



**Competition as a Socially
Desirable Dilemma
Theory vs. Experimental
Evidence**

Christoph Engel





Competition as a Socially Desirable Dilemma Theory vs. Experimental Evidence

Christoph Engel

August 2009

Competition as a Socially Desirable Dilemma Theory vs. Experimental Evidence*

Christoph Engel

Abstract

Cartels are inherently instable. Each cartel member is best off if it breaks the cartel, while the remaining firms remain loyal. If firms interact only once, if products are homogenous, if firms compete in price, and if marginal cost is constant, theory even predicts that strategic interaction forces firms to set the market clearing price. For society, this would be welcome news. Without anti-trust intervention, the market outcome maximises welfare. The argument becomes even stronger if the opposite market side has a chance to defend itself; if imposing harm on the opposite market side is salient; if it is clear that cartels are at variance with normative expectations prevalent in society. There is an equally long list of reasons, though, why such optimism might be unwarranted: capacity is limited; interaction is repeated, and the end is uncertain; firms might be willing to run a limited risk of being exploited by their competitors, hoping that the investment pays. This paper explores the question both theoretically and experimentally. In the interest of capitalising on a rich body of experimental findings, and on the concept of conditional cooperation in particular, the paper offers a formal model that interprets oligopoly as a linear public good.

JEL: A13, C91, D43, D62, H23, H41, K21, L13

Keywords: Cartel, Oligopoly, Bertrand, Cournot, Public Good, Externality, Experiment

I. Research Question	2
II. Competition as a Dilemma	3
1. One Shot Interaction	3
2. Repeated Interaction.....	7
a. Open vs. Defined End.....	7
b. Explicit vs. Implicit Sanctions.....	9
c. Conditional Cooperation.....	11
III. Competition as a Socially Desirable Dilemma.....	16
1. Defense	16
2. Respect for Outsiders.....	17
3. Normativity.....	18
IV. Antitrust Learning from the Lab.....	20
References	22

* Valuable comments by Martin Hellwig and Christian Traxler are gratefully acknowledged.

I. Research Question

You are facing a dilemma – what a mess! Your opponent is facing a dilemma – what a relief! This is the situation of cartel members vs. cartel authorities. For cartel members, keeping prices high, or quantity low for that matter, is a dilemma. Each cartelist is individually best off if the others are loyal to the cartel, while she lowers her price, or while she sells above her quota. If one believes in (admittedly radically simple) theory, in many markets cartel authorities could content themselves with staying at the sideline. Attempts to cartelise the market are anyhow doomed to failure. In this perspective, antitrust authorities might substantially reduce their ex officio enforcement activities. More interestingly, they might target scarce enforcement resources more effectively to markets that theory predicts to be least liable to the erosive forces of dilemma. Ultimately one may even wonder whether there is a need for the legal ban on cartels in the first place. Not by coincidence, some of the claims of the Chicago School come to mind (Bork 1978).

Empirically, dilemmas do not always materialise when theory has predicted them. Since dilemmas are normally perceived as a social problem, not as a social panacea, much energy has been devoted to understanding the conditions under which those facing the dilemma are able to overcome it without governmental intervention (Ostrom 1990; Ledyard 1995; Zelmer 2003). Although our knowledge is certainly still incomplete, it is safe to say that many more dilemmas are successfully managed, or at least assuaged, than a theory of unboundedly prevoiant and perfectly selfish individuals predicts. One major reason is beyond the purview of this paper. In the field, people are not like monads. Their behaviour is embedded in a rich texture of formal and informal institutions, which provides orientation and direction at the same time (Granovetter 1985). This paper focuses on a behavioural explanation. More narrowly even it exclusively capitalises on our experimental knowledge about motivational forces that can help participants overcome dilemma situation without governmental help. The paper wonders whether the same forces that alleviate government's tasks when it comes to, say, protecting the environment, make government's life more troublesome when it sees at workable competition.

Both from a theoretical and from an experimental perspective, it makes a significant difference whether individuals face a dilemma just once, or whether the dilemma repeats itself. Moreover, the structural similarity between socially beneficial and socially detrimental dilemmas neglects a characteristic feature of the latter. If society at large is indifferent, or if it even receives a side benefit when those facing the dilemma overcome it, the situation is free from conflict. This is different if one man's meat is another man's poison, as they say. Those negatively affected might be able to muster some countervailing power. And behaviourally it might matter that overcoming the internal dilemma inevitably implies imposing harm on outsiders. People might dislike the unintended side-effect, and they might be sensitive to normative expectations originating in society at large.

The analysis of incentives in prisoner's dilemma games (e.g. Rapoport and Chammah 1965), public goods (e.g. Cornes and Sandler 1996), and oligopoly (e.g. Tirole 1988) is a rich field. The

theoretical parts of this paper do not intend to survey this literature (for a somewhat more extensive account see Engel 2006). The presentation of theory is confined to those, fairly basic, elements necessary to understand and to appreciate the behavioural findings. The paper also does not speak to the equally rich empirical literature on the success and failure of cartels in the field (for surveys see Slade 1990; Feuerstein 2005; Levenstein and Suslow 2006).

The experimental evidence used to back the behavioural claims is taken from three different papers (Engel and Normann 2008; Engel, Irlenbusch et al. 2009; Engel and Rockenbach 2009) and a meta-study (Engel 2007a) of my own¹.

Section II brackets the fact that collusion is socially detrimental. This section thus adopts an internal perspective. It interprets a cartel as a dilemma of its members. Section III drops the limitation. Adopting an external perspective, it explores, from a behavioural perspective, how the internal dilemma is modified by the fact that overcoming the dilemma is socially undesirable. Section IV concludes with an outlook to antitrust policy.

II. Competition as a Dilemma

1. One Shot Interaction

Starting as simple as possible, let us assume a market with two firms, supplying a homogeneous product. There are high barriers to entry. Each firm would be in a position to serve total demand on the spot. Marginal cost is constant. There is no fixed cost. Demand decreases in price. Demand is fully rational in the sense that it strictly prefers a cheaper over a more costly offer.

This is of course a textbook Bertrand market (Bertrand 1883). Collectively, firms are best off if they coordinate behaviour, i.e. form a cartel, and act like a monopolist. They fix price, or restrict quantity for that matter, such that they maximise producer rent. Producers' success comes at a distributional loss for consumers and, more troubling even, entails a deadweight loss for society at large.

However, for any firm it is not a best response to the other cartel members' strategies to be a loyal member of the cartel. Individually, each firm is best off if the other keeps price high, or quantity low. The first firm then undercuts the collusive price by a trifle. That way it captures the entire demand, and the producer rent that goes with it. The second firm anticipates this, and preempts it by undercutting the collusive price even a bit more. This, in turn, is anticipated by the first firm, and preempted by undercutting by another trifle, all the way down until both firms end up setting the competitive price right from the beginning. Theory thus predicts that both firms offer at marginal cost, and make zero profit. This prediction does not depend on the number of firms in the market.

1 I thank my co-authors on the experimental papers for letting me use our data once more.

Theorists themselves have always felt uneasy with this prediction. One way of surmounting the “Bertrand paradox” (Tirole 1988:211) is dropping the assumption that either firm is able to serve the entire market on the spot. If capacity is an issue, the strategic variable shifts from price to quantity (Edgeworth 1897; Kreps and Scheinkman 1983) (but see Davidson and Deneckere 1986). More precisely, and more importantly, when committing to a quantity, either firm has to take into account that its competitor is simultaneously making the same choice. In such a market, the equilibrium requires these two strategic moves to be coordinated. Technically, the equilibrium is at the intersection of the two reaction curves. This is of course a textbook Cournot market (Cournot 1838). The smaller the number of firms, the further the Cournot equilibrium is away from the competitive, or Walrasian, equilibrium. The smaller the number of firms, the larger their (non-collusive) profit.

In the field, it is often not easy to distinguish Bertrand and Cournot competition. Competitors need not hold coordinated beliefs about their opponents’ ability to adjust supply either. It therefore is attractive to test the effect of competition in price versus competition in quantity under the controlled conditions of a lab experiment. Oligopoly experiments have indeed started more than 50 years ago. Until 2006, a total of 147 experimental papers have been published, most of them comparing different treatments, so that the number of (totally) independent observations is even larger. In a meta-study (Engel 2007a), I have made this evidence comparable by generating two indices, see Figure 1.




Figure 1
 Benchmarks in Meta-Study
 WE Walrasian Equilibrium, NE Nash Equilibrium, CE Collusive Equilibrium

From the formal model used by the respective experimenter, and from the parameters implemented, I have calculated three benchmarks: the competitive Walrasian, the collusive Pareto, and the strategic Nash predictions. Of course, if the strategic variable is price, the Walrasian prediction is smallest, the Pareto prediction is largest and, in the Bertrand case, the Nash prediction

coincides with the Walrasian prediction. By contrast, if the strategic variable is quantity, a large quantity indicates a high degree of competition. My dependent variable is the degree of deviation from the Walrasian or from the Nash predictions. To calculate this number, I express the empirical mean (the dashed line in Figure 1) as a fraction of the distance between the Walrasian (index CW, the WE-CE line in Figure 1) or the Nash predictions (index CN, the NE-CE line in Figure 1) on the one hand, and the Pareto prediction on the other hand. In some experiments, competition has been so fierce that the mean price is below the Walrasian benchmark; then my index is negative. In other experiments, collusive subjects overshoot and fix prices at such a high level that they make less joint profit than they could; then my index goes above 100. Oligopoly experiments are usually motivated by antitrust problems. In some markets, interaction is indeed singular. Tickets for one soccer match are no substitutes for the match next week; firms are bidding for a public project; there is a small time window for selling shares in a company to a raider. In most markets, however, interaction is repeated. This explains why there is only a single published oligopoly experiment where interaction is one shot (Huck and Wallace 2002). It does not provide an empirical test of the Bertrand versus Cournot models if the interaction is one shot. Actually, a direct test would not be easy to design. If there is a larger deviation from the Nash equilibrium when subjects compete in price: is this due to the (qualitative) difference in strategic variables? Does it result from the (quantitative) difference in the Nash predictions? Is the result driven by the absolute, rather than the relative gains from collusion?

To have a clean test, it is preferable to interpret interaction in a one-shot duopoly Bertrand versus Cournot market as two by two games. If one does, a characteristic difference between both markets becomes visible. Defection, i.e. setting the Cournot quantity, is the unique equilibrium when firms compete in quantity. By contrast, in a Bertrand market (with homogeneous products and constant marginal cost), firms are indifferent between playing Nash and letting their competitor exploit them; either way their payoff is zero. Consequently, in a Bertrand market there are multiple equilibria. Note, however, that both firms colluding is the only cell that is not an equilibrium.

In generic notation, in both games actions and payoffs are as in Table 1. If a firm cooperates, it sets the collusive price or the collusive quantity. Defection consists of setting the competitive price if the game is Bertrand, and setting the quantity resulting from the intersection of firm's reaction curves if the game is Cournot². Hence $\pi_{DD}^C > \pi_{DD}^B = 0$, where the first subscript is this firm's action, and the second subscript is the opponent's action. Superscripts stand for Cournot and Bertrand respectively. In both games $\pi_{DC} > \pi_{CC}$: it is individually rational to defect if the opponent cooperates. However, $\pi_{DD}^B = \pi_{CD}^B$, whereas $\pi_{DD}^C > \pi_{CD}^C$. In Bertrand, if the opponent defects, a player is indifferent between cooperation and defection. A Cournot game is a standard prisoner's dilemma. Cooperation is strictly dominated. By contrast, in the Bertrand game the prediction is clear only if one also excludes weakly dominated strategies: while a player is indifferent between cooperation and defection if the other defects, she prefers defection if the other cooperates.

2 Formally, reaction curves result from maximising firm profit, given the opponent's action.

	Cooperate	Defect
Cooperate	π_{CC}, π_{CC}	π_{CD}, π_{DC}
Defect	π_{DC}, π_{CD}	π_{DD}, π_{DD}

Table 1
One Shot Duopoly

This formalisation invites a straightforward experimental test: vary π_{DD} , while holding the remaining payoffs fix. Hans Normann and myself have done this (Engel and Normann 2008), setting $\pi_{DC} = 10, \pi_{CC} = 5, \pi_{CD} = 0$. We have varied π_{DD} in the interval $[0,5]$, in 11 steps, as displayed in Figure 2. The bars indicate the percentage of subjects who have defected in the respective situation. One sees a monotonic relation: the lower the safety payoff, the less defection, or the more cooperation. Behaviour is driven by the difference between the safety payoff and the collusive payoff. If this difference is big, collusion is easy to sustain. Note that, in this experiment, we are not measuring deviations from the Walrasian equilibrium, but deviations from the Nash equilibrium. Hence, while the Nash prediction does better the larger the safety payoff, relative to the collusive payoff, even if this difference is zero (i.e. when the safety payoff is 5), there is still some cooperation.




Figure 2
Degree of Defection in Reaction to Safety Payoff

This finding resonates with a piece of practical wisdom. Cartels are the offspring of penury, as they say. Firms disregard the risk of being exploited by their competitors if they believe collusion to be the only way of making some money at all. Consequently, the risk of cartelisation is most pronounced if sunk costs are high and demand shrinks or is volatile. The quintessential example is the cement industry.

In principle, the theoretical literature invites an alternative interpretation of this finding. The way how we have designed the experimental market fits price competition. Each player is in a position to stop collusion, and to immediately impose the Walrasian result on her competitor. In the parlance of industrial economics: Nash reversion is unhindered. If firms compete in quantity, they at most can impose the Cournot equilibrium, which is the further away from the Walrasian solution the less firms compete, as has been explained above. Therefore the threat to punish a freerider by cutting into her profits in the subsequent period is very credible in the situation we have tested. Since players have more threat power, collusion is easier to sustain (Deneckere and Davidson 1985; Rotemberg and Saloner 1986). Note, however, that we have tested our subjects on a one-shot game, where (strategic) punishment is pointless. Therefore the alternative explanation cannot apply.

2. Repeated Interaction

a. Open vs. Defined End

Non-theorists tend to be surprised if they learn that the theoretical predictions remain unchanged if firms interact repeatedly. Intuitively, the shadow of the future should make collusion easier. Not exploiting your opponent today pays off tomorrow. Why should one be concerned that the opponent would oversee the obvious? Theory assumes perfectly prevoyant actors. In the last round, actors are back to the situation from the one-shot game. Each firm is at the mercy of its opponent. To preempt being the sucker in the last round, each firm has an incentive to defect in the penultimate round. This is anticipated by the respective opponent, and induces her to defect in the antepenultimate round, and so forth, until the entire game unravels. This thought has first been introduced in the tangential situation of the so-called chainstore paradox (Selten 1978), and it has been formalised under the graphic name of the centipede game (Rosenthal 1981).

The theoretical prediction changes if it is uncertain when the game ends. The so-called folk theorem formally proves that, then, cooperation in all periods is an equilibrium. However, any other symmetric or asymmetric, stable or alternating contribution pattern also is an equilibrium. While the dilemma from the one shot game disappears, firms face a serious problem of equilibrium selection (Aumann and Shapley 1994). Theory is thus unable to predict a concrete market outcome. Yet theoretically collusion is impossible when the end of interaction is known *ex ante*, while it is a distinct possibility if the interaction is open ended.

This is a testable proposition. As Figure 3 shows, the prediction is not borne out by the data. Let us first consider index CN, since this is the direct test. Note that the benchmark is taken from the one shot Nash equilibrium. As one sees, if the number of rounds is announced at the beginning, the one-shot Nash equilibrium is a very good predictor, both if experimental subjects compete in price and in quantity. However, if experimenters use some form of randomisation to determine the end of the game, this helps in a Bertrand setting, but it hurts in a Cournot setting. Apparently, if the safe profit is substantial, subjects are more impressed by the uncertainty inherent in the

openendedness of the game, while they see it as a tool to overcome unravelling if the safe profit is zero or small. The regression in Table 2 demonstrates that the finding is significant. Interestingly, there is neither a significant main effect of Bertrand versus Cournot, nor a significant main effect of defined versus open end. Only the interaction term is significant.




Figure 3
Open vs. Defined End in Oligopoly Games

	CN	CW
Cournot	6.088	23.015***
open ended	13.678	-5.873
Cournot*open ended	-46.766**	-19.157**
cons	-3.583	31.140***
N	392	465
p model	.0058	<.001
adj. R ²	.004	.071

Table 2
Explaining Collusion by Strategic Variable and Open Endedness³
OLS, robust standard errors, * p < .05, ** p < .01, *** p < .001
N is smaller with index CN since this index could not always be calculated

From a policy perspective, index CW is more interesting. For ultimately society has to endure a deadweight loss, and consumers lose in distributional terms, whenever price is above, or quantity is below, the Walrasian benchmark. While Figure 3 seems to indicate that there is even more social detriment if the end of the game is defined, regression analysis shows that this difference is insignificant, Table 2. Unsurprisingly, in this regression there is a significant main effect of Cournot. This is what theory predicts as well. However, the significant interaction is not predicted. Again, social detriment is reduced if subjects compete in quantity and if the game is open

3 Note that, in this and many later findings on the CN index, the adjusted R² is very small. This is due to the fact that the dataset pools findings from experiments with very different specifications in multiple dimensions. The fact that the adjusted R² on the CW index is generally much higher shows that this index is less sensitive to the resulting noise.

ended. From a policy perspective, this is welcome news. If the market is not one shot in the first place, it usually is also not determined *ex ante* at which point the market will close. Most markets are therefore openended. At least if firms compete in quantity, i.e. if capacity is an issue, experimental data suggests that this reduces the risk of collusion.

b. Explicit vs. Implicit Sanctions

Intuitively, one might expect collusion to be more sustainable if a firm risks being sanctioned when it breaches the cartel. Even if sanctions are costly, those punishing free riders might see this as an investment into a prosperous joint future. If actors are strictly maximising their profit, these intuitions are wrong. If there are more than two firms in the market, the explanation is straightforward. If one firm exerts effort to discipline a free rider, all other cartel members benefit equally. Each loyal firm is best off if other loyal firms do the job. The problem of free riding repeats itself at the second level (Yamagishi 1986; Heckathorn 1989). In a duopoly, one is back to the distinction between open ended games and games with defined duration. In the latter case, the unravelling result remains unaffected. In the former case, sanctions are just another of the multiple equilibria. From the perspective of profit maximising agents, sanctions only matter in a richer model with discounting. If sanctions are cleverly designed, cooperation becomes more robust with respect to discounting (Abreu 1988).

Moreover, while the enforcement of antitrust is notoriously deficient, the legal ban on formal cartel enforcement is fairly effective. Therefore explicit sanctions are typically unavailable. In principle, this is not a major drawback for cartel implementation. Loyal cartel members can threaten freeriders with setting the competitive price themselves. If they do, they kill the entire advantage from breaking the cartel. This sanction does not even entail an out of pocket cost. In principle, in a cartel, explicit sanctions can thus be substituted by a powerful implicit sanction. Actually, this sanction is even in line with a general recommendation for sustaining cooperation in repeated interaction. If a firm has violated the cartel once, following this recommendation it has to accept that others set the competitive price for one later period. During this period the former perpetrator is thus exploited herself. After this one period, all revert to cooperation for the rest of the game. This is an instance of the tit-for-tat strategy (Axelrod 1984). However, the strategy only works in a duopoly.⁴ In a larger market, if firm A has violated the cartel, even if firm B is happy to sanction, if it reduces price in the subsequent period it inevitably hits all other cartel members equally although they have been loyal. The implicit sanction is thus only a substitute for the explicit sanction if there is just one more firm. In larger markets, firms lack a technology for targeted sanctioning.

The empirical picture is fairly scattered. Again as a check for theory, index CN is more important. There is a clear difference between Bertrand and Cournot markets. In the latter, experimental subjects have a hard time to even end up in the Nash equilibrium. However, in duopoly markets, and only here, they do comparatively better if they compete in quantity. It thus seems that

4 And it is less powerful with Cournot than with Bertrand competition, see above at the end of section 1.

theory does a better job in Cournot markets. Put differently, in Bertrand markets, the general predictions for larger markets might be overshadowed by the fact that the safety payoff is so low. Main effects for market size are somewhat odd though. Deviations from the Nash equilibrium are strongly and significantly above the average of larger markets in markets of 3, while the main effects for market sizes 2 and 4 are insignificant, as is the general main effect for the number of firms in the market. Again, with index CW, the picture is much clearer. Here the main effect of the overall number of firms is significant, as are separate main effects for market sizes 2 and 3, while interactions are again somewhat scattered.

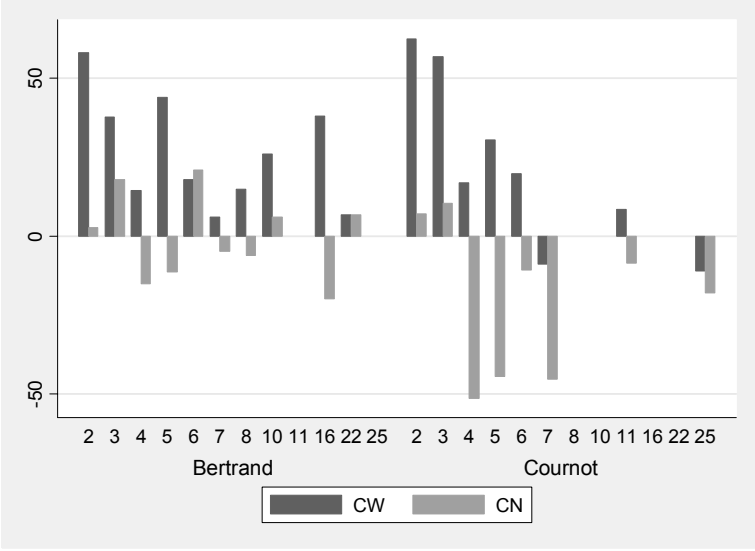


Figure 4
Effect of Market Size

	CN		CW	
size	-.514		-3.572***	
size 2		12.418		41.537***
size 3		27.707**		21.091***
Cournot	3.731	-27.359*	20.388***	.056
size*Cournot	-1.829		-1.067	
size2*Cournot		31.673(*)		4.334
size3*Cournot		19.717		19.168*
cons	.640	-9.732	45.113***	16.611***

N	392	392	465	465
p model	.1146	<.001	<.001	<.001
adj. R ²	-.004	.020	.150	.305

Table 3

Effect of Market Size

OLS, robust standard errors, (*) p < .1, * p < .05, ** p < .01, *** p < .001

c. *Conditional Cooperation*

While the empirical findings do not follow a simple pattern, theory definitely does not do a good job explaining them. One would either expect no deviations from the Nash prediction at all or, if any, they should be much more frequent in markets of two. Actually, there are many such deviations, especially in Bertrand markets. A concept from the theory of public goods seems more promising in explaining these findings: conditional cooperation.

Let us first consider the typical pattern, taken from our own data⁵. In a between subjects design, two different groups of subjects were presented with a linear public good⁶. They played the game in fixed groups of 4 over announced 10 periods. In the Voluntary Contribution Mechanism VCM no additional institutional framework was provided. In the Punishment treatment, participants had the opportunity to punish other group members, at a cost⁷. Initial contributions are much higher in the punishment treatment (Mann Whitney, N = 88, p = .0003), and they do practically not decay over time, whereas decay is pronounced in the VCM treatment. As a statistical check of this, a random effects regression is presented in Table 4. While the main effect of treatment is insignificant, both the main effect of the time trend and its interaction with the punishment treatment are significant. Actually, both regressors practically cancel out, which corroborates the flat trend in the punishment treatment from Figure 5.

5 The data is taken from two baseline treatments, one without and one with punishment. 88 subjects were tested. Per period they had an endowment of 20 tokens. The marginal per capita rate was set at .4.

6 For exact model specifications, see the formalisation below. The VCM data is pooled over two different (baseline) treatments. In the second VCM, the game was over 12 periods. Only data from periods 1-10 is included in the analysis. If one adds the remaining data points, qualitatively the picture does not change, nor do significance tests.

7 Punishment technology was as in (Fehr and Gächter 2000).




Figure 5
Linear Public Good With and Without Punishment

punishment treatment	2.741
period	-.798***
pun treatment*period	.796***
cons	15.728***
N	880
p model	<.001

Table 4
Linear Public Good With and Without Punishment

Random Effects, robust standard error, clustered at the group level, depvar: contribution,
Hausman test insignificant, *** p < .001

These patterns fit a population with a substantial fraction of conditional cooperators. A person is classified as a conditional cooperator if she is willing to contribute to the public good as long as the other group members contribute as well (Fischbacher, Gächter et al. 2001; Fischbacher and Gächter 2008). This can explain cooperation in groups of three or more members, provided initially the fraction of (conditional and unconditional) cooperators is large enough, and provided conditional cooperators are willing to give cooperation a chance at the beginning of the game. Then they update their beliefs about cooperativeness in this group, in light of the experiences they are making. Every single bad experience further deteriorates the overall level of cooperativeness, resulting in the downward trend.

Conditional cooperators are neither stubborn egoists nor unswerving altruists. It however is still largely unclear which motivational forces are exactly at work. One interpretation is exclusively forward looking. Conditional cooperators are willing to take a calculated risk, hoping that they will be rewarded with a higher overall payoff at the end of the game. In the alternative interpretation, conditionally cooperative subjects care about relative profit. They are inequity averse (cf.

Fehr and Schmidt 1999). Their willingness to contribute is conditional on the expected size of the difference between their payoff and the payoff of the group member that has contributed least and therefore made most money.

In both perspectives, punishment is helpful. In the first interpretation, punishment forces freeriders to raise their contributions. Consequently, even if selfish group members have not anticipated punishment, they are likely to adjust their behaviour. This is the forward looking effect. Punishment is a deterrent. In the second interpretation, punishment, on top, has a backward looking effect. The punished individual's income is reduced. Thereby all loyal contributors suffer smaller inequity disutility. Actually there is even a third effect, which is confined to the punishing individual. This individual gains expressive utility from being the one who inflicts harm on the freerider.

The easiest way to show that there is conditional cooperation at all is an experimental procedure that elicits conditional responses. The so-called strategy method does precisely this (Selten 1967). Using this method, it has been shown that, indeed, a substantial fraction of experimental subjects cooperates conditional on other group members' contributions (Fischbacher, Gächter et al. 2001). The strategy method has one disadvantage for the purpose: the condition is defined with certainty. This is not the environment contributors face in a public good game, and even less in real life dilemma situations. This additional uncertainty is captured by a design that has participants play for real but elicits beliefs before choices are made. Interestingly, using the latter design slightly dampens the willingness to contribute. The data is best explained by mixed motives: subjects are willing to contribute conditional on the expected contributions of others, but they also want to have a slightly higher profit than their peers (Fischbacher and Gächter 2008).

One way to show that conditional cooperation is sensitive to relative payoff tests subjects twice. Together with the dilemma, they are assigned the role of dictator in a structurally similar dictator game. If subjects were willing to share freely in the dictator game, but defect in the dilemma game, they must have been afraid of losing in relative terms. Hans Normann and myself have added a dictator game variant to the prisoner's dilemma reported above (Engel and Normann 2008). Again using the strategy method, we asked our subjects to choose between 5 € for sure and a lottery with 5 € as the minimum outcome, and the chance to get a higher outcome that we varied in 11 steps. If they chose the safe outcome, their random counterpart also received 5 €. If they chose the lottery, their random counterpart received nothing. We used this somewhat more complicated procedure, rather than a straightforward dictator game, to control for the uncertainty inherent in the prisoner's dilemma. In neither game, gains from defection are certain. However, only in the prisoner's dilemma must a subject be afraid of being exploited by her counterpart.

The scatterplot in Figure 6 classifies players in both dimensions: cooperativeness in the dictator game (x-axis) versus cooperativeness in the prisoner's dilemma (y-axis). As predicted, there is systematically less cooperation in the prisoner's dilemma (the regression line has a slope smaller than 1). Actually the dots in the upper right corner are unswerving altruists, the dots in the lower left corner are selfish players. Conditional cooperators are below the 45° line. The smaller will-

ingness to cooperate in the prisoner's dilemma is motivated by comparing themselves to their counterpart. A Wilcoxon signed rank test shows that the willingness to cooperate is significantly smaller if cooperation comes at the risk of being exploited oneself ($N = 105$, $p < .001$).




Figure 6
Cooperativeness With and Without Risk of Being Exploited

To the best of my knowledge, conditional cooperation has not been used as an explanation for oligopoly data. Since it has been investigated in linear public goods, it is interesting to know whether an oligopoly can be interpreted as a linear public good of oligopolists. This would be a necessary condition for using the explanation in the oligopoly context. Formally, a linear public good has the following structure:

$$\pi_i = e_i - g_i + \mu \sum_{j=1}^N g_j$$

where π_i is this player's payoff, e_i is her endowment, g_i is the fraction of the endowment she contributes to the project. μ is the marginal per capita rate. It determines how much contributions to the project pay. Finally $j \in N$ is one of the total of group members, including the player under consideration.

In an oligopoly, if a firm assumes all other firms in the industry to be strict profit maximisers, it expects all other firms to set the Nash price or quantity. The best response is to set this price or this quantity itself. The picture changes if a firm expects other firms to be loyal to the cartel with positive probability p . Let us first consider a duopoly. The first firm expects collusive profit π_{CC} with probability p . However, if the other firm is loyal with positive probability, this also implies that this firm has a chance to get the exploitation profit π_{DC} . In a duopoly, this chance

also has probability p ⁸. Hence the firm under consideration compares the payoff from cooperation π_C with the payoff from defection π_D :

$$\pi_C = p\pi_{CC} + (1-p)\pi_{CD}$$

$$\pi_D = p\pi_{DC} + (1-p)\pi_{DD}$$

If we write χ_i for the fact that the firm under consideration is loyal to the cartel, we can combine all terms, to get

$$\pi_i = \chi_i p \pi_{CC} + \chi_i (1-p) \pi_{CD} + (1-\chi_i) p \pi_{DC} + (1-\chi_i) (1-p) \pi_{DD}$$

In oligopoly, the functional equivalent to the endowment is the safety payoff π_{DD} . No firm can be forced to accept a lower payoff. Now in oligopoly payoffs are ranked

$$\pi_{DC} > \pi_{CC} > \pi_{DD} \geq \pi_{CD}$$

We therefore can normalize payoffs with respect to the safety payoff π_{DD} . If a firm is loyal to the cartel, it exposes itself to the risk of being exploited. But this is not the only sacrifice it makes for sustaining cooperation. It also gives up the opportunity to exploit its opponent. Hence, the expected payoff can be written as a function of the decision to cooperate oneself, to get

$$\pi_i(\chi_i) = \pi_{DD} - (1-\chi_i)p(\pi_{DC} - \pi_{DD}) - \chi_i(1-p)(\pi_{DD} - \pi_{CD}) + \chi_i p(\pi_{CC} - \pi_{DD})$$

This reformulation shows two things: one can indeed interpret a duopoly as a linear public good. However, the elements differ in many respects, and these differences are likely to matter in behavioural terms. The analogue of the endowment e_i is the safety payoff π_{DD} . The benefit from cooperation is in the last term. It consists of the difference between the cooperation payoff π_{CC} and the safety payoff. There is no direct equivalent of the marginal per capita rate μ . Instead there is an element not present in a standard public good, $\chi_i p$. The first component indicates that the public good is a “weakest link good” (Cornes and Sandler 1996:186): if this player defects, she herself does also not benefit from the public good. The second component indicates that the contribution of this player is fully lost if the other player defects. In a standard public good, the minimum return is the own contribution, multiplied by $\mu < 1$.

The largest difference, however, concerns contributions. In a standard public good, the contribution is simply g_i . The player gives up a fraction of her endowment. One can either interpret this as an out of pocket cost or as an opportunity cost. In duopoly, the contribution is much more complicated. It consists of two separate terms. The second term $\chi_i(1-p)(\pi_{DD} - \pi_{CD})$ represents an out of pocket cost. By contributing, the firm exposes itself to the risk of being exploited by the opponent. The first term $(1-\chi_i)p(\pi_{DC} - \pi_{DD})$ is an opportunity cost. By contributing, the

8 By the logic explained above, intermediate actions are irrational. Firms only choose between the collusive price/quantity, and the Nash price/quantity.

firm foregoes the opportunity to exploit the opponent. Both terms are qualified by the uncertainty about the other firm's action.

In duopoly, the parallel to a linear public good is technically easier to draw. But in markets of two, collusion could also be explained by a tit for tat logic. The reason to explore the parallel is collusion in larger markets. The mapping of the safety payoff π_{DD} is not sensitive to the number of suppliers⁹. Gains from collusion depend on the willingness of all market participants to cooperate, i.e. the last term becomes $\chi_i p^{N-1} (\pi_{CC} - \pi_{DD})$. Whenever at least one firm defects, all other firms only receive the worst outcome π_{CD} . The size of the loss, compared to the safety payoff, is thus not sensitive to the number of defectors. However, the larger the market, the more it becomes likely that at least one firm violates the cartel. Hence the third term becomes $\chi_i (1 - p^{N-1}) (\pi_{DD} - \pi_{CD})$. The second term is most complicated. Here, not only probabilities matter. Also the size of the gain depends on the number of defectors. If otherwise all firms are equal, defectors share the existing defection gains equally. Consistency requires $\pi_{DC} / N = \pi_{DD}$: if all defect, there is no defection gain to be distributed. To express the opportunity cost generically, one needs binomial expansion, to get

$$\sum_{r=1}^N p^{\frac{(N-1)!}{(N-1-r+1)!(r-1)!}} \frac{\pi_{DC}}{r}$$

where r is the number of defectors. Putting elements together, for larger markets the representation as a linear public good looks like this:

$$\pi_i = \pi_{DD} - (1 - \chi_i) \sum_{r=1}^N p^{\frac{(N-1)!}{(N-1-r+1)!(r-1)!}} \frac{(\pi_{DC} - \pi_{DD})}{r} - \chi_i (1 - p^{N-1}) (\pi_{DD} - \pi_{CD}) + \chi_i p^{N-1} (\pi_{CC} - \pi_{DD})$$

Behaviourally, all qualifications are likely to matter the same way, albeit of course not by the same size, as in duopoly.

III. Competition as a Socially Desirable Dilemma

All of the foregoing has exclusively interpreted oligopoly as an internal problem of oligopolists, neglecting the social harm successful internal cooperation imposes on the opposite market side. To a degree, potential victims are able to protect themselves (1). On top of this, internals might be reluctant to knowingly impose harm on outsiders (2), and they might be sensitive to normative expectations enshrined in antitrust legislation (3).

1. Defense

In the quintessential market, a small number of providers serve a large number of buyers. Not so rarely, the asymmetry turns into one between organised supply and diffuse demand (Olson

9 But the size of π_{DD} is affected if firms compete in quantity, or if they compete in price, but products are heterogenous.

1965), pretty much leaving demand at the mercy of supply. One characteristic element of seller power has been studied very carefully experimentally. If sellers can expose buyers to a take it or leave it offer, they extract a lot of consumer surplus. Collusion drops if there are (balanced) negotiations, and even more so in a sealed bid auction, or in a double oral auction, Figure 7. Regression analysis shows that the difference between the posted offer and the remaining protocols is significant, as is the interaction with the number of firms in the market, Table 5.




Figure 7
Trading Protocol

	CN	CW
posted offer	16.467*	57.972***
number	.399	-.420
poster offer*number	-1.800 ⁺	-3.295***
cons	-8.946 ⁺	3.070
N	395	468
p model	.0075	<.001
adj. R ²	-.000	.289

Table 5
Trading Protocol

OLS, robust standard errors, ⁺ p < .1, * p < .05, ** p < .01, *** p < .001

2. Respect for Outsiders

Arguably, increasing one's own profit is somewhat less attractive if one has to impose harm on outsiders. This is what theories of inequity aversion would predict (Fehr and Schmidt 1999). One finding from oligopoly experiments also points into this direction. Usually, demand is represented by a computer that buys according to a predefined demand function. If experimenters have replaced the computer by real subjects, collusion plummets, Figure 8. Both for the CN in-

dex (Mann Whitney N = 404, p = .0051) and for the CW index, the difference is significant (N = 479, p <.001).




Figure 8
Human versus Computer Buyers

3. Normativity

In the oligopoly data, there is also a proxy for normativity. While most oligopoly experiments deliberately frame the interaction as the supply side of a market, as a check experimenters investigated how results change if the very same incentive structure is presented neutrally. Actually, the difference is pronounced, Figure 9, Table 6.




Figure 9
Neutral versus Market Frame

	CN	CW
market frame	-42.448***	-16.336*
Cournot	-17.611	18.5
market frame*Cournot	14.198	-1.048
cons	39.444	46.667
N	388	460
p model	<.001	<.001
adj. R ²	.003	.056

Table 6
Market versus Neutral Frame

OLS, robust standard errors, * p < .05, ** p < .01, *** p < .001

While these are interesting hints, neither of these findings can serve as proof. This motivated an experiment by Bettina Rockenbach and myself (Engel and Rockenbach 2009). We had groups of 7 interacting over 10 announced rounds, who were randomly split into 4 active and 3 passive players. In the no externality condition, the passive players just kept their endowment. The active players played a standard linear public good, with the same parameters as above. In the externality condition, any point contributed to the public good reduced the period income of each passive player by .2 points. In the second phase, before the 7 players were randomly reassigned roles, they took a non-binding vote about expected contribution levels. The 4 active players of this phase then repeated the game for another 10 rounds. Figure 10 shows that, in line with the findings from the oligopoly literature, the norm reduced contributions in the externality condition. However, in the first 10 periods, contrary to the findings from the oligopoly literature, subjects contributed even more to the public good if this did harm to outsiders.




Figure 10
Externalities and an Endogenous Norm in a Linear Public Good

This is data on 2 times 10 periods of subjects interacting in groups of 4 (7). More than half of the contributions were at the extremes, i.e. either 0 or the full endowment. This is why a random effects Tobit model is appropriate. It demonstrates a significant main effect of the externality manipulation, and a highly significant main effect of introducing the norm, plus a significant interaction effect. I.e. the picture from Figure 10 is fully supported by econometric analysis¹⁰.

externality	6.627*
norm	9.129***
externality*norm	-17.183***
cons	-10.203*
N	1440
p model	<.001

Table 7
 Externalities and an Endogenous Norm in a Linear Public Good
 Random Effects Tobit, group fixed effects (not reported), * p < .05, ** p < .01, *** p < .001

IV. Antitrust Learning from the Lab

This article is published in a volume on the “more economic approach” to antitrust. It encompasses both theoretical and empirical work, including case studies and econometrics. In light of this general receptivity for economic findings and methodology, it is surprising how little attention is paid to the rich body of experimental evidence. For sure, many features of a market an antitrust practitioner might want to understand have not been tested experimentally. For instance to the best of my knowledge markets characterised by economies of scale or scope have not been tested in the lab¹¹. Even worse, some practically relevant features would be very different to implement in the lab, like the presence of a “maverick” firm¹². Most firms are neither owned nor managed by a single individual. While the effects of interacting in a group have been studied extensively in the lab (for a survey Engel 2008), none of these studies have been done within experimental oligopoly markets. Moreover, ad hoc groups impose much less structure than the long term interaction at the interior of a firm. Last but not least, experiments are conducted because the researcher wants to get rid of the identification problem. If the experiment is properly designed, one may be certain about the arrow of causation. Yet the biggest advantage of experiments is also their biggest drawback. To achieve full control, experiments deliberately strip off the situation. If an effect is found in the clean setting of the lab, the same effect can be absent, or different, in the field.

10 Actually, there is no generally acknowledged fixed effects estimator for censored data. As a double check, the Hausman test is performed on the mirror model that ignores censoring. It turns out insignificant.

11 For a list of the lacunae in the experimental literature on oligopoly, see the conclusions of (Engel 2007a).

12 This question featured prominently in the Airtours ruling of the European Court of First Instance. In (Engel 2007b), I am exploring the case from an experimental perspective.

Antitrust practitioners would therefore certainly not want to take action solely based on an experiment with a number of student subjects. But experiments are a valuable complement to other theoretical and empirical methods for understanding competition. They can serve as a double check for effects predicted by theory, or found in the field. They can point antitrust practice to qualitatively new and surprising effects, like the sensitivity of collusion to the size of the safety payoff. And they can solve identification problems.

References

- ABREU, DILIP (1988). "On the Theory of Infinitely Repeated Games with Discounting." Econometrica **56**(2): 383-396.
- AUMANN, ROBERT J. and LLOYD S. SHAPLEY (1994). Long Term Competition - A Game Theoretic Analysis. Collected Papers I. Robert J. Aumann. Cambridge, MIT Press: 395-409.
- AXELROD, ROBERT (1984). The Evolution of Cooperation. New York, Basic Books.
- BERTRAND, JOSEPH LOUIS FRANCOIS (1883). "Théorie mathématique de la richesse sociale par Léon Walras. Recherches sur les principes mathématiques de la théorie des richesses par Augustin Cournot." Journal des savants **67**: 499-508.
- BORK, ROBERT H. (1978). The Antitrust Paradox. A Policy at War with Itself. New York, Free Press u.a.
- CORNES, RICHARD and TODD SANDLER (1996). The Theory of Externalities, Public Goods and Club Goods. Cambridge, Cambridge University Press.
- COURNOT, ANTOINE-AUGUSTIN (1838). Recherches sur les principes mathématiques de la théorie des richesses. Paris, Librairie des sciences politiques et sociales M. Riveáre & cie.
- DAVIDSON, CARL and RAYMOND J. DENECKERE (1986). "Long-Run Competition in Capacity, Short-Run Competition in Price, and the Cournot Model." Rand Journal of Economics **17**: 404-415.
- DENECKERE, RAYMOND J. and CARL DAVIDSON (1985). "Incentives to Form Coalitions with Bertrand Competition." Rand Journal of Economics **16**: 473-486.
- EDGEWORTH, FRANCIS (1897). "La Teoria Pura del Monopolio." Giornale degli Economisti **40**: 13-31.
- ENGEL, CHRISTOPH (2006). Wettbewerb als sozial erwünschtes Dilemma. Recht und spontane Ordnung. Festschrift für Ernst-Joachim Mestmäcker zum achtzigsten Geburtstag. Christoph Engel und Wernhard Möschel. Baden-Baden, Nomos: 155-198.
- ENGEL, CHRISTOPH (2007a). "How Much Collusion? A Meta-Analysis on Oligopoly Experiments." Journal of Competition Law and Economics **3**: 491-549.
- ENGEL, CHRISTOPH (2007b). Tacit Collusion. The Neglected Experimental Evidence <http://ssrn.com/abstract=1020889>.
- ENGEL, CHRISTOPH (2008). The Behaviour of Corporate Actors. A Survey of the Empirical Literature

- ENGEL, CHRISTOPH, BERND IRLENBUSCH, SEBASTIAN KUBE, ANDREAS NICKLISCH, HANS-THEO NORMANN and EMANUEL TOWFIGH (2009). Beware of Broken Windows! First Impressions in Public Good Experiments
- ENGEL, CHRISTOPH and HANS-THEO NORMANN (2008). Greed and Fear in a One-Shot Prisoner's Dilemma. An Experimental Analysis
- ENGEL, CHRISTOPH and BETTINA ROCKENBACH (2009). We Are Not Alone. The Impact of Externalities on Public Good Provision
- FEHR, ERNST and SIMON GÄCHTER (2000). "Cooperation and Punishment in Public Goods Experiments." American Economic Review **90**: 980-994.
- FEHR, ERNST and KLAUS M. SCHMIDT (1999). "A Theory of Fairness, Competition, and Cooperation." Quarterly Journal of Economics **114**: 817-868.
- FEUERSTEIN, SWITGARD (2005). "Collusion in Industrial Economics - A Survey." Journal of Industry, Competition and Trade **5**: 163-198.
- FISCHBACHER, URS and SIMON GÄCHTER (2008). Heterogeneous Social Preferences and the Dynamics of Free Riding in Public Good Experiments
http://home.twi.uni-konstanz.de/files/twi_research/49_No27-08-05-TWI-RPS-Fischbacher-Gaechter.pdf.
- FISCHBACHER, URS, SIMON GÄCHTER and ERNST FEHR (2001). "Are People Conditionally Cooperative? Evidence from a Public Goods Experiment." Economics Letters **71**: 397-404.
- GRANOVETTER, MARC (1985). "Economic Action and Social Structure. The Problem of Embeddedness." American Journal of Sociology **91**: 481-510.
- HECKATHORN, DOUGLAS D. (1989). "Collective Action and the Second-Order Free-Rider Problem." Rationality and Society **1**: 78-100.
- HUCK, STEFFEN and BRIAN WALLACE (2002). "Reciprocal Strategies and Aspiration Levels in a Cournot-Stackelberg Experiment." Economics Bulletin **3/3**: 1-7.
- KREPS, DAVID M. and JOSE ALEXANDRE SCHEINKMAN (1983). "Quantity Precommitment and Bertrand Competition Yields Cournot Outcomes." Bell Journal of Economics **14**: 326-337.
- LEDYARD, JOHN O. (1995). Public Goods. A Survey of Experimental Research. The Handbook of Experimental Economics. J.H. Kagel und A.E. Roth. Princeton, NJ, Princeton University Press: 111-194.
- LEVENSTEIN, MARGARET C. and VALERIE Y. SUSLOW (2006). "What Determines Cartel Success?" Journal of Economic Literature **44**: 43-95.

- OLSON, MANCUR (1965). The Logic of Collective Action. Public Goods and the Theory of Groups. Cambridge, Mass., Harvard University Press.
- OSTROM, ELINOR (1990). Governing the Commons. The Evolution of Institutions for Collective Action. Cambridge, New York, Cambridge University Press.
- RAPOPORT, ANATOL and ALBERT M. CHAMMAH (1965). Prisoner's Dilemma. A Study in Conflict and Cooperation. Ann Arbor,, University of Michigan Press.
- ROSENTHAL, ROBERT W. (1981). "Games of Perfect Information, Predatory Pricing and the Chain Store Paradox." Journal of Economic Theory **25**: 92-100.
- ROTEMBERG, JULIO J. and GARTH SALONER (1986). "A Supergame-theoretic Model of Price Wars During Booms." American Economic Review **76**(3): 390-407.
- SELTEN, REINHARD (1967). Die Strategiemethode zur Erforschung des eingeschränkt rationalen Verhaltens im Rahmen eines Oligopolexperiments. Beiträge zur experimentellen Wirtschaftsforschung. Ernst Saueremann. Tübingen, Mohr: 136-168.
- SELTEN, REINHARD (1978). "The Chain Store Paradox." Theory and Decision **9**: 127-159.
- SLADE, MARGARET EMILY (1990). "Strategic Pricing Models and Interpretation of Price-War Data." European Economic Review **34**: 524-537.
- TIROLE, JEAN (1988). The Theory of Industrial Organization. Cambridge, Mass., MIT Press.
- YAMAGISHI, TOSHIO (1986). "The Provision of a Sanctioning System as a Public Good." Journal of Personality and Social Psychology **51**: 110-116.
- ZELMER, JENNIFER (2003). "Linear Public Goods. A Meta-Analysis." Experimental Economics **6**: 299-310.