

SYMPOSIUM ON EXPERIMENTAL ECONOMICS (i)

On Experimental Research in Oligopoly^{1, 2}

1. INTRODUCTION

Economics has not generally been regarded as an experimental science, although in recent years this view has changed a little. The first experiment in economics of which this writer is aware was conducted by E. H. Chamberlin [2] and appeared in 1948 in the *Journal of Political Economy*. The next effort, appearing in the same journal several years later, is the utility experiment of Mosteller and Noguee [12]. In the past decade the amount of experimental research in economics has increased markedly, although the volume could by no means be called large. Some examples are the experiments in competitive markets of V. L. Smith [15, 16, 17], the utility experiments of Davidson, Suppes and Siegel [3], Dolbear [4] and Yaari [18], the bilateral monopoly experiments of Fouraker and Siegel [6, 14], and the oligopoly experiments of Hoggatt [10, 11], Fouraker and Siegel [6] and Friedman [7, 8, 9].

This paper is concerned with a detailed review and discussion of a group of oligopoly experiments (Hoggatt [10, 11], Fouraker and Siegel [6] and Friedman [7]). These experiments are all concerned with the same economic question: does standard oligopoly theory (e.g., Cournot's) predict behaviour in markets in which the usual textbook assumptions are met? Though this may seem a useless enterprise (why worry about whether an "unrealistic" theory predicts behaviour in an "unrealistic" situation?), it is far from being so. Indeed all theory is unrealistic by nature. Its value lies in organizing the mind in an insightful way and in giving manageable, comprehensible models which can predict behaviour in a complex world. Oligopoly theory is very hard to test empirically, not only because of the complexities of modern industry, but also because much useful data is owned by firms who cherish their right to privacy.

Experimental testing of oligopoly theory is of interest for several reasons. First, in the laboratory an artificial market may be created which satisfies the assumptions of standard theory (e.g. each firm has only its own quantity as a decision variable, no randomness enters the profit functions, a firm knows its own cost and demand functions, etc.). Surely if theory fails to predict behaviour in a market meeting these assumptions, one has strong evidence the theory is not valid. Second, there is an interaction between theoretical and empirical research. Perhaps experimental results which cannot be explained by existing theory are suggestive of new lines of theoretical development. Finally, as a social scientist, the economist has a general interest in understanding decision making by individuals. Experiments involving human subjects making decisions are relevant to this end.

¹ This research was supported by a Junior Faculty Fellowship from Yale University and also by a grant from the National Science Foundation.

² I wish to thank Professor William Brainard, Michael Farrell, Lawrence Fouraker, Austin Hoggatt and an anonymous referee for helpful comments. None of them, however, can be assumed to approve the final result.

The several experiments reviewed below deal with "non-cooperative" oligopoly, that is, oligopoly where the subjects have no verbal or written communication with one another. The only information a subject has when he makes a decision are the previous decisions he and his rivals have made. He is informed of his rivals' decisions for the current period only after he has made his own.

In discussing the experiments, attention will, of course, focus on the economic hypotheses which they test; however, the experimental procedures will also come under close scrutiny. The procedures are of interest because they bear on the validity of the results, and because experimental techniques are not familiar to most economists.

2. DESCRIPTION OF AND COMMENTS ON THE EXPERIMENTS

2.1. General Characteristics

There are a few general characteristics which are identical in all the experiments, and certain others which differ in very specific ways. Concerning the former: (a) Each subject represented a single firm in a market with two, three or four firms. (b) Each firm

TABLE I

	Hoggatt 1	F-S	Friedman	Hoggatt 2
Were subjects paid their actual profits?	No	Yes	Yes	No
How many games did each subject play?	1	1	6 or 9 sequentially	2 simultaneously
In what form was profit information given?	Algebraic profit function	Payoff matrix	Payoff matrix	Algebraic profit function
What information known on profits?	Incomplete	Some complete, some incomplete	Complete	Complete

had exactly one decision variable. In some experiments this was price, in some, output level. (c) Demand functions were always linear, as were marginal cost curves. (d) In all the experiments, subjects played in "games" which lasted for many "periods".¹ In each period, subjects would make their decisions (choose price or output, as the case may be), each in ignorance of what his rivals were choosing. After all had chosen, each would be informed of the decisions taken by the rivals, and the next period would commence. During a single game, a subject knew that his rivals would be the same identical individuals, and the payoff matrix (or other information) he received would be unchanged. Subjects did not know the identity of their rivals. (e) There was no communication allowed between subjects. (That is, no written messages or conversations.) The only information to accumulate to a subject in the course of a game was the past decisions of the rivals.

Systematic differences between experiments were: (a) whether subjects were paid their actual profits, (b) the length of time subjects spent in the experiment, (c) the form in which market information was given and (d) the information state. The two information states are "complete" and "incomplete". Complete information is the case in which the subject knows the profits to each firm in the industry, corresponding to a given set of choices. Incomplete information describes the condition in which the subject knows only the profits to himself. Table I summarizes these differences. Thus, for example, F-S subjects were paid their actual game earnings, each subject was in only one game, information was provided in the form of a payoff matrix, and both complete and incomplete information states were used in their experiments.

¹ Fouraker and Siegel did have some "one shot games" in which subjects had only one decision period which they played for real stakes. None of these is reviewed in the present paper.

2.2. Hoggatt 1

2.2.1. *Description.* Hoggatt's first experiment [10] employed nine of his University of California (Berkeley) business school colleagues, in three triopolies, as subjects. The games were played under incomplete information with each subject being given the industry demand curve and his own cost curve. The demand and total cost functions had the form:

$$p_t = \alpha_0 - \alpha_1(q_{1t} + q_{2t} + q_{3t}) - \alpha_2(q_{1,t-1} + q_{2,t-1} + q_{3,t-1}) - \alpha_3(q_{1,t-2} + q_{2,t-2} + q_{3,t-2}),$$

$$C_{it} = \beta_{i0} + \beta_{i1}q_{it} + \beta_{i2}q_{it}^2, \quad i = 1, 2, 3,$$

where:

p_t is the industry price in period t ,

q_{it} is the output of the i th firm in period t ,

$\alpha_0 > 0, \alpha_1 > \alpha_2 > \alpha_3 > 0, \beta_{ij} > 0, i = 1, 2, 3, j = 0, 1, 2.$

Thus demand is modelled in a way which causes an increase in the output of period t to have a depressing effect on the prices of the next two periods (p_{t+1} and p_{t+2}) as well as on the current price (p_t).

The three yardsticks against which the behaviour of the firm is measured are [10, page 106]:

Generalized Cournot Behaviour Assumption I: The Manager of each firm assumes on day t that the combined outputs of other firms will be the same in the current period as they were in the previous period. Output is set so that profit on day t will be maximized if the assumption is correct.

Generalized Cournot Behaviour Assumption II: The Manager of each firm assumes that the combined outputs of the other firms will be the same in the current and the next two future periods as they were in the previous period. Output level is set on day t so that if it is constant for two future periods, and if the assumption of fixed opponents' outputs is correct, then profit on day $t+2$ will be maximized.

The third yardstick is joint industry profit maximization.

The three games ran, respectively, for 13, 16 and 20 periods, and the firms chose output levels tolerably close to the short-run Cournot levels (Cournot I, above). This experiment is of great interest because it is the first oligopoly experiment to be reported and it is the only one to date in which there is a behavioural hypothesis in which explicit consideration is given to profits in a future period.

2.2.2. *Comments.* The prime weaknesses of the experiment relate to experimental control and the motivation of the subjects. Concerning the first, subjects were not confined to a laboratory and isolated from one another during the experiment; rather, they made decisions every day or few days, leaving them in a mailbox and picking up their results from another mailbox. As a result there is no experimental control insuring that the subjects did not discuss the experiments with one another, although the results are hardly those one would expect from collusion. Even without collusion, the mere fact that subjects *could* talk to one another, and assorted others, during the course of the experiment adds an uncontrolled element (a) the importance of which is very difficult to assess and (b) which could have been easily controlled by a different design.

These subjects were not paid money in proportion to their profits, which raises the question of what objectives they might have had in the game. Under incomplete information, the subject lacks opportunity to formulate an objective function involving the profits of rivals. Furthermore, with their training as professional economists, they are aware of the more popular oligopoly formulations. Both these influences should incline them toward a simple, Cournot-type of behaviour.

In general, to refrain from paying subjects their actual profits is to refrain from building into the experiment a real-world incentive which is easily duplicated in the laboratory. The unwanted effects on behaviour which may result are best illustrated by considering complete information oligopoly—where subjects know their rival's profit functions. Where profits are not paid, the subject will not cherish a marginal "dollar" of profit as if it were a dollar in his pocket. Indeed the absolute level of his profits is without meaning. The subject will be in a parlour game situation where success is judged by how well he does in relation to his rivals. Compare the subject whose profits are low, but are greater than those of any rival, with the subjects whose profits are twice as high, but lower than those of any rival. The first will regard himself a success, the second a failure. Giving real money payoffs to subjects will moderate, perhaps completely overshadow, this parlour game influence.¹

Generally the behaviour which maximizes profit does not maximize the extent to which one's profit exceeds his rivals'. Indeed pursuit of the latter objective requires a sacrifice of one's profit level.

The experimental situation should speak for itself. As the intent of the experiment is to see whether and in what manner subjects react to a profit incentive in an oligopoly setting, the best experiment is one which provides a real profit incentive in its basic design. Then, subjects are merely given instructions which describe the features of the experiment without reference to how they *should* behave; and the experimenter records and analyzes the resulting behaviour.

The "long-run" Cournot behaviour, designed to maximize profits in the second period hence, is a bothersome notion. It appears highly arbitrary and hard to defend. Why maximize tomorrow's profit and pay no attention to today's? A much more natural notion would be that of maximizing the sum of profits over the current and next two periods, with the firm assuming his rivals will continue indefinitely their outputs of the preceding period, and choosing his current and next two output levels so as to accomplish this maximization (with the firm *not* required to assume it must hold its output constant for three periods).²

In summary, it is difficult to gauge the success of this experiment. In part, its success depends upon how its purpose is viewed. As a demonstration that experimental research is possible in oligopoly, it is surely successful despite the shortcomings which have been noted. The importance of paying profits to subjects is minimized by the incomplete information state, and the problem that subjects could have colluded because they were not confined to a laboratory is not of concern here because Hoggatt's subjects were his colleagues, whom he could surely trust. These results form the first experimental support of the hypothesis that the Cournot solution is typical of behaviour in incomplete information games.

2.3. Fouraker-Siegel

2.3.1. *Description.* Fouraker and Siegel (F-S) performed ten oligopoly experiments and half a dozen bilateral monopoly experiments which are described in their second book [6]. Their first book [14] was concerned only with bilateral monopoly. Their experiments are characterized by the virtues of extreme simplicity of design, a high degree of control and excellent documentation of their procedures and results. The subjects were college undergraduate students, who were paid their actual earnings for participating in the experiment.

Among the 10 oligopoly experiments were eight which exhaust all possible combinations of (a) complete and incomplete information, (b) duopoly or triopoly and (c)

¹ No doubt the desired tradeoff between absolute income and one's position *vis-à-vis* rivals will vary from one person to another. This relationship would be a most interesting one to explore experimentally with the aid of a psychologist.

² More sensible still is the assumption that the firm seeks to maximize a discounted stream of profits extending from the current period until the firm will cease to operate.

Cournot or Bertrand market. The crucial difference between Cournot and Bertrand markets is that the latter have discontinuities in demand for the output of a firm and the former do not.¹ Only the four standard Cournot experiments will be discussed here. All the F-S experiments are based on demand functions of the form:

$$\begin{aligned} p_t &= \alpha_0 - \alpha_1(q_{1t} + q_{2t} + q_{3t}) && \text{for triopoly,} \\ p_t &= \alpha_0 - \alpha_1(q_{1t} + q_{2t}) && \text{for duopoly.} \end{aligned}$$

Costs are nil, so profit for a firm is

$$\begin{aligned} \pi_{it} &= p_t q_{it} && i = 1, 2 && \text{for duopoly,} \\ & && i = 1, 2, 3 && \text{for triopoly.} \end{aligned}$$

Information was given to the subjects in the form of payoff matrices which afforded each a choice of 25 output levels (the integers from 8 to 32). Under incomplete information, payoff matrices gave profit to the subject corresponding to any output he might choose and any total output the rivals might choose. Under complete information, the subject's matrix also gave total profit to the rivals corresponding to the choices which might be made.²

There were 16 incomplete information duopoly games, 11 triopoly, 16 complete information duopoly and 11 triopoly. The precise demand functions were, in cents:

$$\begin{aligned} p_t &= 2.4 - 0.04(q_{1t} + q_{2t}) && \text{for duopoly,} \\ p_t &= 2.4 - 0.04(q_{1t} + q_{2t} + q_{3t}) && \text{for triopoly.} \end{aligned}$$

Costs were assumed nil.

F-S singled out three solution concepts as being of special interest. There are (a) the "Cournot solution" which is found by assuming each firm seeks to maximize its own profit with respect to its own decision variable

$$\frac{\partial \pi_i}{\partial q_i} = 0, \quad i = 1, \dots, n;$$

(b) the "joint maximum" which is found by assuming each firm seeks to maximize total industry profits with respect to its decision variable

$$\frac{\partial \sum_j \pi_j}{\partial q_i} = 0;$$

and (c) the "rivalistic solution" which is found by assuming the firm seeks to maximize the excess of its profit over the average profit of its rivals

$$\frac{\partial}{\partial q_i} \left(\pi_i - \frac{1}{n-1} \sum_{j \neq i} \pi_j \right) = 0.$$

¹ In the Cournot experiments, quantities are the decision variables; however, in the Bertrand case prices are chosen by the firms, with the following understanding: (1) Only the firm charging the lowest price has any sales, (2) his sales are given by the industry demand relationship, and (3) if two or three firms tie for low price they share industry sales equally. Obviously for a firm, a discontinuity in the profit function occurs where its price is equal to the lowest of the rivals. At prices above this level, sales are zero and profit equals zero. At the rival's price, profit jumps to half the profit previously enjoyed by the low priced rival and, if price is lowered just below the rival's profit doubles. The original hint for this model is found in Bertrand [1], and further analysis appears in Edgeworth [5] and Shubik [13]. The Bertrand games will not be discussed further.

² There were two additional duopoly experiments which are of a very strange nature. In one of them the subjects are told that he who makes more profit than his rival shall win an extra \$8.00, over and above his regular game winnings. This extra prize so greatly overshadowed the potential ordinary game winnings that a subject interested in simply maximizing his own gain would behave in the game as if he (apart from the \$8.00 prize) wished to maximize the excess of his profit over the profit of his rival. The second experiment of the pair is intended to induce cooperative behaviour by giving a prize of \$4.00 to each subject if the pair attain a joint profit of approximately the joint maximum. Like the other experiment, this induces the desired behaviour in only the most trivial sense—by making it become indistinguishable from non-cooperative behaviour.

Table II gives the output levels which correspond to these solutions for the models used by F-S. All games in the four experiments ran for 25 periods of which the first three were designated "practice", and for which subjects were paid no profits.

The principal hypotheses tested by F-S in these experiments are: (a) industries and individuals under incomplete information will tend to choose the Cournot solution output levels (40 and 20 for duopoly, and 45 and 15 for triopoly), and (b) the variability of responses under complete information is greater than under incomplete.¹ Apparently the rationaliza-

TABLE II
Output quantity

	Duopoly		Triopoly	
	Firm	Industry	Firm	Industry
Rivalistic (R)	30	60	20	60
Cournot (M)	20	40	15	45
Joint maximum (C)	15 *	30	10 *	30

* The joint maximum is achieved so long as total industry output is 30, irrespective of how the total is allocated among the firms.

tion for the latter is that under incomplete information, those subjects with a rivalistic or cooperative bent are prevented from exercising it due to ignorance of their rival's profit functions.

The hypothesis (a) is confirmed. Using as data the output decisions for the penultimate period, the mean observed output levels are: 20.9 for individual duopolists, 16.0 for individual triopolists, 41.8 for duopoly games, and 48.1 for triopoly games. Also a tabulation was made, reproduced in Table III, in which each individual and each game

TABLE III
Observed cooperativeness

	Individual			Game		
	R	M	C	R	M	C
Complete information duopoly	9½	12½	10	3½	7½	5
Complete information triopoly	15	15	3	6	5	0
Incomplete information duopoly	4	26	2	2	14	0
Incomplete information triopoly	9	20	4	2	9	0

was classified as *R*, *M* or *C* according to whether the decision of the next to last period was nearest the *R*, *M* or *C* output level.

The specific hypotheses on Cournot behaviour (item *a*) were two: (1) As the individual Cournot output for duopoly (20) exceeds that for triopoly (15), while the joint Cournot output for duopoly (40) is less than triopoly (45); it was hypothesized that the observed individual mean for duopoly (20.9) would be significantly larger than the triopoly mean (16.0), and the joint duopoly mean (41.8) would be significantly less than the joint triopoly mean (48.1). These hypotheses were confirmed by Student "t" tests. (2) The second hypothesis is that, in the absence of systematic preferences by subjects, there would be a

¹ F-S also tested hypotheses, not reviewed below, relating to subjects' propensity to prefer outcomes in which all have equal profits. These are discussed under the heading "Level of Aspiration" [6, pp. 151-154].

probability of 1/2 that a given game will (in the next to last period) exhibit total output nearer the Cournot level than either the rivalistic or cooperative. At significance levels of 0.002 for duopoly and 0.035 for triopoly (or greater) it is concluded that observed behaviour is not due to chance, and therefore systematic preference is indicated for the Cournot output.¹

A fault common to both the preceding tests is that they employ the wrong data. The analysis is based on the assumption that Cournot maximizing behaviour by an individual is the same as choosing the Cournot equilibrium point output. This is only so if rivals are expected to choose Cournot equilibrium output also. The same applies to rivalistic and cooperative behaviour. For example, consider Cournot maximizing behaviour for firm 1 in a duopoly. His profit function is

$$\pi_{1t} = q_{1t}(2.4 - 0.04q_{1t} - 0.04q_{2t}),$$

and his estimate of his profit for period t is

$$\pi_{1t}^e = q_{1t}(2.4 - 0.04q_{1t} - 0.04q_{2t}^e),$$

where q_{2t}^e is his estimate of the output level his rival will choose in period t . His estimated profit is maximized when:

$$\begin{aligned} \frac{\partial \pi_{1t}^e}{\partial q_{1t}} &= 2.4 - 0.08q_{1t} - 0.04q_{2t}^e = 0 \\ &= 60 - 2q_{1t} - q_{2t}^e \end{aligned}$$

or, equivalently, when

$$q_{1t} = 30 - 0.5q_{2t}^e.$$

Thus, the for Cournot maximizer, the appropriate output level depends on the expected output level of the rival, and, it is the Cournot equilibrium output of 20 if, and only if, the rival is also expected to choose 20.

The F-S procedure for classifying the decisions of individuals implies that whatever type of behaviour a subject decides to exhibit, he expects his rival to do the same *in the current period*. This makes sense if a subject is first presumed to estimate the behaviour he expects his rivals to exhibit and then mimic that behaviour himself. The latter appears an absurd way to behave, unless the subjects are assumed to have an obsession of making equal profits.

One alternative is to assume the subject takes the preceding period output levels of its rivals as his estimate of what they will do in the current period. Then a Cournot response is one which is nearer the Cournot level than the rivalistic or cooperative, with the three points calculated under the relevant output assumptions. Under this procedure Cournot behaviour for duopoly is given by:

$$q_{1t} = 30 - 0.5q_{2,t-1}.$$

Rivalistic behaviour is given by:

$$\frac{\partial(\pi_{1t} - \pi_{2t})}{\partial q_{1t}} = 2.4 - 0.08q_{1t} = 0,$$

or

$$q_{1t} = 30,$$

¹ F-S did not follow the usual procedure of selecting a significance level in advance of performing tests. Instead, they reported, for each test, the significance level for which the null hypothesis would be marginally accepted.

and co-operative by:

$$\frac{\partial(\pi_{1t} + \pi_{2t})}{\partial q_{1t}} = 2.4 - 0.08q_{1t} - 0.08q_{2,t-1} = 0,$$

or

$$q_{1t} = 30 - q_{2,t-1}.$$

Table IV contains the same information as Table III for individual responses, with the classification made on this basis. A marked increase in Cournot behaviour and decrease in cooperative behaviour is seen in Table IV, as compared with Table III.

To test the hypothesis (b), F-S first calculate the standard deviation of total game output, for the next to last period, across all games of an experiment. They then test whether the standard deviation for complete information duopoly exceeds that of incomplete information duopoly, and similarly for triopoly. The hypothesis is confirmed that standard deviations are greater for complete information.

TABLE IV
Alternative classification of observed cooperativeness

	Individual		
	R	M	C
Complete information duopoly	9½	15½	7
Complete information triopoly	14	19	0
Incomplete information duopoly	4	26	2
Incomplete information triopoly	9	23	1 *

* This observation is ambiguous. While the output choice of eight units is nearer M than C, it was the lowest available output level.

2.3.2. *Comments.* Many issues are raised by reflecting on the F-S experiments: (a) They make little use of their data. (b) Are students appropriate as subjects? (c) Can relevant, useful, results be gotten when subjects participate in only one game each? (d) Their behavioural hypotheses are quite crude. These will be discussed in turn.

(a) In each game some twenty-five observations are generated. Most hypotheses used only the next to last. This procedure naturally raises questions: How many periods of play are necessary before a "learning phase" is over and some sort of stable-state equilibrium behaviour begins? Could the experiment shed light on the learning phase? Are there characteristics peculiar to an individual which may be seen in the learning phase and which give an indication of how he will behave in the stable state?

Of course, knowing the length of the learning phase suggests a minimum length for the game if one wants observations on settled behaviour. The minimum differences to be expected between learning phase and later behaviour are: (i) The random component in learning behaviour is likely to be larger because the subject's aims are not as well defined as they will eventually become. (ii) To the extent that learning phase behaviour has regular features, one would expect these to change over time as learning proceeds. One characteristic of stable equilibrium behaviour should be that it is characterized by features which remain constant over time. For example, the quantity choice of firm 1 of a duopoly might be described by a reaction function:

$$q_{1t} = f(q_{2,t-1}, \xi_t).$$

The reaction function gives the decision of firm 1, in the current period, as a function of his rival's decision in the preceding period and a random variable. One would expect, as the learning phase proceeds, that the function, f will change, and the role of ξ_t will diminish in importance.

It would be most interesting to see experiments which distinguished the two phases and tested hypotheses relevant to each.

F-S did not include learning hypotheses among those which they sought to test. As a research strategy, it does make sense to attempt first an understanding of equilibrium behaviour because that is probably the easier task and its solution facilitates the study of learning.

(b) The obvious questions about the student sample (or about Hoggatt's faculty sample) is whether they are drawn randomly and from the "right" population. F-S generally sought male volunteer subjects who were told they "could make some money by participating in a research project."¹ Are students drawn in this way typical of the male student population or are they different because they wanted to earn money or be in a research project? Perhaps these volunteers are more hard working. Are students typical of the population of businessmen? They are both younger and more uniform in age and education. Perhaps there are relevant, singular traits which are typical of businessmen and not of students. Does the businessman's years of experience change him in ways relevant to this sort of experiment? It would be a mistake to pretend to answer these questions here. They are empirical questions to be answered by appropriate, well-designed studies. All the studies reviewed here retain interest, whether or not different samples reveal different behaviour, as studies of human decision making.

(c) Each subject participated in only one game, probably of an hour's duration. As businessmen may be assumed quite familiar and accustomed to their usual circumstances, the appropriate thing is to have data from subjects who are used to theirs—who have got a feeling for the game, as a businessman has a feeling for his business. The experience of the present writer suggests that for games of the approximate complexity of these reviewed here, two to four hours of play must elapse before a subject gets fully accustomed to his circumstances. Striking evidence in [8] shows behaviour changing for approximately three hours and stable for the next seven in an experiment which consisted of two five hour sessions on consecutive nights with the same subjects. The experimenter ought consciously to decide whether he wishes to study the behaviour of novice subjects or subjects who have some feel for the particular experimental environment in which they are placed.

Considering the comments of paragraph (a), above, raises the question whether there may be two sorts of learning relevant to the subject who plays many games. The first is that which gives him general understanding of and feel for the experimental situation, and the second is that which gives him a sense of how his particular rivals of the moment make their decisions. These may be termed "experimental" learning and "within game" learning. Surely that latter proceeds more quickly, the more experimental learning the subject has experienced. Loosely speaking, experimental learning is like acquiring general business experience, while within game learning is like gaining experience in a particular new market.

(d) They admit of three types of behaviour, R , M and C . R and M are polar extremes of cooperativeness, while M , in between, is non-cooperative. Below, in section 2.4, a somewhat less rigid alternative is suggested.

On balance, the F-S experiments are exciting, disappointing and very suggestive. After the very great care with which they designed their experiments, chose and handled their subjects, and reported all they have done, it is a let-down to see the very crude and indirect nature of the hypotheses they test, and to see how little use they make of their

¹ See [6], pages 22 and 14.

data as well as how little attention they pay to determining which data they really want. To call down F-S for these shortcomings is, perhaps, to criticize them for not seeing before they ran their experiments certain things which are much more clearly perceived after seeing what they have done.

Whether or not their experiments are ideal, they are designed to facilitate testing the hypotheses they set out to test in simple and uncluttered ways. Their results provide a further corroboration of the hypothesis that the Cournot solution will prevail under incomplete information. The great dispersion of results under complete information suggests that personal characteristics may play a role in determining equilibrium behaviour, and that the most appropriate solution concept may be one for which the outcome is not unique.

2.4. Friedman

2.4.1. *Description.* The Friedman experiments [7] were designed and executed after F-S and with a knowledge of their work. In many respects F-S techniques and procedures were copied: Student subjects were used and were paid their actual profit. Subjects entered and left the laboratory singly so they would not see fellow subjects. Instructions were given in written form so that the precise instruction each subject received could be fully reported in the research report. Subjects were in single decision variable games which could be represented by payoff matrices.

These experiments differed from F-S in several important respects. They are confined to oligopoly and, although price was the decision variable, the models were not Bertrandese, so there were no discontinuities in the demand relations. All games were the complete information variety, as F-S seemed to have strong results for the incomplete case.

The Friedman experiments were designed to have each subject play in many (6 or 9) games. The games were sequential with the payoff matrices and grouping of subjects into games varying from game to game. While a subject never knew who his rivals were, he did know that he had the same rivals all through a game, and that he would have new ones from game to game. The original intent of having a subject play in many games was to provide sufficient data so that the behaviour of each subject, as a function of the behaviour exhibited by his rivals, could be characterized individually.

The payoff matrices were derived from the following demand and total cost relations:

demand:

$$q_{it} = \alpha_0 - \alpha_1 p_{it} + \frac{\alpha_2}{n-1} \sum_{j \neq i} p_{jt} \quad i = 1, \dots, n;$$

total cost:

$$\begin{aligned} C_{it} &= \beta_0 + \beta_1 q_{it} + \beta_2 q_{it}^2, \\ \alpha_0, \alpha_1, \alpha_2, \beta_1, \beta_2 &> 0, \quad \alpha_1 > \alpha_2, \\ \pi_{it} &= p_{it} q_{it} - C_{it} = \pi_i(p_{it}, \sum_{j \neq i} p_{jt}). \end{aligned}$$

Thus, as with F-S, the firms were symmetric (if each makes the same decision, profits are identical for all). n , the number of firms, was 2, 3 or 4.

The principal analysis was based on an index of cooperativeness: Consider the following for the i th firm:

$$F_{it} = \pi_{it} + \rho_i \sum_{j \neq i} \pi_{jt}.$$

If a firm chose its price so as to maximize F_{it} , ρ_i could be regarded as its index of cooperativeness. An index of zero corresponds to a Cournot, or non-cooperative, maximizer.

Values greater than zero show a positive degree of cooperativeness, with $\rho_i = 1$ corresponding to a firm seeking to maximize industry profits. Negative values correspond to various degrees of rivalistic behaviour, with $\rho_i = \frac{-1}{n-1}$ being the rivalistic behaviour

notion of F-S. When $\rho_i = \frac{-1}{n-1}$, the firm seeks to maximize the excess of its own profit over average profit of its rivals.¹

If it is assumed that a firm believes its rivals will repeat in the current period the prices they charged in the preceding period, then the price the firm should charge in period t is implicitly defined by:

$$\frac{\partial \pi_i(p_{it}, \sum_{j \neq i} p_{j,t-1})}{\partial p_{it}} + \rho_i \sum_{j \neq i} \frac{\partial \pi_j(p_{j,t-1}, p_{it} + \sum_{k \neq i, j} p_{k,t-1})}{\partial p_{it}} = 0.$$

Friedman assumed that the firm might have a different value of ρ_i in each period and that ρ_{it} depended on the $\rho_{j,t-1}$ ($j \neq i$). The degree of cooperativeness exhibited by a subject in period t was assumed to depend on the cooperativeness shown by his rivals in period $t-1$. Clearly, if the $\rho_{j,t-1}$ ($j \neq i$) are known, and it is assumed the firm maximizes in the way described, ρ_{it} is given by:

$$\rho_{it} = \frac{-\partial \pi_i(p_{it}, \sum_{j \neq i} p_{j,t-1})/\partial p_{it}}{\sum_{j \neq i} \partial \pi_j(p_{j,t-1}, p_{it} + \sum_{k \neq i, j} p_{k,t-1})/\partial p_{it}}.$$

The relationship, whose parameters were estimated separately, for each subject, is a variant of:

$$\rho_{it} = \alpha_i + \frac{\beta_i}{n-1} \sum_{j \neq i} \rho_{j,t-1} + \varepsilon_{it}.$$

Where ε is a random disturbance term with zero mean.

The data used were from the ten periods preceding the final period of each game. Earlier periods (the first ten to fifteen) were not used on the premise that the subjects required this time to get used to each new game. The latter assumption appears inconsistent with the validity of the estimated relationship, because it is assumed the parameters α_i and β_i do not depend on how the rivals choose their ρ values. If the α and β of a subject were thought to depend on the α 's and β 's of the rivals, one might wish to exclude the early periods of each game when subjects would be assumed to be sizing up one another and deciding on appropriate parameter values. Thus consistency would appear to demand that Friedman either (a) assume α and β for a subject depend on the α 's and β 's of rivals, and estimate a different α and β for each subject for each game (which justifies using only the later periods of each game to estimate the behavioural relationships) or (b) assume α and β are not dependent on the behaviour of rivals, in which case there is no reason to assume the values of α and β should differ early in a game from their later values. Friedman made the assumption (b), yet did not use the data from the early part of the games.

The justification for Friedman's formulation is not that one firm really cares about the profits accruing to its rival, but, rather, that the firm believes it will make more profits in the long run if it allows its cooperativeness index to be higher, the higher are the indices of its rivals. This could be valid if the firm believed the rivals' indices were positively related to its own. Then it would seem likely that a firm would select its α and β as some function of the α 's and β 's of its rivals.

¹ If $\rho_i = \frac{-1}{n-1}$ for all i , the oligopoly game with the F_i as payoff functions would be a zero sum n -person game. I point this out because zero sum games do not often arise naturally in economics.

The questions raised above about Friedman's analysis are too much for his data to bear. Its proper settlement would require a new, and well thought out, experiment. In part this is because in his experiments, certain parameters of the profit functions, as well as the grouping of subjects into games, were changing from game to game. Thus it is impossible to tell whether α and β for a subject change from game to game due to changes in game parameters or due to differences in the α 's and β 's exhibited by rivals. If α and β did not change from game to game, weak support would be gained for the hypothesis that neither rival behaviour nor game parameters affect the α and β of a subject. This strange conclusion is not borne out, however. Game parameters were introduced in the regression by means of dummy variables, and it was found that they cannot, in general, be removed without significant effect.¹ For each game and each subject, in both experiments, the β 's were usually between zero and one.

Friedman's first experiment involved only duopoly games, with six subjects playing in six games each.² The second utilized nine subjects in 2, 3 and 4 person games. Eight of them were in 3 duopolies, 3 triopolies and 3 four person games, with the ninth subject in 6 duopolies and three triopolies.

A second type of information coming out of his experiments concerns the number of games that eventually reach the joint maximum and, once attained, stay there until the end of the game. In the first experiment it is 6 out of 18 duopolies. In the second, it is 6 out of 15 duopolies, 2 out of 9 triopolies and 0 out of 6 four-person games. The proportion of joint-maximum games appears to decline rapidly as the number of players increases.

2.4.2. *Comments.* In comparing the F-S complete information games to Friedman's, one difference that stands out is the frequency of joint maximum games. There were none among their 16 duopolies or 11 triopolies; while 12 of 33 of Friedman's duopolies were joint maximum, and 2 of 9 triopolies. Only one of these games occurs when subjects were in their first game. It was a duopoly and was among 5 duopoly games which were the first played by subjects. This information is summarized below in Table V. The evidence is not conclusive, but it suggests that there are differences between the mode of play of a subject in his first game, as compared with later games. Additional, and stronger evidence may be found in an experiment not reviewed here [8].

TABLE V
Frequency of joint maximum games

	First game	Later games
Duopoly	1 of 5	11 of 28
Triopoly	0 of 12	2 of 8

Another side to the question of subjects participating in many games is that casual empiricism suggests some may become bored after a while. Boredom may lead to random choices, choices designed to shock one's rivals, mechanical repetition of some pattern of play while the mind dwells on other things, etc. Perhaps the real world of business

¹ *F* tests were performed comparing, for each subject, a regression with several dummy variables present to represent game parameters and a regression with all dummy variables removed in which only two parameters were estimated. For nearly all subjects in the second experiment the two regressions were significantly different at the 5 per cent level. While the reverse was true in the first experiment, there was considerably less data available per subject.

² Actually eight subjects were hired, one quit in the middle of the experiment resulting in complete participation for only six subjects. The seventh was played against dummy strategies in every game after the eighth quit.

involves stakes so large or an environment so interesting and varied that boredom is rare. In any case, it is really not known when and how boredom sets in to the experimental scene, much less whether it occurs in a way different from or similar to the real business world. In the progress of an experiment, a stable equilibrium phase may give way to a boredom phase.

Friedman tried to do too much in his experiments and, as a consequence, somewhat muddled his results. Clearly the more variables one wishes to examine, the more observations are required for good analysis. Friedman gave in to the temptation to let vary almost anything he could, given the number of games and subjects he used. The result is that when all variables are accounted for, nearly none seems significant; however, if, say, half the variables are excluded nearly all the remainder are significant. The half to be used may be chosen arbitrarily.

It was also noted earlier that the theoretical underpinnings of the regression are shaky. The behavioural hypotheses embodied in it do not arise from any well-defined optimization process. They are supported by (somewhat contradictory and fuzzy) *ad hoc* reasoning.

2.5. Hoggatt 2

2.5.1. *Description.* Hoggatt has done a second experiment [11], after those reported earlier in this section, which is, in certain ways, a notable improvement in design over its predecessors; however, it also fails to utilize certain important lessons learned or demonstrated in them. The weak points in this experiment are: (1) using unpaid subjects, and (2) having subjects play only a small number of periods (two games of between 13 and 21 periods each).

Hoggatt had 12 subjects for this duopoly experiment, all of whom were researchers in the School of Engineering of the University of Chile. Like the F-S and Friedman experiments, this one is concerned with the cooperativeness of subjects; however, it is nearer in approach to the latter than the former. The fundamental F-S hypothesis was that behaviour is more uniformly non-cooperative, the less information the subjects possess. Friedman used the index of cooperativeness, ρ , to characterize the behaviour of a subject and estimated the relationship between the index of a subject and that of his rivals' in the preceding period. Hoggatt's hypothesis is the higher the ρ value of the rival, the higher is one's own ρ value.

His experiment is designed in a very simple, yet very powerful and sensible way. His subjects each played in two games simultaneously. They were told they sold in two completely isolated and independent markets with a distinct rival in each. The rivals were in fact robot players, whose modes of play were previously determined by the experimenter. The strategy against which the subject played in one market was $\rho = -0.6$ and in the other $\rho = +0.6$. The precise hypothesis is that in a game with fixed ρ for the robot, the series of ρ_{it} chosen by the subject will go to a limit ρ_i^* . The higher is the value of the robot, the higher is equilibrium value ρ_i^* of the subject.

The precise statement of Hoggatt's hypothesis is [11, pp. 118-119]:

. . . we shall split each time series of outputs by the subject into a "first half" and a "last half," discarding the central term in the event that the number of terms is odd. Similarly, we divide the series of first differences with respect to time into a first half and second half. We say that the path of a particular market shows a tendency toward stability if:

- (a) The mean of the absolute deviations of the last half series of outputs is less than the mean of the absolute deviation of the first half series of outputs for that market and;
- (b) the mean of the first difference with respect to time of the last half series of output

is less than the mean of the first difference with respect to time of the first half series of output.

Seven of the games with robot ρ of -0.6 and eleven with robot ρ of $+0.6$ pass tests (a) and (b) for stability. Table VI below shows the contingency table for the 12 subjects concerning the passing of the tests (a) and (b). A χ^2 test for independence of the entries of Table VI is accepted at the 95 per cent level. On this basis, one cannot conclude for stability.

While one cannot conclude formally for stability, a glance at the data leaves the impression that ρ is higher for a subject, the higher is the ρ played against him. In six of the 12 cases the ρ a subject plays is never lower against the $\rho = 0.6$ robot than the corresponding move against the $\rho = -0.6$ robot. In four cases there are a few instances (6 out of 68 periods) where a lower ρ is played against the 0.6 robot than the -0.6 . The remaining two cases are less clearcut; however, the overall effect leaves the strong impression that Hoggatt's hypothesis concerning the happy effect of cooperation is correct.

TABLE VI

	Does behaviour tend toward stability (pass test (a), (b)) in games where the robot strategy is $+0.6$		
Does behaviour tend toward stability (pass tests (a) and (b)) in games where the robot strategy is -0.6	Yes	Yes	No
	No	7	0
	4		1

2.5.2. *Comments.* Like all the experiments discussed in the present paper, this one is interesting because of some novel insights of design, as well as because the results, themselves, are suggestive. Using unpaid subjects and not allowing them enough experience to get fully used to the game situation has been commented on before. The experimental design is excellent for its simplicity. The parameters allowed to vary are kept down to a reasonable few. The analysis itself utilizes techniques calling for a minimum of underlying statistical assumptions, in contrast with Friedman, who uses ordinary least squares regression.

The really novel feature of this experiment is the use of pre-programmed strategies. The obvious advantages are that subjects may be paired with many types of rival players, and, as the experimenter *knows* the parameters of the robot players, any statistical analysis is simplified. Instead of having a set of parameters to estimate for each subject, with interdependence between the values for them, causality runs only one way. The parameters of a subject are never affected by those of another subject.

Like most good things, the robots come at a price. If their behaviour is very singular and unlike the behaviour of people, it is possible the subjects play against them in ways they would never play with other subjects. For example, Hoggatt's robots have no learning phase. It is possible that the response of a subject to this learning phase may affect the final, stable equilibrium behaviour of his rival. In these circumstances, one could not predict subjects' behaviour in games with other subjects on the basis of their performance with robots. Clearly the ideal is to have robots which behave like people.

Hoggatt's robot strategies were of the form

$$x_{rt} = a + bx_{s,t-1}$$

where x_{rt} is the t th period output of the robot and x_{st} is the t th period output of the subject. Such a simple, rigid pattern may be easily detected by subjects, and spotted as a mechanical strategy. It may be possible to make the robot appear more human by adding a random element to its behaviour rule. In any case, the robot strategies, while potentially very useful, must be understood better and handled carefully.

3. FURTHER COMMENT ON THE EXPERIMENTS

3.1. *Motivation of Subjects*

The first question concerns whether the subjects are likely to have a relevant sort of motivation. All the experiments reported above are aimed at a simple modelling of business behaviour, thus the subjects are supposed to be actuated by the most important forces operating on real businesses. A common list of business motivations would probably include: (a) profit, (b) sales, (c) rate of growth of profit or sales or assets, (d) power, (e) relative share of industry sales going to the firm, (f) level of profit, sales, or growth rate relative to rivals.

Traditionally, economists would argue for (a) and claim that (b) and (c) might be regarded as proxies for future profits while (d), (e) and (f) either play no role or are of secondary importance next to profit. It remains an empirical question of importance to determine whether there are motivations more important than profit, and if there are, after determining what they are, to reconstruct economic theory with these new, correct, motivations in such a way that it gives more insight into and explanation of economic realities than the presently accepted theory. For those who, like the author, believe that profit is unlikely to be dethroned, the experiments which have been described focus on the principal motivation.

3.2. *On the Hypotheses Which are Tested*

A number of behavioural hypotheses have been put forward in the oligopoly literature. In each of the experiments reviewed here, a few of the alternatives are singled out and the experiment is used to choose among them, or a relationship is sought between the mode of behaviour displayed by a subject and the behaviour of his rivals.

Hoggatt's first experiment featured standard Cournot behaviour and a misconceived notion of long-run Cournot behaviour. The F-S experiments focussed on Cournot, rivalistic and joint profit maximizing behaviour. The Friedman and Hoggatt 2 experiments assume firms seek to maximize a weighted sum of their own and their rivals' profits.

All of these modes of behaviour assume single period profit maximization, which is a rather naïve objective for oligopolists. They are presumably aware that their actions will affect the behaviour of their rivals and that there will be many decision periods after the current one. Surely maximizing a discounted profit stream is more sensible. Another difficult assumption is that a firm assumes its rivals will repeat in period t their decisions of period $t-1$. Generally this assumption will be contradicted by the facts, so its maintenance by rational economic agents is untenable.

3.3. *Concluding Comments*

On balance the experiments discussed here are pioneering efforts which go far to demonstrate the usefulness of experimental methods for economic research. They contain much information on and insight into the problems of good experimental design, and they test interesting hypotheses in fruitful ways.

The extreme care with which F-S wrote instructions and handled subjects was a contribution from experimental psychology, and it helped to launch experimental economics at a high level. Perhaps the individual choice circumstances characteristic of oligopoly, bilateral monopoly and utility may require a much longer period of familiarization than the psychologists' subjects need. A longer time is needed than was at first recognized.

Paying subjects is an important piece of realism. When this is done, decisions are not hypothetical. Profit to the subject actually depends on his behaviour, and he need not be asked to pretend it does. Chances are most people are poor at guessing how they would behave in situations with which they are not thoroughly familiar, so anything that can be done to make an experiment real instead of hypothetical is very important.

The advantages of experimental research, when adapted to the needs of economics, are striking. Compared to traditional empirical research in economics the experimenter has considerable control over the design of the markets within which his economic agents operate, and he can exclude some influences from which he wishes to abstract while concentrating on others he wishes to study. Anyone questioning results within the experimental framework is able to replicate an experiment performed by someone else and analyze entirely new and independent data. On the other hand, the world of the experimenter is artificial and of his own design; so the extent to which experimental results may be taken as a clue to behaviour in real life markets is uncertain.

The two conclusions on oligopolistic behaviour which emerge from these studies are: the Cournot solution characterizes behaviour in incomplete information situations, and, under complete information, a subject is more cooperative, the more cooperative is his rival. It is not obvious whether these conclusions will obtain under more intensive examination, or how more comprehensive experiments will find them modified; however, they are interesting results and suggestive of what might be done experimentally.

University of Rochester

J. W. FRIEDMAN

First version received 8.4.68; final version received 25.10.68

REFERENCES

- [1] Bertrand, Joseph. "Review of Walras' *Theorie Mathematique de la Recherche Sociale* and Cournot's *Recherches Sur les Principes Mathematiques de la Theorie des Richesses*", *Journal des Savants*, **68** (1883), 449-508.
- [2] Chamberlin, Edward H. "An Experimental Imperfect Market", *Journal of Political Economy*, **56** (1948), 95-108.
- [3] Davidson, Donald, Suppes, Patrick and Siegel, Sidney. *Decision Making: An Experimental Approach* (Stanford: Stanford University Press, 1957).
- [4] Dolbear, F. Trenergy. "Individual Choice Under Uncertainty: An Experimental Study", *Yale Economic Essays*, **3** (1963), 419-470.
- [5] Edgeworth, Francis Y. *Papers Relating to Political Economy, I* (London: Macmillan, 1925), pp. 111-142.
- [6] Fouraker, Lawrence E. and Siegel, Sidney. *Bargaining Behavior* (New York: McGraw-Hill, 1963).
- [7] Friedman, James W. "Individual Behaviour in Oligopolistic Markets: An Experimental Study", *Yale Economic Essays*, **3** (1963), 359-417.
- [8] Friedman, James W. "An Experimental Study of Cooperative Duopoly", *Econometrica* (October 1967), pp. 379-397.
- [9] Friedman, James W. "Equal Profits as a Fair Division", to appear in *Beitrag Zur Experimentellen Wirtschaftsforschung*, Vol. II (ed. Professor Heinz Sauermann).
- [10] Hoggatt, Austin C. "An Experimental Business Game", *Behavioral Science*, **4** (1959), 192-203.
- [11] Hoggatt, Austin C. "Measuring Behavior in Quantity Variation Duopoly Games", *Behavioral Science*, **12** (1967), 109-121.
- [12] Mosteller, Frederick and Noguee, Philip. "An Experimental Measurement of Utility", *Journal of Political Economy*, **59** (1951), 371-401.

- [13] Shubik, Martin. *Strategy and Market Structure* (New York: Wiley, 1959).
- [14] Siegel, Sidney and Fouraker, Lawrence E. *Bargaining and Group Decision Making: Experiments in Bilateral Monopoly* (New York: McGraw-Hill, 1960).
- [15] Smith, Vernon L. "An Experimental Study of Competitive Market Behavior", *Journal of Political Economy*, **70** (1962), 111-137.
- [16] Smith, Vernon L. "Effect of Market Organization on Competitive Equilibrium", *Quarterly Journal of Economics*, **77** (1964), 181-201.
- [17] Smith, Vernon L. "Experimental Auction Markets and the Walrasian Hypothesis", *Journal of Political Economy*, **73** (1965), 387-393.
- [18] Yaari, Menahem E. "Convexity in the Theory of Choice Under Risk", *Quarterly Journal of Economics*, **79** (1965), 278-290.