

Collusion Among Many Firms: The Disciplinary Power of Targeted Punishment *

Catherine Roux[†] and Christian Thöni[‡]

April 2013

Abstract

We explore targeted punishment as an explanation for collusion among many firms. In a series of Cournot oligopoly experiments with various numbers of firms, we compare output choices with and without the possibility of targeting punishment at specific market participants. We find strong evidence that targeted punishment enables firms to establish and maintain collusion. Furthermore, we find that the collusive effect of targeted punishment is even stronger in markets with more competitors, suggesting a reversal of the conventional wisdom that collusion is easier with fewer firms.

JEL Classification: L13, K21, C91

Keywords: Cournot oligopoly; Experiments; Collusion; Targeted punishment

*We are particularly grateful to Maria Bigoni, Aaron Edlin, Nikos Nikiforakis, Hans-Theo Normann, Daniele Nosenzo, Till Requate, Luis Santos-Pinto, Armin Schmutzler, Andy Skrzypacz, Adriaan Soetevent, Giancarlo Spagnolo for very helpful discussions and valuable comments. We would also like to thank participants at the EALE 2012 in Stockholm, the EARIE 2012 in Rome, the ESA 2012 in New York and various research seminars. Financial support from the Richard Büchner-Stiftung is gratefully acknowledged. The usual disclaimer applies.

[†]University of St.Gallen. E-mail: catherine.roux@unisg.ch

[‡]University of Lausanne. E-mail: christian.thoeni@unil.ch

1 Introduction

It is the conventional wisdom that collusion is easier with fewer firms (Chamberlin, 1933; Bain, 1951; Stigler, 1964). Both theoretical arguments and experimental studies support this assertion. In the theory of collusion in repeated games, deviations from a collusive agreement are punished with a reversion to competitive conditions which hits all cartel members equally (Friedman, 1971; Green and Porter, 1984; Abreu, 1988). The strength of this punishment and, thus, the firms' incentives to collude, decrease as the number of competitors increases. Experimental tests of the tacit collusion model have so far found that while collusion sometimes occurs with two firms, behavior is close to Nash play in markets with three or more firms (Huck et al., 2004).¹

Yet the reality of antitrust enforcement is different: cartels usually involve many firms. Empirical evidence suggests that the median number of cartel members lies between six and ten (Posner, 1970; Hay and Kelley, 1974; Fraas and Greer, 1977). In the recently convicted UK construction cartel the Office of Fair Trading found even 103 firms guilty of bid rigging in the construction industry.² Tacit collusive agreements seem to exhibit a similar pattern. In 1994, for example, about 60 independent Nasdaq market makers systematically followed an implicit collusive scheme when quoting several actively traded stocks (Christie and Schultz, 1994).

Another stark contrast between theory and practice is that cartel members use punishment which is intended to hurt a particular member. A close examination of about 20 recent cartel cases in Europe leads Harrington (2006, p.40) to the conclusion that "typically, though not universally, the threatened aggression would be targeted at the markets and customers of the deviator." That cartel members specifically punish defectors rather than abandon the cartel altogether suggests a richer set of punishment strategies than is taken into account in the theoretical literature. Furthermore, narrative evidence on the means of retaliation firms use to sustain cartel agreements reports that occasional cheating does not typically trigger massive punishment but is rather matched in degree and kind (Genesove and Mullin, 2001).

¹See Haan et al. (2009) for a review of the experimental literature on collusion. For a comparison with collusion in experimental Stackelberg oligopolies see Huck et al. (2001).

²CE/4327-4, Office of Fair Trading.

The present paper investigates targeted punishment as an explanation for collusion among many firms. To this end, we set up a series of Cournot oligopoly experiments where we distinguish between (i) general punishment which consists of increasing output in a market and (ii) targeted punishment intended to hurt a particular firm. We vary the number of firms across treatments and compare their output decisions with and without the possibility of targeting punishment at specific market participants. Our experimental setup is inspired by the literature on public-goods experiments with punishment. The upshot of this literature is that punishment stabilizes cooperation very effectively (Fehr and Gächter, 2000; Masclet et al., 2003; Herrmann et al., 2008). We find strong evidence that targeted punishment enables firms to establish and maintain collusion. We also find that the collusive effect of targeted punishment is even stronger in markets with more firms, suggesting a reversal of the conventional wisdom. Moreover, our data confirm that more severe deviations trigger harsher punishment. Firms fine-tune their retaliation to hit harder those who deviate to a higher degree.

What does targeted punishment in collusive relationships look like? One form of such punishment is guaranteed buy-ins. Guaranteed buy-ins are used to enforce sales-quota schemes whereby each firm is entitled to a certain level of output (Harrington and Skrzypacz, 2011). Firms that produce above their quota have to buy the output from underproducing firms. In the citric acid cartel, for example, Haarmann and Reimer, ahead of its sales quota, had to buy 7000 tons of citric acid from ADM, which fell short of its quota (Harrington, 2006). Defecting firms are typically reluctant to fulfill these purchases and do so only under pressure from the other cartel members. Another form of targeted punishment are the social sanctions to which individual representatives of defecting firms are commonly exposed. In the Nasdaq conspiracy, dealers that narrowed the spread were overtly harassed: “You’re embarrassing and pathetic. . . . You’re breaking the spreads for everybody.” (Paltrow, 1994, p.1). In a meeting of the carbonless paper cartel “[an AWA employee] said quite expressly that he would not tolerate any failure to follow this price increase and that he would ‘personally look after’ anyone who did not ‘play the game’ ”.³

³OJ L 115, 21.4.2004, 104, p. 19.

The distinctive feature of targeted punishment is that its disciplinary power increases with the number of firms that might use it. In the above examples, it is the increasing peer pressure that makes guaranteed buy-ins and social sanctions more effective in larger groups. Hence, larger groups may be able to collude more successfully than smaller ones. At the same time, however, increasing the number of firms may also have the opposite effect: the larger the number of firms that can punish, the more important may become incentives to free ride on others' punishment, especially if punishment is costly. Our findings suggest that this destabilizing effect is of minor importance. In particular, firms hardly adjust their punishment per deviation to different group sizes.

The remainder of this paper is organized as follows. Section 2 describes the experimental design. Section 3 explains the hypotheses. Section 4 discusses our results, and section 5 concludes.

2 Experimental Design and Procedures

We ran two treatments, referred to as *Cournot* and *Punishment*. *Cournot* consists of 20 periods of our basic game in which $n \in \{2, 4, 6, 8\}$ symmetric firms face linear demand and produce at constant marginal cost. The firms simultaneously and independently choose a quantity q_i for the production of a homogeneous good. Their action sets are numbers between 0 and 74 with 0.1 as the smallest increment. The price is determined by the market demand $p = \max\{74 - Q, 0\}$ where $Q = \sum_i q_i$ is the total quantity produced in the market. The marginal costs of production are equal to 2. Firm i 's profit in a given period is

$$\pi_i^1 = \max\{74 - Q, 0\} \cdot q_i - 2 \cdot q_i. \quad (1)$$

At the end of each period, the firms are informed about all individual quantities and profits realized in the market.⁴ In *Punishment*, the game consists of two stages. The first stage is identical to the *Cournot* treatment. In the second stage, all firms have the possibility of punishing the other firms in the market conditional

⁴More detailed feedback about others' play increases quantities in Cournot games (Huck et al., 1999).

on their actions in the first stage. Firms assign punishment points p_{ij} that reduce the payoff of the punished firm j by five units at a cost of one unit to the punishing firm i , with $p_{ij} \in \{0, 1, \dots, 20\}$ and $\sum_j p_{ij} \leq 20$. To cover the costs of potential punishment, each firm receives a per-period endowment of 20 units. Any units not used for punishment contribute to the firm's profits.⁵ Firm i 's profit in a given period is

$$\pi_i = \pi_i^1 + 20 - 5 \cdot \sum_{j \neq i} p_{ji} - \sum_{j \neq i} p_{ij}. \quad (2)$$

At the end of the second stage, the firms learn the total number of punishment points they received from other market participants and their profit of the period. In both treatments, the game is played as a finitely repeated game with 20 periods in a partner matching protocol. All game parameters and the number of periods are commonly known. Assuming subgame perfection and common knowledge of rationality, the punishment option should not affect the outcome of stage one, and all quantities should be chosen as in the Nash equilibrium of the usual Cournot game ($q^N = 72/(n+1)$). Evidence from public-goods experiments, however, suggests that players use the punishment option to enforce quantities lower than q^N . In particular, we hypothesize that the presence of the punishment option allows groups to reach collusive agreements ($q^C = 36/n$). Once the collusive solution is implemented, only little punishment is necessary to maintain collusion, allowing groups to earn supracompetitive profits.

We apply a within-subjects design. In each experimental session, subjects play 20 periods of *Cournot* followed by 20 periods of *Punishment*. Subjects only read the instructions for part two after the first part has ended. To control for order effects, we ran two sessions with two and four-firm oligopolies in the reversed order. As none of the statistical tests for order effects was significant, we pool the data of both orders for the analysis.

The sessions were run in the lab at the University of St.Gallen between February and October 2012 and were programmed in z-Tree (Fischbacher, 2007). Recruitment was carried out with ORSEE (Greiner, 2004). Subjects were randomly

⁵The purpose of this endowment is to avoid systematic wealth effects between the markets with different numbers of firms. Due to the fact that demand is equal in all markets, average profits are lower if there are more firms. To maintain comparability between *Punishment* and *Cournot*, the endowment of 20 was also paid in the latter.

allocated to computer terminals in the lab so that they could not infer with whom they would interact. For the entire experiment, communication was not allowed. We provided written instructions which informed the subjects of all the features of the market, and we read out loud the most important points (see the Appendix). We used a loaded frame similar to other studies on experimental Cournot games (see, e.g., Huck et al. (2004)). The subjects were told that they were representing a manager of a firm in a market with $n - 1$ other firms, and that their job was to make output decisions. Before the start of each treatment, the subjects had to answer a control question to make sure that they understood the instructions and were familiar with the program.⁶ When answering the control question and when choosing the quantity during the game, the subjects had access to a profit calculator allowing them to calculate the payoff of hypothetical combinations of their quantity and the average quantity produced by their competitors.

For the profits during the experiment, we used an experimental currency unit called Guilders, and 1000 Guilders were worth 1.10 CHF for the duopolies. We adapted the exchange rate for the four, six and eight-firm markets so that the Nash earnings would be equalized across markets. The payments consisted of a 10 CHF show-up fee plus the sum of the profits over the course of the experiment. Potential losses were deducted from the show-up fee. None of our subjects went home empty handed. The sessions lasted for about 105 minutes, and the average earnings were about 37 CHF (\sim 39 USD). The 152 subjects were undergraduate students in economics, business, international studies, and law, and had not participated in oligopoly experiments before.

3 Hypotheses

We begin with the conventional wisdom that collusion becomes more difficult with more firms. We hypothesize that this conjecture holds in our *Cournot* condition.

Hypothesis 1. *Without targeted punishment, the total quantities increase with the number of firms.*

⁶We presented the subjects with a randomly generated combination of the own quantity and the average quantity of the other players, and their task was to calculate their payoff. In a methodological note, we checked whether the presented situation influenced the quantities in the first period (Roux and Thöni, 2013). We found no evidence for such an effect.

Next, we consider the impact of targeted punishment on collusion. In particular, we look for a treatment effect in terms of collusive success in each market and conjecture that the possibility of targeted punishment helps to sustain collusion across all markets.

Hypothesis 2. *For all numbers of firms, quantities are lower when targeted punishment is available.*

Finally, we explore the question of whether the availability of targeted punishment can undermine the conventional wisdom. Intuitively, there are two forces that work in opposite directions. On the one hand, the potential harm caused by targeted punishment increases with the number of firms that can target punishment at a defector. On the other hand, the larger the group that can punish, the more likely it is that each individual firm may actually want to free ride on others' punishments. This effect is equivalent to the bystander effect in social psychology first identified by Darley and Latane (1968). The results from public-goods games with punishment suggest that this "second order" free-rider problem does not mitigate the effectiveness of the punishment mechanism (Fehr and Gächter, 2000). Thus, the strength of the individual punishment is presumably invariant to group size, which makes the sum of the punishments targeted at a deviator increasing in the number of firms. If this is the case, then larger groups should find it easier to collude.

Hypothesis 3. *With targeted punishment, the total quantities decrease as the number of firms increases.*

4 Results

We first examine whether the conventional wisdom holds in the *Cournot* condition. Figure 1 shows the median of the total quantities produced across all markets in both treatments.⁷ The lower horizontal line shows the collusive situation which is the total quantity produced in monopoly. The upper horizontal lines show the

⁷Each bar indicates the median of all total quantities produced in a market in a period for a given group size. The figures display the median instead of the mean because the latter is affected by outliers with high quantities. We show the average quantities in Table 1.

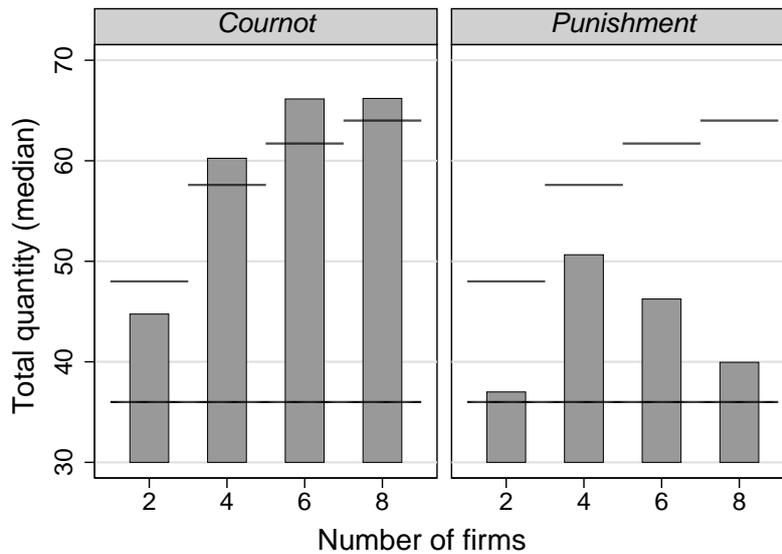


Figure 1: Median of total quantities in *Cournot* and *Punishment* for all numbers of firms. Horizontal lines indicate the theoretical benchmarks: collusion (lower) and Nash equilibrium (higher).

total quantity predicted in Nash equilibrium. From the left panel we can see that the total quantities clearly increase with group size from two to six firms in the *Cournot* condition. The increase in total quantity with increasing numbers of firms is highly significant ($p = .000$, Jonckheere–Terpstra test⁸; all non-parametric tests use independent group averages as the unit of observation). While play in the duopolies lies in between Cournot–Nash and collusion, the oligopolies with four, six, and eight firms tend to produce above the Nash prediction.⁹ These findings support our Hypothesis 1 and are in line with the results reported in Huck et al. (2004).

⁸Similar to the Kruskal–Wallis test, the Jonckheere–Terpstra test is a non-parametric test for whether samples (in our case group averages for the different numbers of firms) originate from the same distribution. Unlike the Kruskal–Wallis test, this test has an ordered alternative hypothesis. As a parametric alternative, we ran a simple linear regression explaining the total quantities by n . The results are similar.

⁹The differences to Cournot–Nash are insignificant with $p > .1$ except for six-firm oligopolies, which produce weakly significantly higher total quantities ($p = .063$, exact Wilcoxon signed-rank test).

Table 1: Quantities: Theory, experiment, and treatment effect

n	Firms	Markets	Quantities				
			Theory		Experiment		p -values
			Nash	Collusion	<i>Cournot</i>	<i>Punishment</i>	
2	32	16	24	18	22.69	20.07	.005
4	52	13	14.4	9	14.96	12.53	.000
6	36	6	10.3	6	11.45	7.61	.016
8	32	4	8	4.5	8.42	5.50	.063

Notes: Number of observations and theoretical and observed firms' average quantities. The last column shows exact p -values from one-sided Wilcoxon signed-rank tests for differences in the firms' quantities between the two treatments *Cournot* and *Punishment*.

Result 1. *Without targeted punishment, the total quantities increase with the number of firms and are close to or above the Cournot-Nash quantities for $n > 2$.*

Turning to the *Punishment* condition, we can see from the right panel in Figure 1 that total quantities substantially decrease for all group sizes. The duopolies produce quantities close to the collusive benchmark, while the oligopolies with four, six, and eight firms produce quantities in between collusion and Cournot-Nash. Table 1 shows the average quantities under the *Cournot* and *Punishment* treatments. For all markets, the average quantities drop drastically, and all treatment differences are significant.¹⁰ We can thus clearly reject the null hypothesis that targeted punishment has no effect on the firms' ability to collude. If we compare the quantities to the theoretical benchmarks, we observe that firms produce significantly less than the Nash quantity in all markets.¹¹

Figure 2 shows the dynamics of the Cournot markets. The quantities are measured on the left vertical axis. The bold lines show the median quantity over the 20 periods of *Cournot* and *Punishment* for each number of firms. For comparison, we show the average quantity as a thin line. The averages are almost always

¹⁰We report the one-sided p -values because we have a directed hypothesis. The fact that the p -value for the eight-firm oligopolies is only weakly significant is mainly due to the use of a conservative test. The corresponding t -test gives $p = .002$.

¹¹The quantities lie above the collusive benchmark. The difference is significant at 5% for all group sizes except eight, even though the median quantity for $n = 2$ is very close to the collusive output. The reason is that the deviations are almost exclusively towards higher quantities, i.e., defecting subjects only rarely undercut the collusive quantity.

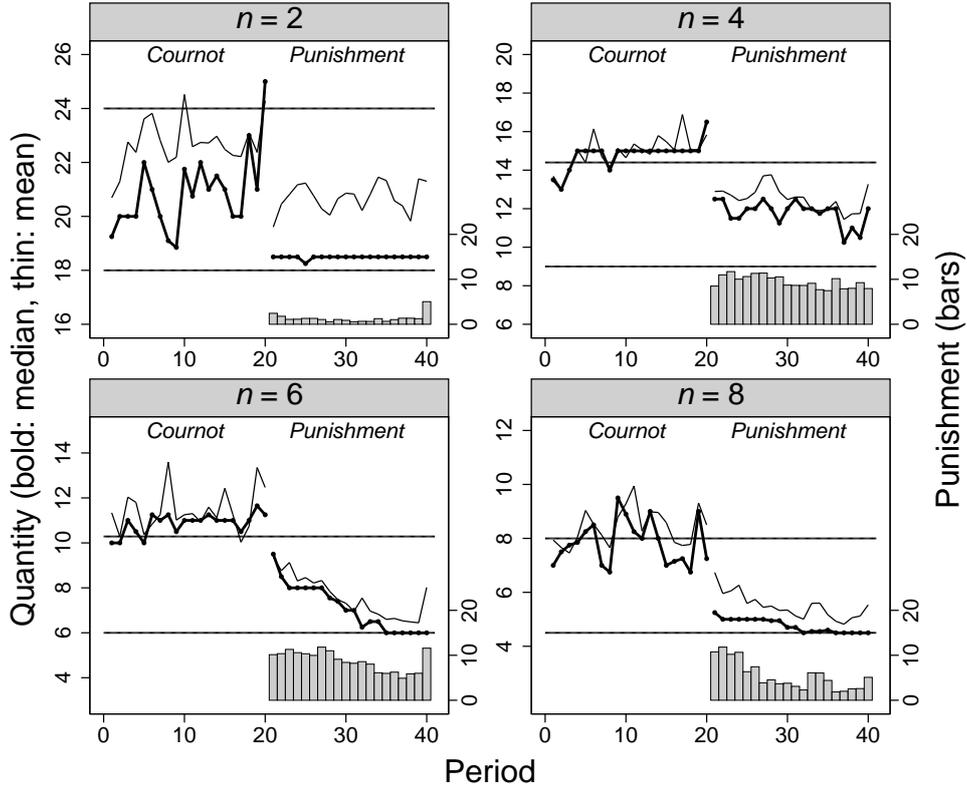


Figure 2: Individual quantities across markets and periods. Bold lines indicate median quantities, thin lines indicate mean quantities. The two horizontal lines indicate the theoretical benchmarks: collusion (lower) and Nash equilibrium (higher). Bars indicate average number of punishment points assigned to other firms in a given period of *Punishment*.

higher than the median quantities. The horizontal thin lines show the theoretical benchmarks: collusion and Nash. The bars show the average number of individual punishment points used to reduce the other firms' payoffs (right vertical axis).

Targeted punishment brings the output closer to the collusive benchmark. In all but the four-firm oligopolies, the collusive quantity is the median choice in the last third of *Punishment*.¹² Figure 2 also shows no tendency towards collusion during the 20 periods of *Cournot*. Thus, it is clearly the introduction of targeted

¹²For $n = 2$, the median quantity is (with one exception) 18.5. Most duopolies coordinated on this quantity and not on the q^C we derived above (18). Due to the fact that we round each payoff to an integer, both situations lead to identical payoffs.

punishment and not experience that allows firms to collude. The bars indicate that the punishment option is effectively used throughout all periods. Establishing collusion comes at a cost. A potential concern could therefore be that the use of targeted punishment lowers firms' profits. This is not the case. For all markets, the average firm profits are higher with the punishment option. The six- and eight-firm oligopolies can almost double their payoff in *Punishment* (increases of 95% and 94%). In the duopolies and the four-firm oligopolies, the increase in profits of 7% and 16%, respectively, is relatively moderate. Overall, these results support our Hypothesis 2.

Result 2. *The availability of targeted punishment substantially reduces quantities for all numbers of firms.*

In Hypothesis 3, we conjecture the reversal of the conventional wisdom in *Punishment*. Figure 1 shows that our evidence is mixed: increasing the number of firms from two to four substantially and significantly increases the total quantity ($p = .003$, exact Wilcoxon rank-sum test). If we consider only the markets with four or more firms, we observe that the total quantity decreases with an increasing number of firms ($p = .037$, Jonckheere–Terpstra test, one-sided¹³). One may argue that the size of the treatment effect, i.e., the difference between the total quantities produced in *Cournot* and *Punishment*, is a more accurate measure for the anticompetitive effect of the punishment option than the total quantity per se. When we relate the difference in total quantities $Q^C - Q^P$ to the number of firms, we observe a very clear positive effect ($p = .000$, Jonckheere–Terpstra test). Overall, we interpret our findings in favor of our third hypothesis, that when targeted punishment is available, collusion is easier with more firms.

Result 3. *The availability of targeted punishment seems to reverse the conventional wisdom.*

Why does the collusive effect of targeted punishment increase with the number of firms? The key to this result lies in how the number of firms affects the disciplinary power of the punishment option. To examine the determinants of targeted

¹³The equivalent p -value in the simple linear regression is .032.

punishment in detail, we ran OLS regressions. We estimated robust standard errors clustered on independent groups.¹⁴ The dependent variable is the sum of punishment points received by firm i from the other $n - 1$ market participants. The independent variables are firm i 's quantity and a number of controls. We rescaled firm i 's quantity to $q_i^r = (q_i - q^C)/(q^D - q^C)$ where q^C denotes the collusive quantity and $q^D = 18(n + 1)/n$ denotes the “optimal deviation,” i.e., the best-response quantity for the situation in which the $n - 1$ other firms choose the collusive quantity. Thus, q_i^r is a linear transformation of q_i such that, independently of the number of firms, $q_i^r = 0$ if firm i chooses the collusive output and $q_i^r = 1$ if firm i chooses the optimal deviation. This rescaling allows a straightforward comparison of the coefficients between markets with different numbers of firms. We give the estimates in Table 2.

In Model (1), we estimate a coefficient of about 15 for q_i^r , which means that, relative to choosing the collusive quantity, the optimal deviation of firm i increases the number of punishment points received by that amount. Recall that each received punishment point reduces the punished firm's payoff by five units. The estimated income reduction is thus 75 units. The effect is highly significant. Model (1) includes the average quantity of the other firms in the market, rescaled by the same procedure as firm i 's quantity, as a control. The coefficient estimate for \bar{q}_{-i}^r is negative, meaning that high quantities of the other firms reduce the amount of punishment. Hence, punishment becomes stronger when the other firms are closer to perfect collusion. We include four dummy variables for the different numbers of firms (and estimate without constant). Controlling for the quantities, firm i receives more punishment in markets with more than two firms than in a duopoly, but the effect does not seem to be monotonic in the number of competitors. With the remaining variables, we control for time and order effects. We include the period number and a dummy for the final period to capture end-game effects in punishment. Punishment decreases over time but, contrary to standard end-game effect arguments, we observe that punishment tends to become harsher in the final period. Finally, the order of the treatments does not seem to affect punishment.

If we want to compare the deterrent effect of punishment for different numbers of firms, then Model (1) is not very informative, because we only estimate the

¹⁴Tobit estimates yielded similar results.

Table 2: Determinants of getting punished: Regression results.

	Dependent variable: Sum of received punishment		
	(1)	(2)	(3)
q_i^r	14.953*** (4.640)		
$q_i^r \times n = 2$		2.742** (1.354)	9.258*** (3.322)
$q_i^r \times n = 4$		19.938*** (5.240)	50.742*** (5.197)
$q_i^r \times n = 6$		66.135*** (8.457)	95.332*** (9.128)
$q_i^r \times n = 8$		90.112*** (9.217)	112.533*** (9.024)
\bar{q}_{-i}^r	-5.978** (2.675)	-4.214** (1.929)	
$n = 2$ (D)	-0.464 (1.176)	0.730 (1.024)	-1.410 (1.168)
$n = 4$ (D)	8.050*** (1.105)	4.328** (1.738)	2.718 (2.041)
$n = 6$ (D)	7.884*** (1.896)	0.510 (1.730)	2.917 (2.120)
$n = 8$ (D)	5.008*** (1.456)	-1.541 (1.608)	-0.708 (1.477)
Period	-0.195*** (0.048)	0.033 (0.062)	0.032 (0.088)
Final period (D)	2.492** (1.144)	-0.177 (0.914)	1.745 (1.150)
Order $C \rightarrow P$ (D)	1.805 (1.215)	1.226 (1.217)	1.854 (1.187)
F -test	44.5	87.9	53.8
Prob > F	0.000	0.000	0.000
R^2 adjusted	0.379	0.593	0.754
N	3040	3040	1330

Notes: OLS estimates, robust standard errors estimated with clustering on group. Dependent variable is the sum of punishment points received by firm i from all other firms in the group. Independent variables are the rescaled quantity q_i^r , where $q_i^r = 0$ if firm i produces the collusive output and $q_i^r = 1$ if firm i produces the best-response quantity to the collusive quantities of the other firms, dummies for the number of competitors and interactions. Further controls are the average rescaled quantity of the other firms in the group (\bar{q}_{-i}^r), Period $\in \{1, 2, \dots, 20\}$, Final period is a dummy for the last period. Order $C \rightarrow P$ is a dummy for the treatment order. Standard errors in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

average amount of punishment across all situations. In Model (2) we estimate the slope of the received punishment in q_i^r separately for each number of firms. The results display huge differences. Comparing firms which produce the collusive quantity ($q_i^r = 0$) to firms which optimally deviate ($q_i^r = 1$), the increase in punishment ranges from merely 2.7 units for duopolies to 90 units for eight-firm oligopolies. Among the other controls, we observe that most dummies for the number of firms become insignificant. Only in the four-firm oligopolies does there seem to be a substantial amount of punishment independent of the quantity choice. All other controls are insignificant with the exception of \bar{q}_{-i}^r , which reaffirms the observation that punishment is stronger if the other firms in the group chose a relatively low quantity. This suggests that punishment is particularly harsh if the punishing firms maintain the collusive quantities.

In Model (3) we restrict our sample to the cases with $|\bar{q}_{-i} - q^C| \leq 1$, i.e., we look only at situations where the other firms' quantities are, on average, close to the collusive quantity.¹⁵ This procedure leaves us with 1330 observations. Model (3) comes closest to estimating the punishment targeted at a firm which unilaterally deviates from a collusive agreement. The coefficients of the slopes are larger than in Model (2), reinforcing the argument that punishment is harshest when the rest of the group remains in a collusive agreement. When we multiply the coefficients by five, we get an estimate for the income reduction an "optimal" deviator faces, namely $\{46.3, 254, 477, 563\}$ for $n = \{2, 4, 6, 8\}$. We have to set these numbers in relation to the maximum profits that can be earned by optimally deviating from a collusive agreement, which are $\{81, 182, 225, 248\}$. While for duopolists deviating from collusion might pay in the particular period, it clearly does not for firms in oligopolies with more than two competitors. Punishment is not only strong on average, it is also quite certain to occur. If we consider the cases where $n - 1$ firms play on average $q^C \pm 1$ (as in Model 3) and calculate the frequency that a deviating firm i with $q_i > q^C + 1$ will be punished, we see that in 59% of the cases, a duopolist receives punishment. In all markets with $n > 2$, there is not a single instance, out of 80 cases, in which the deviator gets away unpunished. When we divide the number of punishment points assigned to the defecting firm by

¹⁵Because we only rarely observe quantities below q^C , this means that the other firms' quantities are relatively similar.

the number of potential punishers, we observe almost no reduction with increased group size {9.26, 16.9, 19.1, 16.1}.¹⁶

Taken together, these results suggest a reason for the reversal of the conventional wisdom in our *Punishment* condition: subjects do not adjust their punishment per deviation to the different group sizes. Consequently, a deviator faces much harsher punishment in markets with more firms. This is remarkable, because the punishment of deviators is a pure public good in markets with more than two firms. Thus, when it comes to enforcing collusive agreements it seems that larger groups are not plagued by the bystander effect. While the OLS models reported here impose a linear relation between the received punishment and q_i^r , a closer inspection of the data reveals that indeed punishment increases gradually in q_i^r . Consequently, the more a defecting firm's quantity exceeds the collusive quantity, the harsher the punishment. This result confirms Genesove and Mullin (2001)'s finding that punishment seems to "fit the crime."

5 Discussion and Conclusion

The conventional wisdom is that collusion is easier with fewer firms. While theories on collusion as well as oligopoly experiments support this assertion, there is abundant evidence on convicted cartels suggesting that firms manage to cooperate also in markets with a large number of competitors. Moreover, firms in these markets target punishment at specific cartel members in order to enforce compliance with the rules of collusion. When cheating occurs, it does not immediately lead to price wars but rather triggers punishment that fits the crime. In a series of Cournot oligopoly experiments with various numbers of market participants, we studied whether targeted punishment explains the occurrence and stability of collusion among many firms.

Without targeted punishment, our results confirm earlier findings in the literature: some collusion occurred in duopolies but behavior was close to the Nash benchmark or even more aggressive in groups of four and more. But when we made

¹⁶On the other hand, the increase between $n = 2$ and $n = 4$ is significant ($p = .040$, F -test). This might be due to the fact that in duopolies, all punishment is targeted, i.e., both an increase in quantity and assignment of punishment points harms the deviating firm exclusively.

targeted punishment available, collusion occurred for all group sizes. Moreover, the collusive effect of targeted punishment was even stronger in markets with more firms, suggesting a reversal of the conventional wisdom. Our analysis of the punishment behavior indicates that the higher the quantity a defecting firm produces, the stronger the punishment imposed by other market participants. This is in line with evidence from the Sugar Institute cartel reported by Genesove and Mullin (2001).

What might drive the strong impact of targeted punishment on collusion? Our analysis shows that, unlike what is predicted using the standard assumptions, targeted punishment is a highly credible threat even in the final round of the game. The literature on non-standard preferences has mainly focused on two explanations for this phenomenon: (i) players have preferences over the distributions of the final payoffs among all players, e.g., inequality aversion, and (ii) players not only consider final outcomes but also the intentions of other players, e.g., reciprocity.¹⁷ The first explanation would clearly imply that in one-deviator situations (as investigated in Model (3) of Table 2), per capita punishment decreases as group size increases. Firms would simply want the deviator's income to be lowered, but prefer others to do the dirty work.¹⁸ If reciprocity is the motive behind targeted punishment, things might look different. A revengeful firm might feel the urge to retaliate, independently of whether others do so or not. In the light of our experimental results and our examples of targeted punishment among oligopolists, we find the reciprocity explanation more compelling.

If reciprocity is the driver behind targeted punishment, it might be interesting for future research to explore situations in which there is ambiguity about the true intention behind an observed market outcome. There might be a multitude of reasons for firms to deviate from the collusive agreement, or the information about market quantities might be distorted.¹⁹

¹⁷See Sobel (2005) for an overview and Doruk and Santos-Pinto (2013) for an application to oligopolies.

¹⁸For a detailed analysis of punishment motivated by distributional preferences, see Thöni (2011).

¹⁹Research from public-goods games suggests that imperfect information reduces the efficacy of the targeted punishment option (Grechenig et al., 2010; Ambrus and Greiner, 2012).

Another extension would be allowing communication. In real-world cartels, members usually meet and talk to coordinate actions.²⁰ As they have the possibility of renegotiating the agreement after a deviation, it is not clear whether deviating from an explicit agreement spurs even stronger punishment than the deviation from collusive situations in our experiment.

Targeted punishment in oligopolies can take many different forms. Typical examples involve guaranteed buy-ins in quota cartels and social sanctions. The driving force in these enforcement mechanisms is peer pressure, which becomes stronger when the number of firms involved in the collusive relationship increases. Research on social networks suggests that friendship between competitors is an important determinant of collusive success.²¹ A wonderful example of social events organized to foster loyalty and friendship in cartel-like operations is the series of “Gary Dinners” named after the Chairman of U.S. Steel, Judge Elbert Gary, in the US steel industry during the years 1907–1911 (Armstrong and Huck, 2010). Further examples of targeted punishment are leniency applications that follow defections (Hinloopen and Soetevent, 2008), exclusive-dealing contracts with suppliers of an essential industry input (Krattenmaker and Salop, 1986), comparative advertising strategies against the defector (Ayres, 1987), or simply the expulsion of the agreement (Dick, 1996).

Our results shed an interesting light on the discussion of coordinated effects in merger decisions. The European Commission’s merger decisions published from 1990 to 2004 indicate that the potential for coordinated effects arise only if the post-merger market would include two remaining firms with fairly identical market shares (Davies et al., 2011). While this perspective, adopted by the European Commission, is in line with the existing theoretical and experimental literature, our findings suggest that it may be too narrow. Our results may also be important in the antitrust evaluation of information exchanges as circumstantial evidence for collusive behavior when a hard-core agreement is absent.²² European case law suggests that whether the information exchanges are considered as unlawful is directly linked to the degree of market concentration.

²⁰Fonseca and Normann (2012) report results from experimental Bertrand games with communication.

²¹See, e.g., Baker and Faulkner (1993); Dobbin and Dowd (1997); Ingram and Roberts (2000).

²²See, e.g., European Commission *UK Agricultural Tractor Registration Exchange*, OJ L 68, 13.3.1992, pp. 19–33.

References

- Abreu, D. (1988). On the theory of infinitely repeated games with discounting. *Econometrica* 56(2), 383–396.
- Ambrus, A. and B. Greiner (2012). Imperfect public monitoring with costly punishment: An experimental study. *American Economic Review* 102(7), 3317–3332.
- Armstrong, M. and S. Huck (2010). Behavioral economics as applied to firms: A primer. *Competition Policy International* 6(1), 3–45.
- Ayres, I. (1987). How cartels punish: A structural theory of self-enforcing collusion. Faculty Scholarship Series 1549, Yale Law School.
- Bain, J. S. (1951). Relation of profit rates to industry concentration: American Manufacturing, 1936–1940. *Quarterly Journal of Economics* 65(3), 293–324.
- Baker, W. E. and R. R. Faulkner (1993). The social organization of conspiracy: Illegal networks in the heavy electrical equipment industry. *American Sociological Review* 58(6), 837–860.
- Chamberlin, E. H. (1933). *The Theory of Monopolistic Competition*. Harvard University Press, Cambridge, MA.
- Christie, W. G. and P. H. Schultz (1994). Why do Nasdaq market makers avoid odd-eighth quotes? *Journal of Finance* 49(5), 1813–1840.
- Darley, J. M. and B. Latane (1968). Bystander intervention in emergencies: Diffusion of responsibility. *Journal of Personality and Social Psychology* 8(4), 377–383.
- Davies, S., M. Olczak, and H. Coles (2011). Tacit collusion, firm asymmetries and numbers: Evidence from EC merger cases. *International Journal of Industrial Organization* 29(2), 221–231.
- Dick, A. R. (1996). When are cartels stable contracts. *Journal of Law and Economics* 39(1), 241–283.
- Dobbin, F. and T. J. Dowd (1997). How policy shapes competition: Early railroad foundings in Massachusetts. *Administrative Science Quarterly* 42(3), 501–529.
- Doruk, I. and L. Santos-Pinto (2013). Tacit collusion under fairness and reciprocity. *Games* 4(1), 50–65.
- Fehr, E. and S. Gächter (2000). Cooperation and punishment in public goods experiments. *American Economic Review* 90(4), 980–994.
- Fischbacher, U. (2007). z-Tree - Zurich toolbox for readymade economic experiments. *Experimental Economics* 10(2), 171–178.
- Fonseca, M. A. and H.-T. Normann (2012). Explicit vs. tacit collusion—The impact of

- communication in oligopoly experiments. *European Economic Review* 56(8), 1759–1772.
- Fraas, A. G. and D. F. Greer (1977). Market structure and price collusion: An empirical analysis. *Journal of Industrial Economics* 26(1), 21–44.
- Friedman, J. W. (1971). A non-cooperative equilibrium for supergames. *Review of Economic Studies* 38(1), 1–12.
- Genesove, D. and W. P. Mullin (2001). Rules, communication, and collusion: Narrative evidence from the Sugar Institute case. *American Economic Review* 91(3), 379–398.
- Grechenig, K., A. Nicklisch, and C. Thöni (2010). Punishment despite reasonable doubt - A public goods experiment with sanctions under uncertainty. *Journal of Empirical Legal Studies* 7(4), 847–867.
- Green, E. J. and R. H. Porter (1984). Noncooperative collusion under imperfect price information. *Econometrica* 52, 87–100.
- Greiner, B. (2004). An online recruitment system for economic experiments. In K. Kremer and V. Macho (Eds.), *Forschung und wissenschaftliches Rechnen* (GWDG Beric ed.), pp. 1–15. Göttingen.
- Haan, M. A., L. Schoonbeek, and B. M. Winkel (2009). Experimental results on collusion. In J. Hinloopen and H.-T. Normann (Eds.), *Experiments and Competition Policy*, pp. 9–33. Cambridge: Cambridge University Press.
- Harrington, J. E. (2006). How do cartels operate? *Foundations and Trends in Microeconomics* 2(1), 1–105.
- Harrington, J. E. and A. Skrzypacz (2011). Private monitoring and communication in cartels: Explaining recent collusive practices. *American Economic Review* 101(6), 2425–2449.
- Hay, G. A. and D. Kelley (1974). An empirical survey of price fixing conspiracies. *Journal of Law and Economics* 17(1), 13–38.
- Herrmann, B., C. Thöni, and S. Gächter (2008). Antisocial punishment across societies. *Science* 5868, 1362–1367.
- Hinloopen, J. and A. R. Soetevent (2008). Laboratory evidence on the effectiveness of corporate leniency programs. *RAND Journal of Economics* 39(2), 607–616.
- Huck, S., W. Müller, and H.-T. Normann (2001). Stackelberg beats cournot: On collusion and efficiency in experimental markets. *Economic Journal* 111(474), 749–765.
- Huck, S., H.-T. Normann, and J. Oechssler (1999). Learning in cournot oligopoly - an experiment. *Economic Journal* 109(454), C80–95.
- Huck, S., H.-T. Normann, and J. Oechssler (2004). Two are few and four are many:

- Number effects in experimental oligopolies. *Journal of Economic Behavior & Organization* 53(4), 435–446.
- Ingram, P. and P. W. Roberts (2000). Friendships among competitors in the Sydney hotel industry. *American Journal of Sociology* 106(2), 387–423.
- Krattenmaker, T. G. and S. C. Salop (1986). Competition and cooperation in the market for exclusionary rights. *American Economic Review* 76(2), 109–113.
- Masclet, D., C. N. Noussair, S. Tucker, and M.-C. Villeval (2003). Monetary and non-monetary punishment in the voluntary contributions mechanism. *American Economic Review* 93(1), 366–380.
- Paltrow, S. J. (1994, October 20). Taking stock of Nasdaq. Los Angeles Times.
- Posner, R. (1970). A statistical study of antitrust enforcement. *Journal of Law and Economics* 13(2), 365–419.
- Roux, C. and C. Thöni (2013). Do control questions influence behavior in experiments? DEEP Working Paper 1303.
- Sobel, J. (2005). Interdependent preferences and reciprocity. *Journal of Economic Literature* 43(2), 392–436.
- Stigler, G. (1964). A theory of oligopoly. *Journal of Political Economy* 72(1), 44–61.
- Thöni, C. (2011). Inequality aversion and antisocial punishment. University of St.Gallen Discussion Paper no. 2011-11.

Appendix

Instructions for *Cournot*

Instructions

You are taking part in an economic study financed by various research foundations. You have the opportunity, depending on your decisions, to earn a fair amount of money. It is therefore important that you read these instructions carefully.

These instructions are solely for your private use. You are not allowed to communicate with the other participants during the entire study. Should you have any questions, please ask us directly. If you violate this rule, you will be dismissed from the study and forfeit all the payments.

During the study we do not speak in terms of CHF (national currency), but in Guilders. Thus, in a first stage, your earnings are determined in terms of Guilders. At the end of the study the total amount of Guilders you have earned are converted to CHF. Prior to the begin of the study the exchange rate (Guilders to CHF) is displayed on the screen. At the end of the study your earnings, plus 10 CHF to cover potential losses, are paid to you in cash.

The study is divided into 20 separate periods. You remain anonymous during the entire study. The participants are divided into groups. Suppose, for example, that your group consists of three, i.e. you and two other, participants. The composition of the groups remains the same for all 20 periods. Hence, you are in a group with the same participants for all 20 periods.

Each participant is the manager of a firm. All the firms in the same group produce a homogenous good and sell this in the same market. Each period, you decide how many units of the good to produce. The other two firms do exactly the same thing. The following rule applies: the larger the production quantity of the firms in the same market, the lower is the market price of the good. The unit cost of production amounts to 2 Guilders, and thus, the per-unit profit equals the market price minus 2 Guilders.

In the following pages we describe the study in detail.

Detailed Information on the Study

At the beginning of each period each participant has to decide how many units of the good to produce. All the firms in a market face the same market demand which is

$$P = 74 - Q.$$

P is the market price and Q is the total quantity of production. For example, if your firm produces with two other firms for the same market, then Q is the sum of the production quantity of your firm and the production quantities of the two other firms.

At the beginning of each period you see the following input screen:

Period			1 out of 20	Remaining time [sec]: 116
What-if calculator			Your decision	
Your quantity	Av. quantity others	Your profit	Your production quantity <input type="text"/>	
Your production quantity <input type="text"/>			OK	
Average production quantity of the other firms <input type="text"/>				
erase table		calculate		

The period number appears in the upper left corner of the screen. In the upper right corner you can see how many more seconds you have to decide on your production quantity.

In the lower left corner of the screen you can see the “What-if-calculator”. You can use this tool to determine your profits from different combinations of your quantity and the average quantity of the other firms in the market.

You can decide on your production quantity by entering a number in the input field (maximum one decimal). Once you have entered your quantity, you must click on the “OK button”. Once you have done this, you can no longer revise your decision for that period.

After all firms have made their decision, the following profit screen shows your profit. It also displays your production quantity and the average production quantity of the other firms.

Period	1 out of 20	Remaining time [sec]: 56
Your profit		
<p>Your production quantity ...</p> <p>Average production quantity of the other firms ...</p> <p>Your profit ...</p>		
<input type="button" value="continue"/>		

Your profit is determined as follows:

$$[74 - (\text{your quantity} + \text{sum of the quantities of the other firms})] \cdot (\text{your quantity}) - 2 \cdot (\text{your quantity}).$$

The profits of the other firms in the market are determined in the same way. The above formula makes clear that, given your quantity, the bigger the sum of the production quantities of the other firms in the market, the smaller your profit. The same is true for the other firms. The larger the quantity you choose, the smaller the profits of the other firms. If the total production quantity exceeds 74 units of the good, then the market price is 0. You can also incur losses.

If you are done before the time is up, you can click on the “Continue button”. Next, the information screen which shows the production quantities and profits of all firms in your market appears.

Period	1 out of 20	Remaining time [sec]: 88
Information production quantities		
You receive additional 20 Guilders		
	Production quantity	Profit
Your firm
Firm 1
Firm 2
		<input type="button" value="continue"/>

On this screen you see the production quantity and the profit of each firm in your market. Your production quantity and your profit are always shown in the first row and are highlighted in blue, while the production quantities and profits of the other firms are listed in the rows below. You observe the production quantities and profits, however, you cannot see the identity of the other firms' managers. At this point, you receive 20 additional Guilders.

Do you have any questions?

Instructions for *Punishment*

Instructions

Now, we are going to repeat this study - call it A - with several modifications. Let us call this new study B. The composition of your group remains the same. Study B is also divided into 20 separate periods. After these 20 periods, the entire study is definitively over, and you receive the total amount of Guilders that you have earned in studies A and B, converted to CHF.

In study B, each period has two stages. In both stages you are again the manager of a firm. The first stage is the same as study A in which you just have participated. In the first stage of each period you thus have to decide how many units of the good to produce. Your profit in the first stage is determined exactly the same way as your profit in study A.

What is New in Study B? New in study B is the second stage which directly follows the profit screen in the first stage. At the start of the second stage you receive 20 additional Guilders. As before, you see the production quantities of the other firms. Moreover, you can now lower the profit of each firm by allocating “deduction point”. The other firms can lower your profit as well. The input screen in the second stage makes this clear:

Period 1 out of 20 Remaining time [sec]: 90

Your decision in stage 2

You receive 20 additional Guilders.

	Production quantity	Profit	Deduction points
Your firm	
Firm 1	<input type="text"/>
Firm 2	<input type="text"/>

OK

On this screen you see the production quantity and the profit of each firm in your market. Your production quantity and your profit are always shown in the first row and are highlighted in blue, while the production quantities and profits of the other firms are listed in the rows below, exactly as in study A. As before, you observe the production quantities and profits, however, you cannot see the identity of the other firms’ managers.

You have to decide whether, and if so, how many deduction points you want to assign to each firm (except to your own). You have to make an entry in all cases. If you do not want to affect the profit of a particular firm you enter 0.

The allocation of deduction points is costly. Deduction points are integers in between 0 and 20. The higher the number of deduction points you assign to a firm, the higher your costs. Each deduction point allocated costs you one Guilder. For example, if you produce with two other firms for the same market, and you assign 9 deduction points to one of the firms, it costs you 9 Guilders. Assigning 0 deduction points to the other firm costs you nothing. You, therefore, have allocated 9 deduction points in total, and your total costs amount to 9 Guilders (9+0).

By assigning 0 deduction points to a particular firm, you do not alter the profit of that firm. If, however, you assign 1 deduction point to a firm, you lower the profit of that firm by 5 Guilders. If you assign 2 deduction points to a particular firm, you lower the profit of that firm by 10 Guilders, and so on. With each deduction point you assign to a firm, the profit of that firm decreases by 5 Guilders.

Whether, and if so, by how much the profit of a firm decreases depends on the total quantity of deduction points received. For example, if a particular firm receives 20 deduction points in total (from the other firms in the same market), its profit from the first stage decreases by 100 Guilders. If a particular firm receives 30 deduction points in total, its profit from the first stage decreases by 150 Guilders. Your total profit earned in both stages is the sum of your profit earned in the first stage in addition to the 20 Guilders minus the total costs of the deduction points allocated as well as received.

After all firms have made their decision, the following profit screen shows your total profit:

Period	1 out of 20	Remaining time [sec]: 87
Total profit		
Your profit from stage 1 ...		
Additional guilders in stage 2 ...		
Assigned deduction points (in sum) ...		
You have received „ Profit reduction due to received deduction points ...		
Your total profit ...		
<input type="button" value="continue"/>		

Do you have any questions?