

Imitation and Belief Learning

in an

Oligopoly Experiment

Theo Offerman, Jan Potters and Joep Sonnemans*

First version: February 1998
Current version: September 2000

Abstract

We examine the force of three types of behavioral dynamics in quantity-setting triopoly experiments: (1) mimicking the successful firm, (2) rules based on following the exemplary firm, and (3) rules based on belief learning. Theoretically, these three types of rules lead to the competitive, the collusive, and the Cournot-Nash outcome, respectively. In the experiment we employ three information treatments, each of which 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, behavioral rules, and outcomes.

JEL codes: L13, C73, 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/CentER
Warandelaan 2
5000 LE Tilburg
Netherlands

Tel: 31-20-5254294
Fax: 31-20-5255283
E-mail: theo@fee.uva.nl
j.j.m.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, Cars Hommes, Georg Kirchsteiger, and Rob Moir, in particular. Very constructive comments were provided by three referees. The research is made possible by a grant from the Netherlands' Organization for Scientific Research (to Theo Offerman) and a fellowship from the Royal Netherlands' Academy of Arts and Sciences (to Jan Potters).

1. INTRODUCTION

Although quantity-setting oligopoly is one of the ‘workhorse models’ of industrial organization (Martin, 1991), empirically there is much ambiguity about its outcome. 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.¹ 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.^{2,3}

In principle, players could spend effort to find the Cournot-Nash equilibrium of the game and play accordingly. Decision costs may prevent them from doing so, however. Given that it is costly to derive best response strategies and to deduce the corresponding Cournot-Nash equilibrium, it may be reasonable to assume that players look for strategies that save on cognitive effort and decision costs.⁴ Deviations from best response behavior could thus explain the deviations from the Cournot-Nash outcome that are observed in empirical and experimental work.

A starting point for our study is a 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

¹ For other stylized games it is often shown that learning may play a decisive role in the process of equilibrium selection (e.g., Brandts and Holt, 1992; Van Huyck, Battalio and Beil, 1990).

² 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.*, 1998. Experimental work on imitation focuses on the role of imitation in decision tasks, see Offerman and Sonnemans, 1998; 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.

³ The 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.

⁴ There is a growing awareness in both empirical and theoretical work that decision costs are a relevant ingredient of decision making (e.g. Abreu and Rubinstein, 1988; Conlisk, 1980; Conlisk, 1988; Conlisk, 1996; Ho and Weigelt, 1996; Pingle and Day, 1996; Rosenthal, 1993).

the other firms. This approach thus provides a dynamic behavioral underpinning for Walrasian equilibrium.

The result is counter-intuitive 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 to mimic the 'villain'. In the social psychological literature, it is also common to hypothesize that cooperative or altruistic acts serve as examples for others (Bandura, 1986; Sarason *et al.*, 1991). In the oligopoly game, the 'exemplary firm' produces the quantity that, if followed by all firms, leads to higher profits than following any of the other firms. The exemplary firm is typically the firm that produces the lowest quantity. We will consider three rules in the 'following class'. The rule 'follow the exemplary firm' assumes that all firms imitate the choice of the exemplary firm. The rule 'follow the exemplary other firm' assumes that firms follow the firm that is exemplary among the other firms. The rule 'follow and guide' assume that the non-exemplary firms follow the exemplary firm, but that the exemplary firm strikes a balance between its own choice and that of the next exemplary firm. Hence, these rules differ (only) in the degree to which the exemplary firm conditions its choices on those of the non-exemplary firms. As will be shown in the next section, the industry will evolve toward the collusive outcome when firms use either of these rules from the following class.

The belief learning approach provides a completely different underpinning of behavior in the quantity-setting oligopoly game. According to this approach firms form beliefs about the choices of other firms on the basis of the assumption that the history of the game provides information about the future actions of the others. Given their beliefs firms myopically choose the quantity with the highest expected profits. In our experiment we choose the structure and parameters of the market such that both best response learning and fictitious play 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. How can the theoretical results be reconciled with the wide range of empirical findings? As Shubik (1975) notes, "... the specific details of communication, information, and the mechanisms of the market have considerable influence on the play when numbers are few". It may be that specific details of the environment trigger a specific dynamic rule of behavior, and that a specific dynamic rule leads to a specific benchmark. In particular, the type of feedback about the choices and performance of other firms may affect the decision costs and hence the use of alternative strategies. 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 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 need information about the quantity chosen by each other firm. Without fully individualized feedback about quantities, firms are unable to detect the choice of either the

successful firm or 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. We refer to this treatment as treatment Q (Q is mnemonic for aggregate information). By excluding the possibility of imitation, we hypothesize this treatment to be relatively favorable for belief learning.

In the other two treatments we provide more feedback information than in treatment Q. This may stimulate the employment of rules based on imitation which are cognitively less demanding than belief learning. In treatment Qq π , firms receive individualized information feedback on both quantities and profits (q is mnemonic for individualized quantity information and π for profit). This information makes comparative profit appraisals available at minimal decision cost. Therefore, we hypothesize this treatment to be relatively favorable for mimicking the successful firm.

In treatment Qq firms receive individualized information about quantities but not about profits. So, players can imitate but it will not be obvious whom to imitate. Mimicking the successful firm is no longer for free and requires some costly deliberations. Thus, relative to treatment Qq π , we hypothesize that treatment Qq encourages the use of rules based on following.⁵

The remainder of the paper is organized as follows. The next section presents the institutions of the market and presents the theoretical results. Section 3 describes the design and procedure of the experiment. Section 4 elaborates on the hypothesized relation between feedback information and dynamic rule of conduct. Section 5 presents the results and, finally, section 6 provides a concluding discussion.

2. MARKET AND THEORY

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 $C(q_i)$:

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

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 $P(Q)$:

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

where a and b denote given demand parameters.

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 :

⁵ 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.

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

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

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

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

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

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

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

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

In the experiment we chose $n=3$, $c=1$, $a=45$ and $b=\sqrt{3}$. For each of the three theoretical benchmarks Table 1 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’ at least three firms should be present. Moreover, with three firms in the industry it is easy to conclude whether a change in the output of the firm with the intermediate production level is better organized by the model that predicts that this firm mimics the successful firm or by a model from the following class. With more than three firms it might be less clear whether a firm with an intermediate production level imitates one of the successful or exemplary firm, or whether it imitates yet another 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 (see Theocharis, 1960).

Table 1
Benchmarks of the game

benchmark	quantities of three firms q_i ($i=1,2,3$)	price $P(Q)$	profits of three firms π_i ($i=1,2,3$)
Walrasian equilibrium	100	15	500
collusive outcome, (a) real quantities:	56.25	22.5	843.75
(b) integer quantities:	56	22.55	843.74
Cournot-Nash equilibrium	81	18	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 (or one of the firms) with the highest profit in the previous period. The successful firm imitates by repeating its former choice.

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

Result 1 is due to Vega-Redondo (1997).⁶ Sketches of the proofs of the results in this section are provided in Appendix A. The result indicates that in the long run, only the Walrasian outcome will be observed a significant fraction of the time. The intuition is as follows. Assume that firms produce different quantities. If the price is above the marginal costs of the high quantity firm, this will be the firm with the highest profit. Imitation then induces the other firms to select this high quantity in the next period. The firms remain at this symmetric, say sub-Walrasian, allocation until one of the firms experiments. If this firm selects a lower quantity, it will earn a lower profit and will return to the other firms in the next period. If, however, it selects

⁶ 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.

a higher quantity it will earn a higher profit (as long as price is above marginal cost) and will be mimicked by the other firms in the next period. Gradually this process moves toward the Walrasian outcome. The Walrasian outcome is more robust than other symmetric outcomes, because a unilateral deviation by one of the firms will not be imitated by other firms.

The other rules based on imitation assume that firms follow the firm that sets the good example from the perspective of industry profit. The exemplary firm is the firm (one of the firms) with the quantity that would give the highest sum of profits if it were followed by all firms. All following rules considered here assume that the non-exemplary firms choose the exemplary quantity. Since following rules are not standard in the literature, we consider three variations to this rule which differ in the degree to which the exemplary firm conditions its choice on those of the non-exemplary firms. According to ‘follow the exemplary firm’ rule the exemplary firm will continue to produce the same output as in the previous period. According to ‘follow and guide’, the exemplary chooses a quantity midway between its own previous choice and that of the exemplary other firm, that is, the firm that is exemplary if its own choice is disregarded. According to ‘follow exemplary other firm’, the exemplary firm follows the choice of the exemplary other firm (just as the non-exemplary firms do). By conditioning its choice on those of the other firms, the exemplary firm can protect itself from being exploited. By applying a mild punishment, it can still send a signal to head toward a more cooperative outcome. The different rules strike a different balance between these two motivations.

Besides the following part, the rules consist of an experimentation part. We will say that a firm uses one of the following rules if it experiments with common and independent probability ϵ and it chooses the predicted quantity with probability $(1-\epsilon)$.

Result 2: If firms choose a rule from the following class, the unique stochastically stable state of the process is (q^c, \dots, q^c) as $\epsilon \rightarrow 0$.

The intuition for the result is easy. After one round of following the non-exemplary firms will follow the firm that was exemplary in the previous period. The exemplary firm itself may have moved away toward a higher quantity (under follow and guide or follow exemplary other), but the next round of following will bring it back to the quantity choices of the other two firms, who are now exemplary. Since the latter two firms are exemplary for each other, they will stay where they are. From this symmetric outcome, they will only move away if one of the firms experiments and chooses a quantity which is closer to the collusive outcome. Eventually, this process will lead to the collusive outcome.

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 it myopically maximizes its expected payoff. 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.

We supplement the belief learning rules with an experimentation part. A firm uses a rule from the belief learning class if it experiments with common and independent probability ϵ and it chooses the predicted quantity with probability $(1-\epsilon)$. Belief learning does not generally converge toward the Cournot-Nash outcome in quantity-setting oligopoly games. Our oligopoly game is dominance solvable, however, which ensures that it does. Experimentation can occasionally take the outcome away from Cournot-Nash, but if the experimentation probability is small the process will spend most of its time at the equilibrium.

Result 3. If firms use a rule from the belief learning class, the unique stochastically stable state of the process in the triopoly market described above is (q^N, \dots, q^N) as $\epsilon \rightarrow 0$.

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 limit states. To that purpose we designed three information treatments, which are discussed in the next section. In section 4 it will be argued that each treatment is particularly conducive to the employment of one of the three dynamic rules.

3. DESIGN AND PROCEDURES

Both the instructions and the experiment were computerized.⁷ A transcript of the instructions is provided in Appendix B. Subjects could read the instructions at their own pace. It was explained to subjects that they made decisions for their ‘own firm’, while two other subjects made the decisions for ‘firm A’ and ‘firm B’. It was also 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.⁸

Table 2 presents the main features of the three treatments. The design is ‘nested’ in the sense that the treatments can be strictly ordered on the amount of information provided to the subjects. In all treatments, the players received feedback information on total quantity, price, private revenue, private cost, and private profit. In fact, in Treatment Q this was all the information that players received in a particular period about the outcome of play in the previous period. In treatment Qq, firms received additional information on the individual quantities produced by the other two firms. Finally, in treatment Qq π , firms were not only provided with individualized information about the quantities but also about the corresponding profits to the other two firms. In the instructions subjects were made familiar with the screens used in the experiment. This was the only point where the three treatments Q, Qq and Qq π and their instructions differed: the screen that popped up after a period had ended, contained different information. Copies of these screens can be found in Appendix B.

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 B contains some copies of these forms of information.

⁷ 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.

⁸ 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). A theoretical advantage of a known final period is that a unique subgame-perfect equilibrium exists for the repeated game. Furthermore, we chose for 100 periods because simulations indicated that this was about the time needed for all the dynamic rules (especially fictitious play) to converge to the corresponding benchmarks.

Table 2
Summary of treatment and predictions

treatment	baseline information	additional information	prediction: favorable for
Q	R_i, C_i, π_i, Q, P	-	belief learning \rightarrow Cournot-Nash
Qq	R_i, C_i, π_i, Q, P	q_i, q_k	following \rightarrow collusion
Qq π	R_i, C_i, π_i, Q, P	q_j, q_k, π_j, π_k	mimic successful firm \rightarrow Walras

Notes: 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.

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 her profit after a period had ended, we did not provide profit tables.⁹ 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.

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.¹⁰ Subjects knew how points were converted into Dutch guilders *before* playing the game. 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: 33 subjects participated in treatment Q, 36 subjects participated in treatment Qq and 33 subjects participated in treatment Qq π . Subjects had no prior experience with directly related experiments, and, of course, no subject participated in more than one experiment. An experiment lasted between 1½ and 2 hours. Average earnings were 55.13 Hfl, which was the Dutch equivalent of about US\$30.

Independent of the present study, Huck *et al.* (1999) conduct an experimental study that also examines the role of imitation in the quantity setting oligopoly game. They carried out five treatments: their treatment BEST is similar to our treatment Q; their treatment FULL is similar to our treatment Qq π . Furthermore, they run three treatments with incomplete information about market conditions. In their

⁹ Of course, subjects had all information necessary to construct such tables themselves if they wanted.

¹⁰ At the time of the experiment 1 dutch guilder was worth about 0.53 US-dollar.

treatment NOIN subjects do not know anything about demand and costs and are only informed about their own realized profits. Their treatment IMIT resembles treatment NOIN, except that subjects receive individualized feedback information about quantities and profits of all firms. Finally, their treatment IMIT+ is the same as their treatment IMIT, except that subjects are informed that demand is downward sloping and that market conditions are identical across sellers and periods.

Apart from 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 latter of which is not considered in Huck *et al.*). In the final section, we discuss how their results complement ours.

4. HYPOTHESES

In this section we elaborate on the hypothesis 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. Also we speculate on how the outcome will be affected if some players in the population will behave more rationally.

First 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 each of the treatments. Treatment Q is conducive to belief learning in the sense that it is the only treatment where firms cannot imitate because they lack the necessary individualized information to do so. As a result we expect that treatment Q will be relatively favorable for outcomes to move toward the Cournot-Nash outcome.

The hypothesis that collusive outcomes will be observed more often in treatment Qq than in treatment Q is not only based on the premise that individualized information on quantities may give firms hints as to what the appropriate quantity may be. More important is that individualized quantity information allows firms to monitor each other's behavior fully, and send and read signals more effectively. This allows for conditional following rules which may facilitate a gradual adjustment toward the collusive outcome. For example, even a small reduction or increase in quantity can be detected as such by other players, whereas this is much more difficult when players receive only aggregate quantity information. As a consequence, we expect relatively

more behavioral strategies aiming for the collusive outcome in treatment Qq than in treatment Q.¹¹

Information-wise, treatments Qq and Qq π are similar in the sense that fully rational agents in treatment Qq could compute the extra information about rivals' profits that is provided in treatment Qq π . Hence, technically speaking, rules based on following as well as mimicking the successful firm can be applied in both treatment Qq π and treatment Qq. However, the cognitive effort required to apply these rules differs between treatment Qq and treatment Qq π . In particular, mimicking the successful firm is a decision rule that can be employed at minimal decision costs in treatment Qq π . Agents receive the information needed for a comparison of profits on a silver plate. In treatment Qq, agents could also mimic the successful firm, but here they would have to spend considerable effort in finding out which of the firms is the successful firm. At the same time, rules based on following are equally costly in treatment Qq and treatment Qq π . In both treatments, firms can detect the exemplary firm at the same cognitive effort. Furthermore, to the extent that it is more natural to evaluate profits at symmetric quantity choices than at asymmetric choices, determining the exemplary firm is cognitively less demanding in treatment Qq than finding the successful firm. Hence, in terms of relative decision costs, treatment Qq is conducive to following and treatment Qq π to mimicking. As result we expect the collusive (Walrasian) outcome to have a relatively strong attraction in treatment Qq (treatment Qq π).

For the theoretical results 1-3 it is assumed that all players follow a specific adaptive rule which is different for each of the three treatments. This assumption provides clear benchmark results, but it seems too extreme to have descriptive relevance. What may be reasonably hypothesized though, is that the probability that a player adheres to a particular rule is inversely related to the decision cost of applying that rule. Belief learning, following, and mimicking can thus be hypothesized to be most relevant in treatment Q, Qq, and Qq π , respectively.

The preceding theoretical arguments assume that *all* players behave according to the same boundedly rational rule. It seems likely though that at least some players will behave more rationally, and that they will try to anticipate rivals' behavioral responses. Now we will tentatively predict what may be the effect when the population consists of a mixture of rational players and decision cost conserving adaptive players (belief learners in treatment Q, followers in treatment Qq, and mimics in treatment Qq π). We will argue that the repeated game aspect and the presence of rational players may reinforce the expected effect toward collusion for treatment Qq. In fact, the repeated nature of the game may to a large extent be an essential ingredient for players' motivation to follow the exemplary firm. On the other hand, the presence of fully rational players may dampen the expected trend toward the Walrasian outcome in treatment Qq π .

First, consider treatment Qq and assume that the rational players anticipate that some players will use one of the following rules as a result of the information provided in this treatment. Rational players may then

¹¹ 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.

find it in their best interest to choose an exemplary (low) quantity to guide the followers toward the collusive outcome. Only in the final periods of the game rational players would then be tempted to defect from the collusive outcome. Thus, to a large extent rational players could decide to play just like followers do. The tendency toward the collusive outcome in treatment Qq may be reinforced by, and may even depend upon, the repeated nature of the game. Second, consider treatment Q and assume that the population consists of a mixture of rational players and myopic belief learners. Rational players would recognize that belief learners cannot be seduced to play more cooperatively. In fact, if a rational player anticipates that she is matched with myopic belief learners, she will find it in her best interest to produce more than the Cournot-Nash quantity, so that the myopic belief learners will be induced to decrease their production. However, a movement toward the collusive outcome cannot be engendered by a rational player in this environment. Finally, also in treatment Qq π a rational player will not be able to exploit the repeated nature of the game if he anticipates that some of his opponents are likely to be mimics of the successful firm. A rational player will be able to draw the outcome somewhat away from the Walrasian outcome by producing less himself, but he cannot induce the mimics to do likewise. In conclusion, if a proportion of the players rationally anticipates the presence and behavior of the others, one may expect that the support for the collusive outcome is reinforced in treatment Qq, but a movement toward collusion cannot be expected in treatments Q and Qq π .

Our repeated game argument in support of collusion in treatment Qq resembles the argument underlying the folk theorem in finitely repeated games with incomplete information (Kreps *et al.*, 1982, Fudenberg and Maskin, 1986). A fraction of irrational players who employ trigger-like strategies (such as tit-for-tat or grim) suffices to support collusive outcomes as an equilibrium. And then only some proportion of such irrational types needs to be present (or believed to be so by the more rational players) in order to induce *all* players to behave in a similar manner. Notice that players who 'follow and guide' or players who 'follow the exemplary other firm' correspond quite closely to irrational types who employ trigger strategies. On the other hand, naive 'followers of the exemplary firm' could be mercilessly exploited by rational players.

Of course, the presence of a small fraction of irrational players using trigger-like strategies could also support collusion in treatment Qq π and treatment Q. After all, trigger strategies can even be employed when there is only aggregate quantity information feedback. As we have argued above, however, in view of the decision costs involved it is more likely that the (irrational) players are of the mimicking type in treatment Qq π and of the belief learning type in treatment Q, rather than of the (conditional) following type. In any case, we can expect that the presence of these other irrational types will hinder the support of collusive outcomes. Hence, if the repeated nature of the game enters the picture because some rational players anticipate the behavior of the boundedly rational others, we hypothesize that the support for the collusive outcome is reinforced in treatment Qq, but not, or at least to a much lesser extent, in treatments Q and Qq π . Whether and to what degree this hypothesis is correct is, of course, an empirical question, and one that is addressed experimentally in the sequel.

5. RESULTS

The experimental results are described in three sections. Section 5.1 focuses primarily on the predictions regarding aggregate and long run outcomes. A first crude dynamic analysis of the data is provided in section 5.2. Section 5.3 contains a detailed dynamic analysis of the various learning rules on the basis of a maximum likelihood procedure.

5.1 Aggregate production levels and long term predictions

If the three treatments trigger different dynamics, as argued in section 4, then one would expect 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, although the difference in production levels between treatments Q and Qq seems to disappear over time.¹² In treatments Q and Qq an end-effect occurs in the last periods. This suggests that at least some subjects were aware of the repeated game aspect.

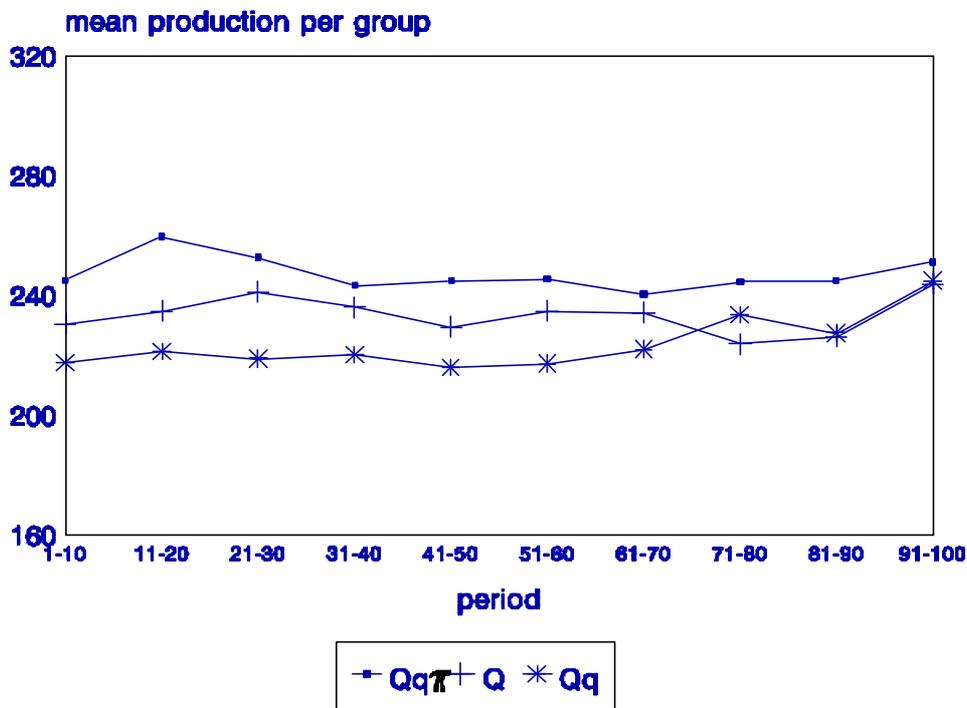


Figure 1.
Average production per block of 10 periods

¹² These results are roughly in line with the results reported by Huck *et al.* (1999). They also find a higher production level in their treatment FULL than in treatment BEST.

Figure 1 does not display much of a time trend in the overall averages. This raises the question whether there is any dynamics in the underlying group data at all. To address this question we examine for each group of three firms by how much the aggregate quantity changes from one period to the next. Figure 2 displays the ‘running frequencies’ of groups' shifts in quantities between successive periods. The picture is roughly the same for each of the treatments so we present them in combination. As would be expected, the production levels change more in the early rounds of the experiment than in the final rounds. For example, about 60% of the groups experience a shift in production by at least 11 units from one period to the next in the early periods of the experiment. By period 90 this is reduced to about 35% of the groups. Moreover, despite the fact that production averages across groups remain more or less constant, there appears to exist substantial adjustments in the production levels of the individual groups, justifying a deeper analysis of the underlying dynamics.

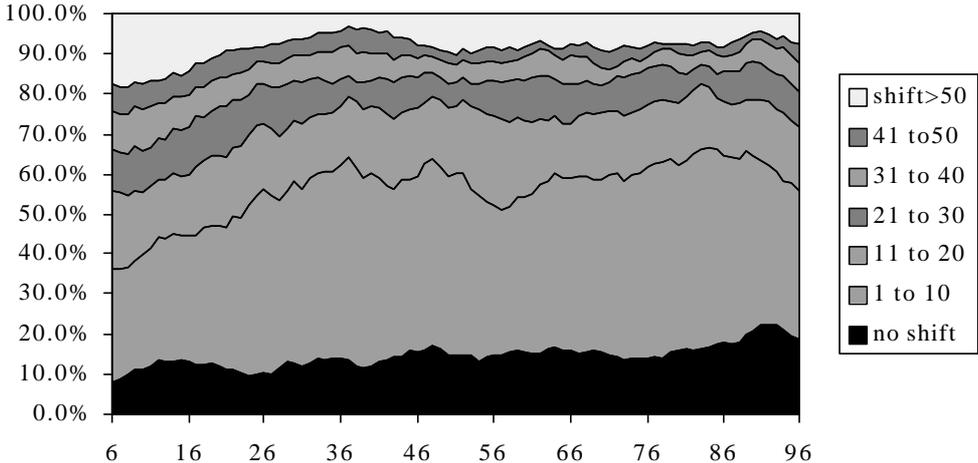


Figure 2.
Running frequencies of groups' shifts between successive periods (all treatments).

notes: For each period t on the horizontal axis, the vertical axis gives the percentage of groups in each category in periods $\{t-4, \dots, t+4\}$.

From figure 1 it appears that although differences in average production levels between the treatments are in the expected direction, they are rather small. Focusing the attention on the averages is not particularly informative in this experiment however, because it masks some interesting patterns in the data. Figure 3 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 $Qq\pi$. As hypothesized, the mode of the distribution is found at the Walrasian outcome of 300 units. The distribution also has a top at the collusive outcome. Despite the fact that the average quantity produced is in the neighborhood of the Cournot-Nash quantity in this treatment, there is no mode at the Cournot-Nash equilibrium!

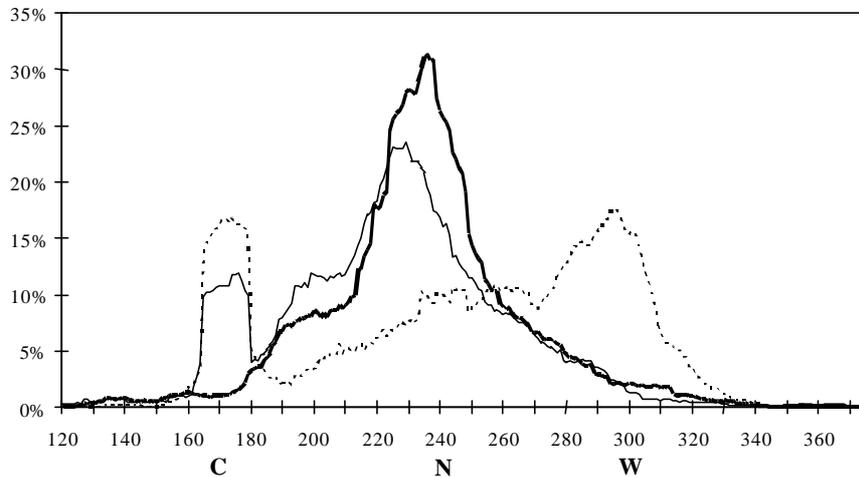


Figure 3.

Running frequencies of aggregate quantities.

notes: For each aggregate quantity Q_{tot} displayed at the horizontal axis the vertical axis reports the percentage of outcomes that falls in the interval $\{Q_{tot} - 7, \dots, Q_{tot} + 7\}$.

Basically, when we divide the data into three blocks of periods [6-35], [36-65] and [66-95] the resulting pictures per block are qualitatively similar to the picture sketched in figure 3 on the basis of all periods: in treatment Q the production-levels remain concentrated around a single peak at the Cournot-Nash outcome; in treatment Qq the production-levels remain bimodal around the collusive outcome and the Cournot-Nash outcome; and in treatment $Qq\pi$ the production-levels remain bimodal around the collusive and the Walrasian outcome. There are two noteworthy trends over time though. First, in treatment Q the peak at the Cournot-Nash outcome grows over time, and the dispersion of the outcomes becomes smaller. Second, the peak at the collusive outcome in treatment $Qq\pi$ starts at a low level of 12% in the first block of periods, and grows via 19% in the second block to 23% in the third block of periods. Perhaps in this treatment some subjects switch to collusion supporting rules once they have experienced the rather disappointing payoffs at the Walrasian outcome.

Another way to look at the data is to compare the relative frequencies of a benchmark across the treatments. In accordance with the predictions, figure 3 shows that the Cournot-Nash outcome is reached more often in treatment Q than in both treatments Qq and $Qq\pi$. These differences are significant even with a conservative testing procedure: for each group of three firms it is counted how often the aggregate production in an industry is in the neighborhood of the Cournot-Nash outcome ($Q_{tot} \in [236, 250]$). There are significantly more Cournot-Nash outcomes in treatment Q than in treatment Qq (Mann-Whitney U test: $n_Q=11$, $n_{Qq}=12$, $p=0.04$). The same test applied to treatments Q and $Qq\pi$ leads to a similar result ($n_Q=11$,

$n_{Qq\pi}=11$, $p=0.04$). Likewise, a similar testing procedure leads to the hypothesized conclusion that the number of Walrasian outcomes ($Q_{tot}\in[293, 307]$) is significantly higher in treatment $Qq\pi$ than in both treatments Q and Qq (Mann-Whitney U tests: $p=0.01$ for $Qq\pi$ versus Q and $p=0.00$ for $Qq\pi$ versus Qq). Finally, as expected the number of collusive outcomes ($Q_{tot}\in[161,175]$) is higher in treatment Qq than in treatment Q , though it is smaller in treatment Qq than in treatment $Qq\pi$. However, neither of these differences is significant at conventional levels.¹³

In the treatments that allows for imitation (Qq and $Qq\pi$) there seems to be more than one focal point. This may (partly) explain the fact that the differences between the averages tend to be small. 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).¹⁴

Table 3 focuses specifically on the long term outcomes. It selects the groups that have roughly stabilized in periods 60-90.¹⁵ A group is defined to have weakly stabilized if on average the mean absolute deviation of each subject's quantity from her average quantity is less than or equal to 8. There are 23 (out of 34) groups that match this weak definition and in the table they are ranked according to the mean group quantity. Walrasian outcomes (W) are only observed in treatment $Qq\pi$. Stable outcomes in the neighborhood of the collusive output (C) are primarily observed in treatments Qq and $Qq\pi$. In treatment Q converging groups tend to be drawn toward the Cournot-Nash outcome (N), although some groups' points of attraction deviate in the direction of the collusive output (C/N). The picture is by and large the same if a stronger definition for stabilization is used (see bold-faced entries).

¹³ All these test results are similar if the attention is restricted to the last 30 periods, except that the first test result comparing the frequency of Cournot-Nash outcomes between treatments Q and Qq ceases to be significant.

¹⁴ The 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.* (1999).

¹⁵ Periods 91-100 are discarded because, as mentioned before, some of the groups display an end round effect in this final part of the experiment.

Table 3
Long term outcomes

% groups converging	treatment Q		treatment Qq		treatment Qq π	
	weakly	strongly	mean production	outcome	mean production	outcome
	63.6%	45.5%			72.7%	36.4%
groupnr:	mean production	outcome	mean production	outcome	mean production	outcome
8					171.3	C
29			175.2	C		
1					177.8	C/N
21	190.3	C/N				
32			213.9	C/N		
30			216.4	C/N		
9					219.0	C/N
15	224.1	C/N				
24			224.8	C/N		
19	227.6	C/N				
31			230.0	C/N		
22	232.3	C/N				
25			235.8	N		
16	239.6	N				
13	239.7	N				
18	240.1	N				
28			243.9	N		
11					248.5	N
27			254.4	N/W		
7					277.0	N/W
10					279.7	N/W
4					293.0	W
5					303.3	W

Notes: A group is defined to weakly (strongly) converge if in periods 61-90 the mean absolute deviation of a subject's production to her average production level averaged over the group is smaller than or equal to 8 (4). In case of weak convergence, the group's average production level in periods 61-90 is displayed, and the outcome is labeled as C(ollusion), N(ash) or W(alras) if it is within 7 units from the collusive (168), Cournot-Nash (243), or Walras (300) outcome, respectively. The entries of strongly converging groups are bold. Outcomes between collusion and Cournot-Nash are labeled C/N, and outcomes between Cournot-Nash and Walras as N/W. The left column displays the group identification number. Groups are ordered on the mean average production of their convergence point.

5.2 Crude dynamic analysis

The previous section suggested which outcomes may actually have functioned as behavioral equilibria in the three markets. It did not discuss the behavioral dynamics in the markets. The goal of the present section is to provide a crude dynamic analysis of the data. 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 a firm in the industry in the previous period is less than or equal to 8 units. Table 4 indicates which outcomes may have served as behavioral equilibria in the three treatments.¹⁶

Table 4
Symmetric rest points

previous own choice in category	treatment Q	treatment Qq	treatment Qq π
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

Notes: The first number in a cell displays the number of symmetric outcomes, that is, outcomes in which the difference between the maximum and the minimum production in an industry is smaller than or equal to 8. 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.

We hypothesized on the basis of theoretical arguments that the Cournot-Nash outcome would primarily be expected in treatment Q, the collusive outcome in treatment Qq and the Walrasian outcome in treat-

¹⁶ Of course, this analysis involves a number of arbitrary choices. The general picture does not depend on these choices, though.

ment $Qq\pi$. The table suggests that the collusive outcome is a stable rest-point in treatments Qq and $Qq\pi$, but not in treatment Q . In the imitation treatments the collusive outcome is reached quite 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 seems to be less stable because firms tend to abandon their previous choice quite often. In treatment $Qq\pi$, on the other hand, the Cournot-Nash outcome is relatively stable but is not reached very often. The Walrasian outcome is only reached quite often in treatment $Qq\pi$. It seems to be a relatively unstable point though, since firms stay at their former choice in only 37% of the cases. It seems that the behavioral dynamics in treatment $Qq\pi$ are such that aggregate play is continuously attracted by the Walrasian outcome. Once firms have actually reached this outcome though, they seem to realize that it is not very profitable. This stimulates experimentation away from the Walras quantity.

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. Table 5 shows that, in line with the hypotheses, middle firms tend to follow the exemplary firm in treatment Qq , whereas they tend to mimic the successful firm in treatment $Qq\pi$.

Table 5
Adaptation process of middle firms

	middle firm decreases production	middle firm production constant	middle firm increases production	N
treatment Qq	41.5%	28.3%	30.2%	643
treatment $Qq\pi$	29.0%	27.7%	43.4%	452

Notes: 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. N represents the number of relevant cases.

5.3 Detailed dynamic analysis

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 $(\mathbf{m}_{t,i}, \mathbf{s})_{[40,125]}$, where $\mathbf{m}_{t,i}$ is set equal to the prediction of a particular model for firm i in period t , and \mathbf{s} is a free parameter. We are not particularly interested whether this assumption itself is accurate, and *a priori* there is no reason to expect that the assumption about the experimentation process will favor one model relative to another.

To reiterate, the model ‘follow the exemplary firm’ predicts that each firm chooses a quantity that is

equal to the quantity produced by the firm that was exemplary in the previous period. The models ‘follow and guide’ and ‘follow exemplary other’ make the same predictions for the firms that were not exemplary in the previous period. The model follow exemplary other predicts that the exemplary firm in the previous period imitates the choice of the most exemplary among the other two firms. The follow and guide model predicts that the exemplary firm chooses the average of its own production in the previous period and that of the exemplary other firm. 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 earned 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 preceding models all assume that subjects learn in one way or another. However, whether subjects learn or not is something to be decided by the data. Therefore, we also include three benchmark models: the ‘collusion plus noise’ predicts that a firm always chooses 56, the ‘Cournot-Nash plus noise’ model predicts constant choices of 81 and the ‘Walras plus noise’ model predicts constant choices of 100. Finally we compare the results with a model that sets a lower bound on the expected performance: the ‘random model’. This model 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 $s \in \mathcal{S}$ in one of the models, the random model is obtained.

Table 6 reports the maximum likelihood results for the models described above. Within the class of collusion supporting models, the learning model follow and guide outperforms both the other learning models and the collusion plus noise benchmark: in both treatments Qq and $Qq\pi$ the follow and guide model yields the better likelihood and the lower estimate for s . By choosing a production level above their own previous level but below the level of the exemplary other firm, an exemplary firm provides a signal that it wants to move toward the collusive outcome but only if at least one other firm also moves in that direction.¹⁷ Within the class of Walras supporting models, the mimic successful firm model organizes the data much better than the Walras plus noise model, both in treatments Qq and $Qq\pi$. In the class of Cournot-Nash supporting models, the fictitious play model is better than the best response model in all three treatments. However, both learning models are outperformed by the static Cournot-Nash plus noise model. This suggests that subjects do not update their beliefs as supposed by either best response or fictitious play models.

In view of these results, the remainder of this section will focus on the follow and guide model as the candidate of the collusion supporting models, the mimic successful firm model as the candidate of the Walras supporting models and the Cournot-Nash plus noise model as the candidate of the Cournot-Nash supporting

¹⁷ Note that the follow and guide rule prescribes the exemplary firm to provide a mild punishment only. We also looked at models that involve a stronger punishment by the exemplary firm. For instance, one such rule prescribes the exemplary firm to choose the average production level of the other firms in the previous period. Such rules perform worse than the follow and guide

models. Each of these three models explains the data significantly better than the random model 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 Cournot Nash plus noise model. If one of the imitation models would provide a better fit of the data than the Cournot-Nash plus noise model in treatment Q, then we could conclude that the Cournot-Nash plus noise model does a very bad job because it is outperformed by a model that cannot be true. However, we find that the Cournot-Nash plus noise model beats the imitation models both on the criterion of likelihood and on the criterion of the estimate \mathbf{s} . In treatments Qq and Qq π firms had in principle sufficient information to behave in accordance with each of the models. The follow and guide model outperforms the mimic successful firm model in treatment Qq both on the likelihood and the estimate of \mathbf{s} . The ordering is reversed in treatment Qq π . There the mimic successful firm model outperforms the follow and guide model on the basis of the likelihood, though the follow and guide model yields the somewhat lower estimate for \mathbf{s} .

These results are by and large in line with the predictions described in Table 2. The main difference is that belief learning does not seem to be the reason for observing the Cournot-Nash outcome in treatment Q. A static Cournot-Nash model organizes the data better than either of the belief learning models.

Table 4 suggested that the imitation treatments Qq and Qq π 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 any of the three models. In the following, let P_{fg} denote the probability that an arbitrary firm uses the follow and guide model; P_{msf} and P_{nsh} denote the probability that a firm uses the mimic successful firm model and the Cournot-Nash plus noise model respectively. We define $P_{fg}+P_{msf}+P_{nsh}=1$. 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:

$$L(q_{i,2}, \dots, q_{i,100}) = P_{fg} * \prod_{t=2}^{100} L(q_{i,t} / fg) + P_{msf} * \prod_{t=2}^{100} L(q_{i,t} / msf) + P_{nsh} * \prod_{t=2}^{100} L(q_{i,t} / nsh),$$

where $L(q_{i,t} / fg)$, $L(q_{i,t} / msf)$ and $L(q_{i,t} / nsh)$ denote the conditional probability that firm i chooses quantity $q_{i,t}$ at period t when it follows and guides, 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 chooses the Cournot-Nash production level in a noisy way, respectively.

rule. Exemplary firms appear to employ relatively mild punishments indeed.

Table 6
Maximum likelihood estimates simple models

	treatment Q n=3267	treatment Qq n=3564	treatment Qq π n=3267
<i>collusion supporting models</i>			
collusion plus noise	-logL=14232.3 σ =36.40 (0.81)	-logL=15382.6 σ =32.49 (0.61)	-logL=14513.9 σ =66.40 (3.90)
follow the 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)
follow and guide	-logL=13936.4 σ =23.12 (0.42)	-logL=15052.1 σ =22.47 (0.37)	-logL=13901.7 σ =23.12 (0.41)
follow exemplary other firm	-logL=13987.4 σ =23.85 (0.45)	-logL=15170.0 σ =23.48 (0.40)	-logL=13949.9 σ =24.16 (0.44)
<i>Walras supporting models</i>			
Walras plus noise	-logL=14501.4 σ =51.95 (2.67)	-logL=15875.4 σ =653.24 (850.57)	-logL=14469.5 σ =45.65 (1.87)
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)
<i>Cournot-Nash supporting models</i>			
Cournot-Nash plus noise	-logL=13736.5 σ =17.16 (0.26)	-logL=15326.8 σ =20.27 (0.35)	-logL=14271.2 σ =24.24 (0.55)
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

Notes: 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 parentheses.

The likelihood functions for three general models that allow for two dynamic rules each are obtained by setting either P_{fg} , P_{msf} or P_{nsh} 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 6. Thus the models are related to each other via a nested structure.

Table 7 reports the maximum likelihood results for these models when applied to the imitation

treatments. We do not estimate the models for treatment Q, because subjects did not have the necessary information to imitate in this treatment. 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 best in treatment Qq π on the criterion of the likelihood. The estimate for \mathbf{s} is a bit smaller for model 3, however. The results for the two-rules models are clearer in treatment Qq. Here, model 2 outperforms the other models both on the criterion of the likelihood and the estimate of \mathbf{s} . These results nicely complement the data reported in table 4.

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 \mathbf{s} . Hence, all three behavioral rules (mimicking, following and guiding, and static equilibrium play) add significantly to explaining observed behavior. As hypothesized though, the degree to which they do varies substantially with the information feedback.

Table 7
Maximum likelihood estimates general models imitation treatments

	treatment Qq n=3564	treatment Qq π n=3267
model 1, 2 rules: (a) follow and guide (b) mimic successful firm ($P_{msf}=1-P_{fef}$)	-logL=14867.5 $\sigma=19.94$ (0.31) $P_{fg}=0.78$ (0.07)	-logL=13357.4 $\sigma=17.67$ (0.27) $P_{fg}=0.45$ (0.09)
model 2, 2 rules: (a) follow and guide (b) Cournot-Nash plus noise ($P_{nsh}=1-P_{fg}$)	-logL=14817.0 $\sigma=18.37$ (0.28) $P_{fg}=0.63$ (0.08)	-logL=13690.3 $\sigma=19.21$ (0.32) $P_{fg}=0.71$ (0.08)
model 3, 2 rules: (a) mimic successful firm (b) Cournot-Nash plus noise ($P_{nsh}=1-P_{msf}$)	-logL=15121.6 $\sigma=19.01$ (0.30) $P_{msf}=0.20$ (0.07)	-logL=13461.7 $\sigma=17.50$ (0.29) $P_{msf}=0.66$ (0.08)
model 4, 3 rules: (a) follow and guide (b) mimic successful firm (c) Cournot-Nash plus noise ($P_{nsh}=1-P_{fg}-P_{msf}$)	-logL=14780.9 $\sigma=18.33$ (0.28) $P_{fg}=0.63$ (0.08) $P_{msf}=0.11$ (0.05)	-logL=13283.8 $\sigma=16.76$ (0.26) $P_{fg}=0.39$ (0.09) $P_{msf}=0.52$ (0.09)

Notes: 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 parentheses.

The results for model 1 are especially interesting when one tries to find out which firm is imitated in

the quantity setting oligopoly game. Providing firms information about individualized quantities without corresponding profits stimulates firms to follow the exemplary firm and guide the non-exemplary firms: in treatment Qq 78% of the population is estimated to follow and guide 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 Qq π 55% of the population is estimated to mimic the successful firm and the remaining 45% follows and guides.

We do not wish to attach too much weight to the exact numbers in these estimation results though. Apart from the fact that the number of subjects is limited and that we only study one particular market structure, we have restricted our attention to a (plausible) subset of the potential learning rules with clear long term predictions. Nevertheless, the estimated differences between the treatments are substantial and they clearly indicate that the type of feedback information affects the behavioral rules that are adopted by the subjects.

6. CONCLUSION

In this paper we set out to test whether different information feedback triggers different behavioral rules which in turn lead to different outcomes in a quantity setting oligopoly experiment. 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. When in addition information on the corresponding individualized profits is provided (treatment Qq π), the frequency distribution is also bi-modal, with 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 rest point in treatment Q. In treatments Qq and Qq π , however, it turns out that the main candidate to serve as a behavioral rest point is the collusive outcome. Although the Cournot-Nash and the competitive outcome are reached quite often in treatments Qq and Qq π , respectively, they are abandoned at a much higher rate than the collusive outcome. A maximum likelihood procedure corroborates the hypothesized relationship between information treatment and dynamic rule of conduct to a large extent. The best fit is given by ‘following the exemplary firm and guiding the non-exemplary firms’ in treatment Qq, and by ‘mimicking the successful firm’ in treatment Qq π . In treatment Q, however, belief learning models are outperformed by static equilibrium play (plus noise).

Mimicking the successful firm could be thought of as a low cognitive effort rule in this game. Following the exemplary firm and guiding non-exemplary firms is a higher cognitive effort algorithm. To follow requires finding out the exemplary firm, which is never part of the immediately available information. And to guide involves at least some basic strategic reasoning. Belief learning and equilibrium play seem to

provide the highest cognitive challenge. These require a firm to compute or estimate the best response function, which is not so easy for the present game. As information conditions deteriorate from treatments $Qq\pi$, to Qq , to Q , subjects' behavior tends to become more sophisticated, from mimicking, to following, to equilibrium play. It is as if subjects try to offset the deterioration of information by more sophisticated behavior. Perhaps the presence of easy-to-use (but suboptimal) clues like information about rivals' profits discourages people to think deeply about the decision problem.

The success of the collusive outcome especially in treatment Qq but also in treatment $Qq\pi$ may, at least partly, be due to the repeated nature of the game. It seems likely that some of the players who follow the exemplary firm and guide the non-exemplary firms are motivated to do so by the fact that the game is repeated. They may hope that following catches on, leading to a movement towards the collusive outcome. In future work it would be interesting to examine the extent to which following depends on the repeated nature of the game. This could be done by running an experiment in which players are rematched with different rivals after each period. Possibly this would increase the success of the other theoretical benchmarks, Walras and Cournot-Nash.

A complementary approach is adopted in Huck *et al.* (1999). They show that the market outcome becomes more competitive when information about demand and cost conditions deteriorates. Thus, an intriguing pattern emerges. On the one hand, when information about the realized profits of competitors deteriorates (like going from our treatment $Qq\pi$ to treatment Qq), the competitive outcome is observed less frequently. On the other hand, when information about market conditions deteriorates (as studied in Huck *et al.*) the opposite effect is observed. This pattern can be explained in terms of the decision costs and availability of alternative behavioral rules. On the one hand, when information about competitors' profits deteriorates, mimicking the successful firm becomes more costly relative to some other rules, like following the exemplary firm and guiding non-exemplary firms. This discourages mimicking of the successful firm, diminishing the force towards the competitive outcome. On the other hand, when information about market conditions deteriorates, cognitively more costly rules (e.g., belief learning or following) are no longer available or become even more costly to implement. This increases the relative attractiveness of imitating the successful firm, thus increasing the force towards competitive outcomes.

Finally, 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. It may then 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 is generally 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 mimics. 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 and guiding non-exemplary firms 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 this perspective it is not so straightforward what 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.
- Abreu, D., and A. Rubinstein (1988): 'The Structure of Nash equilibrium in Repeated Games with Finite Automata', *Econometrica*, 56, 1259-1281.
- 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.
- Brandts, J., and C.A. Holt (1992): 'An Experimental Test of Equilibrium Dominance in Signaling Games', *American Economic Review*, 82, 1350-1365.
- Cheung, Y.W., and D. Friedman (1997): 'Individual Learning in Normal Form Games: Some Laboratory Results', *Games and Economic Behavior*, 19, 46-76.
- Conlisk, J. (1980): 'Costly Optimizers versus Cheap Imitators', *Journal of Economic Behavior and Organization*, 1, 275-293.
- Conlisk, J. (1988): 'Optimization Cost', *Journal of Economic Behavior and Organization*, 9, 213-228.
- Conlisk, J. (1996): 'Why Bounded Rationality?', *Journal of Economic Literature*, 34, 669-700.
- 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.
- Freidlin, M.I., and A.D. Wentzel (1984): *Random Perturbations of Dynamical Systems*. New York: Springer Verlag.
- Fudenberg, D., and E. Maskin (1986): 'The Folk Theorem in Repeated Games with Discounting or with Incomplete Information', *Econometrica*, 54, 533-554.
- Ho, T., and K. Weigelt (1996): 'Task Complexity, Equilibrium Selection, and Learning: An Experimental Study', *Management Science*, 42, 659-679.
- 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 (1999): 'Learning in Cournot Oligopoly - An Experiment', *Economic Journal*, 109, C80-C95.
- Kreps, D.M., P. Milgrom, J. Roberts, and R. Wilson (1982): 'Rational Cooperation in the Finitely Repeated Prisoners' Dilemma', *Journal of Economic Theory*, 27, 245-252.
- Martin, S. (1993): *Advanced Industrial Economics*. Oxford: Blackwell Publishers.
- Milgrom, P., and J. Roberts (1982): 'Predation, Reputation, and Entry Deterrence', *Journal of Economic Theory*, 27, 280-312.
- 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 (1998): 'Quantal Response Models in Step-Level Public Good Games', *European Journal of Political Economy*, 14, 89-100.
- Offerman, T., and J. Sonnemans (1998): 'Learning by Experience and Learning by Imitating Successful Others', *Journal of Economic Behavior and Organization*, 34, 559-575.
- 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.
- Pingle, M. and R. Day (1996): 'Modes of Economizing Behavior: Experimental Evidence', *Journal of Economic Behavior and Organization*, 29, 191-209.
- Rosenthal, R.W. (1993): 'Rules of Thumb in Games', *Journal of Economic Behavior and Organization*, 22, 1-14.
- 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.
- Shubik, M. (1975): 'Oligopoly Theory, Communication and Information', *American Economic Review*, 65,

280-283.

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.

Theocharis, R.D. (1960): 'On the Stability of the Cournot Solution on the Oligopoly Problem', *Review of Economic Studies*, 27, 133-134.

Van Huyck, J.B., R.C. Battalio, and R.O. Beil (1990): 'Tacit Coordination Games, Strategic Uncertainty, and Coordination Failure', *American Economic Review*, 80, 234-248.

Vega-Redondo, F. (1997): 'The Evolution of Walrasian Behavior', *Econometrica*, 65, 375-384.

Young, H. P. (1993): 'The Evolution of Conventions', *Econometrica*, 61, 57-84.

Appendix A

This appendix provides sketches of the proofs of the results in section 2.

Result 1 (for details see Vega-Redondo, 1997).

First note that only symmetric states ($q_1=q_2=q_3$) can be limit states of the imitation dynamics. One round of imitation will bring about a symmetric state from any state with asymmetric quantity choices. When players arrive at a symmetric state, it is obvious that imitation will keep them there. Hence, any symmetric outcome could potentially be a limit state of the imitation process alone. The imitation dynamics, however, is supplemented with mutation dynamics: with a strictly positive probability ϵ , a player does not imitate, but randomly chooses a quantity according to some probability distribution with full support. These mutation dynamics ‘resolve’ the indeterminacy that characterizes the imitation dynamics.

The centerpiece of the proof is that for each $q \neq q^W$ and $1 \leq k \leq n$ ($n=3$ in our experiment), the following inequality holds:

$$P((n-k)q+kq^W)q^W - C(q^W) > P((n-k)q+kq^W)q - C(q) \quad (A1)$$

To see this, first note that:

$$[P(nq^W) - P((n-k)q+kq^W)]q > [P(nq^W) - P((n-k)q+kq^W)]q^W \quad (A2)$$

This follows from the fact that P is decreasing, so that the price difference in the inequality is negative when $q < q^W$ and positive when $q > q^W$. In either case the inequality holds. Now rewrite (A2) as follows:

$$P(nq^W)q + P((n-k)q+kq^W)q^W > P(nq^W)q^W + P((n-k)q+kq^W)q \quad (A3)$$

Subtracting $C(q)+C(q^W)$ from either side and rearranging, we obtain:

$$[P((n-k)q+kq^W)q^W - C(q^W)] - [P((n-k)q+kq^W)q - C(q)] > [P(nq^W)q^W - C(q^W)] - [P(nq^W)q - C(q)] \quad (A4)$$

From the definition of q^W , the Walrasian equilibrium, it follows that the RHS of inequality (A4) is no smaller than zero. This ensures that the LHS of (A4) is larger than zero, which is equivalent to (A1).

Inserting $k = 1$ in inequality (A1) tells us that if one firm produces q^W and the other firms produce some other quantity q , then the firm producing q^W will earn the highest profit. Hence, from any symmetric state in which the firms produce some quantity $q \neq q^W$, it will take only one suitable mutation (plus one round of imitation) to move the state to the Walrasian outcome.

Inserting $k = n-1$ in inequality (A1) tells us that if all firms except one produce the Walrasian quantity, then the firms producing q^W will earn a higher profit than the firm producing some other quantity q . Hence, one round of imitation will bring the latter firm back to the Walrasian quantity. In other words, starting from the Walrasian outcome, one mutation will not be enough to move the process away from this outcome. At least, two mutations are necessary to escape the Walrasian outcome.

We have seen that (a) arriving at a symmetric outcome involves only imitation and no mutation, (b) arriving at q^W from any other symmetric outcome involves only one (suitable) mutation, and (c) escaping

the Walrasian outcome requires at least two mutations. With the help of the graph-theoretic techniques of Freidlin and Wentzel (1984), these facts can be used to show that with $\varepsilon \rightarrow 0$, the Walrasian outcome is the unique stochastically stable state.

Result 2

For the dynamics induced by each of the following rules only symmetric states can be limit states. Starting from any state, after one round (under ‘follow the exemplary firm’) or two rounds (under ‘follow and guide’ or ‘follow the exemplary other firm’) the following dynamics will bring about a symmetric state. At the same time the probability of leaving a symmetric state with following alone is zero. The introduction of a small probability (ε) of mutation singles out (when $\varepsilon \rightarrow 0$) the unique stochastically stable state, namely the collusive outcome in which all firms produce q^C (= 56 in our experiment).

The arguments on which this claim rests are similar to the ones for result 1, but easier to establish. The rules ‘follow the exemplary firm’, ‘follow and guide’ and ‘follow the exemplary other firm’ all assume that the non-exemplary firms follow the exemplary firm. Hence, if at least one firm produces q^C , then - no matter what the other firms produce - the firm (or one of the firms) producing q^C will be the exemplary firm. One round of following will then ensure that at least two firms produce at the collusive output. In the next round they will be exemplary for each other and remain at the collusive outcome, and - should it have moved away - the third firm will return to the collusive output. Hence, only one suitable mutation is needed to arrive at the collusive state. Second, starting from the collusive outcome, suppose one firm mutates and produces some other quantity $q \neq q^C$. Then the firms still producing q^C are the exemplary firms, and the mutant will return to the collusive output. Hence, at least two mutations are needed to escape the collusive outcome. These features ensure that the collusive outcome is the unique stochastically stable state of the following plus mutation dynamics as $\varepsilon \rightarrow 0$.

Result 3¹⁸

Best response dynamics and fictitious play are special cases of the adaptive learning models analyzed in Milgrom and Roberts (1991). Roughly speaking, a learning process is consistent with adaptive learning if each player eventually chooses strategies that are a best response to some beliefs over competitors' strategies, where near zero probability is assigned to strategies that have not been played for a sufficiently long time. Milgrom and Roberts show that a process that is consistent with adaptive learning converges to the serially undominated set, that is, the set of strategies that survive iterated elimination of strongly dominated strategies.

The serially undominated set of the oligopoly game in section 2 is the Cournot-Nash equilibrium. This follows from the fact that the best-reply mapping of the game is a contraction. A sufficient condition for a

¹⁸ The proof was suggested by one of the referees.

contraction (see, e.g., Vives, 1999) is

$$\frac{\partial^2 \mathbf{p}_i}{(\partial q_i)^2} + \sum_{i \neq j} \frac{\partial^2 \mathbf{p}_i}{\partial q_i \partial q_j} / < 0.$$

Since $\partial^2 \pi_i / (\partial q_i)^2 = P'' q_i + 2P' - C''$, and $\sum_{i \neq j} |(\partial^2 \pi_i) / (\partial q_i \partial q_j)| = -2P'' q_i - 2P'$ this condition requires that $-P'' q_i - C'' < 0$, which follows immediately from $P'' > 0$ and $C'' > 0$.

This result cannot be directly applied though, since the subjects in our experiment choose from the finite strategy set $\{40, 41, \dots, 125\}$. It turns out, however, that also on the finite grid the Nash equilibrium is the unique serially undominated strategy profile. The best response to $Q_{-i}=80$ is $q_i=102$ and the best response to $Q_{-i}=250$ is $q_i=60$. Given that the profit function is strictly concave and that $80 \leq Q_{-i} \leq 250$, the first round of elimination of dominated strategies leaves us with the quantity choices $60 \leq q_i \leq 102$. In the second round we can use the restriction $120 \leq Q_{-i} \leq 204$ to leave us with the undominated quantity choices $70 \leq q_i \leq 92$. Repeating this process leaves us with $q_i = 81$ after 8 steps of elimination.¹⁹

Just as in the imitation models we can add noise (experimentation) to the learning process by assuming that in each round with probability ε a player randomly chooses some quantity from the feasible set according to a probability distribution with full support. Theorem 4 in Young (1993) can then be used to establish that the Cournot-Nash outcome is the unique stochastically stable state of the perturbed belief learning process.

¹⁹ A complication, however, is that in the experiment payoffs are rounded to two decimals (cents), as a consequence of which the payoff function is not strictly concave. It turns out that this halts the elimination process after step 7, at which point the surviving set of strategies still consists of $q_i \in \{80, 81, 82\}$. Hence, for practical purposes this is the serially undominated set to which adaptive learning processes will converge.

Appendix B.

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 Qq π , 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.

Translation of the computer screen

<p style="text-align: center;">Period: <input style="width: 50px;" type="text" value="4"/></p> <p style="text-align: center;">Total earnings: <input style="width: 50px;" type="text" value="2356.67"/> points.</p>	<p>Your decision:</p> <p>Production: <input style="width: 50px;" type="text"/></p> <p>Type your production quantity and press Enter.</p>
---	---

The results of the previous period:			
Your production:	76	The price:	19.19
Production of others:	146	Your revenue:	1458.44
Total production:	222	Your costs:	662.55
		Your earnings:	795.89

Results window Q

The results of the previous period:			
Your production:	76	The price:	19.19
Production of others:	146	Your revenue:	1458.44
Total production:	222	Your costs:	662.55
		Your earnings:	795.89

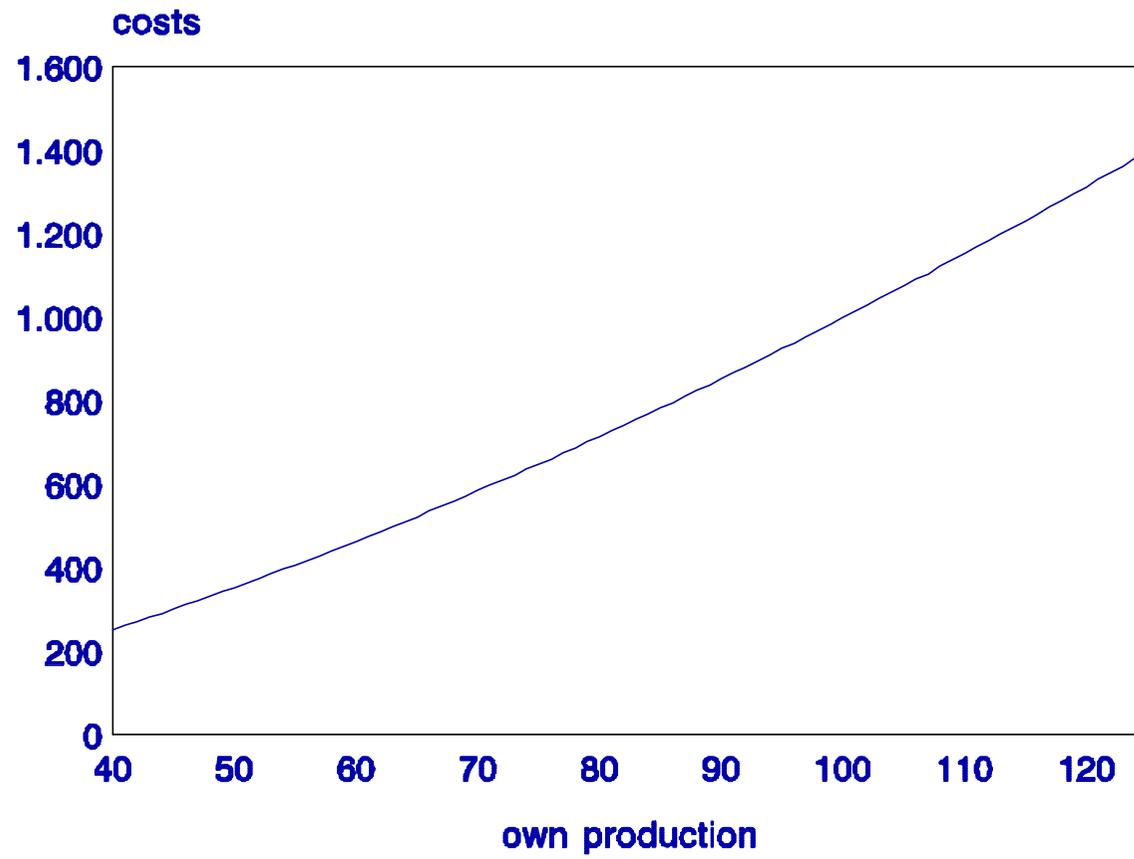
Results window Qq

The results of the previous period:			
Production:		Total production:	222
Your:	76	The price:	19.19
A:	62	Your revenue:	1458.44
B:	84	Your costs:	662.55
		Your earnings:	795.89

Results window Qqπ

The results of the previous period:			
Production:	Earnings:	Total production:	222
Your:	76	The price:	19.19
A:	62	Your revenue:	1458.44
B:	84	Your costs:	662.55
	795.89		
	701.59		
	842.09		

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



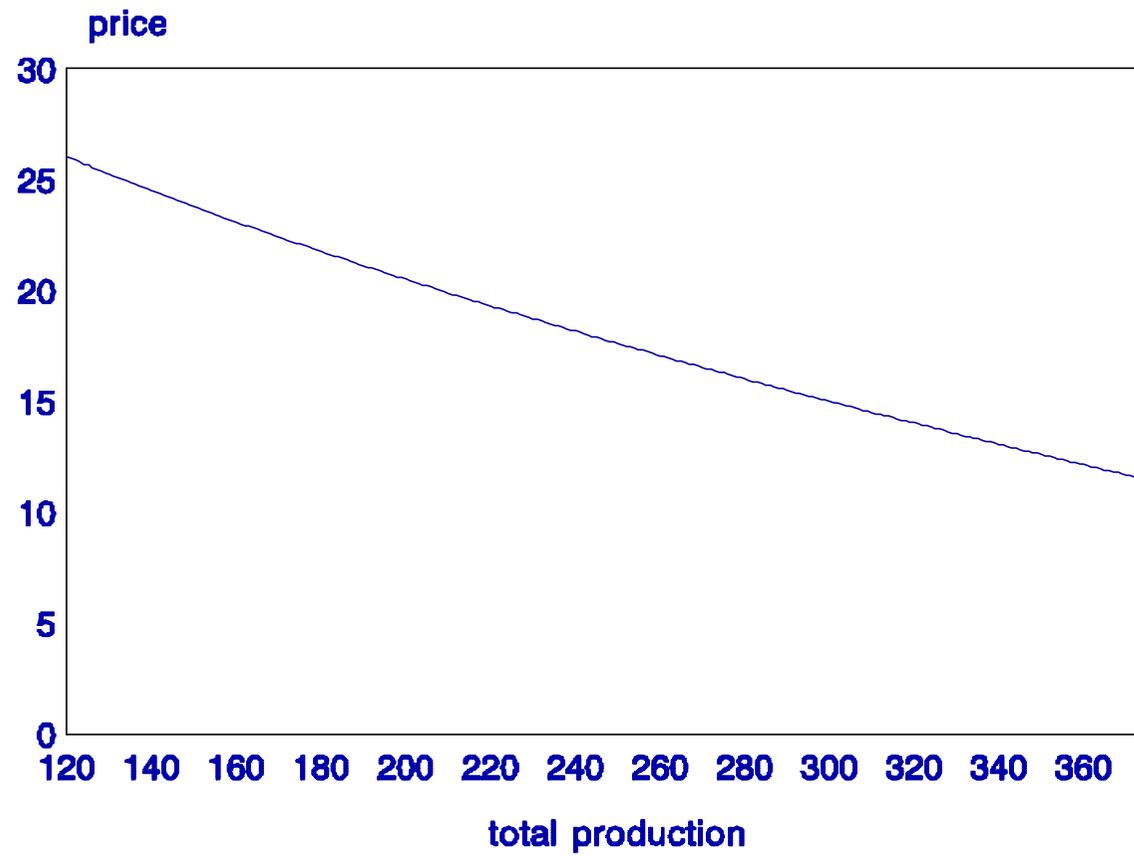
Cost-form: costs of your own production

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.