematching took place exactly 6 times. So, each subject played the indefinitely repeated game exactly 7 times, and this was common knowledge. For the rematching we adopted an absolute stranger protocol meaning that two subjects never met more than once.

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

As discussed in the introduction, in an environment characterized by imperfect monitoring flexibility can have two opposing effects on the sustainability of collusion. On the one hand, flexibility reduces the benefits from defection; on the other hand, it makes it harder to detect defections and punish them, and may trigger punishment even when no deviations took place. In our parametrization, the positive effect is theoretically expected to dominate only at intermediate levels of flexibility, making cooperation sustainable when $\Delta = 2$, but not when $\Delta = 1$ or $\Delta = 3$.

A look at aggregate cooperation frequencies across treatments, reported in Table 3, suggests at first hand that this theoretical prediction is not borne out by our experimental data: there are no significant differences in aggregate cooperation frequencies across the three treatments.[6]

Table 3: 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, forthcoming). 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 con confirmed by the experimental data.

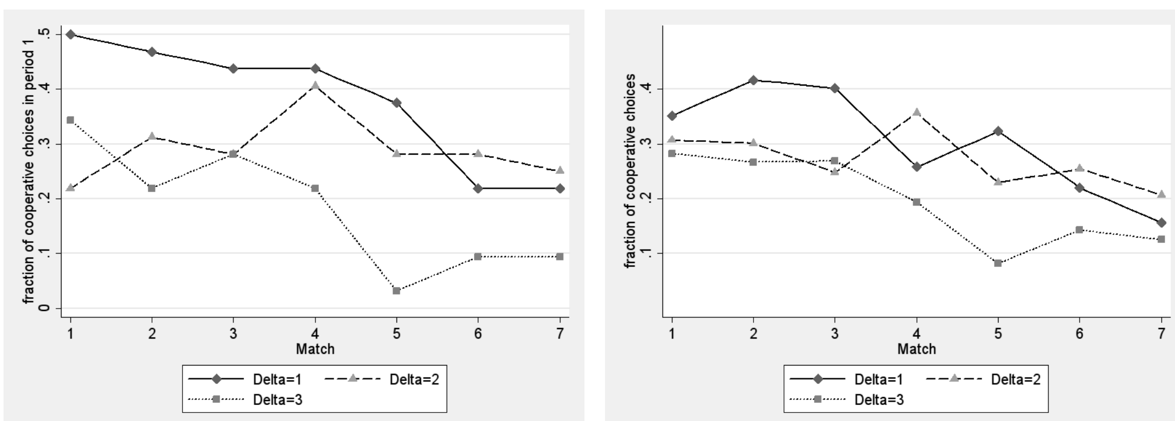

(a) First period.

(b) All periods.

Figure 1: Frequency of cooperation by match.

Figure 1 presents the development of cooperation over the matches for each of the three treatments. The left panel displays the average rates of cooperation in the first period of each match; the right panel gives average rates of cooperation across all periods of a match. In both figures we see that in the early matches cooperation is most frequent in treatment $\Delta = 1$. The decline in cooperation over the matches, however, is more

[6] 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).

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 $\Delta = 3$.[7]

The decline in cooperation rates over matches in treatments $\Delta = 1$ and $\Delta = 3$ is confirmed by the regressions in Table 4. 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, and interactions between the two.[8]

Table 4: 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	
R-squared within	0.052		0.025	
R-squared between	0.073		0.018	
R-squared overall	0.061		0.024	

Notes: 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.

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. 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 will examine how profitable cooperation/collusion 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 an environment in

[7] 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 $\Delta = 2$ and $\Delta = 3$ is not significant though. This holds both for cooperation rates in the first period and for cooperation rates across all periods in a match. The same results is obtained if we look at cooperation rates in only the last two matches.

[8] In this and in the following regressions, standard errors are computed clustering on matching groups. Clustering on sessions would produce qualitatively similar results.

which it is not, and this difference should become more pronounced as subjects gain experience. Figure 2 presents the difference between the average per period profit in a match for subjects who cooperated in the first period of a match and subjects who defected in the first period of a match. It shows that in the early matches cooperation is profitable only in the $\Delta = 1$ treatment. Over time cooperation tends to become more profitable in $\Delta = 2$ while such a trend is not observed for the other two treatments.

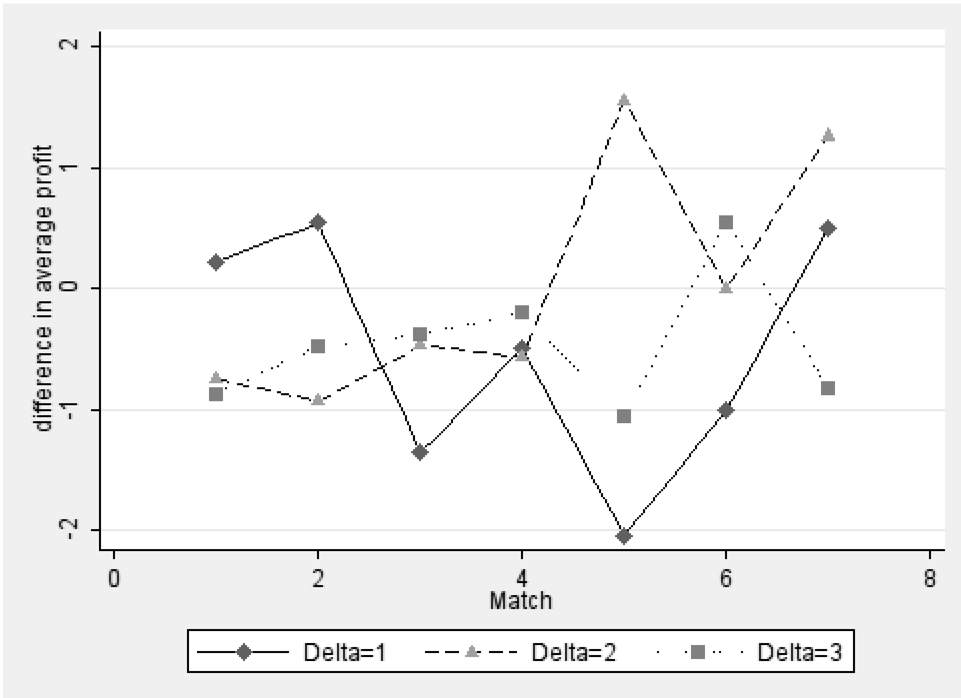

Figure 2: Difference in average profit for first-period cooperators and defectors.

To make these results more precise, Table 5 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. 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.

Results indicate that for treatment $\Delta = 2$ there is a negative relationship between a subject's cooperation rate and his average profits in the early matches, but that there is a strongly significant increase in the profitability of cooperation over the matches. The other two treatments show the reverse trend. The profitability of cooperation decreases over the matches in $\Delta = 1$ and $\Delta = 3$.[9]

Result 2. *The profitability of cooperation increases as subjects gain experience in treatment* $\Delta = 2$*, while it decreases in treatments* $\Delta = 1$ *and* $\Delta = 3$*.*

According to Sannikov and Skrzypacz (2007), when collusion is an equilibrium – as in treatment $\Delta = 2$ of our experiment – it could be supported by a trigger strategy

[9]The estimated coefficient of Match x Coop(t=1) for the reference treatment $\Delta = 2$ is significantly larger ($p < .01$) than those for Match x Coop(t=1) x $\Delta = 1$ and Match x Coop(t-1) x $\Delta = 3$

Table 5: Panel regression of average profits

	Average period profit	
Constant	0.960***	0.346
Coop(t=1)	−1.398***	0.141
$\Delta = 1$	0.271	0.917
$\Delta = 3$	0.312	0.805
Coop(t=1) x $\Delta = 1$	1.297	1.305
Coop(t=1) x $\Delta = 3$	0.713	1.086
Match	−0.135**	0.058
Match x		
Coop(t=1)	0.353***	0.067
$\Delta = 1$	0.016	0.161
$\Delta = 3$	0.077	0.141
Coop(t=1) x $\Delta = 1$	−0.458*	0.271
Coop(t=1) x $\Delta = 3$	−0.298	0.214
Observations	672	
Log-likelihood	−1715.225	

Notes: Standard errors in the second column. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.
Panel regression with random effects at the subject level. Robust standard errors for data clustered on matching groups.

prescribing to collude at the beginning of the game, and to switch to Cournot-Nash behavior when the observed price falls below a cut-off level.

To test for this prediction, 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, the action chosen by the subject in the previous period (cooperation or defection), the average price observed in the last Δ periods, and interactions between these variables. Results are presented in Table 6, while Figure 3 displays the marginal effect of the observed price (price(t-1)) on cooperation, depending on the treatment and the action taken by the subject in the previous period. It shows that, in all treatments, subjects who cooperated in the previous period tend to cooperate more often when they observe a high price, while subjects who defected in the previous period exhibit a weakly negative reaction to the observed level of price.

Results from the regression in Table 6 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 stress an important learning effect. When $\Delta = 2$ the tendency to react positively to prices when subjects are in a cooperative mode increases with experience, and subjects' strategy becomes somewhat more lenient, in the sense that the

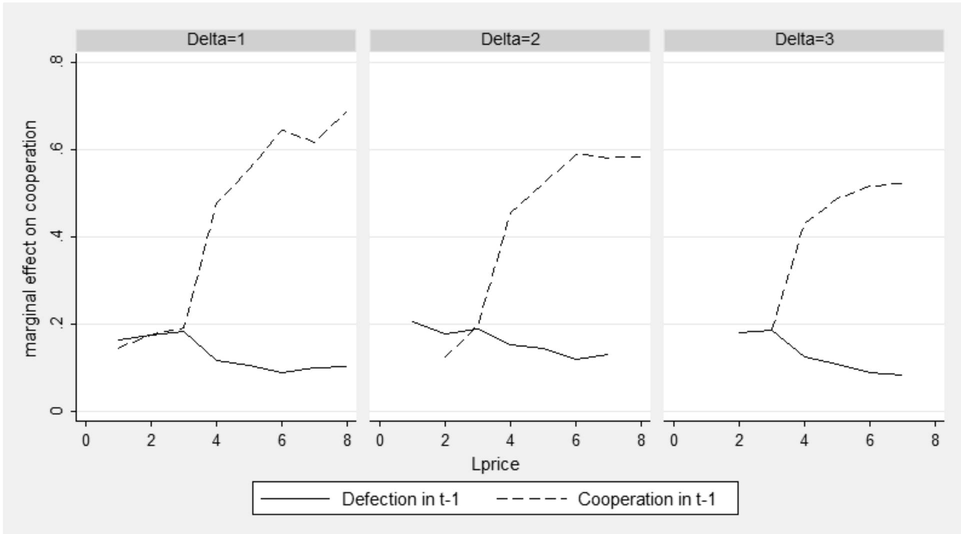

Figure 3: Cooperation rates and prices observed over the previous Δ periods.

negative reaction to prices in the "defective mode" abates in later matches.[10] In contrast, when $\Delta = 1$ the positive effect of observed prices on cooperation tends to disappear and the negative effect of prices on cooperation when subjects are in a defective mode is amplified.

Result 3. *When a subject is in "cooperative mode" (i.e. when he cooperated in the previous period), he tends to cooperate more when he observes a high price, and to cooperate less when he observes a low price.*

[10]This seems to be in line with the results of Fudenberg et al. (2010).

Table 6: Panel regression with random effects.

	Cooperation in all periods > 1	
Constant	0.273^{***}	0.089
price(t-1) x coop. in $t-1$	0.046^{***}	0.005
price(t-1) x defect. in $t-1$	-0.031^{*}	0.018
$\Delta = 1$	-0.021	0.099
price(t-1) x coop. in $t-1$ x $\Delta = 1$	0.040^{***}	0.012
price(t-1) x defect. in $t-1$ x $\Delta = 1$	0.020	0.022
$\Delta = 3$	0.073	0.095
price(t-1) x coop. in $t-1$ x $\Delta = 3$	-0.027	0.020
price(t-1) x defect. in $t-1$ x $\Delta = 3$	-0.013	0.019
Match	-0.026^{**}	0.012
price(t-1) x coop. in $t-1$	0.006^{***}	0.002
price(t-1) x defect. in $t-1$	0.007^{*}	0.004
$\Delta = 1$	0.011	0.014
price(t-1) x coop. in $t-1$ x $\Delta = 1$	-0.010^{***}	0.003
price(t-1) x defect. in $t-1$ x $\Delta = 1$	-0.008^{**}	0.004
$\Delta = 3$	-0.011	0.015
price(t-1) x coop. in $t-1$ x $\Delta = 3$	0.004	0.005
price(t-1) x defect. in $t-1$ x $\Delta = 3$	-0.000	0.004
Observations	3024	
R-squared within	0.078	
R-squared between	0.862	
R-squared overall	0.162	

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

3 Conclusions

Understanding the relationship between flexibility and cooperation is important both theoretically and from a policy perspective. The ‘common wisdom’ follows perfect information models showing that flexibility unambiguously stabilizes cooperation.[11] Abreu, Milgrom and Pierce (1991) and Sannikov and Skrzypacz (2007) suggest however that when monitoring is imperfect more frequent interaction may harm cooperation facilitating deviations and increasing mistakes and the frequency of costly punishments. The interaction between these two forces predicts a non-linear relationship between flexibility under imperfect monitoring. The forces behind the negative effect are subtle however, so one may wonder whether they actually bite in reality. In this paper we examined the relationship between flexibility and collusion under imperfect monitoring in a controlled laboratory environment to understand if these subtle theoretical forces are actually relevant for real world subjects.

Our results support the theoretical prediction once the subjects acquire sufficient experience of play, indicating that policy makers should take these theoretical effects seriously. We find that in treatments with very high flexibility the rate of cooperation sharply decreases, the more the further subjects gain experience. Cooperation rates fall also at very low level of flexibility because gains from defection increase, so we indeed observe the non linear relationship between flexibility and cooperation predicted by theory under imperfect monitoring. Consistent with these main results, we also find that in the intermediate flexibility treatments where collusion is an equilibrium subjects colluding more often earn more on average (the opposite happens when collusion is not an equilibrium); and that subjects appear to react to noisy price observations consistently with the strategies suggested by theory.

[11] As Axelrod writes: “Another 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.” (The Evolution of Cooperation, 1984).

Appendix 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:

$$\delta V(G(2q_C) - G(q_D + q_C)) - (1-\delta)(\pi(q_D, q_C) - \pi(q_C, q_C)) \geq 0. \quad (2)$$

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

$$V = (1-\delta)\pi(q_C, q_C) + \delta V G(2q_C) + 0 * (1 - G(2q_C))$$

, which yields:

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

.

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

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

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 $\Delta = 3$.

Appendix B Instructions

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 (you, other firm)

		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.

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.

Admati, A. & Perry, M. (1991), “Joint Projects without Commitment”, *The Review of Economic Studies* **58**(2), 259--276.

Aoyagy, M., and G. Frechétte (2009): “Collusion as Public Monitoring Becomes Noisy: Experimental Evidence,” *Journal of Economic Theory*, 144, 1135-1165.

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

Blonski, M., Ockenfels, P., and Spagnolo, G. (forthcoming); “Equilibrium Selection in the Repeated Prisoner’s Dilemma: Axiomatic Approach and Experimental Evidence”, *American Economic Journal: Microeconomics*.

Camerer, C., and Weigelt, K. (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.

Dal Bò, P. (2005), “Cooperation under the Shadow of the Future: Experimental Evidence from Infinitely Repeated Games”, *American Economic Review* **95**(5), 1591--1604.

Dal Bò, P. & Fréchette, G. (forthcoming), “The Evolution of Cooperation in Infinitely Repeated Games: Experimental Evidence”, *American Economic Review*.

Duffy, J. and J. Ochs (2009): “Cooperative Behavior and the Frequency of Social Interaction,” *Games and Economic Behavior*, 66, 785-812.

DoJ, “Price Fixing and Bid Rigging - They Happen: What They Are and What to Look For”, *http://www.usdoj.gov/atr/public/guidelines/pfbrprimer.pdf*

Feinberg, R. M. & Husted, T. A. (1993), “An Experimental Test of Discount-Rate Effects on Collusive Behaviour in Duopoly Markets”, *The Journal of Industrial Economics* **41**(2), 153--160.

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 (2010): “A Continuous Dilemma,” Working Paper, UCSD.

Fudenberg, D.; Levine, D.K. (2009), "Repeated Games with Frequent Signals", *Quarterly Journal of Economics*, 124 (2009), 149-195.

Fudenberg, D.; Rand, D. & Dreber, A. (2010), "Turning the other cheek: Leniency and forgiveness in an uncertain world", Technical report, mimeo Harvard University.

Holcomb, J. H. & Nelson, P. S. (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*

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.

Roth, A. E. & Murnighan, J. K. (1978), "Equilibrium Behavior and Repeated Play of the Prisoner's Dilemma", *Journal of Mathematical Psychology* **17**(2), 189 - 198.

Roth, A. E. & Murnighan, J. K. (1982), "The Role of Information in Bargaining: An Experimental Study", *Econometrica* **50**(5), pp. 1123-1142.

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

Sannikov, Y., and A. Skrzypacz (forthcoming): "The Role of Information in Games with Frequent Actions," *Econometrica*.

Sasaki, D. and Strausz, R. (2008), "Collusion and Durability", GESY Discussion Paper 246.

Selten, R. and Stoecker, R. (1988), "End Behavior in Sequences of Finite Prisoner's Dilemma Supergames: A Learning Theory Approach," *Journal of Economic Behavior and Organization*, 7, 47-70.

Shapiro, C. (1983), "Premiums for High Quality Products as Returns to Reputations," *The Quarterly Journal of Economics*, 98, 659-680.

Strausz, R.: "Planned Obsolescence and the Provision of Unobservable Quality," forthcoming in the *Economic Journal*.