ELSEVIER

International Journal of Industrial Organization
18 (2000) 39–57

International Journal of Industrial Organization

www.elsevier.com/locate/econbase

Does information about competitors' actions increase or decrease competition in experimental oligopoly markets?

Steffen Huck, Hans-Theo Normann*, Jörg Oechssler

Department of Economics, Humboldt-Universität zu Berlin, Spandauer Str. 1, 10178 Berlin, Germany

Abstract

This paper investigates the impact the publication of firm-specific data has on the competitiveness of experimental oligopoly markets. We compare two treatments: in one, firms are informed about their rivals' actions and profits. In the other, firms are only given aggregate information about their rivals' actions (average quantities or prices). We find that more information leads to more competition. In the treatment where aggregate information is given, we confirm the theoretical result that Bertrand markets are more competitive than Cournot markets. © 2000 Elsevier Science B.V. All rights reserved.

Keywords: Information; Collusion; Prices versus quantities; Imitation

JEL classification: L13; C92; C72

1. Introduction

The publication of firm-specific data and, in particular, its impact on the competitiveness of markets is still a much debated topic in applied and theoretical work. Competition policies regarding this issue differ. The Commission of the European Union considers the publication of detailed information as anti-competitive. In various decisions the Commission argued that such firm-specific in-

*Corresponding author. Fax: +49-30-2093-5787.
E-mail address: normann@wiwi.hu-berlin.de (H.-T. Normann)

0167-7187/00/$ – see front matter
PII: S0167-7187(99)00033-8

formation would create an ‘artificial transparency’ of the market, leading to less competition. While the Commission allows for the publication of aggregate industry data, the information must not be suited to identify individual actions.[1]

By contrast, US authorities do not consider such information arrangements as a violation of the Sherman Act per se. The arrangements are, however, subject to rule of reason. Relevant cases have focused on the informative versus the conspiratorial functions of individualized information. While in some decisions the publication of firm-specific data has been ruled illegal, in most of those decisions there was additional evidence for collusive behavior (Kühn and Vives, 1994; Scherer and Ross, 1990).

Until only recently, Danish competition policy went even further. The Danish Competition Council explicitly advocated the publication of firm-specific data. In some cases, the Council decided to publish transaction prices of individual firms of certain industries in order to promote market transparency.[2] Though the Danish Competition Council changed the design of its policy recently — presumably to adjust it to the policy of the European Commission — it seems safe to conclude that there is no generally accepted view in competition policy about the publication of firm-specific data.

The theoretical predictions are also mixed. Depending on the model, the publication of individual firm data may either have no effect, increase competition or decrease competition. The first prediction follows from the Nash equilibrium of a static or finitely repeated game of complete information. In such a setting, the publication of individual data does not have an effect since the equilibrium in the repeated game consists simply of repeating the static Nash equilibrium every period.

Since Stigler’s (1964) seminal paper economists have recognized that the publication of private firm-specific data may have an adverse effect on the competitiveness of a market. Stigler argued that the greatest obstacle to collusion is the possibility of secret price cuts. When competitors cannot observe price cuts by rivals, cartels will inevitably break down. On the other hand, if firms are well informed about the actions of their rivals, they can cartelize a market since deviations are immediately detected and punished.

Recently, the contrary argument has been advanced by Vega-Redondo (1997). He shows that the publication of data at the individual firm level may increase the competitiveness of a market. Vega-Redondo works with a game-theoretic learning model based on imitation. He assumes that goods are homogeneous, that firms set quantities and that firms are informed about their competitors’ quantity decisions

[1] See Kühn and Vives (1994) for a detailed discussion of the Commission’s policy.

[2] Albæk et al. (1997) show that this had an adverse effect on the competitiveness of the concrete industry. After the publication of firm-specific transaction prices in October 1993, the industry price level in this industry increased by 15–20%.

and profits in every period. The novel approach of Vega-Redondo's paper is to leave the framework of standard game theory with its best reply logic and instead study the consequences of firms imitating successful behavior. Interestingly, if firms imitate the most successful strategy of the previous period and if there is a small probability that firms mistakenly choose an arbitrary quantity, the market evolves to the competitive outcome.

The intuition for the result is easy to get. If the price is above marginal cost, the firm with the largest quantity has the highest profit, and its quantity is imitated by the other firms. This explains why total quantities increase and why the price is pushed downwards, even below the Cournot–Nash level. On the other hand, if the price is below marginal cost, the firm with the lowest quantity has the highest profit. Vega-Redondo shows that the only long run stable state is where firms produce the competitive outcome, that is, where price equals marginal cost. In the theory section of this paper we extend Vega-Redondo's result to differentiated price and quantity competition.

Given the different theoretical predications as well as the different approaches in competition policy, an empirical investigation should prove to be useful. In this paper we study the problem experimentally. An advantage of experiments is that under laboratory conditions the information available to subjects can be controlled for, which allows to isolate the effects of publication of firm-specific data. Since our experimental markets are finitely repeated games with complete information about the parameters of the model, we cannot attempt to test Stigler's model (or related models like, e.g. Green and Porter (1984) directly.[3] Rather, we test its generalized message, namely that more information about opponents' actions yields less competitive outcomes. However, our design does allow to test the role of imitation, both, on the aggregate and on the individual level.

More specifically, we studied symmetric four firm oligopolies with product differentiation. There were four treatments differing according to the information provided to subjects (aggregate or individual data) and according to the strategic variable (quantity or price competition). In the BASIC treatments subjects received all necessary information about the market structure but only *aggregate* information about their opponent's actions. In the EXTRA treatments subjects additionally received information about *individual* actions and profits of their rivals. The underlying parameters (demand, cost, product differentiation) were the same in across all treatments. Thus, our design also allows us to explicitly test the question whether Cournot or Bertrand is more competitive.[4]

Our findings clearly reject the hypothesis that more information about com-

[3] In particular, we did not implement any uncertainty about demand or cost conditions. For this reason, the literature on information sharing (see, e.g. Vives, 1984) applies only indirectly to our paper. For experimental evidence on information sharing, see Cason and Mason (1998).

[4] Theory predicts that markets are more competitive with Bertrand (see Vives, 1985).

petitors yields a tendency towards collusion. In none of our sessions the outcome was even close to collusion. On the contrary, in the Cournot case we find that more information makes markets significantly *more* competitive, supporting the imitation hypothesis. In the Bertrand case more information also seems to induce slightly more competitive outcomes, though the differences are not significant. In the Basic treatments we find that average actions nearly perfectly match the Nash equilibrium. As a consequence, it turns out that Bertrand markets are indeed more competitive than Cournot markets.

2. Experimental design

In a series of computerized[5] experiments we studied multi-period oligopoly markets with four symmetric firms. We compared two strategic settings, one in which actions were strategic substitutes (a Cournot market) and one in which actions were strategic complements (a Bertrand market). For both strategic settings we have conducted two treatments which varied only with regard to the information given to subjects.

In treatments labeled 'BASIC' subjects were informed about the demand and cost conditions. After each period, they were informed about the average action of the other firms. The information was provided verbally and in the form of a 'profit calculator'. The profit calculator served two functions. When fed with data regarding the other firms (average quantities of the other firms in the Cournot case, average price in the Bertrand case), the calculator allowed to try out the consequences of own actions. Furthermore, it allowed to calculate a best reply against the (hypothetical) actions of the other firms.[6]

Treatments labeled 'EXTRA' provided the same information as BASIC plus additional information about the other firms' individual actions of the previous period and the profits they earned.[7] This 2×2 design is summarized in Table 1.

In all cases we used the following standard demand function for differentiated goods (see e.g. Martin, 1993). For each firm i inverse demand was given by:

$$p_i = \max\left\{a - q_i - \theta \sum_{j \neq i} q_j, 0\right\}, \tag{1}$$

where $\theta \in [0,1)$ denotes the degree of product differentiation. The limiting cases

[5] We thank Abbink and Sadrieh (1995) for letting us use their software toolbox 'RatImage'.

[6] The profit calculator provides essentially the same information as commonly used payoff tables but helps to avoid a possible bias due to limited computational abilities of subjects.

[7] We provided information about profits to make imitation possible. Of course, in a symmetric market, profits could be inferred from actions.

Table 1
Design

Information	Strategic substitutes	Complements
$\{a_i, \Pi_i, \bar{a}_{-i}\}$[a]	Basic Cournot	Basic Bertrand
$\{(a_j, \Pi_j)_{j=1,\ldots,4}, \bar{a}_{-i}\}$	Extra Cournot	Extra Bertrand

[a] Note: a_i denotes the action (quantities in Cournot, prices in Bertrand), $\bar{a}_{-i}$, denote the average action of other firms, and Π_i denotes the profit.

are a homogeneous market for $\theta = 1$, and fully independent markets for $\theta = 0$. The inverse demand functions (1) was used in the Cournot treatments.

The direct demand functions for the Bertrand treatments can be obtained as follows. Solving the system of Eq. (1) for q_i, $(i = 1, \ldots 4)$ and provided that $q_i \geqslant 0$ one obtains:

$$\begin{aligned} q_i &= \frac{1}{1+(n-1)\theta}\left(a - p_i + \frac{n\theta}{1-\theta}(\bar{p} - p_i)\right) \\ &= \frac{1}{1+(n-1)\theta}\left(a - \frac{1+(n-2)\theta}{1-\theta}p_i + \frac{(n-1)\theta}{1-\theta}\bar{p}_{-i}\right), \end{aligned} \tag{2}$$

where n is the number of *active* firms (i.e. firms with $q_i \geqslant 0$), $\bar{p} := 1/n \sum_{\{i:q_i \geqslant 0\}} p_i$ their average price and

$$\bar{p}_{-i} := \frac{1}{n-1} \sum_{\{j \neq i : q_j \geqslant 0\}} p_j$$

the average price of active firms other than i. We had to use this slightly complicated formulation because firms whose prices are so high that a negative quantity would result according to (2) do not enter into the calculation of the average prices $\bar{p}$ or $\bar{p}_{-i}$.

If all firms are active, we can simplify notation by setting $\tilde{a} = a/(1+3\theta)$, $\alpha := \frac{1+2\theta}{(1-\theta)(1+3\theta)}$ and $\beta := \frac{3\theta}{(1-\theta)(1+3\theta)}$. Thus (2) becomes

$$q_i = \max\{\tilde{a} - \alpha p_i + \beta \bar{p}_{-i}, 0\}. \tag{3}$$

Note that $\alpha > \beta$ since $\theta < 1$.

The cost function for each seller was $C(q_i) = cq_i$. Hence, profits were

$$\pi_i = (p_i - c)q_i. \tag{4}$$

In the experiment we used the following parameters

$$a = 300$$
$$c = 2$$
$$\theta = 2/3$$

which implied for the case of four active firms that

$\alpha = 7/3$
$\beta = 2.$

The number of periods was 40 in all sessions and this was commonly known. Actions could be chosen from finite but sufficiently fine grids such that continuous action spaces were approximated.[8] The 96 subjects for this experiment were recruited via posters at Humboldt University, Berlin. Only about half of the subjects were economics or business students. In each session eight or 12 subjects participated. Subjects were randomly allocated to computer terminals in the lab such that they could not infer with whom they would interact in a group of four. For each of our four treatments we had six groups of subjects.

Subjects were paid according to their total profits. Profits as in (4) were denominated in 'Taler'. The exchange rates for German Marks (8000:1 in the Cournot treatments and 4000:1 in the Bertrand treatments) were known. The average payoff was DM 29 across all treatments.

Sessions lasted about 90 min including instruction time. Instructions (see Appendix C) were written on paper and distributed in the beginning of each session. After the instructions were read, we conducted one trial round in which the different windows of the computer screen were introduced and could be tested. When subjects were familiar with both, the rules and the handling of the computer program, we started the first round.

3. Theoretical predictions and hypotheses

In this section we derive four benchmark predictions about the outcomes of our oligopoly games. The standard prediction is, of course, the unique Nash equilibrium of the stage game. In the Cournot oligopoly all firms choose

$$q_i^{\mathrm{C}} = \frac{a-c}{2+3\theta} = 74.5, \quad i \in I,$$

with a corresponding equilibrium price of $p_i^{\mathrm{C}} = 76.5$. In the Bertrand case firms choose a price of[9]

$$p_i^{\mathrm{B}} = \frac{100+\alpha c}{2\alpha - \beta} = 39.25, \quad i \in I,$$

[8] On the Bertrand market prices had to be chosen from a finite grid between 0 and 1000. On the Cournot market quantities between 0 and 400 could be chosen. In both cases 0.01 was the smallest step size.

[9] Note that given the demand function (2) and constant marginal cost of 2, it is never optimal for a firm to choose a price such that $q_i = 0$. Thus, when firms play best replies, they will always be active.

and a corresponding quantity of $q_i^B = 86.92$. Note that these quantities and prices also correspond to the outcome of the unique subgame perfect equilibrium of the repeated game since we have a finite horizon. As is well known (see, e.g. Vives, 1985) the Bertrand setting is more competitive.

Note the informational requirements to play a Nash equilibrium. One needs to know the demand and cost functions and one needs to know how to calculate a best reply. Since we have endowed all subjects with the profit calculator, those requirements are met in all treatments.

If subjects have additional information about individual actions and profits (as in the EXTRA treatments), an alternative outcome becomes possible. Suppose subjects imitate in each period the strategy of the firm which was most successful last period. Furthermore, with some small probability $\varepsilon > 0$ individuals make a mistake and choose some random strategy. This process was analyzed by Vega-Redondo (1997) for a homogenous Cournot market. Here, we prove an equivalent result for differentiated products and for both, quantity and price competition. Let

$$q^I = \frac{(a-c)(n-1)}{2(n-1) + 3\theta(n-2)} = 89.4$$

and the corresponding price $p^I = 300 - 89.4 - \frac{2}{3}\,3(89.4) = 31.8$.

Proposition 1. *Suppose individuals follow an 'imitate-the-best' process with noise ε. Then in both, the Cournot and the Bertrand setting, (q^I, p^I) is the unique long run stable outcome as $\varepsilon \to 0$.*

Proof. See Appendix A. □

There is an alternative interpretation of the imitation outcome.

Remark 1. *If all firms maximize relative profits, then the imitation outcome (q^I, p^I) results as an equilibrium.*

Proof. See Appendix A. □

Finally, another benchmark is the (symmetric) collusive outcome. Again, the outcome is the same in the Cournot and the Bertrand settings and can easily be calculated as

$$q^K = \frac{a-c}{2 + 2\theta(n-1)} = 49.67$$

and a corresponding price of $p^K = 151$. Theoretically, this collusive outcome cannot be an equilibrium of a game with a finite horizon. Behaviorally, the time horizon of 40 periods may be long enough to make repeated game effects possible.

Table 2
Theoretical benchmarks

Outcomes	Indiv. quantities	Prices	Profits
Collusion	49.7	151.0	7400.3
Cournot–Nash	74.5	76.5	5550.3
Bertrand–Nash	86.9	39.3	3237.6
Imitation	89.4	31.8	2664.1

Selten and Stoecker (1986) show that, apart from an end game effect, there was no behavioral difference in their experiment between a treatment with a long finite horizon and one which approximates an infinite horizon by the device of a stopping probability.[10]

Thus we have four benchmark results for our experiment, the collusive outcome, the Cournot–Nash equilibrium, the Bertrand–Nash equilibrium, and the outcome resulting from the imitation process. Table 2 compares the benchmarks for the parameters used in the experiment.

In the BASIC treatments the information available corresponds exactly to what is necessary to play a Nash equilibrium. Imitation is not feasible for subjects in the BASIC treatments. Thus, we have

Conjecture 1. In the BASIC COURNOT treatment behavior converges to the Cournot–Nash outcome. In the BASIC BERTRAND treatment behavior converges to the Bertrand–Nash outcome. In particular, the BASIC BERTRAND treatment is more competitive than BASIC COURNOT.

In the EXTRA treatments subjects had additional information about individual actions of their competitors. Here, the theoretical predictions are contradictory. On the one hand, we have

Conjecture 2A. In the EXTRA treatments behavior will be less competitive than in the BASIC treatments.

On the other hand, we have Proposition 1 based on imitation which would suggest

Conjecture 2B. In the EXTRA treatments behavior will be more competitive than in the BASIC treatments due to imitation.

[10] Selten et al. (1997) point out that such a random stopping device is problematic since an infinite horizon cannot credibly be implemented in the lab.

Table 3
Average quantities and prices

	Rounds	Average quantities		Average prices	
		BASIC	EXTRA	BASIC	EXTRA
COURNOT	Last 20	74.66	83.42	76.30	52.17
		(1.58)[a]	(3.80)	(4.72)	(9.79)
	Last 5	74.20	86.12	77.39	44.05
		(0.70)	(5.32)	(2.11)	(11.41)
BERTRAND					
	Last 20	85.00	86.21	45.23	41.33
		(2.10)	(1.48)	(6.44)	(4.64)
	Last 5	86.03	86.84	41.91	39.49
		(1.03)	(1.66)	(3.08)	(4.98)

[a] Note: Standard deviations in parentheses.

4. Experimental results

Table 3 presents means and standard deviations of quantities and prices, respectively, over all groups and the last 20 and the last five rounds. We chose to report the later rounds to give any possible learning effects enough time to phase out.[11] Since each group counts as a single observation, all means and standard deviations are based on samples of six observations. These group averages are shown in Tables 5 and 6 in Appendix B.

With respect to Conjecture 1 we find that the experimental means are remarkably close to the theoretical predictions. In the BASIC COURNOT treatment the predicted average quantity was 74.5; the average quantities given in Table 3 are 74.66 and 74.20 for the last 20 and the last five rounds, respectively. In the BASIC BERTRAND treatment the predicted average price was 39.25, and the average prices found in the experiment were 45.23 and 41.91, respectively. In all cases the means are within one standard deviation of the theoretical prediction.

Since the Nash prediction holds in both BASIC treatments, the second part of Conjecture 1 is also confirmed. Namely, Bertrand is indeed more competitive than Cournot competition. We tested this by one-tailed Mann–Whitney U-tests and found that average quantities are significantly higher in BASIC BERTRAND than in EXTRA BERTRAND at a p value of 0.0019 for both, the last 20 and the last

[11] However, there was no noticeable time trend in the data. Averages in the first 20 were not significantly different from averages in the last 20 rounds (according to Wilcoxon tests with 5% significance levels). Furthermore, there were no significant endgame effects at the 5% level.

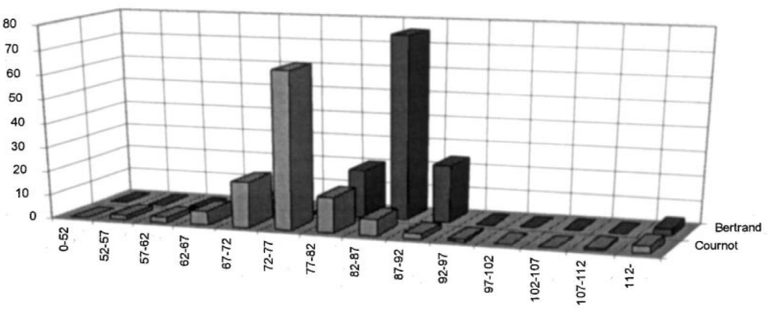

Fig. 1. Frequencies of average quantities in BASIC COURNOT and BERTRAND, last 20 rounds. The Nash quantities are $q_i^C = 74.5$ and $q_i^B = 86.9$.

five rounds. This is also nicely illustrated by Fig. 1. The entire distribution of average quantities is shifted upwards for Bertrand as compared to Cournot.

Thus we have

Experimental Result 1. If subjects have only aggregate information about their rivals' actions, behavior converges to the Nash outcome. As predicted by theory, the Bertrand setting is more competitive than Cournot.

Conjectures 2A and B are examined by testing the appropriate null hypotheses with the one-tailed Mann–Whitney *U*-test. As is immediate from the inspection of Table 3, we do not find any evidence for Conjecture 2A — regardless of whether actions are strategic substitutes or complements.

Experimental Result 2A. Neither in the Bertrand nor in the Cournot case does additional information about rivals' actions and profits facilitate collusion.

With respect to Conjecture 2B we find that behavior is indeed significantly more competitive in Extra in the case of Cournot markets, for both, the last five and the last 20 rounds (MWU tests, each at $p = 0.0019$). The group with the highest average quantity in BASIC COURNOT has still lower quantities than the lowest group in EXTRA COURNOT (see Table 5). Average prices in the Bertrand treatments are somewhat higher in BASIC, but this effect is not significant. Thus, we can state the following result concerning the central question whether more information increases or decreases competitiveness on oligopoly markets.

Experimental Result 2B. In case of strategic substitutes additional information about rivals' actions and profits significantly increases competitiveness. In case of strategic complements additional information does not change the degree of competition significantly.

The difference between strategic substitutes and complements is somewhat puzzling. While there are many documented differences in the IO literature between those two strategic settings, we are not aware of any theoretical result that explains differences with respect to Conjecture 2B.

A closer look at the data might reveal what is going on. Fig. 2 presents histograms of the frequencies of actions in the last 20 rounds on the level of groups. For the Cournot markets the histograms show clearly how behavior shifts towards more competition resulting in the observed increase in mean quantities.

For the Bertrand markets a different picture emerges. As shown above mean prices are not significantly different in BASIC and EXTRA. While in BASIC a clear peak emerges around the Nash price, the histogram reveals that actions are much more dispersed in EXTRA. In fact, it can be seen that although the Nash prediction is very close to the *average* price, specific actions coincide only rarely with Nash. It seems that some groups play more competitive than Nash resulting in the peak at the 35–38 bracket. But there also seem to exist other groups who play less competitive.[12] What accounts for this diverse behavior is an interesting question which should be explored in further work.

One possible behavioral explanation could be the following. In Bertrand markets imitation typically prescribes lowering the current price. Subjects might be reluctant to do so given that it must be immediate to them that this implies lower profits per unit. However, this is different in Cournot markets where imitation typically prescribes increasing ones quantity. Here, the negative effect on prices and profits is less obvious making imitation look less harmful.

It is also instructive to look at individual data, which allows to test explicitly the behavioral hypothesis underlying imitate-the-best behavior.[13] To analyze the individual data we calculated for each round t and for each player i the actions prescribed by imitate-the-best behavior. Then we checked how close each player came to meet this target, T_i^t. The closeness is measured by the following 'hit rates'

$$h_i^t := \left| \frac{T_i^t - a_i^t}{T_i^t} \right| .$$

As a comparison we calculated the same measure for myopic best-reply behavior. Given subjects could observe average quantities or prices, which allowed them to play a best reply against the opponents' action from last period, this behavioral hypothesis seems particularly interesting.

[12] We do not provide a formal test for the difference between those frequency distributions as the observations are not independent.

[13] In principle, we could also test whether deviations from collusion are punished as predicted by the repeated game literature. However, this was impossible as collusive outcomes were never reached in any group.

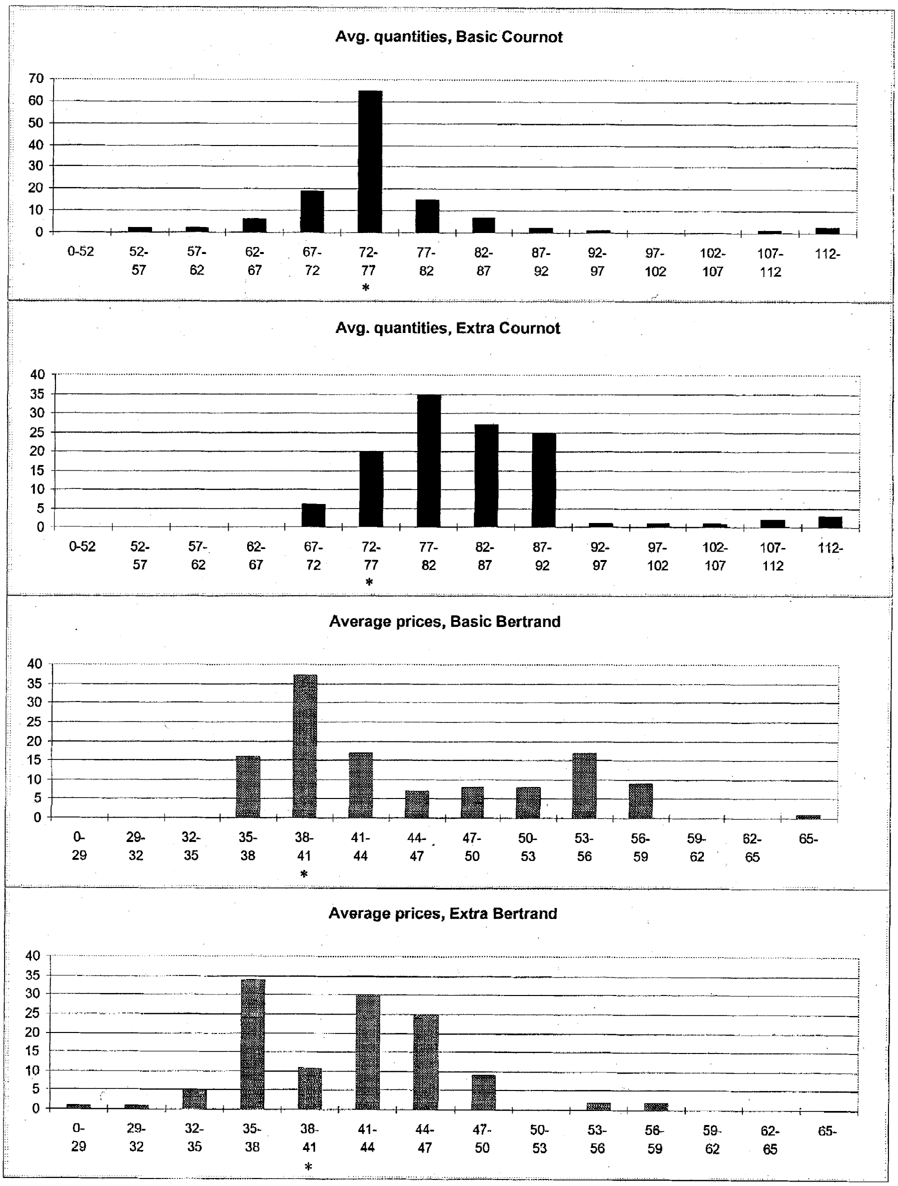

Fig. 2. Frequencies of actions in the last 20 rounds (* Indicates the bracket which includes the Nash outcome).

Table 4
Relative frequency (in %) of hits

	Target	< 0.01	< 0.05	< 0.1	Avg. *h* (S.D.)
EXTRA					
COURNOT	Best reply	3.3	19.9	32.2	0.25 (0.31)
	Imitation	11.2	30.8	50.3	0.14 (0.19)
BERTRAND	Best reply	9.1	39.9	64.2	0.11 (0.18)
	Imitation	7.2	38.2	60.4	0.13 (0.20)
BASIC					
COURNOT	Best reply	10.4	40.0	59.1	0.14 (0.21)
BERTRAND	Best reply	11.3	34.7	56.0	0.13 (0.17)

Table 4 reports the frequency distribution of hit rates in the different treatments. The most striking fact is that the best reply hypothesis performs clearly worse than the imitation hypothesis in EXTRA COURNOT; more than 50% of all decisions were within a 10%-neighborhood of the imitation target. Clearly, this can account for the more competitive behavior in the COURNOT treatment. On the other hand, in EXTRA BERTRAND the best reply hypothesis performs even slightly better than imitation, which could account for our result that prices in EXTRA BERTRAND are relatively close to the Nash outcome. In the BASIC treatments, in which imitation was impossible, the best reply hypothesis performs reasonably well.

Finally, we compare our results with those of other studies. With respect to the Cournot setting there is ample evidence that supports our finding that more information about other players' actions yields more competitive behavior. In the early study by Fouraker and Siegel (1963) there are two treatments on homogeneous Cournot triopoly with incomplete and complete information, which roughly resemble our BASIC and EXTRA treatments, respectively. Fouraker and Siegel found that total quantities were close to Nash with incomplete information but more competitive with complete information, which corresponds nicely to our results.

In our companion paper (Huck et al., 1999), we analyzed learning in homogeneous Cournot oligopolies with four firms. In the treatments comparable to the BASIC and EXTRA treatments in this paper, we found the same result: More information about the competitors yields significantly more competitive markets.

Offerman et al. (1997) as well as Bosch-Domènech and Vriend (1998) have conducted similar experiments. Offerman et al. (1997) analyze triopoly markets. Two of their three treatments are similar to those of Huck et al. (1999). In their experiment, when information about rivals' quantities and profits is provided, average quantities are higher (if however not significantly) than in the treatment where only the rivals' aggregate quantity is available. Offerman et al. (1997), however, observe this effect on a generally lower level of competition as total quantities are always at or below the Cournot level, which does not match our experimental results. They also analyze a treatment in which the firm-specific quantities are given as information, but not the profits. In that treatment, competition was slightly less severe compared to a treatment where also profits were published.[14]

Bosch-Domènech and Vriend (1998) study homogeneous Cournot duopolies and triopolies. While in their experiments the objective amount of available information is always the same (and corresponds to our EXTRA treatment), they vary the presentation of the information and impose different time limits on decisions. Their evidence is slightly mixed. However, there is a noticeable trend towards more competition when the emphasis of the presentation is more on rivals' actions and profits and less on the market structure.

Taken together, these papers on Cournot competition provide robust evidence that collusion is difficult to achieve with more than two firms. While in duopoly experiments (see, e.g. Holt, 1995) collusion occurs frequently, already with three firms cases of successful collusion are rare. With four firms, there is virtually no collusion.

To our knowledge there are only a few experiments about Bertrand markets with product differentiation. Dolbear et al. (1968) investigate collusive behavior in an experimental market similar to ours. Their complete information treatment is similar to our BASIC treatment. But in their incomplete information treatment, firms were not informed about the impact their competitors' price had on their profit. Average prices did not differ in a significant way between treatments. Harstad et al. (1998) also study Bertrand markets with product differentiation but they analyze the effect of non-binding preannoucements of prices.

5. Conclusion

In a series of experiments we have investigated the influence of information about rivals' actions and profits on the competitiveness of oligopolistic markets. In neither strategic setting we find evidence for the hypothesis that additional information facilitates collusive behavior. In the case of strategic substitutes (quantity competition) it renders market outcomes significantly more competitive. Also in the case of strategic complements (price competition) it makes market outcomes more competitive, if however not significantly so. While we do not want to overemphasize the policy implications of our results, they could tentatively be taken as evidence in favor of the publication of individual firm data to foster competition.

[14] See, however, Ruffle (1997) who obtained the opposite result in an experimental posted-offer market.

Acknowledgements

We are grateful to Tim Cason, Eric van Damme, Doug Davis, and participants of the 1998 EARIE conference in Copenhagen and of the 1998 ESA meeting in Mannheim for helpful comments. Two anonymous referees and Stephen Martin helped to improve the paper considerably. The Deutsche Forschungsgemeinschaft, SFB 373, provided financial support.

Appendix A. Proofs

Proof of Proposition 1. The imitation process is analyzed in a finite state space. Thus, in line with the experimental setup we require that quantities or prices must be chosen from a finite grid $\Gamma := \{0, \delta, 2\delta, \ldots, \upsilon\delta\}$, for arbitrary $\delta > 0$, and some $\upsilon \in \mathbb{N}$ large enough. Whenever a firm revises its strategy, it chooses one of those strategies which received the highest payoff last period according to some probability distribution with full support. Furthermore, every period each firm 'mutates' (makes a mistake) with independent probability $\varepsilon > 0$ and chooses an arbitrary action from Γ (all actions are chosen with some strictly positive probability).

Consider first the Cournot case and let $\Delta(q, q')$ be the profit differential between a firm choosing q when all other firms choose q' and a firm choosing q' against $n-2$ firms with q' and one firm with q:

$$\Delta(q,q') := \Pi(q,q') - \Pi\left(q', \frac{n-2}{n-1}\, q' + \frac{1}{n-1}\, q\right)$$

Note that $\max_q \Delta(q,q') \geq 0$ as one can always set $q = q'$. Simple calculations show that

$$q^* = \frac{1}{2(n-1)}\,(a(n-1) - 3(n-2)\,\theta q - c(n-1))$$

is the unique q maximizing $\Delta(q,q')$. Next, define

$$q^{\mathrm{I}} := \min_{q'} \Delta(q^*,q')$$

q^{I} is unique since $\Delta(q^*,q')$ is quadratic in q'. Since $\Delta(q^*,q') \geq 0$, q^{I} is given by setting $q' = q^*$. Thus

$$q^{\mathrm{I}} = \frac{(a-c)(n-1)}{2(n-1) + 3(n-2)\theta}$$

is the unique quantity which when used by every firm is stable against invasion of other quantities. In the following we assume that $q^{\mathrm{I}} \in \Gamma$.

Consider first the Markov process for $\varepsilon = 0$. This process reaches an absorbing state if and only if quantities of all firms are equal. Let Θ denote the set of absorbing states. By standard arguments (see, e.g. Theorem 1 of Samuelson, 1994) only states in Θ can appear in the support of the limit distribution of the Markov process for $\varepsilon \to 0$.

To prove the result is suffices to show that it takes only one mutation to reach $\mathrm{q}^{\mathrm{I}} = (q_1^{\mathrm{I}}, q_2^{\mathrm{I}}, q_3^{\mathrm{I}}, q_4^{\mathrm{I}})$ from any state $\mathrm{q} \in \Theta, \mathrm{q} \neq \mathrm{q}^{\mathrm{I}}$, whereas it takes more than one mutation in the opposite direction.

We will show first that

$$\Delta(q^{\mathrm{I}},q) = \Pi(q^{\mathrm{I}},q) - \Pi\left(q, \frac{n-2}{n-1}\, q + \frac{1}{n-1}\, q^{\mathrm{I}}\right) > 0, \quad \forall q \neq q^{\mathrm{I}}$$

Substituting for q^{I} yields after some calculations

$$[n - 1 + 3\theta(n-2)][(2q - a + c)(n-1) + 3\theta q(n-2)]^2 > 0$$

which is satisfied for all $q \neq q^{\mathrm{I}}$. Thus, one firm mutating from q to q^{I} will suffice to put the process in the basin of attraction of q^{I}. On the other hand one firm's mutation will not suffice to leave q^{I}'s basin of attraction as $\Delta(q,q^{\mathrm{I}}) < 0$ for all $q \neq q^{\mathrm{I}}$ by definition of q^{I}.

The proof for the Bertrand case works equivalently. This can immediately be seen by the following argument. The vector of quantities q^{I} (and the corresponding vector of prices p^{I}) has the property that any deviation by some firm i will yield a negative profit differential for firm i as compared to the profit of firms $j \neq i$. Clearly, it does not matter whether the deviation is in prices or in quantities. Thus, the same stable outcome p^{I}, q^{I} results in the Bertrand setting. □

Proof of Remark 1. Let a_i denote own actions (prices or quantities) and a_{-i} denote the actions of all other firms. Maximization of relative profits, i.e.:

$$\max_{a_i} \frac{\Pi_i(a_i, a_{-i})}{\sum_j \Pi_j(a_j, a_{-j})}$$

yields the first-order condition

$$\frac{\partial \Pi_i}{\partial a_i} \sum_j \Pi_j - \Pi_i\left(\frac{\partial \Pi_i}{\partial a_i} + \sum_{j \neq i} \frac{\partial \Pi_j}{\partial a_i}\right) = 0$$

Because of symmetry we get $(n-1)\,\Pi \frac{\partial \Pi_i}{\partial a_i} = \Pi(n-1)\, \frac{\partial \Pi_j}{\partial a_i}$ and hence:

$$\frac{\partial \Pi_i}{\partial a_i} = \frac{\partial \Pi_j}{\partial a_i}$$

It is easy to check that this condition only holds at the imitation outcome. □

Appendix B. Group data

Table 5
Group average quantities over last 20 and last five rounds in the Cournot games

COURNOT	Rounds	Gr. 1[a]	Gr. 2	Gr. 3	Gr. 4	Gr. 5	Gr. 6
BASIC	Last 20	47.99	76.40	74.84	74.95	71.66	75.12
	Last 5	73.10	74.80	74.65	73.56	74.40	74.70
EXTRA	Last 20	80.60	87.90	81.11	86.50	85.85	78.58
	Last 5	84.08	89.51	82.25	84.00	95.35	81.50

[a] Note that the group number refers to different groups of subjects between treatments.

Table 6
Group average prices over last 20 and last five rounds in the Bertrand games

BERTRAND	Rounds	Gr. 1[a]	Gr. 2	Gr. 3	Gr. 4	Gr. 5	Gr. 6
BASIC	Last 20	38.96	40.85	51.63	49.34	52.05	38.55
	Last 5	39.93	41.83	43.68	40.33	47.09	38.60
EXTRA	Last 20	44.61	40.40	35.32	36.84	43.83	47.06
	Last 5	44.21	34.69	34.58	36.15	45.72	41.60

[a] Note that the group number refers to different groups of subjects between treatments.

Appendix C. Translation of instructions

[These are the instructions for the Cournot games. The instructions for the Bertrand Games were the same except as mentioned below]

Welcome to our experiment. Please read these instructions carefully. In the next 1 or 2 hours you will have to make some decisions at the computer. You can earn some real money. But please be quite during the entire experiment and do not talk to your neighbors. Those who do not follow this rule will have to leave and will not get paid. If you have a question please raise your arm.

You will receive your payment discreetly at the end of the experiment. We guarantee anonymity with respect to other participants and we do not record any information connecting your name with your performance.

You can operate the computer with the keyboard or the mouse. Before the experiment there is enough time to make yourself familiar with the computer in a trial round. Money in the experiment is denominated in ‘Taler’. At the end we exchange your earnings into DM at a rate of 8000 T = 1 DM. [Bertrand: 4000 T = 1 DM] The experiment is divided into several rounds. As said, we start with a trial round. The real experiment starts with round 1.

You represent a firm which produces and sells a certain product. Besides you there are 3 other firms which produce and sell the same product. Production cost are 2T per unit (this is the same for all firms). Your task is to decide how much to sell of your good. [Bertrand: Your task is to decide at which price to sell of your good.]

For this the following important rule holds: The *higher your quantity*, the *lower* is the price you receive for them. Above a certain quantity the price becomes zero. Furthermore, the *higher* the *average price* of the other firms, the *lower* is your price.

[Bertrand: For this the following important rule holds: The *higher your price*, the *fewer* units you sell. Above a certain price you don't sell anything anymore. On the other hand you sell the *more*, the *higher the average price* of the other *active* firms is (a firm is *active* when it sells a positive quantity; it is inactive when its price is so high that it doesn't sell anything. The price of an inactive firm does not enter into the calculation of the average prices).]

Profit per unit is the difference between your price and the cost per unit of 2 T. Note that you make a *loss* if the price is lower than the per unit cost. Your profit in a given round results from multiplying the profit per unit with your supplied quantity.

In each round the quantities [Bertrand: prices] of all firms are recorded and the resulting prices [Bertrand: quantities] and profits are calculated. In each round you will be told your profit. Profits from all periods are added and the sum is paid out to you in cash at the end.

From the second round on you[15] will receive the following information. You are told each firm's last period quantity [Bertrand: price], the average quantity [Bertrand: price] of the other firms [Bertrand: other active firms] last period, and each firm's price [Bertrand: quantity] and profit.

From the second round on you[16] will receive the following information. You are told the average quantity of the other firms last period [Bertrand: the average price of the active firms last period, the number of active firms], and your own profit.

Additionally, you have access to a profit calculator. The profit calculator is shown on the last page of the instructions. It has two functions: (1) It calculates your profit for arbitrary quantity [Bertrand: price] combinations. That is, you can enter two values, an average quantity [Bertrand: price] for the others (button 'A') and a quantity [Bertrand: price] for yourself (button 'I'), and the machine tells you how much you would earn. (2) You can let it calculate for arbitrary average quantities [Bertrand: prices] of others (button 'A') the quantity at which you would make the highest profit (button 'M'). You can use the machine as much as you want before each decision. Before we start you will have enough time to get to know the profit calculator directly at the computer.

Everything we have explained to you holds for the other firms as well. In fact, you are all reading exactly identical instructions.

---

[15] This paragraph only for EXTRA.

[16] This paragraph only for BASIC.

The experiment lasts for 40 periods in total. Afterwards you will receive your payments in DM. We want to reassure you again that all data will be treated confidentially.

References

Abbink, K., Sadrieh, A., 1995. RatImage: Research Assistance Toolbox for Computer-Aided Human Behaviour Experiments, University of Bonn, SFB 303 Discussion Paper B-325.

Albæk, S., Møllgaard, P., Overgaard, P., 1997. Government-assisted oligopoly coordination? A concrete case. Journal of Industrial Economics 45, 429–443.

Bosch-Domènech, A., Vriend, N.J., 1998. Do Boundedly Rational People Imitate, Universitat Pompeu Fabra, Barcelona, Mimeo.

Cason, T.N., Mason, C.F., 1998. Uncertainty, Information Sharing and Tacit Collusion in Laboratory Duopoly Markets. Forthcoming Economic Inquiry.

Dolbear, F.T., Lave, L.B., Bowman, G., Lieberman, A., Prescott, E., Rueter, F., Sherman, R., 1968. Collusion in oligopoly: an experiment on the effect of numbers and information. Quarterly Journal of Economics 82, 240–259.

Fouraker, L., Siegel, S., 1963. Bargaining Behavior, McGraw-Hill, New York.

Green, E., Porter, R., 1984. Non-cooperative collusion under imperfect price information. Econometrica 52, 87–100.

Harstad, R., Martin, S., Normann, H.T., 1998. Experimental tests of consciously parallel behaviour in oligopoly. In: Phlips, L. (Ed.), Applied Industrial Organization, Cambridge University Press, Cambridge.

Holt, C., 1995. Industrial organization: a survey of laboratory research. In: Kagel, J., Roth, A. (Eds.), The Handbook of Experimental Economics, Princeton University Press, Princeton.

Huck, S., Normann, H.T., Oechssler, J., 1999. Learning in Cournot oligopoly – an experiment. Economic Journal 109, C80–C95.

Kühn, K.U., Vives, X., 1994. Information Exchanges Among Firms and their Impact on Competition, European Commission, Report to DG IV.

Martin, S., 1993. Advanced Industrial Economics, Blackwell, Cambridge, MA.

Offerman, T., Potters, J., Sonnemans, J., 1997. Imitation and Adaptation in an Oligopoly Experiment, Amsterdam University, Mimeo.

Ruffle, B.J., 1997. Do Buyers Always Act as Passive Price Takers? Demand Withholding in Experimental Posted-Offer Markets, Ben-Gurion University, Mimeo.

Samuelson, L., 1994. Stochastic stability in games with alternative best replies. Journal of Economic Theory 64, 35–65.

Scherer, F.M., Ross, D., 1990. Industrial Market Structure and Economic Performance, 3rd ed, Houghton Mifflin, Boston.

Selten, R., Mitzkewitz, M., Uhlich, G.R., 1997. Duopoly strategies programmed by experienced players. Econometrica 65, 517–556.

Selten, R., Stoecker, R., 1986. End behavior in finite prisoner's dilemma supergames. Journal of Economic Behavior and Organization 7, 47–70.

Stigler, G.J., 1964. A theory of oligopoly. Journal of Political Economy 72, 44–61.

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.

Vives, X., 1985. On the efficiency of Cournot and Bertrand equilibria with product differentiation. Journal of Economic Theory 36, 166–175.