

The Critical Discount Factor as a Measure for Cartel Stability?

Lisa V. Cwojdzinski[†]

December 2006

Abstract

This paper considers the stability of tacit collusion in price setting duopolies with repeated interaction. The critical discount factor δ^* , which is the minimum discount factor above which tacit collusion can be sustained in a subgame perfect equilibrium, is often used as a measure for stability of tacit collusion. A collusive equilibrium is then assumed to be more difficult to sustain when δ^* increases. The distance between the actual discount factor δ and the critical value δ^* , however, does according to standard theory not matter as long as δ is larger than δ^* . This paper contributes experimental evidence that supports the intuitive idea that a larger critical discount factor makes collusion a less likely outcome, even though the value of δ^* is still below the actual δ .

JEL classification: C12, C71, C91, L13

Keywords: Bertrand duopoly, Sequential game, Collusion, Critical discount factor, Experiment.

[†] Humboldt-University Berlin, Department of Economics, email: lisacwoj@wiwi.hu-berlin.de
I would like to thank Dirk Engelmann, Urs Fischbacher, Ulrich Kamecke, Elmar Wolfstetter and seminar audiences at Paris and Berlin for many helpful comments. Sebastian Schwenen and Wolfram Wilde provided excellent research support.

1 Introduction

We consider the pricing behavior of firms in an infinitely repeated Bertrand oligopoly. Usually explicit price agreements in such a market are prohibited by law, but firms may nevertheless manage to coordinate their pricing strategies with less obvious methods. In this case, where coordination of prices takes place in a more 'mysterious' way, we speak of tacit collusion. The minimum discount factor above which collusion can be sustained in a subgame perfect equilibrium (the critical discount factor δ^*) is often used as a measure for the stability of such tacit agreements. As δ^* increases, a collusive equilibrium is usually assumed to be more difficult to sustain. This characterisation of the impact of a changing critical discount factor is specified by Symeonidis (2002): 'A standard way to examine the impact of exogenous factors on cartel sustainability [...] is to examine the comparative static properties of the critical discount factor δ^* . In particular, a change in any exogenous variable that causes δ^* to increase makes collusion less likely, since collusion is then sustainable for a smaller set of deltas.' There are several experiments on Bertrand as well as Cournot competition, considering various market design variables with respect to their influence on stability of collusive behavior. When analysing the impact of all those factors most of them implicitly presume that the critical discount factor comprises a measure for stability of cooperative behavior in the market.¹

In theory, however, the critical discount factor should only matter for firms' behavior in so far as collusion is a sustainable outcome when their actual discount factor δ is larger than the critical δ^* , but not when it is smaller (see Figure 1). According to the theory, as long as δ is above δ^* , we expect collusion to be stable. Conversely, as soon as δ is below δ^* no collusion should be possible. The size of the distance in between the actual discount factor and its critical value, in contrast, should not matter at all. Either pricing agreements are stable or they are not, but 'a little' or 'less' stability can theoretically not occur.² A change in the critical discount factor δ^* , in literature often treated as a measure of cartel stability, does per se have no impact on the extent of tacit collusion when δ was still larger than δ^* .

¹One empirical example for this statement is the Nestle-Perrier case considered by Compte et al. (2002). Against the background of the merger of two large French suppliers of mineral water they demonstrate how varying capacities through mergers influence the stability of collusion. Within this context an increase in the critical discount factor is considered as a measure of less stability.

²The same could be tested for the opposite case where δ is always below δ^* . Here we would expect cooperation not to be stable in both cases. However, in the case of a higher critical discount factor (with a larger distance to the actual δ) it could be that cooperative attempts are more likely to succeed.




Figure 1: : Stability of tacit collusion depending on discounting.

2 Theoretical Background

The work probably closest to our paper is the study of Feinberg and Husted (1993), who considered the critical delta experimentally in a Cournot environment. However, their experiment differs in various aspects from what we want to show. The main difference is that they keep the critical discount factor constant whereas actual discounting is manipulated in one treatment with a decreasing payment structure and in another with a random termination rule. In the reasoning here presented actual discounting will stay the same while δ^* varies. Second, Feinberg and Husted reduce their δ so drastically that cooperation should collapse even in theory. In our experiment, in contrast, the main interest is the comparison of two settings that theoretically should both be stable, because δ is always above δ^* . Finally, their experiment uses the strong simplification that only two different stages of cooperation and no cooperation exist since they simply consider a repeated prisoner's dilemma. We will extend the choice set and include more options. Similar comments hold for Chincarini 2003.

In order to consider the role of the critical discount factor in more detail, a different experimental design is necessary. To make the interaction as simple as possible, a duopoly framework is selected. The new idea is to compare two similar treatments of a price-setting game that should theoretically both yield the same cooperative outcome. As long as δ^* is sufficiently low (lower than the actual 'patience'-parameter δ) participants' behavior should not differ between the treatments because the game they are playing is independent of the distance between δ and δ^* . If, in contrast, the critical discount factor indeed has an impact on the game played, the treatments should yield different behavior in a way that collusive agreements are less stable when δ^* is higher. If behavior can be manipulated this way, the hypothesis would be confirmed that a higher critical discount factor, although below the actual discount factor, influences equilibrium behavior towards

less cooperation. Intuitively, we can think of the difference $\delta - \delta^*$ as a measure for the 'temptation' to deviate from cooperative behavior. The larger this difference the smaller is the incentive to deviate. For $\delta - \delta^* > 0$ theory predicts completely stable collusion, no matter how large the distance is. However, if δ is only a little larger than δ^* the long run profits of a deviation are nearly equally as large as the long run profits from continuous cooperation. The prospect of a high short-run gain might then suffice to induce a deviation. This attractiveness of the high short-run profit is more tempting if the difference $\delta - \delta^*$ is small and, thus, when δ^* is large.

Looking for a theoretical frame for the experiment, the classical Bertrand supergame is not the most appropriate model because it exhibits a problem with multiplicity of equilibria. Depending on the strategy formation, each equilibrium profit between the Walrasian outcome (Bertrand-Nash-equilibrium with prices equal to marginal costs) and joint profit maximisation is a possible equilibrium in the repeated game. Since participants are free to choose any price they want and since they are playing for several rounds, the possible set of equilibria is overwhelming (see, for example, Shapiro 1989). But without a clearly defined strategy assumption, discount factors cannot even be calculated. Theory usually simplifies the oligopoly situation to the two-choice prisoner's dilemma game and considers only trigger strategies, where cooperation means joint profit maximisation and non-cooperation yields the strongest punishment possible, the pure Bertrand Nash Equilibrium with zero profits. If the game is repeated infinitely, of course, all outcomes in between those extremes are also possible equilibria. To calculate a unique δ^* in a repeated game, however, strategies and resulting equilibria must be clearly defined.

To reduce multiplicity of equilibria, Maskin and Tirole (1988) offer an interesting approach. They present a sequential move framework, which has some useful properties the Bertrand supergame lacks. In sequential games, the key results concerning critical discount factors in infinitely repeated games are the same as in simultaneous Bertrand competition (Tirole 1988). However, sequential play allows restricting participants' behavior to Markov strategies and therefore reduces the possible equilibria to a smaller set. Intuitively, firms do not have to build beliefs about the other firm's future price decision in the following period to optimise their behavior because they are currently the only one adjusting the price.³ According to the Markov restriction Maskin and Tirole impose, players' strategies depend only on profit-relevant state variables. In a sequential game the state can simply be defined as the opponent's price in the previous period because this

³This is similar to the inertia introduced in the game by Huck et al. (1999). The idea of sequential moves in a Bertrand model is in fact not very revolutionary since there are already experiments that introduced inertia into pricing oligopolies. From inertia, it is only one step further to come to actual sequential moves.

is the only profit relevant information determining a firm's pricing decision. Consequently, firms' reaction to a certain price has always to be the same. The Markov assumption in a sequential game makes equilibrium analysis much easier. It is neither necessary to care about firms' beliefs about their opponent's behavior nor does history more than one period ago matter. And the multiplicity of equilibria decreases drastically, since the set of possible profits can be restricted to a tighter interval (see Propositions 3 and 7 in the Maskin and Tirole paper). Of course, Markov is a somewhat restrictive assumption. The experimental test of the model will show that it is reasonable, though.

The basic idea of the Maskin and Tirole model is as follows: in an infinite horizon sequential duopoly game, two firms ($i = 1, 2$) compete in prices. In odd-numbered periods, firm 1 chooses its price, whereas in even-numbered periods firm 2 sets prices. Thus, each firm is limited to change its price in every second period and to leave it unchanged for the subsequent period. The total industry profit in each period depends on price p , costs c , and demand $D(p)$:

$$\Pi(p) \equiv (p - c) \cdot D(p).$$

Consumers buy from the cheapest supplier, thus individual profits are determined as a function of prices.

$$\Pi^i = \Pi(p_t^1, p_t^2) = \begin{cases} \Pi(p_t^i), & \text{if } p_t^i < p_t^j \\ \Pi(p_t^i)/2, & \text{if } p_t^i = p_t^j \\ 0, & \text{if } p_t^i > p_t^j \end{cases}$$

Future profits are discounted with factor δ . Intertemporal profit is therefore

$$\sum_{s=0}^{\infty} \delta^s \cdot \pi^i(p_{t+s}^1, p_{t+s}^2).$$

Markov Perfect Equilibria are within this framework derived as pairs of dynamic reaction functions (R^1, R^2) for all prices \hat{p} . R^1 gives the best reply if firm 1 itself is about to move, whereas the distribution of R^2 is (from firm 1's point of view) the distribution of best reply strategies of firm 2. Both reaction functions are based on valuation functions for all possible price decisions. These valuations are for firm 1

$$V^1(\hat{p}) = \max_p [\pi^1(p, \hat{p}) + \delta \cdot W^1(p)]$$

if it is firm 1's turn to move and firm 2's current price is \hat{p} and

$$W^1(\hat{p}) = E_p[\pi^1(\hat{p}, p) + \delta \cdot V^1(p)],$$

if last period firm 1 played \hat{p} and firm 2 is about to move.

$V^1(\hat{p})$ gives firm 1's present discounted profit of choosing p and $W^1(\hat{p})$ is firm 1's valuation of the current situation if firm 2 is about to move while firm 1 is still bound to the price it set in the period before. For firm 2 symmetric conditions hold.

As long as the discount factor is sufficiently high two different types of equilibria are possible within the model, Edgeworth cycles and focal price equilibria. Within a focal price equilibrium, the pair of strategies (R^1, R^2) results in pricing behavior, where it is optimal for both firms to collude at the focal price (for example, the monopoly price). At any price above, the focal price is a best reply, whereas at any price below, firms sooner or later revert to collusion, probably after some 'punishment' phases. The other class of Markov perfect equilibria comprises cyclical price patterns where firms, starting from a very high price level, successively undercut each other until the price war becomes too costly and then return to a price above the monopoly level. Strategies of this type are called Edgeworth cycles.

Additionally to the Markov restriction to firms' behavior Maskin and Tirole require a stable equilibrium to be renegotiation-proof and, thus, to suffice stronger conditions than pure subgame-perfection. A Markov perfect equilibrium is therein called renegotiation-proof, if there exists no other equilibrium that Pareto-dominates it⁴. In our tacit collusion framework, this idea refers to situations in which firms might have an incentive to reconsider the strategies they committed to, for example after a defection from cooperation. If we require renegotiation-proofness we do not allow strategies where firms might be tempted to change their mind at some point of time. The requirement of renegotiation-proofness disqualifies most equilibria in the sequential price game. Neither focal price equilibria besides monopoly pricing nor Edgeworth cycles are

⁴Van Damme (1989) already introduced renegotiation-proofness into a repeated prisoners' dilemma game. He applies a similar concept of renegotiation-proofness and shows that the multiplicity of equilibria can be reduced to a smaller set. However, his approach is a different one than that of Maskin and Tirole. Both have in common that they name a pair of strategies renegotiation-proof if it automatically leads to a payoff path where at no stage the payoff is dominated by another reaction than prescribed by the strategy. In the simultaneous-move framework of Van Damme this requires that all intermediate equilibria on the payoff path have to be renegotiation-proof, too. There are many different possible equilibria, all of which are renegotiation-proof.

robust against temptation to renegotiate because they are Pareto-dominated by strategies that lead to the monopoly price as focal price. Thus, Maskin and Tirole show that the simple monopoly price equilibrium is the only one that is renegotiation-proof⁵ so that neither firm would prefer to establish any other level of collusion (see Proposition 6 of their paper). Of course, the unique renegotiation-proof equilibrium strategy for completeness has to contain a reaction to all other possible prices, too, in case they nevertheless occur. In detail, it comprises the following pair of symmetric strategies:

- Set the monopoly price p^m if the other firm did so in the previous period
- If the other firm deviates, set a price of \underline{p}
- If the other firm set \underline{p} , return to p^m

\underline{p} is determined in a way that it satisfies $4 \cdot \Pi(\underline{p}) \geq \Pi(p^m) > 4 \cdot \Pi(\underline{p} - k)$ with k being the (sufficiently fine) price grid.

The renegotiation-proof focal price equilibrium concept of Maskin and Tirole is appropriate to illustrate tacit collusion. First, it contains a unique renegotiation-proof equilibrium outcome, namely collusion at the monopoly price. Thus, it precisely predicts equilibrium behavior in a way that upcoming deviations, which are in the original centre of interest, can be clearly observed. Second, it nevertheless contains a detailed prediction of the reaction in case of a deviation. The collusive equilibrium can be reached again if one firm for whatever reason started a price war because both firms profit even from relenting themselves. However, none of the firms has an incentive to deviate from the focal price since for a sufficiently high discount factor the best reply to the monopoly price is also the monopoly price⁶. Edgeworth cycles will be ignored for the moment since the concept has been shown not to generate renegotiation-proof equilibria. Thus, the model by Maskin and Tirole fits very well for an experiment which test whether a change in the critical discount factor influences the stability of tacit collusion.

⁵Van Damme's definition is a weaker one, because it requires only internal consistency. That is, in his model a renegotiation-proof equilibrium can be dominated by another equilibrium and nevertheless be renegotiation-proof. The dominance-criterion of Maskin and Tirole includes also external dominance so that for a renegotiation-proof equilibrium any other equilibrium that Pareto-dominates the renegotiation-proof one does not even exist.

⁶From a behavioral point of view, the idea of renegotiation-proofness is a reasonable assumption because the outcome it predicts is supported by other experimental research concerning efficiency. As long as only the supplier side of the market is modelled and no consumers suffer from higher prices in our experiment, collusion at the monopoly price is in our framework the most efficient outcome. As earlier experiments have shown (Engelmann and Strobel 2004; Charness and Rabin 2002), participants in experiments attach high value to maximisation of total payoff. Thus, if people are concerned about efficiency in general, their optimal strategy will lead them to monopoly pricing, too.

3 A numerical example

For the implementation of the experiment a numerical example from the Maskin and Tirole model is derived. In order to enable participants in the experiment to get an overview over their choice set in reasonable time, the price set is restricted to the integers $[0;10]$ ⁷. From the assumptions in Maskin and Tirole industry profits have to be strictly concave. The profit function of the experiment is based on linear demand $D(p) = 2 \cdot (10 - p)$. Costs are assumed to be zero. Thus, industry profits depending on prices are the following:

Price	0	1	2	3	4	5	6	7	8	9	10
Industry Profit	0	18	32	42	48	50	48	42	32	18	0

Table 1: Industry Profits depending on prices

Products are homogeneous so that consumers always buy from the cheaper firm. If both firms charge the same price, demand (and therefore aggregate profit) is shared equally. Although only one firm at a time is allowed to adjust its price, profits in each round are calculated according to a Bertrand situation. For the 'waiting' firm, its price decision from the previous round still holds. In our example, the unique renegotiation-proof equilibrium mentioned above consists in strategies of the two firms that lead to coordination at $p = 5$. Equilibrium strategies are illustrated by the symmetric reaction functions $R(p)$ in Table 2.

Price	0	1	2	3	4	5	6	7	8	9	10
Best Reply	5	1 with prob α 5 with prob $1 - \alpha$	1	1	1	5	5	5	5	5	5

Table 2: Reaction functions depending on prices

Suppose the first moving firm 1 chooses a price of 5. In the second period, firm 2 can decide whether it wants to cooperate or to start a price war. In the case of successful tacit collusion, both firms set a price of 5 and share the monopoly profit of 50 equally. Doing this forever and discounting their profits with δ , their intertemporal profit is equal to $25 \cdot \frac{1}{1-\delta}$. If firm 2 for whatever reason decides to start a price war in period 2, it underbids firm 1 by choosing a price of 4 and, thus, faces a profit of 48. In the next period, firm 1's best reaction is to undercut its opponent drastically by setting a price of 1 (see appendix for a proof of this reaction being optimal) so that firm 1 receives 18, whereas firm 2 receives zero profit. Intuitively,

⁷Admittedly, this price grid is not fine enough to support all properties of the underlying theoretical model. In particular, it undermines the proof of existence of a lower bound for equilibrium focal prices for pure strategies. The restriction of the price set is nevertheless justifiable since the unique renegotiation-proof equilibrium, this work will concentrate on, does not need this particular proof.

this drastic reaction is reasonable because it shortens the time without cooperation. Otherwise, firms would spend some rounds with relatively low prices while sequentially undercutting each other which results in lower profits than the immediate reversal to cooperation. This idea is similar to the 'punishment' threat in Bertrand supergames that deters a firm's opponent from undercutting. Suppose now the firms arrive at $p = 1$. Further undercutting is no profitable option⁸, since in this case the undercutting firm would face a profit of zero instead of 9 when continuing with a price of 1. It can also return to the monopoly price, receiving zero profit in the next round, but afterwards gaining from collusion again. For a smaller δ the return to monopoly pricing becomes less attractive. In this case, both firms would still like to collude again but each firm would prefer the other one to return first. Therefore, it may happen that firms do not revert immediately but after a short price war. Maskin and Tirole illustrate the war of attrition that is played in such a situation. If one firm could credibly commit itself in the first place never to revert to monopoly pricing, this would make the other firm concede. However, this announcement would not be credible since - if the other firm is actually not impressed by the threat - the first firm would be tempted to reconsider its commitment one period later. Therefore, the only renegotiation-proof strategy in this situation is playing mixed strategies. Each firm relents (sets the monopoly price) with probability $1 - \alpha$ and continues the price war with α , where α is selected in such a way that the *other* firm is just indifferent between relenting and not relenting itself. Thus, the expected payoffs to a firm from relenting and further low-pricing have to be equal. Continuing the price war yields a firm a profit of $9 + \delta \cdot W_1$, whereas returning to the monopoly price leads to a profit of $25 \cdot \frac{\delta}{1-\delta}$. For indifference, profits resulting from each decision have to be equal:

$$25 \cdot \frac{\delta}{1-\delta} = 9 + \delta \cdot W_1,$$

where W_1 is the expected payoff of a firm if the opponent plays mixed strategies,

$$W_1 = \alpha \cdot [9 + 0 \cdot \delta + 25 \cdot \frac{\delta^2}{1-\delta}] + (1 - \alpha) \cdot [18 + 25 \cdot \frac{\delta}{1-\delta}].$$

With probability α the opponent continues the price war in the next period. In this case, the first firm is again indifferent between relenting and further low-pricing and we can assume that the first firm will relent itself in the next but one period. Thus,

⁸This can be derived from the condition that \underline{p} has to satisfy $4 \cdot \Pi(\underline{p}) \geq \Pi(p^m) > 4 \cdot \Pi(\underline{p} - k)$, which here turns out as $4 \cdot 18 > 50 > 4 \cdot 0$. For this condition to hold here the used price grid is fine enough.

$$25 \cdot \frac{\delta}{1-\delta} = 9 + \delta \cdot W_1$$

$$25 \cdot \frac{\delta}{1-\delta} = 9 + \delta \cdot \{\alpha \cdot [9 + 25 \cdot \frac{\delta^2}{1-\delta}] + (1-\alpha) \cdot [18 + 25 \cdot \frac{\delta}{1-\delta}]\}$$

From this equation α can be derived as a function of delta with

$$\alpha(\delta) = \frac{9-7\delta}{9\delta+25\delta^2}.$$

For high values of δ (when future profits are highly valued) α is low because the incentive to re-cooperate is large, whereas for a decreasing discount factor the probability to continue low-pricing is higher. For δ below 0.36 firms would never revert to p^m .




Figure 2: Probability of continuing the price war depending on δ

We now come back to the initial decision problem of a colluding firm whether to deviate from collusive monopoly pricing or not. Suppose, firm 2 has to choose between undercutting to $p = 4$ and continuing collusion with $p = 5$. Collusion yields half the monopoly profit forever and therefore

$$\pi_c^i = 25 \cdot \frac{1}{1-\delta},$$

whereas undercutting results in a profit of

$$\pi_u^i = 48 + 0 \cdot \delta + 0 \cdot \delta^2 + 25 \cdot \frac{\delta^3}{1-\delta}.$$

Thus, cooperation is profitable if

$$25 \cdot \frac{1}{1-\delta} \geq 48 + 25 \cdot \frac{\delta^3}{1-\delta}$$

which holds for approximately $\delta > 0.58$.

To investigate the effect of a higher δ^* , firm's payoffs are modified such that the δ^* of the game changes. This is basically done by introducing another treatment where the firm with a higher price than the other (which formerly faced zero profit) receives a positive payment instead of zero. Table 3 summarises the payoff modification in the relevant area of prices between zero and 5.

		Price A						
		0	1	2	3	4	5	...
Price B	0	0	1	4	7	10	13	
	1	0	9	1	4	7	10	
	2	0	18	16	1	4	7	
	3	0	18	32	21	1	4	
	4	0	18	32	42	24	1	
	5	0	18	32	42	48	25	
	⋮							
	⋮							

Table 3: Modified payoff table

A rough intuition for such a payoff modification is a situation with demand inertia. The critical discount factor that is necessary for successful tacit collusion increases due to the modification. In theory, however, this modification should not influence firms' behavior as long as the actual discount factor is sufficiently high. Cooperation remains the optimal payoff-dominant renegotiation-proof strategy and deviations should theoretically not occur in both treatments. The best reply strategies for the new situation remain exactly

the same as before. Gains from colluding and undercutting respectively can be calculated as before. Individual collusive profit π_c^i in each period remains the same whereas the profit from undercutting is now given by

$$\pi_u^i = 48 + \delta \cdot 7 + \delta^2 \cdot 10 + 25 \cdot \frac{\delta^3}{1-\delta}.$$

This has two implications. First, it makes the war of attrition disappear because α , the probability of continuing a price war, is below zero for all possible value of δ . Thus, $\alpha(\delta) = \bar{\alpha} \equiv 0$ and firms return to monopoly pricing immediately from a price level of $p = 1$. Second, deviating becomes more attractive. Cooperation is now profitable if

$$25 \cdot \frac{1}{1-\delta} \geq 48 + \delta \cdot 7 + \delta^2 \cdot 10 + 25 \cdot \frac{\delta^3}{1-\delta},$$

which only holds approximately for $\delta > 0.78$. Thus, the critical discount factor that is necessary for successful collusion is higher than in the basic treatment. We will call the basic situation with $\delta^* = 0.58$, where a duopolist with a higher price than the other faces zero profit, treatment 'LOW'. The second treatment, where payoffs are modified according to table 4, is called 'HIGH'. The actual discount factor in both treatments is assumed to be 1. This is reasonable since participants in our experiment receive their total payment at the end of the experiment. Thus, the actual discount factor δ is higher than δ^* in both treatments. We expect, however, that the temptation to deviate from cooperative monopoly pricing increases in the modified treatment and, thus, that more deviations occur in this treatment. The intuition for the higher temptation to deviate becomes clear when we compare the costs of a deviation in both treatments taking the profits from successful collusion in each round as a reference point. In the LOW treatment, the expected costs of a deviation for the deviating firm depend on the duration of the war of attrition. The expected value of the costs of a deviation is

$$\begin{aligned} \pi(dev) - \pi(col) &= \underbrace{48 - 25}_{p_1=4; p_2=5} + \underbrace{0 - 25}_{p_1=4; p_2=1} + (1-\alpha) \underbrace{(10 - 25)}_{p_1=5; p_2=1} + \alpha \underbrace{(9 - 25)}_{p_1=1; p_2=1} + (1-\alpha) \underbrace{(18 - 25)}_{p_1=1; p_2=5} + \alpha \underbrace{(9 - 25)}_{p_1=1; p_2=1} + \dots \\ &= 23 - 25 - 25(1-\alpha)(1 + \alpha^2 - \alpha^4 + \dots) - 16(\alpha + \alpha^2 + \alpha^3 + \dots) - 7(1-\alpha)(\alpha + \alpha^3 + \alpha^5 + \dots) \\ &= 23 - 25 - 25 \frac{1}{1+\alpha} - 16 \frac{\alpha}{1-\alpha} - 7 \frac{\alpha}{1+\alpha} \end{aligned}$$

For $\delta = 1$ the probability to continue a price war is equal to $\alpha(\delta = 1) = \frac{1}{17}$ so that the expected costs of a deviation are equal to $\pi(dev) - \pi(col) = -27$ in the LOW treatment. In the HIGH treatment, there is no war of attrition ($\alpha = 0$) and the costs of a deviation are straightforward: Compared to the profits of collusive monopoly pricing a deviation yields a cost of

$$\pi(dev) - \pi(col) = \underbrace{48 - 25}_{p_1=4;p_2=5} + \underbrace{7 - 25}_{p_1=4;p_2=1} + \underbrace{10 - 25}_{p_1=5;p_2=1} = -10$$

We see that the costs of a deviation are much smaller in the HIGH than in the LOW treatment and therefore the temptation to deviate is harder to resist in the HIGH treatment with $\delta^* = 0.78$

4 Experimental Design

In the experiment, two treatments of the sequential price game are compared. In both treatments the actual δ is above δ^* so that collusion should be stable both times. The difference between δ and δ^* , however, is larger in the first treatment than in the second one. It will be analysed whether this difference in the critical discount factor leads to a change in stability of collusion.

In the theoretical derivation of the game an infinite time horizon has been chosen. For practical reasons, it is very convenient to approximate this by a long finite horizon of the experiment, because it is definitely impossible to have a forever-lasting experiment.⁹ Theoretically a finite horizon implies inability to collude due to a backward induction argument: The incentive to coordinate disappears in the final round and, anticipating this, firms fail to coordinate from the very beginning. However, experience has shown that people do not behave as predicted by the backward induction argumentation (and they also do not expect their fellow subjects to do so) but do actually cooperate in finite games almost as in infinite ones (see Rosenthal 1981; Selten and Stoecker 1986). Normann and Wallace (2005) show that the choice of the termination rules has no significant impact on cooperation in repeated prisoner's dilemma experiment except for an endgame-effect given the known finite horizon. Transferring their conclusion our experiment is conducted with a long finite horizon of 50 rounds assuming that behavior is the same as in an infinitely repeated game until some rounds before the announced end of the game. Consequently, we will divide our data for the analysis into two parts. In the first part of the experiment we expect participants to behave as if they face a quasi-infinite horizon. The final rounds during the second part of the experiment are treated separately because we expect endgame effects to determine the results.

The experiment is programmed and conducted with the software z-Tree (Fischbacher 1999). In the first

⁹The construction of a random stopping rule to simulate infinity, which has often been used in experiments, is not appropriate because participants nevertheless know that the laboratory experiment has a limited duration (see Selten, Mitzkewitz and Uhlich 1997).

treatment, the duopoly participants face the original payoff table which implies $\delta^* = 0.58$. In the second treatment, a shift in the critical discount factor to 0.78 takes place whereas the actual delta of 1 remains the same. Remember that in theory this variation in the critical discount factor should not cause a difference in participants' behaviour. Coordination at the focal price should be fully stable and no deviations should occur at all since the actual discount factor is larger than the critical discount factor. If a higher δ^* (in contradiction to the theory) actually leads to more deviations from the collusive outcome, the initial hypothesis of an observable impact of the critical discount factor's value would be confirmed.

Hypothesis:

A higher critical discount factor δ^* leads to a higher temptation to deviate from collusion. Therefore, deviations from the collusive outcome $p = 5$ occur more often if δ^* increases.

A deviation is therein defined in the following way:

Definition 1:

A price decision of player i in period t is called deviation if the following conditions are satisfied:

- (a) In period $t - 1$ (at least) player j set the monopoly price (or above).
- (b) In period t player i chooses a price below 5.

In condition (a) above-monopoly pricing is included because in some cases participants use such prices to signal their willingness to cooperate. It is not necessary that both players are cooperating the period before because due to sequential play it is also possible that one player deviates immediately after the other finished a price war. Condition (b) takes into account that some deviations were not in the expected form of a price of 4 but below, which obviously has to be treated as deviation, too. To compare the average frequency of deviations in both treatments, we simply count how many deviations occur within each duopoly.

Participants are called type 'A' and 'B' to distinguish which of them is allowed to adjust his price in odd and even periods, respectively. Type A and B are determined randomly. Both treatments are conducted with overall 82 participants. All of them are business and economics students but inexperienced with oligopoly experiments.¹⁰ 38 out of 82 participate in 'LOW', 44 in 'HIGH'. The group size varies between 8 and 12.

¹⁰Since subjects are randomly assigned to both of the treatments, biases through their education should not matter because

Participants can only take part in one of the sessions. In each session, subjects are randomly matched with one of their fellow participants in a partner design. Of course, they only know that it is someone sitting in the same room as they do, but they do not know which one it is. The initial price of B is predetermined to zero so that both firms start with zero profits to avoid that firm A, given that it is allowed to set its price in the very first, round has a small advantage.

5 Results

The unique renegotiation-proof equilibrium strategy describes behavior quite well. In principle, participants cooperate at the monopoly price of 5 and they never set a price of zero.¹¹ Sporadically prices above monopoly pricing can be observed. This may be due to initial difficulties of participants understanding the structure of the game. The questionnaires filled in by the participants after each session furthermore indicate that some participants considered a price above 5 as an appropriate measure to induce the opponent to cooperate. For most participants, the assumption of Markov strategies seems reasonable, since pricing structure definitely follows an almost predictable pattern and participants usually react in the same way to a certain price. There is only one remarkable difference to what the model predicts. If subjects deviated to price 4, the reaction of the opponent does several times not follow the strategy that Maskin and Tirole would have predicted. Instead of reacting with a price decrease to 1, they underbid each other sequentially in steps of 1. This does partially question the applicability of renegotiation-proof strategies. For the experiment, however, this has no significant impact because under this strategy pair all former considerations regarding discounting still hold although the difference in critical discount factors between the two treatments becomes smaller. Finally, the occurrence of deviations at all is, of course, not predicted by the theoretical model. In fact, deviations occur because the temptation of short-term profit maximisation is strong but it has to be kept in mind that deviations are strictly speaking irrational.

The finite horizon does - as in Normann and Wallace (2005) - not seem to hinder cooperative behavior in

we are only interested in the relative difference between the two treatments. Furthermore, Dufwenberg, Gneezy, and Rustichini (2004) have shown that neither gender nor education causes a significant difference in price competition behavior.

¹¹One observation in the HIGH case differs completely from all other observations. The pricing in this certain experimental duopoly is definitely not cooperative because the price level is always far below monopoly pricing. In a certain way, this supports our hypothesis of less cooperation if the critical discount factor is high, because when participants do not expect collusive attempts to be successful they behave rational if they do not try to cooperate at all. However, this noncooperative behavior does not result in a high number of deviations according to our definition because none of the duopolists ever reverted to a price of 5. Thus, no deviation from monopoly pricing according to our definition can occur, because they never reach the monopoly price where one firm could deviate from. If an exception from the definition was made and any initial underbidding was counted as a deviation, this would result in 8 deviations. We decided to stick to our definition but to exclude this observation from regular analysis. The results of all statistical tests, of course, do not change if we nevertheless include it and assume the number of 8 deviations within 50 periods for this specific duopoly.

			Reaction (in %)										
		# cases	0	1	2	3	4	5	6	7	8	9	10
Price	0	33	48.5	3.0	0.0	0.0	0.0	42.4	0.0	0.0	0.0	0.0	6.1
	1	222	3.2	65.3	0.9	0.9	4.1	12.2	6.3	1.4	2.3	1.4	2.3
	2	163	1.2	46.0	33.7	1.2	0.6	3.1	9.2	1.2	0.6	1.2	1.8
	3	177	2.3	1.1	60.5	26.6	0.0	2.8	5.1	0.0	0.6	1.1	0.0
	4	158	1.3	0.0	3.2	69.0	21.5	3.2	0.0	1.9	0.0	0.0	0.0
	5	439	0.0	0.0	0.0	0.7	23.0	76.3	0.0	0.0	0.0	0.0	0.0
	6	36	0.0	0.0	0.0	5.6	11.1	83.3	0.0	0.0	0.0	0.0	0.0
	7	8	0.0	0.0	0.0	0.0	25.0	75.0	0.0	0.0	0.0	0.0	0.0
	8	6	0.0	16.7	0.0	0.0	33.3	50.0	0.0	0.0	0.0	0.0	0.0
	9	5	0.0	0.0	0.0	0.0	0.0	60.0	40.0	0.0	0.0	0.0	0.0
	10	8	0.0	0.0	0.0	0.0	25.0	75.0	0.0	0.0	0.0	0.0	0.0

Table 4: Frequency distribution of price reactions.

general. Many participants seem pretty well to be able to establish cooperation except for a few final periods of the game, where end-game effects occur. Defining the last deviation within a duopoly as the starting point of the respective endgame, we find the surprising result that the endgame starts earlier in the LOW than in the HIGH treatment. In the LOW treatment, the last deviation occurs 5.38 periods before the end whereas in the HIGH treatment cooperation continues until 1.95 rounds before the end of the game. This is surprising because it seemingly contradicts our hypothesis that collusion should be more stable in the LOW treatment, since we might have expected that more stable collusion also implies a later occurrence of endgame-effects.¹² However, we have to be very careful when analysing the final rounds of the experiment. The remarkable long time horizon of 50 periods also leads to other forms of puzzling behavior due to boredom because of the simple and highly repetitive task in the experiment: In the questionnaire after the experiment a question was included whether and if that was the case, when participants became bored during the experiment. Many participants indicated that they indeed were bored from a certain point of time onward, which was located between rounds 25 and 40. Round 33 seems a good guess about the average beginning of boredom. Consequently, participants' intention to 'make something happen' induced them to deviate from established cooperation even in the LOW treatment. In our experimental data the frequency of deviations in both treatments converged by the end of the 50-period-game and makes the observed overall

¹²In particular, we could have expected that participants reach cooperation only because each of them pretends to be willing to cooperate, similar to the argumentation in Kreps et al. (1982). If both firms send such signals to make the other firm trust them and cooperate, we should also observe collusive behavior until short before the end of the game. For this explanation to hold, however, we would expect that collusion breaks down earlier in the HIGH than in the LOW treatment. The opposite is the case. Thus, our data do not support the hypothesis that such a signalling story is an explanation for cooperation in the finite horizon game.

differences between them less significant. When evaluating the data in the statistical tests, we therefore concentrate for the statistical tests on round 1-33. The restriction to 33 rounds will show clearer results, but we can replicate all results for the whole 50 period time horizon, although sometimes on a lower significance level.

To test the main hypothesis that a higher critical discount factor in HIGH leads to less cooperation than in LOW, the frequency of deviations in both treatments is compared. In the low- δ^* -treatment, on average 2.89 deviations occur within 33 rounds. In the other treatment with a higher critical discount factor this number increases to 3.95 deviations. Statistical significance is tested using a Wilcoxon-Mann-Whitney test. The null hypothesis 'no difference in the average number of deviations' is tested against the alternative 'more deviations in HIGH'. The null hypothesis can be rejected at the 5%-level. Thus, there are significantly more deviations from collusion in the treatment with a higher critical discount factor. Repeating the test for all 50 periods basically confirms the result, but only at a 10%-level (LOW: 4.69 deviations, HIGH: 5.67 deviations).




Figure 3: Frequency distribution of the total number of deviations per pair within all 33 rounds

Another aspect of different stability of cooperative behavior is the difference in the average duration of mutual cooperation in both treatments. More stability of collusion is not only indicated by fewer deviations within a certain period of time but also by longer duration of cooperative monopoly pricing. A higher stability of tacit collusion in LOW is reflected by a longer average duration of mutual cooperation $DC(LOW)=6.67$ than in HIGH ($DC(HIGH)=5.41$). In the repetition of the test for the overall 50 rounds these values increase to 9.59 in LOW and 7.58 in HIGH respectively. This difference, however, is statistically not significant due to high variances in the length of the sequences.

Duration of mutual cooperation			
33 rounds		50 rounds	
LOW	HIGH	LOW	HIGH
6.67	5.41	9.59	7.58
(8.50)	(8.01)	(11.19)	(10.97)
p-value = 0.27		p-value = 0.24	

Table 5: Duration of mutual cooperation. Standard errors in brackets.

Average prices, in contrast, cannot serve as a measure for stability of collusion. At first, we could think that average prices combine both measures, frequency of deviation and duration of cooperative periods, so that they would provide an appropriate measure of stability of collusion. However, the duration of periods of defection, which is also included in average prices, diminishes their explanatory power. The longer defection lasts the lower average prices are over time. In a numerical example we have seen in the numerical example that the average duration of mutual defection is longer in the LOW than in the HIGH treatment. Thus, low average prices can likewise be caused by frequent deviations in the HIGH treatment and by long duration of defection in the LOW treatment, which indicates low stability of collusion in the first case and high stability in the second case.

The stronger temptation to deviate in the HIGH treatment is presumably driven by the less costly price war after a deviation we have just mentioned. Theory predicts no war of attrition in the case of a high δ^* which means a faster return from low-level-pricing to monopoly pricing. We have seen that the model predicts at $p = 1$ a war of attrition in the original LOW but not in the modified HIGH treatment. Thus, we expect that the duration of mutual defection is definitively longer in LOW than in HIGH. On average, mutual defection lasts in the LOW treatment for 2.33 rounds and in HIGH for 1.67 rounds. This difference is significant at the 1%-level in a Wilcoxon-Mann-Whitney test. Analysis of all 50 rounds confirms the results, with a mean duration of mutual defection of 2.71 in LOW and 1.77 in HIGH. Thus, the war attrition is definitely shorter in the HIGH than in the LOW treatment. It is, however, also longer than theory would predict.

This supports our consideration that the higher number of deviations in the HIGH treatment is caused by a higher temptation to deviate because the consequences are less severe than in the LOW treatment. The price war lasts shorter and the reversal to monopoly pricing is cheaper because the reverting firm receives a positive profit in the HIGH treatment. Anticipating that the cost of a deviation are smaller the incentive to deviate in the first place becomes larger. The shorter duration of mutual defection in the LOW treatment

Average duration of mutual defection			
33 rounds		50 rounds	
LOW	HIGH	LOW	HIGH
2.33	1.67	2.71	1.77
(1.21)	(1.14)	(2.48)	(1.19)
p-value = 0.002		p-value = 0.0005	

Table 6: Duration of mutual defection. Standard errors in brackets.

can also partially explain the higher number of deviations. As long as firms set their prices competitively none of them can deviate from cooperation because they just do not cooperate at that time. Thus, if firms return faster to monopoly pricing, it takes less time until there is another opportunity to deviate. However, if the difference in the total number of deviations was completely explained by the shorter duration of periods of defection, the duration of cooperative periods (the length of a sequence of mutual monopoly pricing), in contrast, would last equally long in both treatments. This is not the case.

If we jointly examine all three parts of the former analysis, we can see their dependence more clearly. Between two deviations we usually find several stages. First, undercutting in the $(p_1 = 5, p_2 = 4; p_1 = 1, p_2 = 4)$ -pattern, second, a period of mutual defection, third, reversal to monopoly pricing $(p_1 = 5, p_2 = 4)$, and, finally, a period of mutual cooperation. If we sum up the duration of these four stages, and multiply the sum by the number of deviations, we should end up with the total duration of the game. Or, if we divide the total duration by the number of deviations and subtract the joint duration of mutual defection and mutual cooperation, we find the average duration of the in-between phases. We restrict this analysis to the data for the first 33 periods (see Table 7), because for 50 periods we cannot consider full cycles due to the endgame-effect. We see that the duration of the in-between phases is surprisingly small - 2.42 in LOW and 1.27 in HIGH - compared to the expected time it takes to move from one kind of mutual behavior to the other. This may be due to the fact that the first cycle often starts with an immediate price war so that the first period of mutual cooperation is shorter than the average.

	Duration cooperation	Duration defection	# Deviations	Duration 'in-between'
$\delta^* = 0.58(\text{LOW})$	6.67	2.33	2.89	$\frac{33}{2.89} - (6.67 + 2.33) = 2.42$
$\delta^* = 0.78(\text{HIGH})$	5.41	1.67	3.95	$\frac{33}{3.95} - (5.41 + 1.67) = 1.27$

Table 7: Duration of phases in between mutual cooperation and defection.

Finally, a closer look into individual behavior is worthwhile since the presented aggregate data arises from

quite different patterns of individual willingness to cooperate. Participants' strategies can roughly be classified into more and less cooperative-minded attempts. Since one 50-round-price-sequence always is the result of the interplay of two strategies, pricing sequences can be grouped into four different clusters. If two cooperative strategies meet each other (cluster 'all_cooperate'), a long sequence of cooperative behavior arises after some initial coordination time, sometimes even over 50 rounds from the beginning to the very end of the session. If a peaceful strategy is confronted with a less cooperative one, the typical price pattern exhibits several deviations followed by forgiving attempts of the other participant to re-establish cooperation that are again exploited by ongoing deviations ('defect_forgive'). If the deviating strategy is, in contrast, matched with a non-forgiving strategy, this leads to the deviator himself making the first move to reach the higher mutual price level again ('defect_relent'). In those two clusters, defect_forgive and defect_relent, it is mainly one firm that deviates several times, but the reaction of the other strategy differs. In the defect_relent cluster, the other participant is willing to cooperate but does not accept being exploited. In defect_forgive cases, in contrast, the goal to reach cooperation is worth so much that even several deviations by the opponent are being forgiven. Finally, if two noncooperative strategies are matched, frequent deviations of both firms occur ('all_defect'). Out of all 41 markets in both treatments, 32% fit into all_cooperate, 22% into defect_forgive, 17% into defect_relent and 29% into all_defect patterns. The following pictures are examples for all clusters. They show pricing behavior of four selected experimental duopoly markets over 50 rounds.

No significant difference of the cooperativeness of strategies can be observed between the HIGH and LOW treatment. Thus, the ex-ante distribution of cooperative-minded people is equal in both treatments and different observations in both treatments arise only from the difference in the critical discount factor.

6 Conclusion

At the beginning, we highlighted the fact that many papers implicitly make the assumption that the critical discount factor provides an appropriate measure to analyse the stability of a tacit cartel. This practice, however, contradicts standard economic theory. The experiment in this paper solves this conflict by providing experimental evidence that the critical discount factor's distance from the actually applied discount factor indeed plays a role for stability of tacit cartel agreements. For both LOW and HIGH treatment, the actual discount factor lies above its critical value which means that we should observe perfect cooperation. However, in both treatments we see cooperation to be unstable which contradicts the theoretical prediction in the sense that deviations from tacit collusion occur. Most interestingly, instability is higher in the modified HIGH

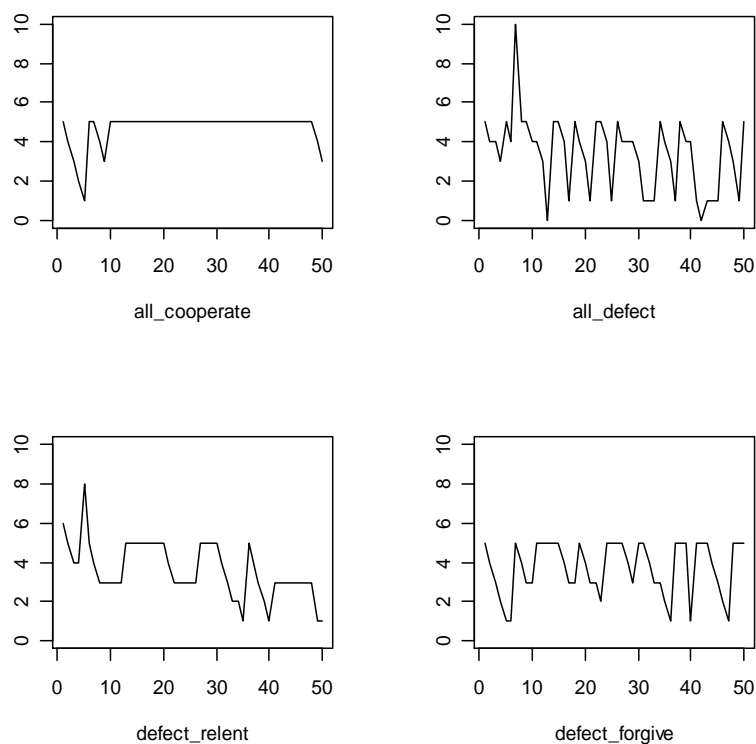


Figure 4: Patterns of play.

treatment, in which the critical δ is higher than in the basic LOW treatment. First, there are significantly more deviations from the collusive outcome within a certain time period in the high- δ^* -treatment. Second, the average duration of periods of mutual cooperation is shorter. Both results support the hypothesis that an increase in δ^* indeed leads to less tacit collusion. Furthermore, the individual ex-ante willingness to cooperate varies widely across participants. For future research, it might be interesting to further investigate the role that initial cooperative-mindedness plays for successful profit maximisation.

7 Appendix

To see whether it is optimal to react to a deviation with an immediate price cut to $\underline{p} = 1$, we have to compare expected future profits from all different possible reactions. The proof is here only presented for the basic treatment, for the payoffs in the modified treatment all results become even clearer. First, we show that $p = 1$ yields a higher payoff than $p = 2$ and next that $p = 1$ also yields a higher profit than $p = 3$ and $p = 4$. For the calculation of the profits we can assume that each firm reverts to monopoly pricing $p = 5$ immediately as soon as price $p = 1$ is reached. This is, because the respective other firm mixes its strategy in such a way that the deciding firm is just indifferent between relenting and continuing the price war. Thus,

once $p = 1$ is reached, profits of further price war continuation have to be equal to profits of an intermediate return to monopoly pricing.

$$\pi(1) = 18 + \alpha(\delta \cdot 9 + \delta^2 \cdot 0 + \frac{\delta^3}{1-\delta} \cdot 25) + (1 - \alpha)(\delta \cdot 18 + \frac{\delta^2}{1-\delta} \cdot 25)$$

if we insert $\alpha(\delta) = \frac{9-7\delta}{9\delta+25\delta^2}$ this can be simplified to

$$\pi(1) = 9 + \delta \cdot 25 + \frac{\delta^2}{1-\delta} \cdot 25.$$

$$\pi(2) = 32 + \delta \cdot 0 + \delta^2 \cdot 0 + \frac{\delta^3}{1-\delta} \cdot 25$$

Thus, $\pi(1) > \pi(2)$, if

$$9 + \delta \cdot 25 + \frac{\delta^2}{1-\delta} \cdot 25 > 32 + \frac{\delta^3}{1-\delta} \cdot 25,$$

which holds for $\delta > 0.6$.

$$\pi(3) = 42 + \delta \cdot 0 + \delta^2 \cdot 0 + \frac{\delta^3}{1-\delta} \cdot 25$$

Thus, $\pi(1) > \pi(3)$, if

$$9 + \delta \cdot 25 + \frac{\delta^2}{1-\delta} \cdot 25 > 42 + \frac{\delta^3}{1-\delta} \cdot 25,$$

which holds for $\delta > 0.75$.

Staying at $p = 4$ is according to the following calculations definitely no stable option, because the other firm's best reply would be $p = 1$:

$$\pi(4) = 24 + \delta \cdot 0 + \delta^2 \cdot 0 + \frac{\delta^3}{1-\delta} \cdot 25$$

Thus, $\pi(1) > \pi(4)$, if

$$9 + \delta \cdot 25 + \frac{\delta^2}{1-\delta} \cdot 25 > 24 + \frac{\delta^3}{1-\delta} \cdot 25,$$

which holds for $\delta > 0.35$.

References

- Brock, W.A. and J.A. Scheinkman (1985), Price Setting Supergames with Capacity Constraints, *Review of Economic Studies* 52, 371-382.
- Charness, G. and M. Rabin (2002), Understanding Social Preferences with Simple Tests, *Quarterly Journal of Economics* 117 (3), 817-869.
- Chincarini, L. (2003), Experimental Evidence of Trigger Strategies in Repeated Games, Rydex Working paper.
- Compte, O., F. Jenny and P. Rey (2002), Capacity Constraints, Mergers and Collusion, *European Economic Review* 46, 1-29.
- Dal Bó, P. (2005), Cooperation under the Shadow of the Future: Experimental Evidence from Infinitely Repeated Games, *American Economic Review* 95 (5), 1591-1604.
- Dufwenberg, M., U. Gneezy and A. Rustichini (2004), Price Competition: The Role of Gender and Education, in: S. Huck (ed.) *Advances in Understanding Strategic Behaviour: Game Theory, Experiments, and Bounded Rationality*, Basingstoke: Palgrave.
- Engelmann, D. and M. Strobel (2004), Inequality Aversion, Efficiency, and Maximin Preferences in Simple Distribution Experiments, *American Economic Review* 94 (4), 857-869.
- Feinberg, R.M. and T.A. Husted (1993), An Experimental Test of Discount-Rate Effects on Collusive Behaviour in Duopoly Markets, *Journal of Industrial Economics* 41 (2), 153-160.
- Fischbacher, Urs (1999), z-Tree - Zurich Toolbox for Readymade Economic Experiments - Experimenter's Manual, Working Paper Nr. 21, Institute for Empirical Research in Economics, University of Zurich.
- Huck, S., Normann, H. and Oechssler, J. (1999), Learning in Cournot Oligopoly: An Experiment, *The Economic Journal*, 109 (454), C80-95.
- Maskin, E. and J. Tirole (1988), A Theory of Dynamic Oligopoly, II: Price Competition, Kinked Demand Curves and Edgeworth Cycles, *Econometrica* 56 (3), 571-599.
- Noel, M.D. (2004), Edgeworth Cycles and Focal Prices: Computational Dynamic Markov Equilibria, Working Paper, University of California.
- Normann, H. and B. Wallace (2005), The Impact of the Termination Rule on Cooperation in a Prisoner's

Dilemma Experiment, Working Paper, Royal Holloway.

Rosenthal (1981), Games of Perfect Information, Predatory Pricing and the Chain-Store Paradox, *Journal of Economic Theory* 25, 92-100.

Selten, R., M. Mitzkewitz and G.R. Uhlich (1997), Duopoly Strategies Programmed by Experienced Players, *Econometrica* 65 (3), 517-555.

Selten, R. and R. Stoecker (1986), End Behavior in Sequences of Finite Prisoner's Dilemma Supergames, *Journal of Economic Behavior and Organization* 7, 47-70.

Shapiro, C. (1989), Theories of Oligopoly Behavior, Ch. 6 in: R. Schmalensee and R.D. Willig (eds.) *Handbook of Industrial Organization*, Volume I, New York: North-Holland.

Symeonidis, G. (2002), Cartel Stability with Multiproduct Firms, *International Journal of Industrial Organization* 20, 339-352.

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

Van Damme, E. (1989), Renegotiation-Proof Equilibria in Repeated Prisoners' Dilemma, *Journal of Economic Theory* 47, 206-217.