



UvA-DARE (Digital Academic Repository)

Imitation and belief learning in an oligopoly experiment.

Offerman, T.; Sonnemans, J.; Potters, J.

Publication date
1997

[Link to publication](#)

Citation for published version (APA):

Offerman, T., Sonnemans, J., & Potters, J. (1997). *Imitation and belief learning in an oligopoly experiment*. (Tinbergen Institute discussion paper; No. 97-116/1). Tinbergen Institute.

General rights

It is not permitted to download or to forward/distribute the text or part of it without the consent of the author(s) and/or copyright holder(s), other than for strictly personal, individual use, unless the work is under an open content license (like Creative Commons).

Disclaimer/Complaints regulations

If you believe that digital publication of certain material infringes any of your rights or (privacy) interests, please let the Library know, stating your reasons. In case of a legitimate complaint, the Library will make the material inaccessible and/or remove it from the website. Please Ask the Library: <https://uba.uva.nl/en/contact>, or a letter to: Library of the University of Amsterdam, Secretariat, Singel 425, 1012 WP Amsterdam, The Netherlands. You will be contacted as soon as possible.

Imitation and Belief Learning in an Oligopoly Experiment

Theo Offerman, Jan Potters and Joep Sonnemans*

Current version: October 1997

Abstract

We examine the force of three types of behavioral dynamics in quantity-setting triopoly experiments (1) mimicking the successful firm, (2) following the exemplary firm, and (3) belief learning. Theoretically, these three rules of dynamic conduct lead to the competitive, the collusive, and the Cournot Nash outcome, respectively. In the experiment we employ three information treatments. Each of these treatments is hypothesized to be conducive to the force of one of the three dynamic rules. To a large extent, the results are consistent with the hypothesized relationships between treatments, dynamic rules, and outcomes.

JEL codes: H41, C91

Key-words: quantity-setting oligopoly game, imitation, belief learning, experiments

Theo Offerman and Joep Sonnemans
University of Amsterdam
Department of Economics /CREED
Roetersstraat 11
1018 WB Amsterdam
Netherlands

Jan Potters
Tilburg University
Department of Economics
Warandelaan 2
5000 LE Tilburg
Netherlands

Tel: 31-20-5254294
Fax: 31-20-5255283
E-mail: theo@fee.uva.nl
potters@kub.nl
joeps@fee.uva.nl

* We are thankful for comments by participants of the Amsterdam Workshop on Experimental Economics and the meeting of the Economic Science Association, and by Rachel Croson, Charles Holt, Cas Hommes, and Georg Kirchsteiger in particular. The first author gratefully acknowledges financial support by the Netherlands' Organization for Scientific Research.

1. INTRODUCTION

Although quantity-setting oligopoly is one of the 'workhorse models' of industrial organization (Martin, 1991), there is much ambiguity about its outcome, both empirically and theoretically. A recent survey by Slade (1994) indicates that most empirical studies reject the hypothesis that the outcome is in line with the Cournot-Nash equilibrium of the corresponding one-shot game. Interestingly, however, outcomes on both sides of the Cournot-Nash outcome are found. In the experimental literature, a similar state of affairs obtains. Many experimental oligopoly games result in higher than Cournot-Nash production levels, some result in lower production levels (Holt, 1995).

Theoretically, three main benchmarks for the quantity setting oligopoly game exist: the Walrasian equilibrium, where each firm's profits are maximized given the market clearing price; the Cournot Nash equilibrium, where each firm's profits are maximized given the other firms' quantity choices; and the collusive outcome, where aggregate profits are maximized. To motivate these benchmarks, one may search for dynamic underpinnings of these benchmarks. In this paper we present an experimental investigation of such dynamic underpinnings based on belief learning and imitation. Both belief learning and imitation have received support in other experimental settings.^{1 2}

A starting point for our study is a recent result by Vega-Redondo (1997) on the effect of a rule based on imitation in symmetric oligopoly. He shows that if firms tend to mimic the quantity choice of the most successful firm, then there will be a tendency for the industry to evolve toward the competitive outcome. The successful firm is the firm with the highest profit. Typically, this firm produces a higher quantity than the other firms. This approach thus provides a dynamic behavioral underpinning for Walrasian equilibrium.

The result is paradoxical in some sense. By imitating the choice of the most profitable firm, the industry as a whole develops toward an outcome which is very unprofitable. Therefore, as suggested by Sinclair (1990), in the presence of a partial conflict between the individual's and the group's interest, it may be reasonable for imitators to follow the 'saint' rather than the 'villain'. In the social psychological literature, it is also common to hypothesize that cooperative or altruistic acts serve as examples for

¹ For experimental work on belief learning see Boylan and El-Gamal, 1993; Cheung and Friedman, 1997 ; El-Gamal *et al.*, 1994; Mookherjee and Sopher, 1994; Offerman *et al.*, 1997. Experimental work on imitation focuses on the role of imitation in decision tasks, see Offerman and Sonnemans, 1997; Pingle, 1995. Of course, there are other dynamic rules - such as reinforcement learning (Roth and Erev, 1995) - that have found experimental support, but here we focus on imitation and belief learning.

²The discussion about the appropriateness of the three benchmarks has long centered around the so-called conjectural variations. An advantage of this approach is that it encompasses all three benchmarks within a unified framework. A problem, however, is that conjectures are used to reflect the manner in which firms react to each others' choices, whereas the models are essentially static in nature (Daughety, 1985). As suggested by Fraser (1994), it would be worthwhile to search for dynamic underpinnings of particular conjectures.

others (Bandura, 1986; Sarason *et al.*, 1991). In the oligopoly game, the 'exemplary' firm produces such a quantity that if all firms would have produced exactly that quantity, the profit would have been higher for all firms than if they would have followed any of the other firms. The exemplary firm is often the firm that produces the lowest quantity. As will be shown in the next section, the industry will evolve toward the collusive outcome when firms tend to follow the firm that sets the good example from the perspective of industry profit.

The belief learning approach provides a completely different underpinning of behavior in the quantity-setting oligopoly game. According to this approach decision makers form beliefs about the choices of other decision makers on the basis of the assumption that the history of the game provides valuable information about the future actions of the others. Given their beliefs decision makers myopically choose the outcome with highest expected value (utility). It is well known that representatives of the class of belief learning rules, like best response learning or fictitious play, can lead to cycles of behavior in quantity-setting oligopoly games. However, if play converges, it converges to the Cournot-Nash equilibrium (see Milgrom and Roberts, 1991). In our experiment we choose the structure and parameters of the market such that both best response learning and fictitious play would converge to the Cournot-Nash equilibrium independent of initial play.

Hence, each of the three benchmarks in the quantity-setting oligopoly can be founded on a different behavioral rule. The theoretical results could then be reconciled with the wide range of empirical findings in the following way. It may be that different environments trigger different dynamic rules of behavior and that different dynamic rules lead to different benchmarks. Especially the type of feedback about the choices and performance of other firms may play an important role here. Some feedback may be conducive to mimicking the most successful firm and may direct the industry toward the Walrasian equilibrium. Other feedback may be conducive to either following the exemplary firm or belief learning and thus lead to the collusive or the Cournot-Nash outcome, respectively. Experimental research has the important advantage that the type of feedback can be precisely controlled. Thus the force of this potential explanation for the range of empirical findings can be examined. Such an examination is the main goal of our paper.

For imitation to be possible, firms (at least) need information about the quantity chosen by each other firm. Without fully individualized feedback about quantities, firms are unable to mimic the successful firm or to follow the exemplary firm. This simple observation leads to the base-line treatment of our experiment. In this treatment the information that firms receive about other firms is restricted to the sum of the quantities produced in the previous period. By excluding the possibility of imitation, we hypothesize this treatment to be relatively favorable for belief learning. We refer to this treatment as treatment Q (Q is mnemonic for aggregate information).

In the other two treatments we provide more information than in treatment Q. In principle, firms

could still adapt their choices on the basis of belief learning in these treatments, but now imitation also becomes available as potential source of adaptation. In treatment Qq , besides information about the sum of the quantities in the previous period (Q), also individualized information about the quantities produced is made available (q is mnemonic for individualized quantity information). This individualized information enables firms to find out which firm set the good example in the previous period and to follow this firm if they so wish. In principle, they could also compute the profit of each of the other firms and mimic the most successful firm. However, we think that it is relatively more likely that firms mimic the successful firm in treatment $Qq\pi$, where firms not only receive individualized information feedback on quantities but also on profits (π is mnemonic for profit). We expect that this individualized profit information will induce firms to make comparative profit appraisals.^{3 4}

The remainder of the paper is organized as follows. The next section presents the institutions of the market and formulates the models and hypotheses. Section 3 describes the experimental design and procedures. Section 4 presents the results. Section 5 provides a concluding discussion.

2. MARKET, MODELS, AND HYPOTHESES

In our market n firms producing a homogeneous commodity form an industry. Firms use similar production methods: each firm i faces the following total cost function $C(q_i)$:

$$C(q_i) = cq_i^{\frac{3}{2}},$$

$$P(Q) = a - b\sqrt{Q}; \quad Q = \sum_{j=1}^n q_j,$$

where c denotes a given cost parameter and q_i denotes the quantity produced by firm i ($1 \leq i \leq n$). The price for the commodity is determined by the inverse demand function $P(Q)$:

where a and b denote given demand parameters.

³ In a public goods experiment, Sell and Wilson (1991) find that the provision of individualized information on contribution levels allows subjects to achieve higher levels of cooperation than in case feedback is restricted to aggregate contributions. This result is in line with our hypothesis about the effect of individualized information on the production level in the oligopoly game.

⁴ The availability and dissemination of information about the actions of individual members of an industry has also been a matter of some concern to antitrust authorities (Scherer and Ross, 1991, Ch. 9). It is feared that 'tacit collusion' is easier when other members can be closely monitored and, if necessary, punished. Note, however, that from a game-theoretical perspective individualized feedback information is not necessary for the employment of trigger strategies in quantity-setting oligopoly. Deviations can be detected on the basis of aggregate information feedback.

The three benchmarks for this game are defined as follows. First, the Walrasian equilibrium is obtained if and only if each firm produces quantity q^w :

$$P(nq^w)q^w - C(q^w) \geq P(nq)q - C(q) , \quad \forall q \geq 0 ;$$

$$\Rightarrow q^w = \frac{4a^2}{(2b\sqrt{n} + 3c)^2} .$$

Second, the collusive outcome is reached if and only if each firm produces quantity q^c :

$$P(nq^c)q^c - C(q^c) \geq P(nq)q - C(q) , \quad \forall q \geq 0 ;$$

$$\Rightarrow q^c = \frac{4a^2}{(3b\sqrt{n} + 3c)^2} .$$

Third, the Cournot-Nash equilibrium is reached if and only if each firm produces quantity q^N :

$$P(nq^N)q^N - C(q^N) \geq P((n-1)q^N + q)q - C(q) , \quad \forall q \geq 0 ;$$

$$\Rightarrow q^N = \frac{4a^2}{\left(\frac{b(2n+1)}{\sqrt{n}} + 3c\right)^2} .$$

In the experiment we chose $n=3$, $c=1$, $a=45$ and $b=\sqrt{3}$. For each of the three theoretical benchmarks Table I shows the corresponding values of quantities, prices and profits. The following considerations played a role for the choice of functional forms and parameter values. First, if we wish to allow for both 'mimicking the successful firm' and 'following the exemplary firm' at least three firms should be present. Moreover, with three firms in the industry it is relatively easy to see whether a firm mimics the successful firm or follows the exemplary firm. With more than three firms it might not be so easy to conclude that a firm imitates one of these two firms, or whether its choice is affected by the choice of any other firm. Second, we wanted to ensure that even at the Walrasian equilibrium subjects made positive profits. Therefore we chose $c > 0$. To prevent subjects from making losses, we restricted the production set of each firm to (integer) values of the interval [40,125]. Third, we wanted to separate the benchmarks as much as possible. At the same time we wanted the difference between the Walrasian quantity and the Cournot-Nash quantity to be more or less equal to the difference between the Cournot-Nash quantity and the collusive quantity. Finally, we wanted to employ an environment in which both fictitious play and best-response dynamics ('Cournot learning') would converge to the Cournot-Nash equilibrium. This implied that we could not use a purely linear specification.

Table I
Benchmarks of the game

benchmark	quantities of three firms			price	profits of three firms		
	q_1	q_2	q_3	$P(Q)$	π_1	π_2	π_3
Walrasian equilibrium	100	100	100	15	500	500	500
collusive outcome,							
(a) real quantities:	56.25	56.25	56.25	22.5	843.75	843.75	843.75
(b) integer quantities:	56	56	56	22.55	843.74	843.74	843.74
Cournot-Nash equilibrium	81	81	81	18	729	729	729

For the remainder of the paper we will restrict our attention to the case where three firms interact repeatedly from period $t=1,2,\dots$ onward in a stationary environment. Then it is possible to relate the benchmarks to different dynamic rules.

The first dynamic rule that we examine hypothesizes that firms mimic the choice of the firm that was most successful in the previous period. Besides an imitation part, the rule consists of an experimentation (randomization) part. We will say that a firm "mimics the successful firm" if it experiments with common and independent probability ϵ and it imitates the successful firm with probability $(1-\epsilon)$. When a firm experiments, it chooses from the feasible set according to a probability distribution with full support. When a firm imitates the successful firm, it chooses the same quantity as was produced by the firm (one of the firms) with the highest profit in the previous period.

Result 1: If all firms mimic the successful firm, the unique stochastically stable state of the process is (q^w, \dots, q^w) as $\epsilon \rightarrow 0$.

In the long run, only the Walrasian outcome will be observed a significant fraction of the time. Result 1 is due to Vega-Redondo (1997).⁵ The intuition of the result is simple. Assume that firms produce different quantities. Usually the firm that produces the high quantity is the firm that earns the highest profit. Then, the imitation part of the rule induces the other firms to select this high quantity in the next period. By mimicking the successful firm, firms will arrive at a symmetric, say sub-Walrasian quantity.

⁵ The result holds as long as there are identical cost functions and a decreasing demand function. Firms are assumed to choose from a common finite grid containing the Walrasian equilibrium. It is allowed that a firm does not change its quantity with positive probability.

The process remains at this allocation until one of the firms experiments. If a firm experiments by selecting a lower quantity, it will earn a lower profit. If it mimics the more successful firms in the next period, the industry will return to the last symmetric allocation. On the other hand, if the firm experiments by selecting a higher quantity in the direction of the Walrasian quantity, it will earn the higher profit. The next period its choice will be mimicked by the other firms. Gradually this process moves toward the Walrasian allocation. The Walrasian allocation is relatively stable, because a unilateral deviation of one of the firms will not be imitated by the other firms.

The other rule based on imitation assumes that firms follow the firm that sets the good example from the perspective of industry profit. Besides the following part, the rule consists of an experimentation part. We will say that a firm "follows the exemplary firm" if it experiments with common and independent probability ϵ and it chooses the exemplary quantity with probability $(1-\epsilon)$. The exemplary quantity is the quantity (one of the quantities) produced in the previous period that would give the highest sum of profits if it were chosen by all firms.

Result 2: If firms follow the exemplary firm, the unique stochastically stable state of the process is (q^c, \dots, q^c) as $\epsilon \rightarrow 0$.

The dynamic adjustment process when firms follow the exemplary firm resembles the adjustment process when firms mimic the successful firm. The crucial difference between the two forms of imitation lies in the fact that they develop in opposite directions. To prove result 2, the proof provided by Vega-Redondo for result 1 needs only slight adaptation: obviously, all references to Walrasian quantities need to be replaced by references to collusive quantities. The essential but trivial change needs to be made in the proof of lemma 1. It is sufficient to prove that q^c is the exemplary choice in case one firm produces q^c while all other firms produce q . This follows directly from the definition of q^c .⁶

An alternative dynamic approach hypothesizes that firms use belief learning processes to adapt their behavior. Belief learners trust the stability in the pattern of choices made by others. The simplest version of belief learning is known as the Cournot rule or best response rule. According to this rule each firm believes that the aggregate quantity produced by the other firms in the previous period will be produced again in the present period. A firm best responds if and only if it myopically maximizes its expected payoff. As already noted, best response dynamics can lead to cycles. However, for the

⁶ Vega-Redondo's results assume that the prospective outcome is an element of the finite set from which firms choose their quantities. This is not the case for the collusive outcome (56.25, 56.25, 56.25). It is not difficult to see though that in our case following the exemplary firm will lead to the unique stochastically stable state (56, 56, 56).

particular market structure and parameterization described above, convergence to the Cournot-Nash equilibrium can be proven.

Result 3a: For the triopoly market described above, best response dynamics converge to the Cournot Nash equilibrium (81,81,81) independent of the choices made in the first period.

Given that each of the three firms in the industry chooses an integer quantity from the interval [40,...,125], there exist $86 \times 86 \times 86$ different possible states for the first period. With the help of a computer-program we checked the development of play for each potential state of the world. Best response dynamics converge with high speed to the Cournot-Nash equilibrium: on average convergence is obtained in 4.06 steps (where each initial state is weighted equally). One of the worst cases is the initial state where two firms produce 125 units each and one firm produces 64 units. This initial state requires 6 steps for convergence.⁷

Fictitious play is another representative of the class of belief learning models. This rule looks further back than one period only. According to fictitious play, firms believe that others select a quantity with probability equal to the observed empirical frequency of that quantity in past play. One has to make (rather arbitrary) assumptions about firms' prior distribution before they enter the game. Given its updated beliefs a firm chooses the quantity that (myopically) maximizes its expected payoff. Because the production set of each player is rather large in the present game, fictitious play only becomes meaningful after a considerable length of play. Therefore, we consider an adapted version of fictitious play that is easier to implement. According to adapted fictitious play each firm chooses a best response to the average aggregate quantity produced by the other firms in all previous periods. In the remainder of the paper we will simply refer to fictitious play when we have this adapted version in mind.

Result 3b: For the triopoly market described above, fictitious play converges to the Cournot-Nash equilibrium (81, 81, 81) independent of the choices made in the first period.

Again, we proved result 3b with the help of a computer-program. It turns out that fictitious play converges slower to the Cournot-Nash equilibrium than best response dynamics. On average firms converge in 36.49 steps to the Cournot-Nash equilibrium (where each initial state is weighted equally). The worst case is the state where two firms produce 40 units each and one firm produces 125 units in

⁷ The computer-program is written in Turbo-Pascal. It is available from the authors on request.

the first period. This initial state requires 102 steps for convergence.⁸

The main aim of our paper is to investigate whether the information feedback that is available to the players affects the degree to which they employ the above mentioned dynamic rules of conduct (mimicking, following, belief learning), and whether this, in turn, affects the degree to which the three theoretical benchmarks (Walras, collusion, Cournot-Nash) serve as states of attraction. To that purpose we designed three information treatments, each of which is hypothesized to be particularly conducive to the employment of one of the three dynamic rules. The design is 'nested' in the sense that the treatments can be strictly ordered on the amount of information provided to the subjects.

Table II presents the main features of the three treatments. In all treatments, the players receive feedback information on total quantity, price, private revenue, private cost, and private profit. In fact in Treatment Q this is all the information that players receive in a particular period about the outcome of play in the previous period. In treatment Qq, firms receive additional information on the individual quantities produced by the other two firms. Finally, in treatment Qq π , firms are not only provided with individualized information about the quantities but also about the corresponding profits to the other two firms.

The hypothesis is that treatments Q, Qq, and Qq π are conducive to the employment of belief learning, following the exemplary firm, and mimicking the successful firm, respectively. Note that in all treatments a firm receives information about total production in the previous period. In principle, it is therefore possible to behave in accordance with belief learning in all treatments. Treatment Q is conducive to belief learning in the sense that it is the only treatment

⁸Note that we propose stochastic rules for the imitation approach, while we propose deterministic rules for the belief learning approach. It is also possible to derive convergence results for stochastic versions of the belief learning rules.

Table II
Summary of treatment and predictions^a

treatment	baseline information	additional information	prediction: favorable for
Q	R_i, C_i, π_i, Q, P	-	belief learning
Qq	R_i, C_i, π_i, Q, P	q_j, q_k	follow exemplary firm
Qq π	R_i, C_i, π_i, Q, P	q_j, q_k, π_j, π_k	mimic successful firm

^a All information concerns the previous period. Subscript i refers to the firm itself; subscripts j and k concern the other two firms in the industry. R denotes revenue; C denotes costs; π denotes profit; Q denotes aggregate production; P denotes price.

where firms cannot imitate, because they lack the necessary individualized information to do so. Similarly, it can be argued that mimicking is possible in treatment Qq and that following is possible in treatment Qq π . For both following and mimicking, individualized feedback information is sufficient to determine the relevant profits and adapt behavior accordingly. Treatment Qq π , however, is favorable for mimicking in the sense that no profit calculations need to be made; the information needed for a comparison of profits is presented on a silver plate. In short, our hypothesis should be read in terms of treatments being "relatively conducive" to one of the three rules.

Mimicking the successful firm could be thought of as a low cognitive effort algorithm in this game. Following the exemplary firm is a higher cognitive effort algorithm. For each output produced in the previous period a firm has to calculate or estimate the profit that would have resulted when all firms had chosen this output. To apply this rule, a firm should at least have some idea about which firm 'contributes the most' to the aggregate profit. This has to involve some reasoning, because it is not part of the immediately available information. Belief learning seems to be the highest cognitive effort algorithm. It requires a firm to think about the quantity that would have yielded the highest profit given the quantity of others. To do this, a firm has to compute or estimate the best response function, which is not so easy for the present game.

Independent of the present study, Huck *et al.* (1997) conduct an experimental study that is related to ours. They also examine the role of imitation in the quantity setting oligopoly game by manipulating the information that is available to the firms. Besides many smaller differences in the details of the design and procedure (e.g., number of firms, number of rounds, cost and demand functions, the use of profit calculators), there are two major differences. First, contrary to Huck *et al.* (and also Fouraker

and Siegel, 1963), we provide subjects with full information about demand and cost functions in all treatments. Hence, in principle our subjects always have enough information to calculate any benchmark that they like. Second, we employ a feedback information treatment (Qq) which is in between full individualized feedback of quantities and profits and no individual feedback. This, we believe, gives a sharp view on the relative importance of mimicking the successful firm and following the exemplary firm (the latter of which is not considered in Huck *et al.*).

3. DESIGN AND PROCEDURES

Both the instructions and the experiment were computerized. A transcript of the instructions is provided in Appendix A. Subjects could read the instructions at their own pace.⁹ Subjects were explained that they made decisions for their 'own firm', while two other subjects made the decisions for 'firm A' and 'firm B'. It was explained that they would interact with the same two other subjects for the whole experiment, but that they could not know the identity of these two subjects. Subjects were also informed that the experiment would consist of 100 periods.¹⁰

Each period, firms had to decide simultaneously how much to produce. They could only choose integer values between and including 40 and 125. It was emphasized that all firms had symmetric circumstances of production, *i.e.*, all firms had the same cost function. All firms in an industry received the same price for each commodity produced. Both the relationship between own production and costs and between aggregate production and price (*cf.* section 2) was communicated to subjects in three ways: via a table, a figure and the formula. It was explained that all three forms contained exactly the same information, and that subjects could make use of the form that they liked best. Our impression is that most subjects used the tables. We also projected the cost and price tables on the wall, in order to induce common knowledge of the market structure. Appendix A contains some copies of these forms of information.

We explained that a firm's profit in a period was its revenue (own production \times price) minus its costs. Although we provided a subject with the information about its profit after a period had ended, we

⁹ The program is written in Turbo Pascal using the RatImage library. Abbink and Sadrieh (1995) provide documentation of this library. The program is available from the authors.

¹⁰ The disadvantage of a known final period is that subjects may anticipate the end of the experiment and then an end-effect may occur. The alternative is to try to induce an infinite horizon, but such a procedure has its own disadvantages, especially on the point of credibility (see also, Selten *et al.*, 1997). We chose for 100 periods because simulations indicated that this was about the minimum time for the dynamic rules to converge to the benchmarks.

did not provide profit tables.¹¹ Neither did we give away best responses, as is done in some experiments. However, subjects were given regular calculators to reduce computational problems.

Before the experiment started, subjects had to correctly answer some questions testing their understanding before they could proceed with the experiment. Also, subjects were made familiar with the screens used in the experiment. This was the only point where the instructions for the three treatments Q, Qq and Qq π differed: the screen that popped up after a period had ended, contained different information, as explained in Table II. Copies of these screens can be found in Appendix A.

A subject's profits were added up for all 100 periods. During the experiment subjects generated experimental points. At the end of the experiment the experimental points were exchanged for Dutch guilders at an exchange rate of 1300 experimental points = 1 Hfl.¹² The subjects filled in a questionnaire, asking for some background information, before they were privately paid their earnings.

Subjects were recruited through an announcement on a bulletin board at the University of Amsterdam. A total of 102 subjects participated in 6 sessions. In treatment Q 33 subjects participated divided over 2 sessions (18+ 15). In treatment Qq 36 subjects participated divided over 2 sessions (21+ 15) and in treatment Qq π 33 subjects participated divided over 2 sessions (15+ 18). Subjects had no prior experience with directly related experiments, and, of course, no subject participated in more than one session. Sessions lasted between 1½ and 2 hours. Average earnings were 55.13 Hfl, which is the Dutch equivalent of about US\$30.

4. RESULTS

The experimental results are described in two sections. A general overview of our results is presented in section 4.1. Section 4.2 contains a comparison of the explanatory power of the models on the basis of a maximum likelihood procedure.

4.1 General results

If the three treatments trigger different dynamics, as argued in section 2, then one would expect

¹¹ Of course, subjects had all information necessary to construct such tables themselves if they wanted.

¹² At the time of the experiment the 1 dutch guilder was worth about 0.53 US-dollar.

that firms produce most in treatment $Qq\pi$ with its Walrasian benchmark and least in treatment Qq with its collusive benchmark. The results are in the direction of the predictions. Figure 1 displays production levels averaged per ten periods. The average production per group of three firms appears to be pretty stable throughout the experiment in each treatment. In treatments Q and Qq an end-effect occurs in the last periods. The difference in production levels between treatments Q and Qq seems to disappear over time.

[Figure 1. Average production per block of 10 periods]

Table III presents some experimental outcomes. In treatment $Qq\pi$ results are more competitive than the Cournot-Nash outcome and in treatment Qq results are more collusive than the Cournot-Nash outcome.¹³

Table III
Average quantities, prices and profits per treatment^a

treatment	quantity	price	profit in points
Q	77.84 (16.19)	18.59 (1.76)	731.43 (109.09)
Qq	74.63 (17.27)	19.17 (2.05)	748.73 (116.47)
$Qq\pi$	82.43 (20.06)	17.90 (2.72)	667.08 (147.12)

^a Standard deviations between parentheses. For individual quantities and profits, the number of observations is 3300 in treatments Q and $Qq\pi$, and 3600 in treatment Qq . For prices the number of observations is 1100 for treatments Q and $Qq\pi$, and 1200 for treatment Qq .

Whether the differences in quantities and profits between the three treatments are significant depends on the way the data are analyzed. When all individual choices count as data-points, the differences in quantities, prices and profits are very significant for each possible pair of treatments (Mann-Whitney U-tests, with $n_Q = n_{Qq\pi} = 3300$, $n_{Qq} = 3600$). On the other hand, when each industry of three firms yields only one data-point, most comparisons miss the significance level of 5% (Mann-Whitney U-tests, $n_Q = n_{Qq\pi} = 11$, $n_{Qq} = 12$).

Within each treatment that allows for imitation (Qq and $Qq\pi$) there seems to be more than one

¹³These results are in line with the results reported by Huck *et al.* (1997). They also find a higher production level in treatment $Qq\pi$ than in treatment Q (treatments 'full' and 'best' in their terminology). The y did not carry out treatment Qq .

focal point. This may (partly) explain the fact that many differences miss the significant level when each industry yields only one average data-point. There exists more variation in the production-levels between groups within each imitation treatment than in the treatment that does not allow for imitation (Q). Figure 2 shows the 'running frequencies' of the aggregate production of the three treatments when the data are pooled over all periods. In treatment Q the production-level is concentrated around the Cournot-Nash outcome of 243 units. In treatment Qq the distribution of the production-levels is bimodal: as expected, there is a top at the collusive outcome of 168 units. However, the mode of the distribution is at the Cournot-Nash outcome. Again, the distribution of production-levels is bimodal in treatment Qq π . As hypothesized, the mode of the distribution is found at the Walrasian outcome of 300 units. Interestingly, the distribution also has a top at the collusive outcome.¹⁴

[Figure 2. Frequencies of group-outcomes]

Figure 2 suggests which outcomes may actually have functioned as (behavioral) equilibria in the three markets. However, this figure is silent about the behavioral dynamics in the markets. In the following, we call an allocation a rest-point if firms tend to repeat their choices at that point. We focus on symmetric rest-points for a theoretical and an empirical reason. The theoretical reason is that all benchmarks defined in section 2 imply that all firms produce the same quantity. The empirical reason is that firms more often repeat their decision of the previous period when all firms made a symmetric choice (49.5%) than they do when this is not the case (31.8%). We say that firms make a symmetric choice if the difference between the maximum and the minimum production of that industry in the previous period is less than or equal to 8 units. Table IV provides some insight into which allocations may have served as behavioral equilibria in the three treatments.

The table suggests that the collusive outcome is a stable rest-point in treatments Qq and Qq π , but not in treatment Q.¹⁵ In the imitation treatments the collusive outcome is reached quite

¹⁴The finding that the variability of aggregate quantities increases when subjects are supplied with more feedback information is also reported by Fouraker and Siegel (1963, p. 143) and Huck *et al.* (1997).

¹⁵ Of course, this analysis involves a number of arbitrary choices. The general picture does not depend on these choices, though.

Table IV
Behavioral equilibria or rest points^a

own choice in category	treatment Q	treatment Qq	treatment Qqπ
40-48	15	16	3
49-57 (C)	12	354 (79%)	485 (85%)
58-66	161 (27%)	166 (51%)	136 (43%)
67-75	239 (41%)	332 (34%)	121 (31%)
76-84 (N)	361 (55%)	282 (27%)	138 (60%)
85-93	16	28 (7%)	140 (26%)
94-102 (W)	0	9	326 (37%)
103-111	0	1	33 (18%)
112-125	0	9	1

^a The first number in a cell displays the number of symmetric choice situations; for more than 25 symmetric choice situations the second number between parentheses displays the percentage of choices that are exactly the same as the choice in the previous period. A choice situation is symmetric if the difference between the maximum and the minimum production in an industry is smaller than or equal to 8.

often, and given that it is reached firms tend to stay there. The Cournot-Nash outcome is a rest-point for treatment Q. It is also reached quite often in treatment Qq, but there it does not seem to be very stable because firms tend to abandon their previous choice quite often. The Walrasian outcome is only reached quite often in treatment Qqπ. It does not seem to be a very stable rest-point, because firms stay at their former choice in only 37% of the cases.

Another dynamic aspect relevant for the comparison of the relative explanatory power of the two explanations based on imitation is the adaptation process of "the middle firms" in the imitation treatments. The middle firm is the firm that produced (strictly) more than the exemplary firm but (strictly) less than the successful firm in the previous period. In treatment Qq the middle firms decrease their production in 41.5% of the 643 relevant cases, as they are expected to do when they follow the exemplary firm. They increased their production in 30.2% of the cases, and they did not change their previous production quantity in the remaining 28.3% of the cases. In treatment Qqπ the middle firms increased their production in 43.4% of the 452 relevant cases, as they are expected to do when they mimic the successful firm. They decreased their production in 29.0% of the cases and did not change their previous production level in the remaining 27.7% of the cases. Middle firms tend to follow the exemplary firm in treatment Qq, whereas they tend to mimic the successful firm in treatment Qqπ.

4.2 Comparison of dynamic models

A more encompassing comparison of the dynamic models requires one to be more specific about how firms experiment. We make a simple assumption about the experimentation process: a firm's choices are assumed to be i.i.d. generated by a truncated normal distribution $(\mu, \sigma)_{[40, 125]}$, where μ is set equal to the prediction of a particular model and σ is a free parameter. In the following we are not particularly interested whether our assumption about the experimentation process itself is accurate. We are interested in finding the best dynamic model describing behavior in the quantity setting oligopoly game. A priori it is not clear that the assumption about the experimentation process will hurt one model more than another.

The prediction of the model "follow the exemplary firm" is that each firm chooses a quantity that is equal to the quantity produced by the firm that was exemplary in the previous period. The prediction of the model "mimic the successful firm" is that each firm chooses a quantity that is equal to the quantity produced by the firm that had the highest profit in the previous period. The prediction of the model "best response" is that each firm chooses a best response to the aggregated quantity produced by the other two firms in the previous period. The prediction of the model "fictitious play" is that each firm chooses a best response to the average aggregated quantity produced by the other two firms in all previous periods. The models are compared with the "random model", a benchmark model that predicts that each of the 86 possible choices (integers from 40 to 125) is chosen with probability $1/86$. Note that the random model is nested in each of the other models: if $\sigma \rightarrow \infty$ in one of the dynamic models, the random model is obtained.

Table V reports the maximum likelihood results for the models described above in the three treatments. It is clear that one of the two models of the belief learning approach dominates the other: all three treatments yield a better likelihood and a lower estimate for σ for the fictitious play model than for the best response model. Therefore, in the remainder of this paper we will present fictitious play as the candidate of the belief learning approach. Each of the remaining three models explains the data significantly better than the random model does in all treatments (Likelihood Ratio test at 1% level).

Remember that in treatment Q it is impossible to imitate because firms lack the necessary information to do so. Nevertheless, the two imitation models are also estimated for this treatment to obtain a sort of lower-bound on the expected performance of the fictitious play model. The fictitious play model beats the imitation models both on the criterium of likelihood and on the criterium of the estimate σ (lower). In treatments Qq and Qq π firms had in principle sufficient

Table V
Maximum likelihood estimates simple models^a

	treatment Q n= 3267	treatment Qq n= 3564	treatment Qqπ n= 3267
follow exemplary firm	-logL= 13949.3 σ= 23.87 (0.44)	-logL= 15078.5 σ= 23.27 (0.39)	-logL= 13927.2 σ= 23.78 (0.43)
mimic successful firm	-logL= 14118.6 σ= 26.10 (0.54)	-logL= 15461.7 σ= 28.42 (0.59)	-logL= 13842.2 σ= 23.81 (0.42)
best response	-logL= 14084.5 σ= 21.37 (0.40)	-logL= 15772.1 σ= 34.98 (1.29)	-logL= 14548.9 σ= 86.58 (16.35)
fictitious play	-logL= 13837.6 σ= 18.09 (0.29)	-logL= 15687.2 σ= 28.60 (0.79)	-logL= 14516.2 σ= 45.11 (2.73)
random	-logL= 14552.4	-logL= 15875.3	-logL= 14552.4

^a The likelihood is computed on the basis of all choices from period 2-100 of all subjects in a treatment . The models are explained in the main text. The standard error of an estimated parameter is displayed between n parentheses.

information to behave in accordance with each of the models. The follow exemplary firm model outperforms the mimic successful firm model in treatment Qq both on the likelihood and the estimate of σ . The result is opposite for treatment Qqπ: there the mimic successful firm model outperforms the follow exemplary firm model on the basis of the likelihood, though both models yield an approximately equal estimate for σ . These results are in line with the predictions described in Table II.

Table IV suggested that the imitation treatments Qq and Qqπ contain more than one rest-point each. This may be the result of heterogeneity in the population. Perhaps one part of the population uses one dynamic rule while another part uses another. To investigate this possibility, more general models have to be formulated. The most general model considered here allows a firm to use either of the three models. In the following, let P_{ef} denote the probability that an arbitrary firm uses the follow exemplary firm model; P_{msf} and P_{fp} denote the probability that a firm uses the mimic successful firm model and the fictitious play model respectively. We define $P_{ef} + P_{msf} + P_{fp} = 1$. Let $q_{i,t}$ denote the quantity produced by firm i in period t . Then, the unconditional likelihood function of firm i 's choices from period 2-100 $L(q_{i,2}, \dots, q_{i,100})$ is given by:

$$L(q_{i,2}, \dots, q_{i,100}) = P_{fef} * \prod_{t=2}^{100} L(q_{i,t}|fef) + P_{msf} * \prod_{t=2}^{100} L(q_{i,t}|msf) + P_{fp} * \prod_{t=2}^{100} L(q_{i,t}|fp) ,$$

where $L(q_{i,t}/fef)$, $L(q_{i,t}/msf)$ and $L(q_{i,t}/fp)$ denote the conditional probability that firm i chooses quantity $q_{i,t}$ at period t when it follows the exemplary firm, the conditional probability that firm i chooses quantity $q_{i,t}$ at period t when it mimics the successful firm and the conditional probability that firm i chooses quantity $q_{i,t}$ at period t when it best responds to the average quantity produced by the two others in the previous periods, respectively.

The likelihood functions for three general models that allow for two dynamic rules each are obtained by setting either P_{fef} , P_{msf} or P_{fp} equal to 0 in the equation describing the likelihood function for the three rules. By setting two of these three probabilities equal to 0, each of the three remaining simple models can be obtained for which the results were described in Table V. Thus the models are related to each other via a nested structure.

Table VI
Maximum likelihood estimates general models imitation treatments^a

	treatment Qq n= 3564	treatment Qqπ n= 3267
model 1, 2 rules:	-logL= 14905.5	-logL= 13371.6
(a) follow exemplary firm	σ= 20.70 (0.33)	σ= 17.93 (0.28)
(b) mimic successful firm ($P_{msf}= 1-P_{fef}$)	$P_{fef}= 0.78$ (0.07)	$P_{fef}= 0.45$ (0.09)
model 2, 2 rules:	-logL= 14908.8	-logL= 13785.3
(a) follow exemplary firm	σ= 19.70 (0.32)	σ= 20.60 (0.35)
(b) fictitious play ($P_{fp}= 1-P_{fef}$)	$P_{fef}= 0.67$ (0.08)	$P_{fef}= 0.76$ (0.08)
model 3, 2 rules:	-logL= 15326.1	-logL= 13532.8
(a) mimic successful firm	σ= 22.91 (0.43)	σ= 18.36 (0.29)
(b) fictitious play ($P_{fp}= 1-P_{msf}$)	$P_{msf}= 0.63$ (0.08)	$P_{msf}= 0.70$ (0.08)
model 4, 3 rules:	-logL= 14845.6	-logL= 13316.3
(a) follow exemplary firm	σ= 19.37 (0.30)	σ= 17.14 (0.26)
(b) mimic successful firm	$P_{fef}= 0.66$ (0.08)	$P_{fef}= 0.39$ (0.09)
(c) fictitious play ($P_{fp}= 1-P_{fef}-P_{msf}$)	$P_{msf}= 0.19$ (0.07)	$P_{msf}= 0.52$ (0.09)

^a The likelihood is computed on the basis of all choices from period 2-100 of all subjects in a treatment . The models are explained in the main text. The standard error of an estimated parameter is displayed between n parentheses.

Table VI reports the maximum likelihood results for these models when applied to the imitation treatments.¹⁶ Each of the two-rules models (models 1, 2 and 3) yields a significantly better fit than either of the two simple models that are nested in it (Likelihood Ratio test at 1% level) in both treatments. Of the two-rules models, model 1 performs unequivocally best in treatment Qqπ. Both the estimate for σ and the likelihood are better for this model than for the other two two-rules models. Thus, firms either seem to mimic the successful firm or follow the exemplary firm in this treatment. The results for the two-rules models are less clear in treatment Qq. Here, model 1 and model 2 perform about equally well (model 1 has the best likelihood, but model 2 has the best estimate for σ).

From a statistical point of view, the results seem most favorable for the most general model allowing for three rules. This model outperforms each of the three two-rules models significantly (Likelihood Ratio test at 1% level) in both treatments. This model also yields the lowest estimates for σ .

The results for model 1 are especially interesting when one tries to find out which firm is imitated in the quantity setting oligopoly game. As hypothesized in section 2, providing firms information about individualized quantities without corresponding profits stimulates firms to follow the exemplary firm: in treatment Qq 78% of the population is estimated to follow the exemplary firm and the remaining 22% is estimated to mimic the successful firm. On the other hand, providing firms with information about individualized quantities and corresponding profits stimulates firms to mimic the successful firm: in treatment Qqπ 55% of the population is estimated to mimic the successful firm and the remaining 45% follows the exemplary firm.

Using the estimates of model 4 and the choices made by a subject, it is easy to compute her α his posterior probability distribution for behaving in accordance with either of the three simple models contained in the general model.¹⁷ It is remarkable that for most subjects one of the posterior probabilities is almost 1. In treatment Qq for 32 of the 36 subjects one of the posterior probabilities is estimated to be higher than 99%. In treatment Qqπ for 28 of the 33 subjects one of these probabilities is estimated to be higher than 99%. This suggests that the idea that subjects follow either of the three rules makes sense. Tables B1 and B2 in Appendix B report the posterior probability distributions for

¹⁶ We do not estimate the models for treatment Q, because subjects did not have the necessary information to imitate in this treatment.

¹⁷ For example, the posterior probability that subject i uses rule *fef* (follow the exemplary firm) in treatment Qq is computed as:

$$P(model_i = fef | q_{i,2}, \dots, q_{i,100}) = \frac{0.66 * \prod_{t=2}^{100} L(q_{i,t} | fef \wedge \sigma=19.37)}{0.66 * \prod_{t=2}^{100} L(q_{i,t} | fef \wedge \sigma=19.37) + 0.19 * \prod_{t=2}^{100} L(q_{i,t} | msf \wedge \sigma=19.37) + 0.15 * \prod_{t=2}^{100} L(q_{i,t} | fp \wedge \sigma=19.37)}$$

treatment Qq and for treatment $Qq\pi$, respectively.

As was to be expected, the 21 players who are estimated to follow the exemplary firm, on average produce less and earn more than the 33 players who are estimated to mimic the successful firm. At the same time, if we make comparisons not over all groups, but only within each group of three firms, then the players who mimic the successful firm do relatively good. In particular, a mimicker always earns more than a follower in the same group of three firms.

5. CONCLUSION

The main goal of this paper was to give a possible account for the wide range of outcomes observed, empirically and experimentally, in quantity setting oligopoly games. It was hypothesized that the type of feedback information that is available to the players may be an important clue in this puzzle. We investigated whether different feedback information can trigger different behavioral heuristics which, in turn, can lead to different outcomes.

Our results can be summarized as followed. As hypothesized, firms produce less when the feedback is conducive to following the exemplary firm (treatment Qq) than when the feedback is conducive to belief learning (treatment Q). They produce most when the feedback is conducive to mimicking the successful firm (treatment $Qq\pi$). Average production quantities do not deviate very much from the Cournot-Nash outcome. However, underlying the moderate differences in average production quantities are remarkable differences in the frequency distributions of the outcomes. With only feedback information about the aggregate quantities produced (treatment Q), the frequency distribution is unimodal and more or less symmetric around the Cournot-Nash equilibrium. When feedback information about individual quantities is available (treatment Qq), there are two peaks; one at the Cournot-Nash equilibrium, and one at the collusive outcome. In treatment $Qq\pi$ the frequency distribution is also bi-modal, with equal-sized peaks around the collusive and the competitive outcome. In the latter case, Cournot-Nash seems to have lost all of its attraction.

A qualitative look at the underlying dynamics permits us to sharpen these results. The Cournot-Nash outcome is the only candidate behavioral equilibrium in treatment Q . In treatments Qq and $Qq\pi$, however, it turns out that the main candidate to serve as behavioral equilibrium is the collusive outcome. Although the Cournot-Nash and the competitive outcome are reached quite often in treatments Qq and $Qq\pi$, respectively, they are abandoned at a much higher rate than the collusive outcome. The Walrasian outcome only attracts aggregate behavior in treatment $Qq\pi$. It does not appear to be a very stable behavioral equilibrium, though.

A quantitative investigation of the dynamics corroborates the hypothesized relationship between

information treatment and dynamic rule of conduct. Maximum likelihood estimates of single rule models indicate that in treatment Q (Qq) [Qq π], the best fit is given by fictitious play learning (following the exemplary firm) [mimicking the successful firm]. The fit increases significantly if we allow for a heterogeneous population of firms, employing different dynamic rules. The picture remains similar though. As long as individualized feedback information is restricted to quantities (treatment Qq), then a majority of the subjects is estimated to follow the exemplary firm. If comparative profit appraisals are easily made (treatment Qq π), a majority of the subjects is estimated to mimic the successful firm. In both of these imitation treatments, the part of the population adapting in accordance with belief learning is estimated to be the smallest of the three.

Whether the better fit of the data by a more general model is sufficient to compensate for the loss in parsimony can of course not be decided on the basis of a criterium of statistical significance only. If one searches for simple elegant theories, then it might be concluded that a different type of environment elicits a different dynamic rule. Firms imitate if they have the opportunity to do so. However, whether they follow the exemplary firm or mimic the successful firm depends on the type of individualized feedback. If firms do not have the possibility to imitate, fictitious play seems to organize the data reasonably well.

On the other hand, if one wants to give a description of what's actually happening in the experiment, then it might be concluded that there is substantial heterogeneity in the dynamic rules used by subjects. If treatments allow for all dynamic rules, then all dynamic rules seem to be used. However, the proportion of the population using a rule is systematically affected by the feedback provided. There are even signs in the data that some of the subjects use more sophisticated strategies than the ones we have allowed for. For example, there are cases in which a firm for some periods produces the collusive quantity of 56, switches to the maximum quantity of 125 as others do not follow, and then switches back to a quantity in the neighborhood of 56. Although these switches could be interpreted as experimentation, it seems to be more likely that subjects switching so wildly are actively trying to direct the others toward the collusive outcome. In the model, firms are the exemplary firm by accident; in the experiment, at least some are so by intent (cf. Selten *et al.*, 1997).

It is useful to recall that Vega-Redondo's motivation for a model of imitation of the successful firm is not based on bounded rationality concerns as much as on evolutionary considerations. According to the evolutionary approach it is relative rather than absolute performance that matters for success. In that case it may be a sensible strategy for a firm to increase its quantity, even if total production is at or above the Cournot-Nash equilibrium. By doing so a firm will hurt itself, but it will hurt the other firms even more. Consistent with this suggestion, within each group of three firms the one producing the highest average quantity generally was the one earning the highest profit. However, this advantage is more than offset by the fact that groups of followers perform better than groups of mimickers. In our

experiment, the Pearson correlation coefficient between firms' average production and their average profit is significantly negative ($r = -0.65$, $p = 0.001$, $n = 102$). If the relative performance of firms is determined across industries and not merely within industries, then following the exemplary firm is a more successful rule of conduct than mimicking the successful firm. As already argued in the evolutionary theory of cultural transmission by Boyd and Richerson (1991, 1993), intergroup rivalry may serve as a potential explanation for a tendency to imitate cooperative behavior. From an evolutionary perspective it is not so straightforward which would be the most successful strategy of imitation: mimic the villain or follow the saint?

References

- Abbink, K., and A. Sadrieh (1995): 'RatImage - Research Assistance Toolbox for Computer-Aided Human Behavior Experiments', *Discussion Paper* No. B-325, University of Bonn.
- Bandura, A. (1986): *Social Foundations of Thought and Action. A Social Cognitive Theory*. New Jersey: Prentice-Hall.
- Boyd, R., and P. Richerson (1991): 'Culture and Cooperation', in: *Cooperation and Prosocial Behavior*, ed. by R. Hinde and J. Groebel. Cambridge: Cambridge University Press.
- Boyd, R., and P. Richerson (1993): 'Rationality, Imitation and Tradition', in: *Nonlinear Dynamics and Evolutionary Economics*, ed. by R. Day and P. Chen. Oxford: Oxford University Press.
- Boylan, R.T., and M.A. El-Gamal (1993): 'Fictitious Play: A Statistical Study of Multiple Economic Experiments', *Games and Economic Behavior*, 5, 205-222.
- Cheung, Y.W., and D. Friedman (1997): 'Individual Learning in Normal Form Games: Some Laboratory Results', *Games and Economic Behavior*, 19, 46-76.
- Daughety, A. (1985), 'Reconsidering Cournot: The Cournot Equilibrium is Consistent', *Rand Journal of Economics*, 16, 368-379.
- El-Gamal, M.A., R.D. McKelvey, and T.R. Palfrey (1994): 'Learning in Experimental Games', *Economic Theory*, 4, 901-922.
- Fouraker, L.E., and S. Siegel (1963): *Bargaining Behavior*. New York: McGraw-Hill.
- Fraser, C. (1994): 'Conjectural Variations', in: *Current Issues in Industrial Organizations*, (ed.) J.Cable. Hampshire: Macmillan.
- 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: Princeton University Press, pp. 349-443.

- Huck, S., H.T. Normann, and J. Oechssler (1997): 'Learning in Cournot Oligopoly - An Experiment', *Discussion paper*, Humboldt University Berlin.
- Martin, S. (1993): *Advanced Industrial Economics*. Oxford: Blackwell Publishers.
- Milgrom, P., and J. Roberts (1991): 'Adaptive and Sophisticated Learning in Normal Form Games', *Games and Economic Behavior*, 3, 82-100.
- Mookherjee, D., and B. Sopher (1994): 'Learning Behavior in an Experimental Matching Pennies Game', *Games and Economic Behavior*, 7, 62-91.
- Offerman, T., A. Schram, and J. Sonnemans (1997): 'Quantal Response Models in Step-Level Public Good Games', forthcoming in *European Journal of Political Economy*.
- Offerman, T., and J. Sonnemans (1997): 'Learning by Experience and Learning by Imitating Successful Others', forthcoming in *Journal of Economic Behavior and Organization*.
- Partow, Z., and A. Schotter (1993): 'Does Game Theory Predict Well for the Wrong Reasons?', *Discussion Paper*, RR 93-46, New York University.
- Pingle, M. (1995): 'Imitation Versus Rationality: an Experimental Perspective on Decision Making', *Journal of Socio-Economics*, 24, 281-315.
- Potters, J., and F. Van Winden (1996): 'Comparative Statics of a Signaling Game: An Experimental Study', *International Journal of Game theory*, 25, 329-353.
- Roth, A., and I. Erev (1995), 'Learning in Extensive Form Games: Experimental Data and Simple Dynamic Models in the Intermediate Term', *Games and Economic Behavior*, 8, 164-212.
- Sarason, I.G., B.R. Sarason, G.R. Pierce, and E.N. Shearin (1991): 'A Social Learning Approach to Increasing Blood Donations', *Journal of Applied Social Psychology*, 21, 896-918.
- Scherer, F., and D. Ross (1990), *Industrial Market Structure and Economics Performance*. Boston: Houghton Mifflin.
- Sell, J., and R. Wilson (1991): 'Levels of Information and Contributions to Public Goods', *Social Forces*, 70, 107-124.
- Selten, R., M. Mitzkewitz and G.R. Uhlig (1997): 'Duopoly Strategies Programmed by Experienced Players', *Econometrica*, 65, 517-555.
- Sinclair, P. (1990), 'The Economics of Imitation', *Scottish Journal of Political Economy*, 37, 113-144.
- Slade, M. (1994), 'Empirical Games: the Oligopoly Case', *Canadian Journal of Economics*, 26, 368-402.
- Vega-Redondo, F. (1997): 'The Evolution of Walrasian Behavior', *Econometrica*, 65, 375-384.

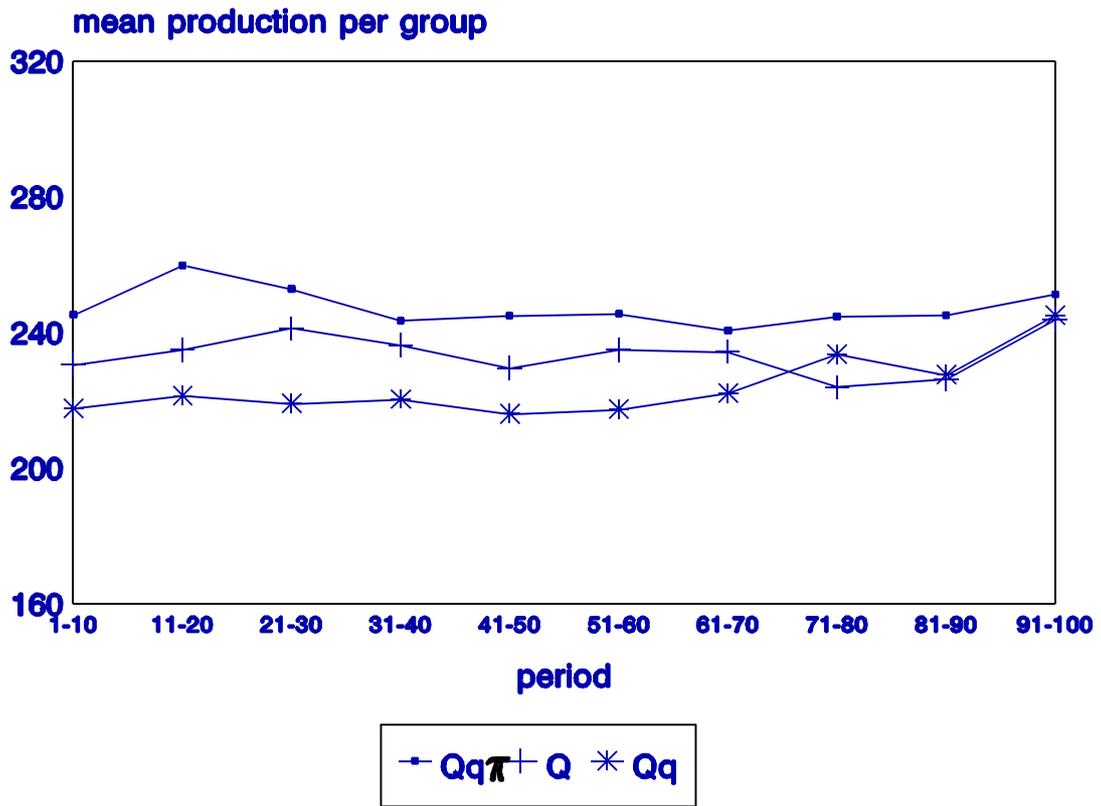


Figure 1. Average production per block of 10 periods

Figure 2. Frequency distributions of aggregate quantities in treatments Q, Qq and Qq π

Appendix A.

This appendix presents the instructions, the figures and tables that were used to communicate the market structure, and prints of the computer screens that were used in each of the three treatments.

Instructions

After a short welcome, subjects read the computerized instructions at their own pace. They could go forward and backward in the program. At the start of the instructions it was explained how they should use the computer. A translated version of the Dutch instructions runs as follows:

"Experiment

You will make decisions for a firm in this experiment. You will be asked repeatedly to determine the quantity that your firm will produce. Your payoff will depend on your production and the production of two other firms. The decisions for these two other firms will be made by two other participants of the experiment. For convenience we will refer to these other firms as firm A and firm B.

You will be matched with the same two other firms throughout the experiment. The same participants will make the decisions for these two firms. You will not know with whom you will be matched, like others will not know with whom they will be matched. Anonymity is ensured.

Own production and costs (table)

The experiment will last for 100 periods. Each period you will decide how much your firm produces. Your production must be greater than or equal to 40 units and smaller than or equal to 125 units. You may only choose integer numbers. For example, it is not allowed to produce half units.

Producing involves costs. There is a cost-form on your table. On one side of this form the cost table is displayed. Each time you pick a lower row of the table, your production increases by 10 units. Each time you pick a column more to the right of the table, your production increases by 1 unit. For example, if you want to know the costs of producing 79 units, you descend in the left column until you have reached the row where 70 units are produced. Then you shift nine columns to the right to find the cell containing the costs of a production-quantity of 79 units. A production of 79 units costs 702 points.

Own production and costs (graph)

The other side of the cost-form displays graphically the relationship between quantity produced and costs. You find the costs of a certain production-quantity by searching this quantity on the horizontal axis and by determining the level of the costs at this quantity on the vertical axis.

Above the graph the formula is displayed which is used to compute the costs. Both the table and the graph are constructed with the help of this formula. Therefore, the table and the graph convey exactly the same information. You may decide for yourself which representation you want to use in the experiment.

Question

What will be your costs if you produce 87 units?

[Subjects had to type in the correct answer before they could continue with the instructions. If they did not answer correctly, the relevant part of the instructions was repeated in different terms before the question was posed again.]

Total production and price (table)

The price that you will receive for each unit produced depends on the total production. The total production is the sum of the production of firm A, the production of firm B and your production. Since firm A and firm B each will also produce between 40 and 125 units, total production will not be smaller than 120 units and not be greater than 375 units.

There is a price-form on your table. One side of this form shows the price-table. This table gives the prices for each possible total production. The table works the same way as the cost-table. Each time you pick a lower row of the table, total production increases by 10 units. Each time you pick a column more to the right of the table, total production increases by 1 unit. For example, if total production consists of 272 units, you descend in the left column until you have reached the row where 270 units are produced. Then you shift two columns to the right to find the cell containing the price of a total production-quantity of 272 units.

Total production and price (graph)

The other side of the price-form displays graphically the relationship between total quantity produced and price. You find the price of a certain production-quantity by searching this quantity on the horizontal axis and by determining the level of the price at this quantity on the vertical axis.

Again, above the graph the formula is displayed which is used to compute the price. Both the table and the graph are constructed with the help of this formula. You may decide for yourself which representation you want to use in the experiment.

Question

What will be the price if you produce 87 units, firm A produces 66 units and firm B produces 102 units?

[Again, subjects had to type in the correct answer before they could continue with the instructions.]

Production and profit

Each period your revenue is equal to your produced quantity multiplied by the price. Your profit is your revenue minus the costs of your production.

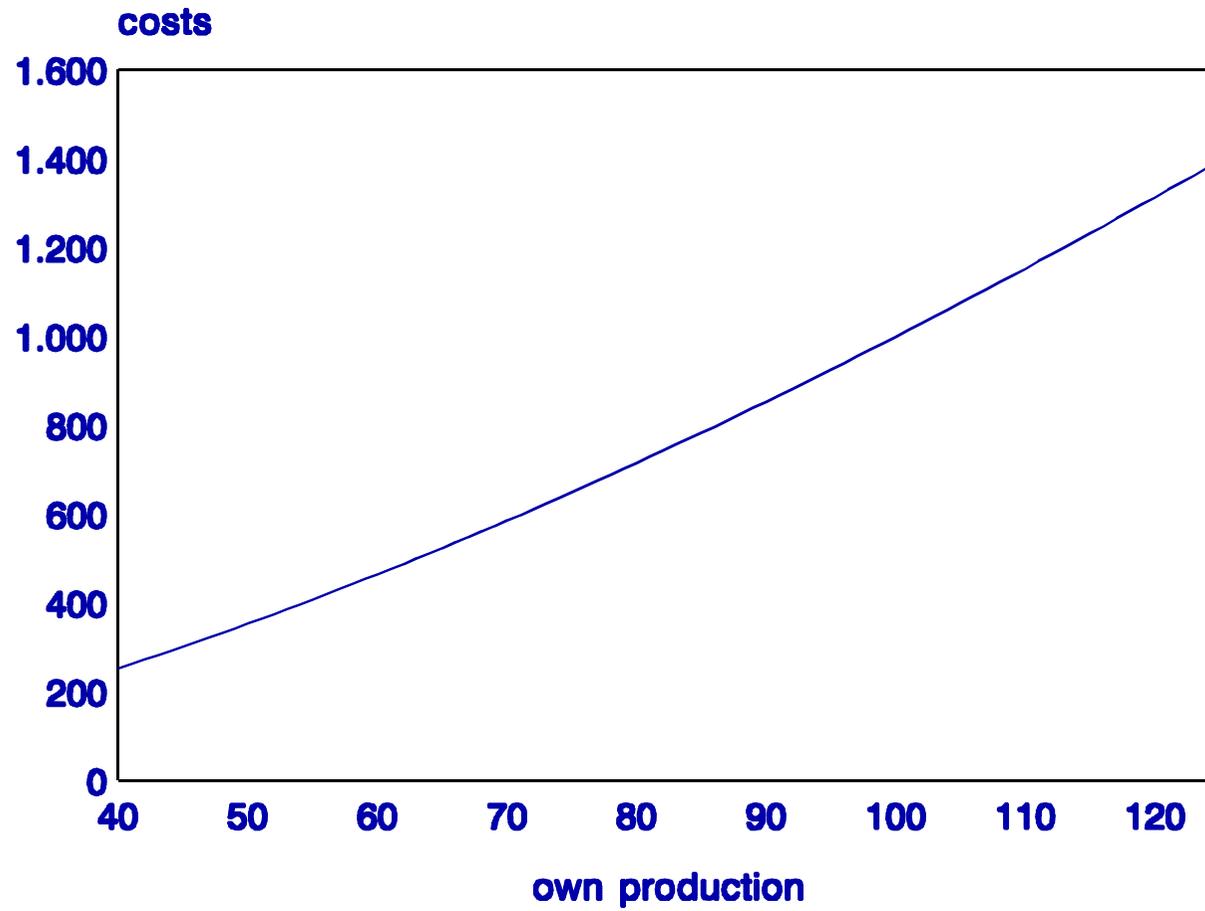
The circumstances of production are exactly the same for firm A and firm B as they are for you. The costs of their production is determined in a similar way as your production. They receive the same price for each unit they produce as you receive.

The experiment lasts for 100 periods. Each period the circumstances of production will be the same as described above. Your payoff in the experiment is equal to the sum of your profits in each of the 100 periods. These profits are denoted in points.

At the end of the experiment your points will be exchanged for real money. For each 1300 points you will receive 1 guilder."

Then subjects were made familiar with the screens that they would see in the experiment. The instructions were exactly the same for treatment Q, treatment Qq and treatment Qqr, except for the screen where information about the results of the previous period was communicated. Before the experiment started, a handout was given to the subjects with a summary of the instructions and the cost and price-tables were projected on the wall.

costs own production-quantity (q) = $q\sqrt{q}$;



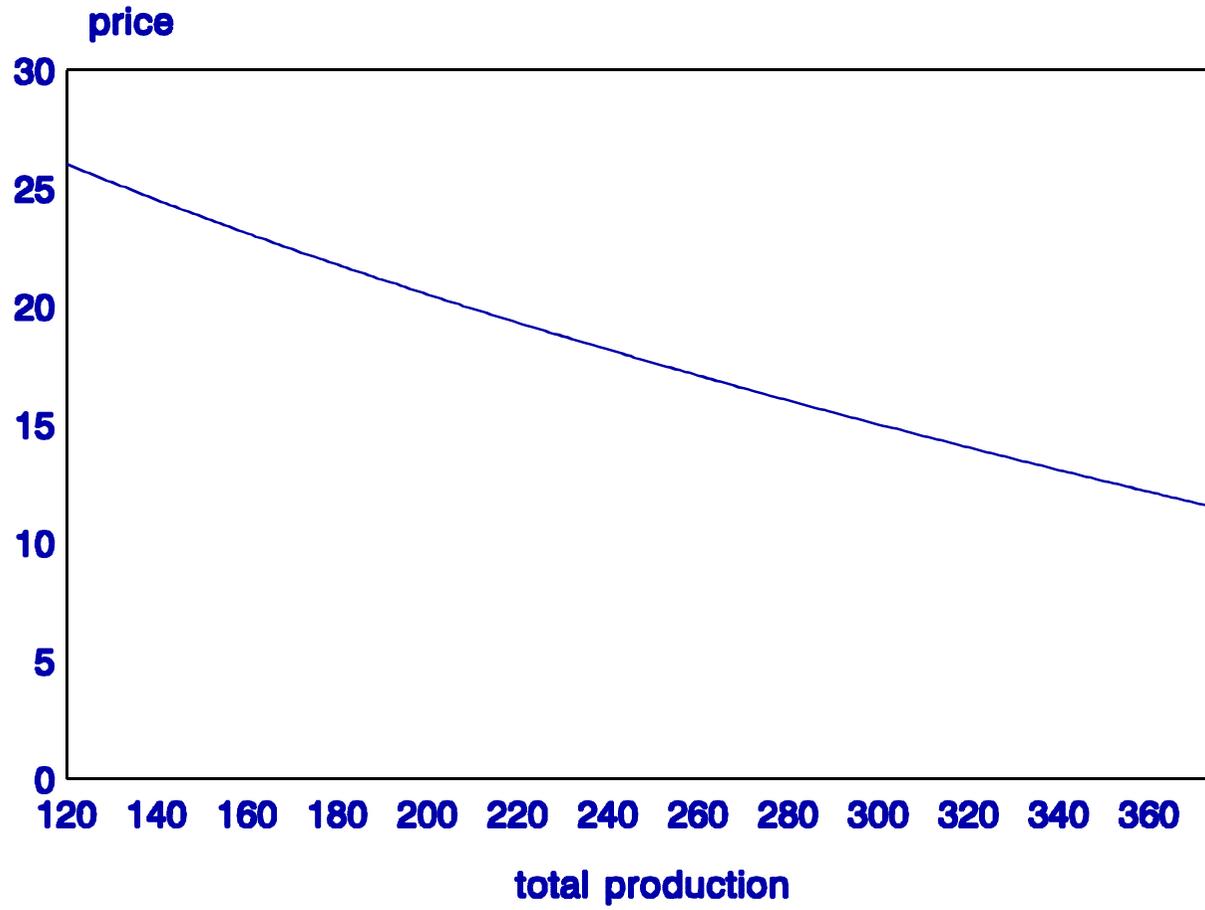
Cost-form: costs of your own production

Production	Costs	0	1	2	3	4	5	6	7	8	9
40		253	263	272	282	292	302	312	322	333	343
50		354	364	375	386	397	408	419	430	442	453
60		465	476	488	500	512	524	536	548	561	573
70		586	598	611	624	637	650	663	676	689	702
80		716	729	743	756	770	784	798	811	826	840
90		854	868	882	897	911	926	941	955	970	985
100	1000	1015	1030	1045	1061	1076	1091	1107	1122	1138	
110	1154	1169	1185	1201	1217	1233	1249	1266	1282	1298	
120	1315	1331	1348	1364	1381	1398					

Example

When you want to know the costs of producing 83, you look at the row displaying **80**, and then at the number below **3**: 756.

*price total production-quantity (Q) = $45 - \sqrt{3} * \sqrt{Q}$;*



Price-form (total production = production of all three firms added up)

Total production	Price 0	1	2	3	4	5	6	7	8	9
120	26.03	25.95	25.87	25.79	25.71	25.64	25.56	25.48	25.40	25.33
130	25.25	25.18	25.10	25.03	24.95	24.88	24.80	24.73	24.65	24.58
140	24.51	24.43	24.36	24.29	24.22	24.14	24.07	24.00	23.93	23.86
150	23.79	23.72	23.65	23.58	23.51	23.44	23.37	23.30	23.23	23.16
160	23.09	23.02	22.96	22.89	22.82	22.75	22.68	22.62	22.55	22.48
170	22.42	22.35	22.29	22.22	22.15	22.09	22.02	21.96	21.89	21.83
180	21.76	21.70	21.63	21.57	21.51	21.44	21.38	21.32	21.25	21.19
190	21.13	21.06	21.00	20.94	20.88	20.81	20.75	20.69	20.63	20.57
200	20.51	20.44	20.38	20.32	20.26	20.20	20.14	20.08	20.02	19.96
210	19.90	19.84	19.78	19.72	19.66	19.60	19.54	19.49	19.43	19.37
220	19.31	19.25	19.19	19.14	19.08	19.02	18.96	18.90	18.85	18.79
230	18.73	18.68	18.62	18.56	18.51	18.45	18.39	18.34	18.28	18.22
240	18.17	18.11	18.06	18.00	17.95	17.89	17.83	17.78	17.72	17.67
250	17.61	17.56	17.51	17.45	17.40	17.34	17.29	17.23	17.18	17.13
260	17.07	17.02	16.97	16.91	16.86	16.81	16.75	16.70	16.65	16.59
270	16.54	16.49	16.44	16.38	16.33	16.28	16.23	16.17	16.12	16.07
280	16.02	15.97	15.91	15.86	15.81	15.76	15.71	15.66	15.61	15.56
290	15.51	15.45	15.40	15.35	15.30	15.25	15.20	15.15	15.10	15.05
300	15.00	14.95	14.90	14.85	14.80	14.75	14.70	14.65	14.60	14.55
310	14.50	14.46	14.41	14.36	14.31	14.26	14.21	14.16	14.11	14.07
320	14.02	13.97	13.92	13.87	13.82	13.78	13.73	13.68	13.63	13.58
330	13.54	13.49	13.44	13.39	13.35	13.30	13.25	13.20	13.16	13.11
340	13.06	13.02	12.97	12.92	12.88	12.83	12.78	12.74	12.69	12.64
350	12.60	12.55	12.50	12.46	12.41	12.37	12.32	12.27	12.23	12.18
360	12.14	12.09	12.05	12.00	11.95	11.91	11.86	11.82	11.77	11.73
370	11.68	11.64	11.59	11.55	11.50	11.46				

Example

When you want to know the price at a total production of 267, you look at the number in row **260** and column **7**: 16.70.

Appendix B.

Tables B1 and B2 report the posterior probabilities that a subject uses each of the three rules, given the maximum likelihood estimates for model 4 (table VI), along with quantities and profits averages over the 100 periods.

Table B1. Posterior probabilities, quantities and profits in treatment Qq

group	subject	P_{ref}	P_{msf}	P_{fp}	quantity	profits
1	1	1.000	0.000	0.000	65.85	785.6
1	2	1.000	0.000	0.000	70.60	792.5
1	3	1.000	0.000	0.000	65.82	787.5
2	4	0.000	1.000	0.000	79.53	762.7
2	5	0.999	0.000	0.001	75.17	748.1
2	6	1.000	0.000	0.000	73.66	734.5
3	7	1.000	0.000	0.000	74.84	761.9
3	8	0.004	0.774	0.222	78.73	770.4
3	9	1.000	0.000	0.000	74.11	746.7
4	10	1.000	0.000	0.000	73.35	653.4
4	11	0.000	1.000	0.000	89.25	723.0
4	12	0.000	1.000	0.000	85.85	698.7
5	13	0.106	0.000	0.894	78.66	670.5
5	14	0.000	1.000	0.000	88.45	712.1
5	15	0.801	0.051	0.148	82.35	678.9
6	16	0.000	0.000	1.000	88.90	731.6
6	17	0.013	0.000	0.987	77.53	677.7
6	18	0.002	0.000	0.998	79.12	683.0
7	19	1.000	0.000	0.000	60.97	807.1
7	20	1.000	0.000	0.000	59.56	805.3
7	21	1.000	0.000	0.000	63.09	829.7
8	22	0.000	1.000	0.000	77.33	797.3
8	23	0.000	1.000	0.000	76.02	799.1
8	24	1.000	0.000	0.000	65.80	729.7
9	25	0.998	0.002	0.000	76.02	767.7
9	26	0.999	0.001	0.000	75.76	760.2
9	27	0.999	0.000	0.001	76.55	769.0
10	28	1.000	0.000	0.000	73.08	779.3
10	29	1.000	0.000	0.000	73.45	780.0
10	30	1.000	0.000	0.000	72.28	774.3
11	31	1.000	0.000	0.000	68.61	775.9
11	32	1.000	0.000	0.000	70.68	785.1
11	33	1.000	0.000	0.000	71.06	804.1
12	34	0.000	0.000	1.000	80.71	690.7
12	35	1.000	0.000	0.000	72.49	715.9
12	36	1.000	0.000	0.000	71.50	679.7

Table B2. Posterior probabilities, quantities and profits in treatment $Qq\pi$

group	subject	P_{fef}	P_{msf}	P_{fo}	quantity	profits
1	1	0.997	0.003	0.000	70.82	758.1
1	2	1.000	0.000	0.000	69.71	762.1
1	3	0.000	1.000	0.000	74.38	777.0
2	4	0.320	0.680	0.000	74.31	750.5
2	5	1.000	0.000	0.000	69.06	727.7
2	6	1.000	0.000	0.000	73.04	747.2
3	7	0.000	1.000	0.000	101.72	601.4
3	8	0.000	0.000	1.000	78.80	565.5
3	9	0.000	1.000	0.000	97.09	574.8
4	10	0.000	1.000	0.000	94.72	540.1
4	11	0.000	1.000	0.000	95.66	546.8
4	12	0.000	1.000	0.000	98.86	543.5
5	13	0.000	1.000	0.000	98.55	525.4
5	14	0.000	1.000	0.000	98.02	517.5
5	15	0.000	1.000	0.000	95.97	513.3
6	16	0.000	0.000	1.000	76.94	709.5
6	17	1.000	0.000	0.000	64.55	693.5
6	18	0.000	1.000	0.000	82.40	777.1
7	19	0.000	1.000	0.000	92.83	656.4
7	20	0.000	1.000	0.000	90.93	658.3
7	21	0.840	0.000	0.160	81.50	601.9
8	22	0.819	0.181	0.000	58.39	832.7
8	23	1.000	0.000	0.000	57.30	833.8
8	24	0.819	0.181	0.000	58.90	835.3
9	25	0.014	0.986	0.000	80.49	689.1
9	26	0.000	1.000	0.000	82.70	700.6
9	27	0.000	1.000	0.000	84.08	685.5
10	28	1.000	0.000	0.000	87.65	564.6
10	29	0.000	1.000	0.000	103.84	591.6
10	30	1.000	0.000	0.000	90.01	573.6
11	31	0.000	0.001	0.999	86.24	772.3
11	32	0.995	0.000	0.005	75.20	682.3
11	33	1.000	0.000	0.000	75.31	704.7