Imitation and Belief Learning in an Oligopoly Experiment

THEO OFFERMAN

University of Amsterdam

JAN POTTERS

Tilburg University

and

JOEP SONNEMANS

University of Amsterdam

First version received September 1998; final version accepted November 2001 (Eds.)

We examine the force of three types of behavioural 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, behavioural rules, and outcomes.

1. INTRODUCTION

Although quantity-setting oligopoly is one of the “workhorse models” of industrial organization (Martin, 1993), 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

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)).

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 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 behaviour 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 the other firms. This approach thus provides a dynamic behavioural 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 a very unprofitable outcome. Therefore, as suggested by Sinclair (1990), in the presence of a partial conflict between individual and group 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, Sarason, Pierce and Shearin (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”, which all propose that the non-exemplary firms choose the exemplary quantity. 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 behaviour. 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. Given their beliefs firms myopically choose the quantity with the highest expected profits. For some market structures, like the one we implement in this paper, 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 behavioural 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

2. For experimental work on belief learning see Boylan and El-Gamal (1993), Cheung and Friedman (1997), El-Gamal, McKelvey and Palfrey (1994), Mookherjee and Sopher (1994) and Offerman, Schram and Sonnemans (1998). Experimental work on imitation focuses on the role of imitation in decision tasks, see Offerman and Sonnemans (1998) and 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 centred 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, 1988, 1996), Ho and Weigelt (1996), Pingle and Day (1996), Rosenthal (1993)).

rule of behaviour, 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 following the exemplary firm or to 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 favourable 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 $Qq\pi$, firms receive individualized feedback information on both quantities and profits (q is a 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 favourable 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 $Qq\pi$, we hypothesize that treatment Qq encourages the use of rules based on the following.⁵

In a related paper, Huck, Normann and Oechssler (1999) independently examine the role of imitation in an experimental quantity setting oligopoly game. The emphasis of their paper is on the impact of information *prior* to the experiment, while our paper is primarily concerned with the impact of feedback information provided *during* the experiment. In the concluding section, we will explain how the results of Huck *et al.* complement ours.

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 produce a homogeneous commodity. Firms face identical cost functions:

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

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

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

5. The availability and dissemination of information about the actions of individual members of an industry have also been a matter of some concern to antitrust authorities (Scherer and Ross, 1990, Chapter 9). It is feared that "tacit collusion" is easier when other members can be closely monitored and, if necessary, punished.

TABLE 1
Benchmarks of the game

Benchmark	Quantities of three firms $q_i (i = 1, 2, 3)$	Price $P(Q)$	Profits of three firms $\pi_i (i = 1, 2, 3)$
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

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 :

$$\begin{aligned}
 P(nq^w)q^w - C(q^w) &\geq P(nq^w)q - C(q), \quad \forall q \geq 0; \\
 \Rightarrow q^w &= \frac{4a^2}{(2b\sqrt{n} + 3c)^2}.
 \end{aligned}$$

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

$$\begin{aligned}
 P(nq^c)q^c - C(q^c) &\geq P(nq)q - C(q), \quad \forall q \geq 0; \\
 \Rightarrow q^c &= \frac{4a^2}{(3b\sqrt{n} + 3c)^2}.
 \end{aligned}$$

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

$$\begin{aligned}
 P(nq^N)q^N - C(q^N) &\geq P((n-1)q^N + q)q - C(q), \quad \forall q \geq 0; \\
 \Rightarrow q^N &= \frac{4a^2}{\left(\frac{b(2n+1)}{\sqrt{n}} + 3c\right)^2}.
 \end{aligned}$$

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.⁶ 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.

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 probability ε and it imitates the successful firm with probability $(1 - \varepsilon)$. When

6. 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 firms, 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)).

a firm experiments, it chooses some arbitrary quantity level from the feasible set.⁷ 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 $\varepsilon \rightarrow 0$.

Result 1 is due to Vega-Redondo (1997).⁸ 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 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 firm 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 probability ε and it chooses the predicted quantity with probability $(1 - \varepsilon)$.

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 $\varepsilon \rightarrow 0$.

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),

7. Here and in what follows the Results 2 and 3, it is assumed that the probability that a firm experiments is independent across players and time. Furthermore, if a firm experiments, it chooses a quantity level according to a given stationary probability distribution with full support (see Vega-Redondo (1997), Young (1993)).

8. 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.

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 behaviour. Belief learners trust the stability in the choice pattern of 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 an adapted version of 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.⁹

We supplement the belief learning rules with an experimentation part. A firm uses a rule from the belief learning class if it experiments with probability ε and it chooses the predicted quantity with probability $(1 - \varepsilon)$. 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 $\varepsilon \rightarrow 0$.

3. DESIGN AND PROCEDURES

Both the instructions and the experiment were computerized.¹⁰ 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.¹¹

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,

9. According to the “true” version of 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.

10. 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.

11. 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, Mitzkewitz and Uhlig (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 treatments

Treatment	Baseline information	Additional information
Q	R_i, C_i, π_i, Q, P	—
Qq	R_i, C_i, π_i, Q, P	q_j, q_k
$Qq\pi$	R_i, C_i, π_i, Q, P	q_j, q_k, π_j, π_k

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.

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 $Qq\pi$, firms were not only provided with individualized information about the quantities but also about the corresponding profits to the other two firms. This was the only point where the three treatments Q , Qq and $Qq\pi$ and their instructions differed: the screen that popped up after a period had ended, contained different information (see 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. Copies of these forms of information will be sent on request.

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.¹² 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. 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 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 $Qq\pi$. Subjects had no prior experience with directly related experiments, and, of course, subjects participated in the experiment only once. An experiment

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

lasted between 1 1/2 and 2 h. Average earnings were 55.13 Hfl, which at the time was the Dutch equivalent of about U.S.\$30.

4. HYPOTHESES

In this section we elaborate on the hypothesis that treatments Q , Qq and $Qq\pi$ 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 favourable 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 based on the premise that individualized information on quantities may give firms hints as to what the appropriate quantity may be. Unlike in treatment Q , firms have the possibility to employ either unconditional or conditional rules that aim to identify and follow the exemplary firm. This treatment may especially foster the use of the more plausible conditional following rules, as the individualized quantity information allows firms to send and read signals more effectively, thereby decreasing the likelihood that they will be unilaterally exploited. 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 behavioural strategies aiming for the collusive outcome in treatment Qq than in treatment Q .¹³

Information-wise, treatments Qq and $Qq\pi$ 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 $Qq\pi$. Hence, technically speaking, rules based on following as well as mimicking the successful firm can be applied in both treatment $Qq\pi$ and treatment Qq . However, the cognitive effort required to apply these rules differs between treatment Qq and treatment $Qq\pi$. Mimicking the successful firm is a decision rule that can be employed at minimal decision costs in treatment $Qq\pi$. 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 identifying the successful firm. At the same time, rules based on following are equally costly in treatment Qq and treatment $Qq\pi$. In both treatments, firms can detect the exemplary firm at the same cognitive effort. 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 $Qq\pi$ to mimicking. As a result, we expect the collusive (Walrasian) outcome to have a relatively strong attraction in treatment Qq (treatment $Qq\pi$).

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

13. 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.

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 $Qq\pi$, respectively.

There is a possibility that at least some players will behave rationally, and that they will try to anticipate and guide rivals' behavioural 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 $Qq\pi$).

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 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. A movement toward the collusive outcome cannot be engendered by a rational player in this environment. Finally, also in treatment $Qq\pi$ 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 behaviour 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 $Qq\pi$.¹⁴ Whether and to what degree this hypothesis is correct is, of course, an empirical question, and one that is addressed experimentally in what follows.

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.

14. 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, Milgrom, Roberts and Wilson (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. 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. If in treatments Q and $Qq\pi$ the (irrational) players are of the mimicking type and of the belief learning type, a similar argument *cannot* be made to support the collusive outcome.

FIGURE 1
Average production per block of ten periods

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. 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 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. The production levels change more in the early rounds 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.

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 neighbourhood of the Cournot–Nash quantity in this treatment, there is no mode at the Cournot–Nash equilibrium!

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 $\{t - 4, \dots, t + 4\}$

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 $\{Q_{\text{tot}} - 7, \dots, Q_{\text{tot}} + 7\}$. C, N and W refer to the collusive outcome, the Cournot–Nash outcome and the Walrasian outcome, respectively

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. 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 out-

come is reached 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 neighbourhood of the Cournot–Nash outcome ($Q_{\text{tot}} \in [236, 250]$). 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 ($Q_{\text{tot}} \in [293, 307]$) is significantly higher in treatment $Qq\pi$ than in both treatments Q and Qq (Mann–Whitney U -tests: $p = 0.01$ for $Qq\pi$ vs. Q and $p = 0.00$ for $Qq\pi$ vs. Qq). Finally, as expected, the number of collusive outcomes ($Q_{\text{tot}} \in [161, 175]$) is higher in treatment Qq than in treatment Q , though it is smaller in treatment Qq than in treatment $Qq\pi$. Neither of these differences is significant at conventional levels.¹⁵

In the treatments that allow 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).¹⁶

Table 3 focuses specifically on the long-term outcomes. It selects the groups that have roughly stabilized in periods 60–90.¹⁷ 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 neighbourhood 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).

5.2. Crude dynamic analysis

The previous section suggested which outcomes may actually have functioned as behavioural equilibria in the three markets. It did not discuss the behavioural 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 behavioural equilibria in the three treatments. This analysis involves a number of arbitrary choices but the general picture is not affected.

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

15. 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.

16. 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).

17. 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 Weakly Strongly	Treatment Q 63.6% 45.5%		Treatment Q 66.7% 25.0%		Treatment $Qq\pi$ 72.7% 36.4%	
	Mean production	Outcome	Mean production	Outcome	Mean production	Outcome
Group No.						
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 labelled 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 labelled 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 production of their convergence point.

TABLE 4
Symmetric rest points

Previous own choice in category	Treatment Q	Treatment Qq	Treatment $Qq\pi$
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.

TABLE 5
Adaptation process of middle firms

	Middle firm decreases production (%)	Middle firm production constant (%)	Middle firm increases production (%)	<i>N</i>
Treatment <i>Qq</i>	41.5	28.3	30.2	643
Treatment <i>Qqπ</i>	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.

Walrasian outcome in treatment *Qqπ*. The table suggests that the collusive outcome is a stable rest-point in treatments *Qq* and *Qqπ*, 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 *Qqπ*, 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 *Qqπ*. 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 behavioural dynamics in treatment *Qqπ* 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 *Qqπ*.

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 $(\mu_{i,t}, \sigma)_{[40,125]}$, where $\mu_{i,t}$ is set equal to the prediction of a particular model for firm *i* in period *t*, and σ 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 favour one model relative to another.

To reiterate, the models in the collusion supporting class all predict that the non-exemplary firms choose the quantity produced by the exemplary firm in the previous period. The “follow the exemplary firm” model makes the same prediction for the exemplary firm itself. 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 model “mimic the successful firm” predicts each firm to choose a quantity equal to the quantity produced by the firm that earned the highest profit in the previous period. The prediction of the “best response” model is that each firm chooses a best response to the aggregated quantity produced by the other two firms in the previous period. The “fictitious play” model predicts each firm to choose 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 $\sigma \rightarrow \infty$ 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 σ . 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.¹⁸ 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 treatments. However, the static Cournot–Nash plus noise model outperforms both learning models. This suggests that subjects do not update their beliefs as supposed by either best response or fictitious play models.

As a result, the remainder will focus on the follow and guide model, the mimic successful firm model and the Cournot–Nash plus noise model as the candidates of the collusion supporting, the Walras supporting and the Cournot–Nash supporting models, respectively. Each of these models fits 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 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 σ . In treatments Qq and $Qq\pi$ 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 σ . The ordering is reversed in treatment $Qq\pi$. 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 a somewhat lower estimate for σ .

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 $Qq\pi$ contain more than one rest-point each. This may reflect heterogeneity in the population. Perhaps one part of the population

18. 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 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 $Qq\pi$ $n = 3267$
<i>Collusion supporting models</i>			
Collusion plus noise	$-\log L = 14,232.3$ $\sigma = 36.40 (0.81)$	$-\log L = 15,382.6$ $\sigma = 32.49 (0.61)$	$-\log L = 14,513.9$ $\sigma = 66.40 (3.90)$
Follow the exemplary firm	$-\log L = 13,949.3$ $\sigma = 23.87 (0.44)$	$-\log L = 15,078.5$ $\sigma = 23.27 (0.39)$	$-\log L = 13,927.2$ $\sigma = 23.78 (0.43)$
Follow and guide	$-\log L = 13,936.4$ $\sigma = 23.12 (0.42)$	$-\log L = 15,052.1$ $\sigma = 22.47 (0.37)$	$-\log L = 13,901.7$ $\sigma = 23.12 (0.41)$
Follow exemplary other firm	$-\log L = 13,987.4$ $\sigma = 23.85 (0.45)$	$-\log L = 15,170.0$ $\sigma = 23.48 (0.40)$	$-\log L = 13,949.9$ $\sigma = 24.16 (0.44)$
<i>Walras supporting models</i>			
Walras plus noise	$-\log L = 14,501.4$ $\sigma = 51.95 (2.67)$	$-\log L = 15,875.4$ $\sigma = 653.24 (850.57)$	$-\log L = 14,469.5$ $\sigma = 45.65 (1.87)$
Mimic successful firm	$-\log L = 14,118.6$ $\sigma = 26.10 (0.54)$	$-\log L = 15,461.7$ $\sigma = 28.42 (0.59)$	$-\log L = 13,842.2$ $\sigma = 23.81 (0.42)$
<i>Cournot–Nash supporting models</i>			
Cournot–Nash plus noise	$-\log L = 13,736.5$ $\sigma = 17.16 (0.26)$	$-\log L = 15,326.8$ $\sigma = 20.27 (0.35)$	$-\log L = 14,271.2$ $\sigma = 24.24 (0.55)$
Best response	$-\log L = 14,084.5$ $\sigma = 21.37 (0.40)$	$-\log L = 15,772.1$ $\sigma = 34.98 (1.29)$	$-\log L = 14,548.9$ $\sigma = 86.58 (16.35)$
Fictitious play	$-\log L = 13,837.6$ $\sigma = 18.09 (0.29)$	$-\log L = 15,687.2$ $\sigma = 28.60 (0.79)$	$-\log L = 14,516.2$ $\sigma = 45.11 (2.73)$
Random	$-\log L = 14,552.4$	$-\log L = 15,875.3$	$-\log L = 14,552.4$

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.

uses one dynamic rule while another part uses another. To investigate this possibility, more general models are 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

$$\begin{aligned} L(q_{i,2}, \dots, q_{i,100}) = & P_{fg} * \prod_{t=2}^{100} L(q_{i,t} \mid fg) + P_{msf} * \prod_{t=2}^{100} L(q_{i,t} \mid msf) \\ & + P_{nsh} * \prod_{t=2}^{100} L(q_{i,t} \mid nsh), \end{aligned}$$

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

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

TABLE 7
Maximum likelihood estimates general models imitation treatments

	Treatment Qq $n = 3564$	Treatment $Qq\pi$ $n = 3267$
Model 1, 2 rules:	$-\log L = 14,867.5$	$-\log L = 13,357.4$
(a) follow and guide	$\sigma = 19.94(0.31)$	$\sigma = 17.67(0.27)$
(b) mimic successful firm ($P_{msf} = 1 - P_{fef}$)	$P_{fg} = 0.78(0.07)$	$P_{fg} = 0.45(0.09)$
Model 2, 2 rules:	$-\log L = 14,817.0$	$-\log L = 13,690.3$
(a) follow and guide	$\sigma = 18.37(0.28)$	$\sigma = 19.21(0.32)$
(b) Cournot–Nash plus noise ($P_{nsh} = 1 - P_{fg}$)	$P_{fg} = 0.63(0.08)$	$P_{fg} = 0.71(0.08)$
Model 3, 2 rules:	$-\log L = 15,121.6$	$-\log L = 13,461.7$
(a) mimic successful firm	$\sigma = 19.01(0.30)$	$\sigma = 17.50(0.29)$
(b) Cournot–Nash plus noise ($P_{nsh} = 1 - P_{msf}$)	$P_{msf} = 0.20(0.07)$	$P_{msf} = 0.66(0.08)$
Model 4, 3 rules:	$-\log L = 14,780.9$	$-\log L = 13,283.8$
(a) follow and guide	$\sigma = 18.33(0.28)$	$\sigma = 16.76(0.26)$
(b) mimic successful firm	$P_{fg} = 0.63(0.08)$	$P_{fg} = 0.39(0.09)$
(c) Cournot–Nash plus noise ($P_{nsh} = 1 - P_{fg} - P_{msf}$)	$P_{msf} = 0.11(0.05)$	$P_{msf} = 0.52(0.09)$

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.

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 1–3) 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 $Qq\pi$ on the criterion of the likelihood. The estimate for σ 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 σ . These results nicely complement the data reported in Table 4.

From a statistical point of view, the results seem most favourable for the most general model allowing for three rules. This model outperforms each of the three two-rules models significantly (likelihood ratio test at 1% level) in both treatments. This model also yields the lowest estimates for σ . Hence, all three behavioural rules (mimicking, following and guiding, and static equilibrium play) add significantly to explaining observed behaviour. As hypothesized though, the degree to which they do varies substantially with the feedback information.

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 $Qq\pi$ 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 feedback information affects the behavioural rules that are adopted by the subjects.

6. CONCLUSION

In this paper we test whether different feedback information triggers different behavioural 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 $Qq\pi$), 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 behavioural rest point in treatment Q . In treatments Qq and $Qq\pi$, however, it turns out that the main candidate to serve as a behavioural rest point is the collusive outcome. Although the Cournot–Nash and the competitive outcome are reached quite often in treatments Qq and $Qq\pi$, respectively, they are abandoned at a much higher rate than the collusive outcome. 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 $Qq\pi$. 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’ behaviour 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 behaviour. 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). 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 between their and our approach. 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 treatment (Qq) which provides feedback information in between complete feedback of individual quantities and profits and no individualized 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.*). Huck *et al.* show that the market outcome becomes more competitive when information about demand and cost conditions deteriorates. Thus, an intriguing pattern emerges: 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 and 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 behavioural 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, it is more costly or even impossible to implement cognitively more expensive rules like belief learning and following. 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 behaviour. From this perspective it is not so straightforward what would be the most successful strategy of imitation: mimic the villain or follow the saint?

APPENDIX A. 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, "resolve" this indeterminacy of the imitation dynamics.

The centrepiece 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:

$$P((n-k)q + kq^W)q^W - C(q^W) > P((n-k)q + kq^W)q - C(q). \quad (\text{A.1})$$

To see this, first note that

$$[P(nq^W) - P((n-k)q + kq^W)]q > [P(nq^W) - P((n-k)q + kq^W)]q^W. \quad (\text{A.2})$$

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 (A.2) as follows:

$$P(nq^W)q + P((n-k)q + kq^W)q^W > P(nq^W)q^W + P((n-k)q + kq^W)q. \quad (\text{A.3})$$

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

$$[P((n-k)q + kq^W)q^W - C(q^W)] - [P((n-k)q + kq^W)q - C(q)] > [P(nq^W)q^W - C(q^W)] - [P(nq^W)q - C(q)]. \quad (\text{A.4})$$

From the definition of q^W , the Walrasian equilibrium, it follows that the R.H.S. of inequality (A.4) is no smaller than zero. This ensures that the L.H.S. of (A.4) is larger than zero, which is equivalent to (A.1).

Inserting $k = 1$ in inequality (A.1) 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 (A.1) 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 3¹⁹

Best response dynamics and fictitious play are special cases of the adaptive learning models analysed in Milgrom and Roberts (1991). Roughly speaking, a learning process is consistent with adaptive learning if each player eventually

19. The proof was suggested by one of the referees.

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

$$\frac{\partial^2 \pi_i}{(\partial q_i)^2} + \sum_{i \neq j} \left| \frac{\partial^2 \pi_i}{\partial q_i \partial q_j} \right| < 0.$$

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 eight steps of elimination.²⁰

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.

APPENDIX B

This appendix presents the instructions and prints of the computer screens that were used in each of the three treatments. The figures and tables that were used to communicate the market structure will be sent on request. 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 firms A and 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.

20. 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.

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 firms A and 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 firms A and 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 one 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 $Qq\pi$, 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
Upper part screen (same for all conditions)

<p style="text-align: center;">Period: <input style="width: 50px; text-align: center;" type="text" value="4"/></p> <p style="text-align: center;">Total earnings: <input style="width: 80px; text-align: center;" type="text" value="2356.67"/> points.</p>	<p>Your decision:</p> <p>Production: <input style="width: 100px;" type="text"/></p> <p style="margin-top: 20px;">Type your production quantity and press Enter.</p>
--	---

Lower part screen: Results window Q

The results of the previous period:															
<table style="width: 100%; border-collapse: collapse;"> <tr> <td style="width: 60%;">Your production:</td> <td style="text-align: right;">76</td> </tr> <tr> <td>Production of others:</td> <td style="text-align: right;">146</td> </tr> <tr> <td>Total production:</td> <td style="text-align: right;">222</td> </tr> </table>	Your production:	76	Production of others:	146	Total production:	222	<table style="width: 100%; border-collapse: collapse;"> <tr> <td style="width: 60%;">The price:</td> <td style="text-align: right;">19.19</td> </tr> <tr> <td>Your revenue:</td> <td style="text-align: right;">1458.44</td> </tr> <tr> <td>Your costs:</td> <td style="text-align: right;">662.55</td> </tr> <tr> <td>Your earnings:</td> <td style="text-align: right;">795.89</td> </tr> </table>	The price:	19.19	Your revenue:	1458.44	Your costs:	662.55	Your earnings:	795.89
Your production:	76														
Production of others:	146														
Total production:	222														
The price:	19.19														
Your revenue:	1458.44														
Your costs:	662.55														
Your earnings:	795.89														

Lower part screen: Results window Qq

The results of the previous period:																			
<table style="width: 100%; border-collapse: collapse;"> <tr> <td colspan="2" style="padding-bottom: 5px;">Production:</td> </tr> <tr> <td style="width: 60%;">Your:</td> <td style="text-align: right;">76</td> </tr> <tr> <td>A:</td> <td style="text-align: right;">62</td> </tr> <tr> <td>B:</td> <td style="text-align: right;">84</td> </tr> </table>	Production:		Your:	76	A:	62	B:	84	<table style="width: 100%; border-collapse: collapse;"> <tr> <td style="width: 60%;">Total production:</td> <td style="text-align: right;">222</td> </tr> <tr> <td>The price:</td> <td style="text-align: right;">19.19</td> </tr> <tr> <td>Your revenue:</td> <td style="text-align: right;">1458.44</td> </tr> <tr> <td>Your costs:</td> <td style="text-align: right;">662.55</td> </tr> <tr> <td>Your earnings:</td> <td style="text-align: right;">795.89</td> </tr> </table>	Total production:	222	The price:	19.19	Your revenue:	1458.44	Your costs:	662.55	Your earnings:	795.89
Production:																			
Your:	76																		
A:	62																		
B:	84																		
Total production:	222																		
The price:	19.19																		
Your revenue:	1458.44																		
Your costs:	662.55																		
Your earnings:	795.89																		

Lower part screen: Results window $Qq\pi$

The results of the previous period:																	
<table style="width: 100%; border-collapse: collapse;"> <tr> <td style="width: 60%;">Production:</td> <td style="width: 40%;">Earnings:</td> </tr> <tr> <td>Your:</td> <td style="text-align: right;">76 795.89</td> </tr> <tr> <td>A:</td> <td style="text-align: right;">62 701.59</td> </tr> <tr> <td>B:</td> <td style="text-align: right;">84 842.09</td> </tr> </table>	Production:	Earnings:	Your:	76 795.89	A:	62 701.59	B:	84 842.09	<table style="width: 100%; border-collapse: collapse;"> <tr> <td style="width: 60%;">Total production:</td> <td style="text-align: right;">222</td> </tr> <tr> <td>The price:</td> <td style="text-align: right;">19.19</td> </tr> <tr> <td>Your revenue:</td> <td style="text-align: right;">1458.44</td> </tr> <tr> <td>Your costs:</td> <td style="text-align: right;">662.55</td> </tr> </table>	Total production:	222	The price:	19.19	Your revenue:	1458.44	Your costs:	662.55
Production:	Earnings:																
Your:	76 795.89																
A:	62 701.59																
B:	84 842.09																
Total production:	222																
The price:	19.19																
Your revenue:	1458.44																
Your costs:	662.55																

Acknowledgements. The authors 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).

REFERENCES

- ABBINK, K. and SADRIEH, A. (1995), "RatImage—Research Assistance Toolbox for Computer-Aided Human Behavior Experiments" (Discussion Paper No. B-325, University of Bonn).
- ABREU, D. and RUBINSTEIN, A. (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 RICHERSON, P. (1991), "Culture and Cooperation", in R. Hinde and J. Groebel (eds.) *Cooperation and Prosocial Behavior* (Cambridge: Cambridge University Press).
- BOYD, R. and RICHERSON, P. (1993), "Rationality, Imitation and Tradition", in R. Day and P. Chen (eds.) *Nonlinear Dynamics and Evolutionary Economics* (Oxford: Oxford University Press).
- BOYLAN, R. T. and EL-GAMAL, M. A. (1993), "Fictitious Play: A Statistical Study of Multiple Economic Experiments", *Games and Economic Behavior*, **5**, 205–222.
- BRANDTS, J. and HOLT, C. A. (1992), "An Experimental Test of Equilibrium Dominance in Signaling Games", *American Economic Review*, **82**, 1350–1365.
- CHEUNG, Y. W. and FRIEDMAN, D. (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., MCKELVEY, R. D. and PALFREY, T. R. (1994), "Learning in Experimental Games", *Economic Theory*, **4**, 901–922.
- FOURAKER, L. E. and SIEGEL, S. (1963) *Bargaining Behavior* (New York: McGraw-Hill).
- FRASER, C. (1994), "Conjectural Variations", in J. Cable (ed.) *Current Issues in Industrial Organizations* (Hampshire: Macmillan).
- FREIDLIN, M. I. and WENTZEL, A. D. (1984) *Random Perturbations of Dynamical Systems* (New York: Springer).
- FUDENBERG, D. and MASKIN, E. (1986), "The Folk Theorem in Repeated Games with Discounting or with Incomplete Information", *Econometrica*, **54**, 533–554.
- HO, T. and WEIGELT, K. (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 J. H. Kagel and A. E. Roth (eds.) *The Handbook of Experimental Economics* (Princeton: Princeton University Press) 349–443.
- HUCK, S., NORMANN, H. T. and OECHSSLER, J. (1999), "Learning in Cournot Oligopoly—An Experiment", *Economic Journal*, **109**, C80–C95.
- KREPS, D. M., MILGROM, P., ROBERTS, J. and WILSON, R. (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 ROBERTS, J. (1982), "Predation, Reputation, and Entry Deterrence", *Journal of Economic Theory*, **27**, 280–312.
- MILGROM, P. and ROBERTS, J. (1991), "Adaptive and Sophisticated Learning in Normal Form Games", *Games and Economic Behavior*, **3**, 82–100.
- MOOKHERJEE, D. and SOPHER, B. (1994), "Learning Behavior in an Experimental Matching Pennies Game", *Games and Economic Behavior*, **7**, 62–91.
- OFFERMAN, T., SCHRAM, A. and SONNEMANS, J. (1998), "Quantal Response Models in Step-Level Public Good Games", *European Journal of Political Economy*, **14**, 89–100.
- OFFERMAN, T. and SONNEMANS, J. (1998), "Learning by Experience and Learning by Imitating Successful Others", *Journal of Economic Behavior and Organization*, **34**, 559–575.
- PARTOW, Z. and SCHOTTER, A. (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 DAY, R. (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 EREV, I. (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., SARASON, B. R., PIERCE, G. R. and SHEARIN, E. N. (1991), "A Social Learning Approach to Increasing Blood Donations", *Journal of Applied Social Psychology*, **21**, 896–918.
- SCHERER, F. and ROSS, D. (1990) *Industrial Market Structure and Economic Performance* (Boston: Houghton Mifflin).
- SELL, J. and WILSON, R. (1991), "Levels of Information and Contributions to Public Goods", *Social Forces*, **70**, 107–124.
- SELTEN, R., MITZKEWITZ, M. and UHLIG, G. R. (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., BATALLIO, R. C. and BEIL, R. O. (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.
- VIVES, X. (1999) *Oligopoly Pricing: Old Ideas and New Tools* (Cambridge, MA and London: MIT Press).
- YOUNG, H. P. (1993), "The Evolution of Conventions", *Econometrica*, **61**, 57–84.