

Experimental Results on Collusion

The role of information and communication

Marco A. Haan Lambert Schoonbeek Barbara M. Winkel*

June 8, 2005

Abstract

We review the experimental literature on collusion, focusing in particular on the roles of information and communication. We confront the results with the theoretical literature and discuss the policy implications. The main insights from the experimental literature are the following. In an environment in which firms are unable to communicate, experimental firms have little success in achieving collusion. The only mildly collusive results are achieved when just two firms are active. Increasing the amount of available information in general has little effect on the likelihood of collusion. The ability to communicate among sellers has a strong and positive effect on the ability to collude. This effect may be mitigated when buyers can communicate with sellers by asking for discounts.

1 Introduction

Collusion is a major issue in competition policy. One important objective of competition authorities is to prevent firms that are supposed to compete amongst each other from achieving prices that are higher than the prices that would result from free and fair competition. Another important objective is to track down and punish those firms that do engage in price- or quantity-fixing agreements.

It is not easy to achieve these goals, especially in a world in which competition authorities have only limited funds at their disposal, and need to make choices with respect to

*All authors: Department of Economics, University of Groningen, P.O. Box 800, 9700 AV Groningen, The Netherlands. e-mail: m.a.haan@rug.nl, l.schoonbeek@rug.nl, b.m.winkel@rug.nl. The authors thank participants of the ENCORE workshop "Experiments for Antitrust Policies", in particular Eric van Damme, for useful comments and discussion.

the firms and industries they investigate. Many important yet difficult questions have to be answered. When exactly are prices too high? Is collusion always detrimental, or are there circumstances in which it can actually be beneficial? How can one establish that an industry engages in collusion? Under what circumstances is collusion more likely to occur? Which market structures and which market institutions are susceptible to collusion, and therefore deserve closer scrutiny from antitrust authorities? Many related questions readily spring to mind.

In trying to answer these questions, policy makers can draw upon a large and growing body of economic research. Economic theory provides a framework to understand what collusion exactly entails, why it is harmful, and which market data are consistent with price fixing. The framework is that of repeated games, pioneered by Friedman (1971) and described in more detail in e.g. Tirole (1988) or Motta (2004). Empirical economics provides tools that are helpful in evaluating whether certain markets are collusive. Experimental economics provides further insights into which institutional settings are most susceptible to collusion. In this paper, we survey those insights. Our aim is thus to give an overview of the experimental results on collusion. In doing so, we try to emphasize how these experimental methods and results can be of help to policy makers.

We are not the first to survey at least some of the experimental literature on collusion. Holt (1995) is an excellent survey of experimental work in the field of Industrial Organization in general, and also gives an extensive discussion of collusion. Wellford (2002) is a more recent survey of similar issues. Huck et al. (2004) provide a useful overview of experimental studies with repeated oligopolies that have Cournot competition as a stage game, and that study the effect of the number of firms on cartel stability. Our focus is therefore not so much on the experimental details of the papers that we survey. Rather, we emphasize the relevance for our understanding of real-world markets. As a final caveat, we do not do full justice to all papers that we discuss, in the sense that we do not always give a full description of the main research questions that these papers address, and do not always evaluate the papers in the light of their original research question. Rather we only discuss the aspects of these experiments that can teach us something about collusion.

Arguably, there are two broad categories of experimental work regarding collusion. On the one hand, we have the experiments that are primarily inspired by theory, and that mainly aim to test that theory. On the other hand, we have experiments that are more loosely related to theoretical work, and that mainly aim to identify the factors that facilitate or hinder collusion in the real world. Needless to say, the latter category is of particular interest to policy makers. Yet, the first category is also relevant, as it aims to test the theory that ultimately underlies competition policy towards collusion. We therefore survey both categories.

In the next section we first give a short introduction into the theory of collusion. As noted, the insights obtained in theoretical work have been a major source of inspiration for experimental work. Therefore, a short discussion of that theory is indispensable. We will especially discuss the theoretical work that has relevance for either experiments or policy. Section 3 gives some general background to experiments in this field. Section 4 addresses the question to what extent experimental firms are able to collude. For that purpose, we survey the experiments that are as close to the standard theory as possible. The remainder of the paper studies institutional environments that may hinder or facilitate collusion. Factors concerning the amount of information that market participants have, are discussed in section 5. The work surveyed in that section is primarily inspired by theory. In section 6, we discuss the experimental evidence concerning the effect of communication between market participants. The work surveyed in that section is primarily inspired by policy issues rather than theory.

Our paper ends with a conclusion, in which we sketch the main lessons that can be drawn from experiments, and their relevance for real markets. We also discuss how experiments may further aid to shape policy in this important area.

2 A short overview of the standard theory

The typical textbook treatment of oligopoly behavior starts with static competition models. In such models firms meet only once. The noncooperative Nash equilibrium has each

firm setting a price or quantity that maximizes its profits, given the action taken by its competitors. Profits of each firm would be higher if the firms could agree to either restrict output, or set a higher price. Yet, in a one-shot game such an agreement is impossible to make. Every firm would have an incentive to defect from the agreement. Doing so would increase profits, at the expense of the other firms. Moreover, a defection cannot have any future repercussions, as there simply is no future in a one-shot model.

In a seminal paper, Friedman (1971) introduced a dynamic model. In his paper, firms play a supergame which consists of a simple static stage game that is repeated an infinite number of times. This changes the picture entirely. In the noncooperative subgame perfect equilibrium of a repeated game, firms *are* able to achieve profits that are higher than the Nash profits in a one-shot game, provided that the discount factor is high enough. Even the cooperative equilibrium of the one-shot game can be achieved as the noncooperative equilibrium of a repeated game. In other words, in a repeated game firms may be able to maximize their joint profits. Moreover, they may be able to do so without any of the firms having an incentive to defect and increase its profits. A firm that defects now faces the possibility of adverse effects in the future, as this will likely cause a breakdown of the cartel. Hence, a firm will only defect if the short-term benefits of doing so outweigh the long-term costs caused by the breakdown of the cartel. Importantly, firms can even do so without making an explicit agreement. In the literature, this is referred to as *tacit collusion*, contrary to *explicit collusion*, where an explicit agreement is made.

The simplest possible version of this model is as follows. Suppose that we have n identical firms that form a cartel. The Nash equilibrium profits of the stage game are denoted π^N . Per-period profits of each firm in the cartel are denoted π^k , with $\pi^k > \pi^N$. Defecting from the agreement yields a one-time profit of $\pi^d > \pi^k$. Suppose that the firms have the implicit understanding that, as soon as one of them defects from the cartel agreement, they will play the one-shot equilibrium forever after. This is a *grim trigger strategy*. In each period, every firm then faces a tradeoff. By sticking to the cartel agreement, it can earn π^k in every single period. By cheating, it will earn π^d in this period, but only π^N in every future period. When the discount factor is denoted by δ , each firm will thus stick to

the agreement if $\pi^k + \delta\pi^k + \delta^2\pi^k + \dots > \pi^d + \delta\pi^N + \delta^2\pi^N + \dots$. After some rearranging, this implies

$$\delta > \frac{\pi^d - \pi^k}{\pi^d - \pi^N}. \quad (1)$$

If this condition is satisfied, the cartel is said to be stable. If the condition is not satisfied, every firm knows in advance that nobody will adhere to the agreement, which in turn implies that such an agreement will never be made.

This model of collusion is simple and attractive, and also readily yields some interesting comparative statics results (see Tirole 1988 or Motta 2004). First, it can be shown that an increase in the number of firms makes it less likely that collusion will occur: there is a smaller range of δ 's for which collusion is stable. Also, when firms meet less often, or can observe each other's behavior less often, collusion becomes less likely. The amount of time between two market periods then becomes longer, which implies that the effective discount factor δ becomes smaller.

However, this model has one major drawback, which is the multiplicity of equilibria. Note that in the discussion above we have not pinned down π^k in any manner. Indeed, the theory only states that *any* agreement that yields cartel profits $\pi^k > \pi^c$ can be sustained as a Nash equilibrium of the repeated game, provided that (1) is satisfied. But the multiplicity of equilibria is even worse. So far, we have assumed grim trigger strategies. But also the punishment phase is not uniquely determined. Abreu (1986) shows that other forms of punishment may yield cartel stability for a larger range of δ s. He focuses on *stick and carrot strategies*, where firms earn a negative profit in the period after the deviation, and return to collusive behavior in the subsequent period.

To reflect the multiplicity of equilibria, we will refer in the remainder of this paper to a *collusive outcome* as *any* outcome of the market process in which realized profits are higher than the Nash equilibrium profits of the stage game. We will refer to *perfect collusion* as a situation in which firms manage to maximize their joint profits. We thus focus on an economic rather than a legal definition of collusion. For example, legal action often also requires that there is some explicit agreement between firms. Following the economic

literature, we merely require that "prices are higher than some competitive benchmark" (Motta 2004, pg. 138), and hence that profits also are,

One crucial assumption in the theory set out above is that the game is repeated an infinite number of times. As soon as the number of periods is finite and known in advance, the entire logic breaks down. Firms then know that in the final period they will all defect from any collusive agreement they could possibly have, as there is no future left in that period. Backward induction implies that they will also defect in the penultimate period, knowing that they will play the Nash equilibrium of the stage game in the final period. Carrying through this logic implies that collusion can never be stable, and firms play the static Nash equilibrium in every single period. In the next section, we will discuss how experiments deal with this problem.

As argued above, the model is often interpreted as one of tacit collusion, i.e. a situation in which firms do not have to formally agree on some collusive arrangement. But that interpretation becomes hard to swallow with such a multiplicity of equilibria. It seems that in practice firms at least need to communicate in order to make some agreement with respect to which of the infinitely many equilibria they are going to play.

This issue becomes even more relevant when firms are asymmetric with respect to e.g. their cost function, or their individual demand. In the real world, this is obviously the case. In an idealized world where firms are perfectly symmetric and demand is known, there may be an infinite number of equilibria, but there is one obvious focal equilibrium, which is the one that maximizes joint profits of the firms. Yet, when firms are asymmetric, such a focal equilibrium does not exist. Hence, coordination then becomes even more difficult.

The basic model assumes that firms can readily observe the price or quantity their competitors set, and can react accordingly. In reality, this will often not be the case. Stigler (1964) already argued that if firms cannot observe each other's actions, collusion is less likely, as there is scope for secret price cuts. Green and Porter (1984) assume that there is uncertainty regarding total market demand. A simplified version of their model, which due to Tirole (1988), is as follows. Suppose that in each period demand can be either high or low. Each firm can only observe its own demand. It cannot observe the

price set by its competitor, and neither can it observe total market demand. Suppose that in a given period a firm observes that its own sales are low. This may be due to low market demand, but it may also be due to the other firm cheating on the cartel agreement by setting a low price. Hence, when a firm's sales are low, it is clearly not a good idea to play the Nash equilibrium of the one shot game in all future periods. Such a reaction would imply a fierce punishment, whereas it may well be the case that no one has cheated. Green and Porter show that under certain conditions the best strategy is a punishment phase that lasts for a finite number of periods. After that punishment phase, firms go back to the collusive agreement. Such a strategy yields the optimal tradeoff between on the one hand taking into account the possibility of low market demand, and on the other hand preventing defections from the collusive agreement. Needless to say, in such a framework, collusion becomes much more difficult to sustain than in the standard case.

Summarizing, the main results from the theory are as follows. First, when firms meet repeatedly, collusive outcomes can be sustained in a subgame perfect equilibrium. Second, collusion becomes more difficult as the number of firms increases, as firms meet more often, or as firms observe each other's behavior less often. Third, many collusive outcomes can be sustained in equilibrium, and different punishment strategies can be part of an equilibrium. Hence, especially when firms cannot communicate, it may be hard for them to coordinate on some collusive outcome. Fourth, asymmetric firms will find it harder to coordinate on a collusive outcome. Fifth, when firms cannot fully observe each other's actions, collusion may still be feasible, but is much less likely to occur.

3 Some experimental issues

It is not always straightforward how to implement theories regarding collusion in economic experiments. Many choices can be made, regarding e.g. the trading institution, the mode of competition, and the way the demand side of the market is implemented. Implementing the possibility of collusion in an experiment inevitably requires one to make choices regarding the set-up of that experiment. In this section, we highlight and discuss some possible

choices.

Note that the theory describes a situation in which the same group of firms play a repeated game for an infinite number of periods. For that reason, we mostly restrict attention to experiments that use the same group composition for the entire duration of the experiment, rather than e.g. reshuffling groups after each trading period. Of course it is impossible to play an infinitely repeated game in the laboratory. It can be shown, however, that this is analytically equivalent to having a fixed probability of the experiment ending after each period, with the discount factor equal to the probability of continuation. Experiments therefore often use such a fixed probability of continuation after each period, provided that a certain minimum number of periods has been played. Anyhow, there is probably relatively little harm in using a fixed horizon in collusion experiments, provided that that horizon is sufficiently long. Indeed, Selten and Stoecker (1986) note that the observed behavior in a treatment with a long finite horizon is very similar to behavior observed in an infinitely repeated game, apart from an end-game effect.

After determining the number of periods and group composition, an important issue is how to implement the actual market. A first crucial choice is whether to use real buyers, or to implement a simulated market using some implicit demand function. With real buyers, the next choice is the trading institution to use. With a demand function, the next choice is the mode of competition to use.

First consider the trading institution. Holt (1995; Section 5) gives an exhaustive description of all possibilities. For our purposes, the following institutions deserve attention. Consider a market in which both buyers and sellers are active. In a *posted offer auction*, each seller independently selects a price, and buyers are called on in random order to make purchase decisions. In a *posted bid auction*, each buyer selects a price, and sellers are called on in random order to make selling decisions. In a *double auction* buyers and sellers are treated symmetrically, in the sense that both can actively post and accept prices. The choice of trading institution can have a crucial influence on the outcome of the experiment. For example, Smith (1981) compares these institutions using a monopolist seller. In the last session of each of his experiments, he calculates the *monopoly effectiveness index*,

defined as the difference between the actual profit and the competitive profit, divided by the difference between the theoretical monopoly profit and the competitive profit. Thus, when a monopoly can perfectly exploit its monopoly power, the index will be equal to one. When the monopolist only manages to set a price equal to marginal costs, the index equals zero. Smith (1981) finds that with a posted offer monopoly the average value of the index in his experiments is 1. Yet, with a posted bid monopoly it is only 0.15, whereas it equals 0.36 with a double auction monopoly. These results suggest that the results of collusion experiments with active buyers should be treated with particular care. If an oligopoly is not able to reach the perfectly collusive outcome, then this may not say too much about the possibility of collusion when in that particular trading institution even a monopoly is not able to achieve the monopoly outcome.

Second, consider the mode of competition. Some experimentalists prefer the use of Cournot models, whereas others prefer Bertrand models. Bertrand competition can be studied using homogenous products. In that case, competition has a winner-take-all character: the firm that sets the lowest price will capture the entire market. Alternatively, Bertrand competition can be studied with heterogenous products. If that is the case, firms that do not set the lowest price may still earn positive profits. In many experiments, subjects are not given the explicit demand, cost, and profit function. Rather, they face a payoff table that gives their payoffs for all possible choices of themselves and their competitors. Such a payoff table may consist of many possible choices, but there are also cases in which subjects can choose between only three options; for example, a high, an intermediate, or a low price.

The choice of whether to use Bertrand or Cournot competition is a contentious issue. Holt (1995), for example, argues against the use of Cournot competition, since it implies the use of a rather mechanical market-clearing assumption. Of course, following Kreps and Scheinkman (1983) one could defend the use of Cournot competition with the argument that if firms first choose capacities and then set prices, the outcome will be the Cournot equilibrium.

Unfortunately, the experimental evidence for this is weak. Davis (1999) considers tri-

opoly markets. There are two treatments: the first is a posted offer market, the second is a posted offer market with advance production. Without advance production, prices decline slowly toward the competitive level, where prices equal marginal costs. With advance production prices are somewhat higher and quantities somewhat lower. However, there is no convergence to Cournot levels. Muren (2000) also considers triopolies with advance production. In sessions with inexperienced participants quantities are higher than the Cournot level. In sessions with experienced participants outcomes are close to the Cournot level. Anderhub et al. (2003) look at duopoly markets with heterogeneous goods where sellers first choose a capacity, then a price. Chosen capacities are fixed for 5 periods. In the first 10 periods, quantities are exogenously fixed. In contrast to Davis (1999) and Muren (2000) capacity constraints are not strict: sellers can sell more at a cost. Subjects set prices at or close to the equilibrium level most of the time, given the capacity choices. Capacity choices are clustered around the competitive level, and to a lesser extent around the Cournot level. The average lies between the two. So, Anderhub et al. (2003) also do not find that capacity-then-price competition leads to Cournot outcomes.

Experiments regarding collusion also differ in more subtle ways, for example with respect to the amount of information that experimental firms receive regarding the market environment, and regarding the actions and performance of their competitors. The effects that such variables have on the extent of collusion is an issue that we address in the remainder of this paper.

4 Experimental tests of the standard theory

In this section, we first try to answer the question as to whether experimental firms manage to collude in the first place. In other words, we consider experiments with a set-up that is as close as possible to the canonical theoretical model outlined in section 2: firms meet each other repeatedly, and no communication is possible. Are such experimental firms able to collude? That is, are they able to achieve profits that are consistently above the Nash equilibrium profits of the stage game? And if so, are they able to reach the price that

maximizes their joint profits?

Huck et al. (2004) survey the evidence for Cournot models. All the experiments they consider have no communication. After each round, firms only receive aggregate information about the behavior of other firms. There is complete information about the own payoff function, and firms are symmetric. Huck et al. (2004) find that, in studies with two firms, total output falls short of the Cournot prediction by some 7% on average. Hence, industries with two firms manage to collude to some extent; total quantity is lower than in the Cournot model and hence, profits are higher. Yet, these firms come nowhere near perfect collusion, in which joint industry profits are maximized. With a linear demand function, which is commonly used in experiments, perfect collusion would imply that total output falls short of the Cournot prediction by 25%.

With more than two firms, however, the effect disappears entirely. On average, total output in markets with more than two firms slightly exceeds the Cournot prediction. Tentatively, this effect seems to become slightly stronger as the number of firm increases, but this is hardly significant. Hence there seems to be a discontinuity when moving from 2 to 3 firms. Experimental industries with two firms have at least some success in achieving collusion. Experimental industries with more than two firms are not successful in doing so at all. In fact, this result was already found in the seminal work by Fouraker and Siegel (1963). They find industry outputs often below the Cournot level with a duopoly, but above the Cournot level in about two-thirds of the cases with a triopoly.

So far, we have only discussed aggregate outcomes. Individual differences between markets can be huge. Nevertheless, from his reading of the literature, Holt (1995) notes that "with Cournot duopolies, outcomes fall on both sides of the Cournot prediction, and some cases of near collusion occur [...] with more than two sellers, outcomes are often more competitive than the Cournot prediction" (pg. 404). He also concludes that increasing the number of sellers results in more competitive behavior. But with more than three sellers, there "seems to be little or no evidence for a pure-numbers effect that is measured by changing the number of sellers in a way that does not alter the incentive structure".

In experiments with Bertrand pricing, results are similar to those with Cournot com-

petition. From her review of the literature, Wellford (2002),⁷ concludes that experimental price-setting duopolies are sometimes able to achieve collusive outcomes, but that there is often fierce competition. With more than 2 firms, competitive outcomes are again the norm.

Some papers have tried to gain more insights into how exactly experimental firms behave. Mason and Phillips (2002) investigate whether trigger strategies can explain observed behavior in infinitely repeated experimental Cournot duopolies with perfect information. Using a comparative statics approach, they conclude that observed behavior is generally consistent with the use of such strategies. They also find evidence that observed behavior is better explained by grim trigger strategies than by stick and carrot strategies (see section 2). Selten et al (1997) however, seem to contradict this. They elicit strategies from experienced players in a two-player Cournot game using the strategy method, in which subjects have to specify in advance how they will react to different possible choices of their opponents. The authors find that players use strategies that are similar to the tit for tat strategies in Axelrod (1984).

It is an interesting question why in an experiment, at least with Cournot competition, two firms are able to collude, whereas three firms are never able to do so. Holt (1995) gives an intriguing interpretation. The difference may be due to the inability of triopolists to punish one noncooperative rival without harming the other. The standard model predicts that, after a defection from some implicit cartel agreement, firms will "punish" other firms forever by effectively abandoning the cartel. Yet, such punishment is crude and often ill-directed. In the case of two firms, the punishment is still relatively straightforward. When one firm defects, the other reacts by setting a much higher quantity in all future periods. This indeed hurts the other firm. But now consider a set-up in which more firms are active. Suppose again that one firm defects. When one of the other firms reacts by increasing its quantity, it will not only punish the defector, but will also adversely affect the other firms. It is likely that experimental subjects will be much more reluctant to carry out such a punishment. First, they often cannot observe which firm exactly did the defection. Second, even if they could observe this, they are not able to single out this firm

for punishment. Indeed, much of the experimental literature suggests that subjects are driven by some sense of fairness, over and above the objectives that are routinely assumed by economists. And once experimental firms are reluctant to punish a defection from some experimental agreement, it becomes much more attractive to defect in the first place. This in turn makes it much less likely that cartels will be formed.

An important question then becomes how relevant these results are for the real world. Indeed, experimental subjects may be very reluctant to punish cartel defection when this also hurts innocent firms. But it is likely that firms in the real world will be much less reluctant to punish in that manner. One may argue that real firms will be much less concerned about the well-being of their competitors than experimental firms are. If that is indeed the case, then real markets with more than two firms are more likely to collude than experimental markets with more than two firms. Then, the discontinuity in the success of collusion when moving from 2 to 3 firms may just be an artefact of experimental work.

Summarizing, economic experiments that test the standard *tacit* collusion model find the following. Industries with two firms are able to collude to some extent. Yet, they are not able to achieve perfect collusion at all: on average, total output is still much closer to Cournot equilibrium output than it is to monopoly output. Industries with more than two firms are not able to collude at all. Yet, this may be an artefact of experiments. A plausible explanation is that experimental firms are reluctant to implicitly punish firms that did not defect from the collusive outcome. Arguably, real-world firms are much less reluctant to do so.

It is important to note that in all the experiments considered so far, firms were not able to communicate. Hence, these were all tests of really tacit collusion. As we argued in our theory section, one may indeed expect that firms will find it hard to collude in such a world. The model allows for a wide range of possible equilibria and without an exchange of information and communication it may be hard to coordinate on one of these. Therefore, in section 5 we study the role of information in experiments. In section 6 we turn to experiments that allow some form of communication. Interestingly, in a survey paper on competition policy enforcement against collusion, Kühn (2001)? focuses on these

exact same issues.

In the remainder of this paper we will not so much focus on the extent to which experimental firms are successful in colluding. Rather, we will stress the comparative statics of collusion. In other words, we will study the extent to which changing some aspect of the environment in which experimental firms operate, has an effect on their ability to collude. We will focus on recent work.

5 The role of information

5.1 Introduction

In the standard experiments described in the previous section, firms only receive information about the aggregate actions of their competitors in the previous round. Firms know their own payoff functions, and they know that firms are symmetric, so they can also infer their competitor's payoff function. In this section we review the literature that studies the effects of changing the information regime. Such effects are important: if more information indeed facilitates collusion then trade associations, which collect and disseminate information among their members, can be deemed a facilitating practice. In section 5.2, we consider information regarding the state of the world. That state may concern the competitor's payoff function, or parameters of the demand function. In section 5.2.1 we study experiments in which it is exogenously given whether or not that information is available. Section 5.2.2 studies a more subtle issue. In some cases, firms have private information, but may decide to share that information with their competitor. Section 5.3 studies the effect of information regarding the competitor's actions and realized profits. A recent strand of the theoretical literature suggests that such information may induce firms to imitate the firm that has achieved the highest profits in the past. If true, this would yield the perfectly competitive outcome in the long run, with prices equal to marginal costs. In section 5.3, we review experiments testing this hypothesis.

5.2 The state of the world

5.2.1 Exogenous information

Dolbear et al. (1968) are among the very first to study the role of information in experiments regarding collusion. They implement what effectively boils down to a Bertrand model with differentiated products. The authors choose a specification in which demand for a single firm only depends on its own price and the average price of its competitors. They study cases with 2, with 4, and with 16 firms. There are 15 market periods in each session, unknown to the subjects in advance. After each period, firms are also informed about the price their competitor has chosen. The authors use two set-ups: one with incomplete information, and one with complete information. With incomplete information, each firm receives a payoff table which gives its profits for different values of its own price and a range of possible values of its own demand. With complete information each firm has a payoff table which gives its profits for different combinations of the firm's own price and the average rivals' price. Moreover, in this case each firm knows that all other firms have the same payoff table.

The authors find that information does not have a significant effect on average price and profit levels. Yet, keeping the number of firms fixed, more information significantly increases the dispersion of per firm prices and profits among sessions. More information also increases price stability within given markets, as measured by the variation of the average price in given markets during the periods 8 through 12 out of a total of 15 periods. Hence, more information increases the variability across markets, but reduces the variability within one market over time. Dolbear et al. (1968) find that the number of firms adversely affects the extent of collusion.

Mason and Phillips (1997) also investigate the effect of incomplete information regarding the competitor's type. In their experiment, the uncertainty concerns the rival's cost function. Their set-up is based on a Cournot model with homogeneous products. They only study markets with two firms. For each firm, marginal costs are either high or low. As a result there are markets where firms are symmetric, with both firms having either high or

low costs, and markets in which firms are asymmetric, with one having low costs, while the other has high costs. There are two information treatments. With complete information both firms also know their competitor's marginal costs (or rather, their competitor's payoff table). With incomplete information, they only know their own marginal costs. Firms know that they play against the same randomly chosen rival during an unknown number of periods. Sessions turn out take between 35 and 46 periods. At the end of each period, firms are informed about their rival's choice.

The authors find the following. First, in symmetric markets average relative outputs (as a fraction of the Cournot outputs) are significantly lower with full information than with incomplete information. This effect does not occur in asymmetric markets, which suggests that information about one's competitor facilitates collusion, but only when firms are symmetric. With incomplete information, outputs are close to Cournot levels. Second, in both symmetric and asymmetric markets, information about one's competitor increases convergence towards steady-state outputs levels. Third, in asymmetric markets, more information significantly raises the steady-state output share of the most efficient firm. Mason and Philips (1997) conclude on the basis of their experiment that a rule-of-reason approach towards information sharing between rivals is the most appropriate. A critical stance might be deemed necessary towards information sharing between firms with symmetric costs. However, a more relaxed approach might be justified towards firms with asymmetric costs. More information then has no pro-collusive effects and can be welfare enhancing. Of course, it is hard to see how such a policy could be implemented. It seems to require that an antitrust authority already knows the firms' marginal costs, and on the basis of that can decide whether firms are allowed to share that information.

Mason and Phillips (1997) note that their set-up with incomplete information is related to the Green and Porter (1984) model with random demand, that we discussed in our theory section. Indeed in both models each firm is unaware of its rival's payoff. However, for the results in Green and Porter (1984), it is crucial that the rival's payoff function may fluctuate in every period. In this experiment, the rival's payoff function is fixed.

Feinberg and Snyder (2002) is more closely related to Green and Porter (1984). Their

experiment has symmetric price-setting duopolists. Importantly, firms can only choose among three prices: a collusive price, an undercutting price and a punishment price. Obviously, the prices are not labelled in this manner in the experiment. They study three treatments. Treatment 1 is the benchmark treatment. Demand is given and known in each period and after each period each firm can observe the price of its rival. In treatment 2 the prices set by the competitor are no longer revealed. Also, known to the firms, negative demand shocks occur in 10-15% of all periods. A demand shock leads to low payoffs irrespective of the rival's choice, and is revealed immediately after it occurs. Treatment 3 is identical to treatment 2, the only difference being that demand shocks are not revealed. Some experimental subjects play treatments 1 and 2, while others play treatments 1 and 3. After each treatment, they are rematched with a different competitor. The parameters of the experiment are chosen in such a manner that in all treatments tacit collusion can be sustained with perfectly-collusive prices by means of trigger-strategies with a one-period punishment phase.

In treatment 1, the authors find some convergence to the collusive outcome. In treatment 2 the outcomes are similar to those of treatment 1, with no effect on the ability to collude. However, in treatment 3 there is a sharp decline in collusion. Uncertainty about both rival's actions and demand turns out to dramatically impede collusion. This is an interesting finding since the set-up has no effect on the *theoretical* ability to collude. Hence, the anti-collusive effect of demand uncertainty and the possibility of secret price cuts is starker than suggested by Green and Porter (1984). In the experiment, punishments are used by most firms one period after the rival has chosen an undercutting price. Most of the time the rival also reverts to the collusive price after such a punishment period. All this is in line with the theory. Also the occurrence of a negative demand shock does not often lead to less collusive behavior in treatments 2 and 3. But, in treatment 3 this is merely due to the fact that there was little collusion to start with.

5.2.2 The decision to share information

In the experiments described above, the possible sharing of information is merely something that is imposed by the experimenter. Some experimental papers investigate whether voluntary information sharing might facilitate collusion. A body of theoretical research studies whether such arrangements are desirable. Typically, papers in this area study one-shot duopoly models that consist of two stages. In the first stage non-cooperative firms decide whether or not they will exchange information about uncertain demand or cost conditions. In the second stage they compete on the output market. The theoretical literature shows that the decision to share information depends on whether competition is in prices or quantities, whether uncertainty is about demand or costs, and whether goods are substitutes or complements. For instance, Vives (1984) demonstrates that with quantity competition, demand uncertainty and substitute goods, the equilibrium has firms concealing information from their rivals. Based on such results, Clarke (1983) argues that information-sharing arrangements must be considered as direct evidence of collusive behavior. However, further theoretical research has shown that non-cooperative firms might exchange information in equilibrium if costs are sufficiently quadratic or if firms are sufficiently risk averse.

Cason (1994) conducts an experiment to test the predictions of this literature. He implements a single-period duopoly model with differentiated products where firms compete in prices. There may be uncertainty regarding either demand (the common intercept of the linear demand functions of both firms) or marginal costs, assumed constant for both firms. Only firm 1 receives accurate information, and has to decide whether to share this information with its rival. The timing is as follows. First, firm 1 decides whether it will share information. Second, the uncertainty is resolved for firm 1. Third, the information is shared if firm 1 decided to do so. Fourth, the firms compete on the output market. Theory predicts that firm 1 will conceal the information with demand uncertainty and complementary products, and with cost uncertainty and substitute products. Firm 1 will reveal the information with demand uncertainty and substitute products, and with cost uncertainty and complementary products.

Observed behavior in the experiment strongly supports the theory in cases with cost uncertainty. With demand uncertainty, the theory did much worse. The author suggests that that may be caused by small differences in firm 1's payoff function under sharing and non-sharing of information in that case. Some collusive pricing was observed in the case of cost uncertainty, complementary goods and information sharing between the firms. In all other cases, prices were close to Nash equilibrium levels. Hence, the experiments with single-period duopolies provide only weak support for the often-expressed claim that information sharing between firms enhances the possibility of collusive behavior.

However, that picture may change entirely if we study this issue in the framework of a repeated game. Cason and Mason (1999) do exactly that. In their experiment, they study a duopoly in which firms compete in quantities and sell homogeneous products. Demand (the vertical intercept of the linear demand function) is unknown. Cost functions are quadratic. In their treatment 1, the timing is as follows. First, each firm decides whether it will share information with its rival. Second, each firm receives some signal about the true demand. Demand can be either high, medium or low, and it varies across periods. If *both* firms have agreed to share information, then combining the two signals allows one to infer the true state of demand. Without information sharing, each firm remains only imperfectly informed about true demand. Third, both firms compete in the output market. At the end of each period firms are informed about their own payoff as well as the rival's output choice and payoff. The authors set the parameters such that in the theoretical equilibrium of a single-period game each firm chooses to conceal its information. In this way it becomes possible to focus on the question whether sharing of information might arise as a strategy in a repeated game context.

The results are as follows. First, firms very often share information. Second, in the cases where firms do share, they reduce their output significantly below the Cournot output in situations with high demand, and even more so in situations with low demand. This seems in line with the theoretical predictions of Rotemberg and Saloner (1986) who argue that in states with low demand collusion is more effective. Yet, in the experiment, the perfectly collusive output is not attained. In cases where firms choose not to share information,

output in low-demand states is significantly *larger* than the Cournot output. It seems that in those cases firms conceal information in an attempt to punish non-collusive behavior of rivals in previous periods.

In a different treatment, Cason and Mason (1999) change the model parameters such that information sharing can no longer be explained by collusive motivations, by drastically reducing the difference between the perfectly-collusive payoffs and the Cournot payoffs. In that treatment, firms still share information at the same rate as they do in the original experiment. Hence, it seems that the decision to share information cannot be explained as a strategy to deliberately facilitate collusion. Instead, the authors argue that risk aversion is the most plausible explanation for information sharing. The authors conclude that information sharing only has a moderate collusive effect in states of low demand but might be valuable for firms since it reduces their uncertainty. They go on to argue that the results are supportive of a relaxed policy towards information-sharing arrangements by trade and professional associations.

Yet, one may argue about the conclusions the authors draw from their experiment. First, they note that information sharing only has moderately collusive effects. Yet, the outputs that are found are still significantly below Cournot outputs, so there *is* collusion. From an experiment, it is difficult to gauge how important this effect is in the real world. Second, even though the authors suggest that collusion is not the main *reason* that firms share information, it is still an important *effect* of such information sharing. Third, this experiment may not be the most appropriate way to evaluate the effects of trade associations. When firms *unilaterally* benefit from revealing information to their competitors, then they may simply choose to do so, and they do not need a trade association. Yet, there may be situations in which disclosing information is a prisoners' dilemma: each individual firm does not have an incentive to reveal information, but all firms would be better off if they could all commit to do so. The relevant policy question then becomes not whether firms should be allowed to disclose information, but rather whether they should be allowed to commit to exchange information.

5.3 Actions and performances

Stigler (1964) argues that the availability of firm-specific information facilitates collusion as it hinders secret price cuts. Vega-Redondo (1997) challenges that argument. He considers a homogeneous symmetric Cournot model. After each period, each firm receives information about the outputs and realized profits of each of its rivals. Vega-Redondo assumes that each firm imitates the output chosen by the firm with the highest profit in the previous period. Yet, with a small probability a firm mistakenly selects some arbitrary output. Under those assumptions, he shows that the market converges towards a Walrasian equilibrium with price equal to marginal cost. Intuitively, if price is larger than marginal cost, the firm with the highest output earns the highest profits. Imitation of the most successful firm thus induces an increase of outputs towards the Walrasian level. Similarly, if price is lower than marginal cost, the firm with the lowest output makes the highest profits. Imitation then yields a lower total output. Again, price moves towards the Walrasian equilibrium. Hence, under this behavioral assumption, the use of firm-specific information leads to a competitive outcome rather than to a collusive one.

A number of experimental studies have examined whether the outcome predicted by Vega-Redondo (1997) is also observed in the laboratory. Huck et al. (1999) examine Cournot markets of a homogeneous good in which fixed groups of four identical firms select outputs in a sequence of 40 periods. Because the endpoint of the sequence was known in advance to all firms, we do not have an infinitely repeated game here, and as a consequence the theoretical results for supergames discussed in Section 2 do not apply. However, as we argued in section 3, this need not matter that much. In treatment 1 each firm only knows that it is active on a market with three rivals and that all firms set quantities. After each period each firm only learns its own profit. In treatment 2 each firm is fully informed about its demand function and cost function, and receives the information that all firms are identical. In addition, in each period each firm learns the outputs and profits of each rival, as well as the average rivals' output. In other treatments the amount of information provided are in between these two extreme cases.

In all treatments outcomes are very competitive in the sense that average outputs are always larger than the Cournot equilibrium. In one treatment, the average output almost coincides with the Walrasian outcome, providing support for Vega-Redondo (1997). More generally, the market outcome is less competitive if more information is provided about market conditions. Note that this is roughly in line with Mason and Phillips (1997), discussed in section 5.2.1. The market outcome is more competitive if more information is provided about rivals' choices and profits. At the individual firm level, there is quite some behavioral variation across markets.

In a companion paper, Huck et al. (2000) discuss a similar experiment concerning a market of a differentiated good. In all treatments firms know the demand and the identical cost functions. The treatments differ with respect to the mode of competition (Cournot or Bertrand) and the amount of information firms receive. In treatment 1 each firm only learns the average choice of its rivals in previous period. In treatment 2 each firm receives firm-specific information about its rivals' choices and profits. If firms only receive information about the average action of their rivals, the outcome converges to the Nash equilibrium of the stage game, regardless of the mode of competition. Firm-specific information leads to more intense competition with output choices, again supporting Vega-Redondo (1997). It has no effect with price competition. Again, at the individual firm level, no uniform behavioral pattern emerges.

Summarizing, Huck et al. (1999, 2000) provide at the average output level some support for the theoretical prediction of Vega-Redondo (1997), and thus suggest that publication of firm-specific data about actions and profits might enhance competition on markets. However, the authors are cautious to draw strong policy implications from their experimental results. Also note that both papers do not provide any evidence of successful collusion, reinforcing the general conclusion that collusion is seldomly encountered in experimental markets with more than two firms.

Results related to these are obtained by other experimental studies of repeated Cournot markets with a homogeneous good. Rassenti et al. (2000) report that in a market with five firms, where firms have private information about their costs, total output converges to the

Cournot level, irrespective of whether firms received aggregate or firm-specific information about their rivals in that period. Offerman et al. (2002) consider markets with three firms and do not find a significant and lasting effect on total output if firms receive firm-specific rather than aggregate data about their rivals in any period. Yet, in their experiment total output is always smaller than or equal to the Nash equilibrium level of the stage game. In that sense their outcomes are more collusive than those of Huck et al. (1997, 2000).

Bru et al. (2002) challenge these studies. They conduct an experiment in which four symmetric firms always have complete information about the demand and cost function. In one treatment, each firm always receives the entire history of each rival's outputs and profits. In the other treatment, this information is restricted to the last period of play, as we have in the studies discussed above. When firms have information about the entire history of play, total outputs are more often and persistently lower and thus *more* collusive. Yet, outcomes vary widely across markets. It appears that behavior on a given market is highly influenced by the specific history of that market; for instance, by the ability of firms to understand the strategic interactions at stake, and the patience they have with deviating firms.

If anything, these experiments serve as an illustration that one should be very careful when drawing conclusions from experiments for markets in the real world. In most of the experiments discussed in this section, firms only receive information about their competitor's performance in the previous period. Hence, firms are also likely to only use *that* information. Yet, it is hard to imagine that firms that operate in the real world behave in such manner, disregarding everything that has happened before the final period. Real firms will gather, interpret and use information in a manner far more careful and elaborate than experimental subjects who take their decisions in a matter of seconds. In that sense, the experiment in Bru et al. (2002) seems closer to the real world than the other experiments discussed in this section, which would also imply that in the real world, firms are more likely to achieve collusive outcomes than most of these experiments suggest.

5.4 Summary

Evidence on the effect of information regarding the competitor's type on the ability to collude is mixed. More information regarding the competitor's type does make a given market more stable, but also less predictable. Demand shocks have no effect on the ability to collude – provided that the demand shock is revealed afterwards. In a repeated game, the voluntary sharing of information has a collusive effect, even though collusion is not the main reason that firms choose to share information. In experiments, the effect of revealing information regarding past performance of rivals is ambiguous. Some experiments seem to suggest that revealing such information makes markets more competitive. Yet, revealing information about performance in all previous periods seems to facilitate collusion.

6 The role of communication

6.1 Introduction

As we argued in section 4, many experiments test whether experimental firms are able to collude *tacitly*. One strand of the literature though, does allow communication between firms about future conduct. This may make it much easier for them to coordinate on some collusive outcome among the infinite number of possible collusive equilibria. In this section, we review those experiments. Early experiments already studied the effect of communication on collusion. This goes back as far as Friedman (1967). He finds that communication through written messages between sellers has a large and positive effect on prices in experimental posted-offer markets. Isaac et al. (1984) find that the same is true with face-to-face-communication. Section 6.2 discusses the more recent literature. The possibility to communicate may be a double-edged sword, however. If sellers are able to communicate not only amongst themselves, but also with buyers, then this may undermine a tentative collusive agreement. Sellers may then be able to give secret discounts. That possibility is likely to undermine a potential collusive agreement. Experiments with the possibility of secret discounts are also discussed in section 6.2.

However, direct communication between firms about their future prices is routinely for-

bidden by antitrust laws. Firms that do communicate about prices are therefore vulnerable to antitrust action. Therefore, they may be able find more subtle ways to communicate their intentions, and to communicate on an equilibrium. One way to do so is through price announcements. Section 6.3 discusses experimental work allowing for these.

6.2 Direct communication: explicit collusion

Davis and Holt (1998) study experimental markets that consist of three buyers and three sellers. Participants do not know each other's costs or values. Sellers post prices that are readily observable to buyers. In treatment 1, with no communication, prices converge to the competitive level. In treatment 2, sellers can communicate; buyers have to leave the room before the actual market starts, and sellers are then able to communicate freely. In this treatment, prices are close to monopoly levels. Some sellers even devise elaborate collusive schemes, in which they took turns in being the lowest pricing seller. In treatment 3, sellers are also able to communicate with buyers, in the sense that they can give secret discounts to buyers who ask for these. Buyers have search costs in the sense that they are charged a small amount every time they approach a seller to ask for a discount. In this treatment, list prices are still close to monopoly levels. But actual transaction prices are much lower, approaching competitive levels in 5 out of 6 sessions. Apparently, the colluding firms cannot resist the temptation to defect from the collusive price when explicitly asked to do so. In treatment 4, actual sales quantities of each firm are revealed after each period. This information is given before negotiations start for the next round, and is meant to represent data such as a trade organization might release. In 5 out of 6 sessions, prices again reach the collusive level.

Thus, also in this experiment, the possibility of sellers to communicate yields collusive outcomes. But that is mitigated by the possibility of buyers to also communicate with sellers, and secure discounts. Yet, when firms can observe each other's quantities, and hence can observe whether there is cheating, prices may well return to the collusive level. The authors conclude that “[t]he implication is that antitrust hostility to agreements to limit or coordinate discounts is well founded, and that trade association activities should

receive careful scrutiny in markets where collusion is suspected.” (pg. 755).

Note that in Isaac and Plott (1981) face-to-face communication also does not help to coordinate on a collusive price in double auctions. Hence, it appears that in institutions where buyers have a more active role than in standard posted-offer markets, pro-collusive effects of direct communication between firms are mitigated.

Davis and Holt (1994) also study the possibility of secret discounts. But in their set-up, sellers are not able to communicate amongst each other. Two firms compete in prices. There are three buyers. The authors consider two different treatments. In treatment 1 sellers compete in a standard posted offer market. In treatment 2 sellers post prices, after which buyers can approach a seller and ask for a discount. The seller can either give the discount or refuse it. The buyer then decides whether or not to buy the good. If he decides not to, he can switch to the other seller without any costs. Discounts cannot be observed by other buyers and sellers. Davis and Holt (1994) use 6 sessions for each treatment. Sessions run for at least 15 periods, followed by a random termination rule.

The results are as follows. In the normal posted offer markets prices are close to the theoretical prediction, even though sellers do not randomize over prices as predicted by theory. In the sessions where discounts are possible there are two different results: if sellers choose to compete both with list prices and transaction prices then transaction prices are low. If sellers choose to post uninformatively high list prices and compete only with transactions prices then transaction prices are high.

6.3 Indirect communication: signalling intentions

As we have argued repeatedly, one major obstacle for firms to achieve collusion in the real world is the sheer multiplicity of collusive equilibria. Firms may have difficulty in coordinating which of these equilibria they will play. Any device that allows firms to achieve such coordination is likely to be a major aid in achieving collusion. One such device is direct communication. But there may also be more subtle ways in which firms can communicate their intentions. One device could be the use of non-binding price announcements, see e.g. Blair and Romano (2002). Here, firms announce some price that they plan to set in

the future. They can then observe the reactions of other firms to their announcements. Suppose for example that a firm wants to coordinate on some high collusive price. It can announce to set that price in the future. If other firms follow by announcing that same price, the firm simply goes through with its announcement. If other firms do not follow, the firm retracts its announcement and sets a lower price. Since no trade takes place at the announced prices, firms risk losing anything from these announcements. Of course, these price announcements may serve purposes different from collusion. Firms using them often claim that they are helpful for their customers, as they reduce uncertainty regarding future prices, and therefore serve as an aid in their customers' planning.

In an oft-cited study, Grether and Plott (1984) examine the effect of such announcements. Their study is inspired by the Ethyl-case in the late 1970s in the US, where the Federal Trade Commission (FTC) filed a complaint against a number of domestic producers of additives that raise the octane level in gasoline. The allegation of the FTC was that price competition on this market was significantly reduced by the joint presence of the following practices: (i) advance notice of price changes and price announcements, (ii) most favoured nation clauses, and (iii) delivered pricing. The advance notice provision implies that a firm gives at least thirty-days in advance public notice of price changes to its rivals and customers. Non-binding price announcements were often made during a short period just before the thirty-day deadline. With a most favoured nation clause a firm guarantees each of its customers a price that is not higher than the price charged to its other customers. A delivered pricing clause implies that each firm uses a uniform list price for all its customers, regardless of their location. Thus, net prices can only differ due to differences in transportation costs. From a theoretical point of view, one can argue that the alleged practices reduce the uncertainty about market prices and can facilitate coordination on higher prices. The defense put forward by the four accused firms was that the industry was as competitive as could be expected given its structure. Moreover, the firms argued that the practices had no effect on the competition in the industry.

To shed some light on this controversy, Grether and Plott (1984) study an experimental market of a homogeneous good with four firms and nine buyers. Demand and supply

conditions are designed to reflect the relevant industry. In treatment 1 all three practices are present. In treatment 2 none of them are. Average prices are significantly higher in treatment 1. Yet, the average price level was not very high: somewhat between the static competitive price level and the price level associated with Cournot output. Moreover, Grether and Plott (1984) focus on the effect of the joint presence of the facilitating practices, and do not analyze the marginal effect of either one of them.

Holt and Davis (1990) do consider the isolated effect of non-binding advance price announcements. Groups consisting of three firms take decisions in a sequence of periods with unknown endpoint. First, 15 periods are played under a standard posted-offer rule. After that, there is a number of periods with price announcements. These periods consist of two stages. In the first stage, there is a tightly structured process of price announcements. One of the sellers can send one price to its two rivals. Both rivals can then react with either ‘A’, ‘L’ or ‘H’, denoting agreement, disagreement and a lower price, and disagreement and a higher price respectively. Both the price announcement and the two reactions are non-binding, and all information is displayed to all firms. In stage 2, each firm selects a binding price. Buyers are represented by a demand function.

Initially, price announcements have a large and positive effect on prices. However, the effect appears to be only temporary. Soon prices converge downwards towards the same level as in the environment without price announcements. Thus, this experiment suggests that non-binding price announcements only have a temporary effect on the ability to collude.

Cason (1995) and Cason and Davis (1995) also consider the effects of non-binding advance price announcements. Both experimental studies are inspired by a US antitrust case in the 1990s. Eight airlines used the Airline Tariff Publishing (ATP) system, an electronic network through which they could make non-binding price announcements. The system enabled them to continuously announce, modify and withdraw non-binding prices of future tickets, and all changes could be made at a very low cost. Potential customers could not observe these announcements. The airlines were accused of illegally raising their ticket prices by means of the ATP-system. In 1994 the case was settled, and the airlines

stopped using the ATP-system.

Cason (1995) investigates whether firms in an experiment are able to coordinate on higher prices if they can continuously announce non-binding prices in an entirely unstructured manner, i.e. without a formal communication structure such as the one imposed by Holt and Davis (1990). In Cason (1995), five firms are active in a repeated posted-offer market with simulated buyers. The firms play 10 periods without the option of price announcements and 10 periods with such an option. In half of the sessions the price announcement game is played first, in the other half the non-price announcement game is. All sessions last at least 20 periods. In the treatment with price announcements, each firm can continuously post non-binding prices during one minute. Each firm is free to enter, change or cancel any price announcement at no cost, and all announcements are displayed to all sellers. The standing price announcement at the end of the one-minute window is binding. Demand and supply conditions are such that nearly all the surplus at the Nash equilibrium of the stage game is captured by the buyers. Hence, firms have much to win from collusion. The results show that the option to announce prices has a weak positive but transitory effect on prices. After some time, the announcements lose their significance and actual prices converge to competitive levels. Hence, the results in this experiment are qualitative identical to those in Holt and Davis (1990). In a treatment in which sellers have all bargaining power, announcements have no influence whatsoever.

Cason and Davis (1995) observe that an important feature of the airline industry is that most firms are active on different markets. Therefore, they extend the experiment in Cason (1995) to one with multiple markets, where three sellers simultaneously operate on three different markets. They find that prices are persistently higher when firms have the opportunity to continuously make non-binding price announcements. However, price paths varied largely between sessions, and it appears that firms are not able to use price announcements in combination with trigger-strategies as predicted by theory.

The experimental results thus only provide weak and somewhat mixed evidence for anti-competitive effects of non-binding price announcements. Still, there is a widespread belief that the ATP-system has facilitated collusive pricing in the airline industry in the

real world.

In a somewhat related paper, Harstad et al. (1998) test whether the concern of competition authorities about parallel pricing is justified. There are four subjects in each session, two with a ‘high quality’ product and two with a ‘low quality’ product. In treatment 1, players are independent price setters. In treatment 2, each period has two stages. In the first stage, players may announce as many prices as they like. At the end of that stage, they set their final price. In treatment 3, players can choose whether or not to match announced price changes. Finally, in treatment 4, subjects that have played in treatment 3 are now interacting in the setting of treatment 2 to see whether they will match price changes even if they are not constrained to do so. Prices in markets where firms are constrained to match competitors’ price changes are higher than in markets in which firms can make price announcements and then pick any price. This is true to a lesser extent in markets when firms have experience with parallel pricing. These results suggest that competition authorities are somewhat justified in their concerns that parallel pricing may lead to more collusive outcomes.

6.4 Summary

The possibility of communication between sellers has a strong and positive effect on the ability to collude. Yet, this effect is mitigated when buyers are able to ask for secret discounts from sellers and when sellers cannot observe each other’s quantities. The experimental evidence for the collusive effect of nonbinding price announcement is only temporary.

7 Does history matter?

In the previous section we discussed how successful communication can be in coordinating on a collusive outcome. Yet, there may still be another device through which such coordination can be achieved. The history of an industry may have an influence on the feasibility of collusion. Firms that have colluded in the past may also be more successful in achieving collusion in the future. Firms that have never colluded may find it hard to start colluding.

Indeed, one may argue that in the Netherlands, one reason that the construction sector was so successful in achieving collusion was probably that they were simply allowed to collude in the past, and that the institutions and mindsets that facilitate collusive practices were still in place. If history is important, and an industry has colluded in the past, then it may be particularly difficult to break down such a cartel once and for all.

Also from a theoretical perspective, past experience may affect the future ability to collude. In our theory section we argued that the multiplicity of equilibria may make it difficult for firms to coordinate on some collusive equilibrium, especially when there is no focal point. Whatever has happened in the past may exactly provide such a focal point. Theoretical work by Kreps and Spence (1985) also suggests that market history might have a large impact on the ability of firms to reach a collusive outcome.

Phillips et al. (1987; as cited by Holcomb and Nelson, 1997) provide experimental evidence for this suggestion by examining repeated quantity-setting and repeated price-setting duopolies. The former are more conducive to collusion than the latter. Players start out in one treatment, and are then put in the other. They authors find that players that start out in a more conducive environment are also more successful in achieving collusion in the less conducive environment. Similarly, players that start out in the less conducive environment are also less successful in achieving collusion in the more conducive environment. So, indeed, history does matter in the ability to collude.

Holcomb and Nelson (1997) qualify this result by looking at the role of information (or monitoring, as they call it) in a repeated duopoly. They have a symmetric Cournot set-up with differentiated products. First, at least 20 periods were played under complete information, where after each period a firm learns the choice made by its rival. During the first 5 periods, firms are also allowed to communicate by means of written messages. After the periods with complete information, the experiment is put on hold and the firms are told that they will enter a new environment with incomplete information. With incomplete information, firms also receive information about their competitor's choice after each period. Yet, known to them, with probability 1/2 this information is not the true choice but rather some randomly chosen number. After a sequence of at least 20 periods with

imperfect information, the experiment is put on hold again and returns to the environment with complete information.

In 17 out of 20 markets, the firms reach the perfectly collusive outcome during the first 20 periods with complete information. After incomplete information is introduced output significantly increases towards the Cournot output in 16 of these markets, with the one remaining market sticking to the perfectly collusive outcome. Yet, when firms switch back to the complete information environment, all pairs again set collusive output levels.

The lessons that can be drawn from this experiment are numerous. First, at one level it appears that here history does not matter for the ability to collude. A switch to incomplete information destroys collusion, while the switch back to complete information reinstates it just as easily. This is exactly what the authors claim. But on the other hand, one could argue that the possibility to communicate at the beginning of the game is crucial in achieving collusion – as firms in very similar experiments that are not allowed to communicate fall short in achieving full collusion by far. In fact the effect of communication is so strong that, even after a spell of non-collusion, firms are still able to get back to the fully collusive outcome that results from that early communication. Hence, in that sense, history matters a great deal.

8 Conclusion

The main insights we gain from the experimental literature on collusion are the following. First, in an environment in which firms are unable to communicate, experimental firms have little success in achieving collusion. The only mildly collusive results are achieved when only two firms are active. In that case, per-period profits are somewhat higher than in the Nash equilibrium of the stage game. With more than two firms, this is no longer the case. Second, increasing the amount of available information concerning the pay-off function of the competitor facilitates collusion only mildly, if at all. It does, however, have a positive effect on market stability. Third, voluntary information sharing has some collusive effect, even though collusion is not the main reason that firms choose to do so. Fourth, the

effect of revealing information regarding past performance of rivals is ambiguous. Fifth, the ability to communicate among sellers has a strong and positive effect on the ability to collude. This effect is mitigated when buyers can communicate with sellers and sellers cannot observe each other's output levels. Sixth, nonbinding price announcements are collusive, but only temporarily. Seventh, history seems to have some effect on the ability to collude.

Taken at face value, the evidence thus suggests that collusion, if it occurs at all, is most likely to occur in industries where a limited number of firms is active, where firms are able to communicate with each other, and where there is no scope for secret discounts. There is no unambiguous experimental evidence that trade associations, where firms are able to exchange information regarding their sales, prices and profits, are detrimental. Of course, it is an entirely different story when these trade associations also facilitate communication among firms.

As a policy implication, this seems to suggest that in order to avoid collusion, antitrust authorities should first and foremost prevent firms from communicating about future prices and quantities. From a theoretical point of view, this is somewhat surprising, as one could argue that communication is just cheap talk that can have no effect as long as it does not lead to binding agreements. Yet, cheap talk can still be useful to select an equilibrium among the many possible equilibria in a repeated game model. The experimental evidence also suggests that antitrust authorities do not have to worry about tacit collusion. Coordinating the actions of an oligopoly does require explicit communication. Tacit collusion just won't work.

At the same time, we feel that one should be extremely careful in drawing general real-world lessons from experiments in collusion. A situation in which inexperienced students have to make decisions in a matter of seconds is fundamentally different from one in which multinationals can spend months pondering the strategic decisions to be made. Still, even when acknowledging this important caveat, we feel that experiments are potentially able to tell more about antitrust policy towards collusion than they have done so far. Many experiments are primarily designed to test the existing theory. They are conducted by

experimentalists that put an emphasis on the importance of control in the laboratory. This emphasis is justified if the aim of an experiment is to test some theoretical results, but much less so if the aim is to test how collusion occurs in the real world, and how an antitrust authority may be able to curb collusion. For that purpose, the situation in the laboratory should resemble the real world as closely as possible, rather than resembling the idealized world assumed by the theory. The ideal experimental set-up would be one with human buyers and human sellers, that have some possibility to interact with each other. Davis and Holt (1998), for example, is a first step towards that ideal. Also, it may be fruitful to have much more case-based experiments, that are inspired by specific antitrust cases, the classic example being Grether and Plott (1984).

From real-world cartels, we also know that it is hard to draw general conclusions (see e.g. the survey by Suslov and Levenstein, 2002). This is also true in experimental work on cartels. In some experimental markets, firms are able to collude, whereas in others they are not, even if these markets are subject to the exact same experimental treatment. Something similar seems to hold in the real world. In real cartels, the challenge for firms is often to find robust cartel agreements that survive external shocks (again, see Suslov and Levenstein, 2002). A particularly robust agreement is one in which firms divide the market along some geographical lines. Also, bargaining over the rents earned by the cartel are an important issue. To our knowledge, there are no experiments that focus on these issues.

To be useful for antitrust policy, economic experiments should study issues that are close to real-world situations, and that theory is not or hardly able to address. Bringing experiments closer to practice could yield some new insights. As an example, experiments show that firms are much more successful in colluding with free communication than they are without communication. But in the real world, neither of these two extreme cases occurs. A world in which firms are not able to communicate whatsoever, is impossible to imagine. But a world in which firms can explicitly talk about future prices also does not exist, as antitrust authorities do not allow it. Hence, it would be interesting to conduct an experiment in which firms are able to communicate — but not about future conduct. More

generally, it seems worthwhile to have experiments with an explicit role for an antitrust authority. Other issues that have not been studied include e.g. firm size asymmetries, the possibility of a geographical division of the market, price leadership and information sharing among buyers.

If experiments about collusion teaches us anything, it is that small institutional details can matter a great deal for market outcomes. We feel that this lesson also holds for the real world. Institutional details in a certain industry may determine whether firms in that industry are able to collude successfully. Yet, if small institutional details are indeed that important, then that also implies that it is extremely difficult to use experiments in order to evaluate whether real-world industries are susceptible to collusion.

References

- ABREU, D. (1986): “Extremal equilibria of oligopolistic supergames,” *Journal of Economic Theory*, 39, 191–225.
- ANDERHUB, V., W. GÜTH, U. KAMECKE, and H.-T. NORMANN (2003): “Capacity choices and price competition in experimental markets,” *Experimental Economics*, 6, 27–52.
- AXELROD, R. (1984): *The Evolution of Cooperation*. Basic, New York.
- BLAIR, R. D., and R. E. ROMANO (2002): “Advance price announcements and antitrust policy,” *International Review of Law and Economics*, 21, 435–452.
- BRU, L., R. GOMEZ, and J. M. ORDOÑEZ (2002): “Information in repeated experimental Cournot games,” in *Experiments investigating market power*, ed. by R. M. Isaac, and C. A. Holt, pp. 191–215. JAI Elsevier, Amsterdam.
- CASON, T. N. (1994): “The impact of information sharing opportunities on market outcomes: an experimental study,” *Southern Economic Journal*, 61, 18–39.

- (1995): “Cheap talk price signaling in laboratory markets,” *Information Economics and Policy*, 7, 183–204.
- CASON, T. N., and D. D. DAVIS (1995): “Price communications in a multi-market context: an experimental investigation,” *Review of Industrial Organization*, 10, 769–787.
- CASON, T. N., and C. F. MASON (1999): “Information sharing and tacit collusion in laboratory duopoly markets,” *Economic Inquiry*, 37, 258–281.
- CLARKE, R. N. (1983): “Collusion and the incentives for risk sharing,” *Bell Journal of Economics*, 14, 383–394.
- DAVIS, D. D. (1999): “Advance production and Cournot outcomes: an experimental investigation,” *Journal of Economic Behavior and Organization*, 40, 59–79.
- DAVIS, D. D., and C. A. HOLT (1994): “The effects of discounting opportunities in laboratory posted offer markets,” *Economics Letters*, 44, 249–253.
- (1998): “Conspiracies and secret price discounts in laboratory markets,” *Economic Journal*, 108, 736–756.
- DOLBEAR, F. T., L. B. LAVE, G. BOWMAN, A. LIEBERMAN, E. C. PRESCOTT, F. H. RUETER, and R. SHERMAN (1968): “Collusion in oligopoly: an experiment on the effect of numbers and information,” *Quarterly Journal of Economics*, 82, 240–259.
- FEINBERG, R. M., and C. SNYDER (2002): “Collusion with secret price cuts: an experimental investigation,” *Economics Bulletin*, 3, 1–11.
- FOURAKER, L. E., and S. SIEGEL (1963): *Bargaining Behavior*. McGraw-Hill, New York.
- FRIEDMAN, J. W. (1967): “An experimental study of cooperative duopoly,” *Econometrica*, 35, 379–397.

- (1971): “Non-cooperative equilibrium for supergames,” *Review of Economic Studies*, 38, 1–12.
- GREEN, E. J., and R. H. PORTER (1984): “Noncooperative collusion under imperfect price information,” *Econometrica*, 52, 87–100.
- GREETHER, D., and C. PLOTT (1984): “The effects of market practices in oligopolistic markets: An experimental examination of the ethyl case,” *Economic Inquiry*, 22, 479–507.
- HARSTAD, R., S. MARTIN, and H.-T. NORMANN (1998): “Intertemporal pricing schemes: Experimental tests of consciously parallel behaviour in oligopoly,” in *Applied Industrial Organization*, ed. by L. Philips. Cambridge University Press, Cambridge.
- HOLCOMB, J. H., and P. S. NELSON (1997): “The role of monitoring in duopoly market outcomes,” *Journal of Socio-Economics*, 26, 79–93.
- HOLT, C. A. (1995): “Industrial organization: a survey of laboratory research,” in *The handbook of experimental economics*, ed. by J. H. Kagel, and A. E. Roth. Princeton University Press, Princeton, NJ.
- HOLT, C. A., and D. D. DAVIS (1990): “The effects of non-binding price announcements on posted offer markets,” *Economics Letters*, 34, 307–310.
- HUCK, S., H.-T. NORMANN, and J. OECHSSLER (1999): “Learning in Cournot oligopoly: an experiment,” *Economic Journal*, 109, C80–C95.
- (2000): “Does information about competitors’ actions increase or decrease competition in experimental oligopoly markets?,” *International Journal of Industrial Organization*, 18, 39–57.
- (2004): “Two are few an four are many: number effects in experimental oligopolies,” *Journal of Economic Behavior and Organization*, 53, 435–446.

- ISAAC, R. M., and C. PLOTT (1981): “The opportunity for conspiracy in restraint of trade,” *Journal of Economic Behavior and Organization*, 2, 1–30.
- ISAAC, R. M., V. RAMEY, and A. WILLIAMS (1984): “The effects of market organization on conspiracies in restraint of trade,” *Journal of Economic Behavior and Organization*, 5, 191–222.
- KREPS, D., and J. SCHEINKMAN (1983): “Quantity precommitment and Bertrand competition yield Cournot outcomes,” *Bell Journal of Economics*, 14, 326–337.
- KREPS, D., and A. M. SPENCE (1985): “Modelling the role of history in industrial organization and competition,” in *Issues in contemporary microeconomics and welfare*, ed. by G. R. Feiwel, Albany, NY. State University of New York Press.
- MASON, C., and O. PHILLIPS (2002): “In Support of Trigger Strategies: Experimental Evidence from Two-Person Noncooperative Games,” *Journal of Economics and Management Strategy*, 11, 685–716.
- MASON, C. F., and O. R. PHILLIPS (1997): “Information and cost asymmetry in experimental duopoly markets,” *Review of Economics and Statistics*, 79, 290–299.
- MOTTA, M. (2004): *Competition policy: theory and practice*. Cambridge University Press, Cambridge.
- MUREN, A. (2000): “Quantity precommitment in experimental oligopoly,” *Journal of Economic Behavior and Organization*, 41, 147–157.
- OFFERMAN, T., J. POTTERS, and J. SONNEMANS (2002): “Imitation and belief learning in an oligopoly experiment,” *Review of Economic Studies*, 69, 973–997.
- PHILLIPS, O. R., R. C. BATTALIO, and J. H. HOLCOMB (1987): “Duopoly behavior with market history,” Working paper, University of Wyoming.

- RASSENTI, S., S. S. REYNOLDS, V. L. SMITH, and F. SZIRADOVSKY (2000): "Adaptation and convergence of behavior in repeated experimental Cournot games," *Journal of Economic Behavior and Organization*, 41, 117–146.
- ROTEMBERG, J. J., and G. SALONER (1986): "A supergame-theoretic model of price wars during booms," *American Economic Review*, 76, 390–407.
- SELTEN, R., M. MITZKEWITZ, and G. R. UHLICH (1997): "Duopoly Strategies Programmed by Experienced Players," *Econometrica*, 65(3), 517–55.
- SELTEN, R., and R. STOECKER (1986): "End behavior in finite prisoner's dilemma supergames," *Journal of Economic Behavior and Organization*, 7, 47–70.
- SMITH, V. L. (1981): "An Empirical study of decentralized institutions of monopoly restraint," in *Essays in contemporary fields of economics in honor of E.T. Weiler (1914-1979)*, ed. by J. Quirk, and G. Horwich, pp. 83–106, West Lafayette. Purdue University Press.
- STIGLER, G. J. (1964): "A theory of oligopoly," *Journal of Political Economy*, 72, 44–61.
- SUSLOW, V. Y., and M. LEVENSTEIN (2002): "What Determines Cartel Success?," Working Paper 02-001, University of Michigan Business School, Ann Arbor.
- TIROLE, J. (1988): *The Theory of Industrial Organization*. MIT-Press, Cambridge, Mass.
- VEGA-REDONDO, F. (1997): "The evolution of Walrasian behavior," *Econometrica*, 65, 375–384.
- VIVES, X. (1984): "Duopoly information equilibrium: Cournot and Bertrand," *Journal of Economic Theory*, 34, 71–94.