WORKING PAPERS

AN EXPERIMENTAL TEST OF THE CONSISTENT - CONJECTURES HYPOTHESIS*

Charles A. Holt Jr.

WORKING PAPER NO. 72

August 1982

BUREAU OF ECONOMICS
FEDERAL TRADE COMMISSION
WASHINGTON, DC 20580

Abstract

Recently, the notion of a "consistent conjecture" has been proposed as a way of avoiding the indeterminacy of conjectural variations models of oligopoly behavior. This paper reports the results of a laboratory experiment designed specifically to discriminate between the consistent-conjectures equilibrium and other commonly used equilibrium concepts. The consistent conjectures equilibrium does not provide a good prediction of subjects' behavior for the particular cost and demand parameters used in this experiment. The static Nash/Cournot equilibrium provides a more accurate prediction, although subjects in some markets managed to collude tacitly.

AN EXPERIMENTAL TEST OF THE CONSISTENT-CONJECTURES HYPOTHESIS*

June 1982

Charles A. Holt Jr.

Bureau of Economics
Federal Trade Commission
and
Department of Economics
University of Minnesota

A common way of analyzing multiperiod oligopoly models without dynamic interactions in the payoff structure is to compute a Nash equilibrium for each period taken separately. Many economists believe that behavior in a repeated market game cannot be predicted accurately with a period by period sequence of such "static" Nash equilibria, but an explicitly dynamic analysis can be extremely difficult unless the class of feasible dynamic strategies is restricted.[1]

There is an embarrassing multiplicity of alternative oligopoly "solutions" that are computationally less complex than game-theoretic approaches to multiperiod games. Many of these alternative solutions can be classified as conjectural variations models in which firms are assumed to conjecture that changes in their own decisions will induce reactions by other firms. These reactions are typically assumed to be characterized by functions that are locally linear. Almost any configuration of decisions can be an equilibrium for some conjectured reaction functions, so these models have little empirical content unless the reaction functions themselves are determined endogenously.

Timothy Bresnahan (1981) has proposed a consistency condition that can often be used to determine specific conjectured reactions. Martin Perry (1982, p. 197) provides a clear explanation of this consistency condition in the context of a duopoly in which firms' decisions are output quantities:

> Each firm's first-order condition defines its profit-maximizing output as a reaction function on (1) the output of the other firm and (2) the conjectural variation about the other firm's response. Thus a conjectural variation by one firm about the other firm's response is consistent if it is equivalent to the derivative of the other firm's reaction function with respect to the first firm's output at equilibrium.

Many economists have found this notion of consistency to be appealing; Perry cites a large number of recent working papers on the theoretical properties of consistent-conjectures equilibria.

Although not explicitly dynamic, the consistent-conjectures equilibrium (CCE) approach initially seemed plausible to me because it predicts deviations from a static Nash equilibrium that are qualitatively consistent with the data reported in several published laboratory experiments with student subjects. These experiments, however, were not designed to provide a clear distinction between the CCE and other equilibrium concepts. This paper reports the results of an experiment designed specifically to test the consistent-conjectures hypothesis.

In section I, the computation of a consistent-conjectures equilibrium is explained in the parametric context that is used to construct the experiment. Section II contains a discussion of how the payoff structures used in the previous laboratory experiments must be modified to permit a good test of the consistent-conjectures hypothesis. In section III, I report the results of an experiment in which the theoretical predictions of the static

Nash and consistent-conjectures equilibria are quite different. The data are clearly inconsistent with the CCE hypothesis. A related experiment is discussed in section IV, and section V contains a conclusion.

I. THE CONSISTENT-CONJECTURES HYPOTHESIS

The notion of a consistent-conjectures equilibrium is easily explained for a homogeneous-product duopoly in which variable costs are zero and industry demand is linear: $p = A - B(x_1 + x_2)$, where $A > 0$, $B > 0$, p denotes price, and x_i denotes the output of firm _i_. The profit function for firm _i_ is: $x_i(A - Bx_1 - Bx_2)$.

The first-order condition for the profit-maximization problem for firm _i_ is:

(1) $A - Bx_j - 2Bx_i - Bx_i\lambda_j = 0$, $(i = 1, 2;\ j \neq i)$,

where $\lambda_j \equiv dx_j/dx_i$. The conjectural variation λ_j is assumed to be a constant.[2]

The two equilibrium outputs cannot be determined from the two equations in (1) unless the λ_j conjectural-variation parameters can be determined. To do this, Bresnahan uses a consistency condition that the actual profit-maximizing reaction of the _i_th firm's output to a change in x_j must be equal to the λ_i conjecture that characterizes the beliefs of firm j. Suppose that x_j changes by an amount of dx_j. Then Bresnahan computes the _i_th firm's profit-maximizing response to this change by totally differentiating equation (1) to obtain

(2) $- Bdx_j - 2Bdx_i - B\lambda_j dx_i = 0$, $(i = 1, 2;\ j \neq i)$.

Dividing (2) by dx_j and using the definition of λ_i, one can express (2):

$$(3) \qquad - B - 2B\lambda_i - B\lambda_j\lambda_i = 0, \qquad (i = 1, 2;\ j \neq i).$$

It follows from the two equations in (3) that $\lambda_i = \lambda_j = -1$. Then (1) implies that $x_i + x_j = A/B$, so price and profits are zero for the consistent-conjectures equilibrium in this example. Note that the industry output equals A/B, but the consistent-conjectures equilibrium outputs need not be equal in this example. This is because the graphs of the reaction functions that satisfy the consistency requirement in the example are overlapping straight lines.

The consistent-conjectures equilibrium concept can be applied when decision variables are prices and there are more than two firms. When the demand is linear, the product is homogeneous, and all firms have the same constant average total cost, it can be shown that the CCE price equals average cost and profits are zero regardless of the number of firms and regardless of whether the decision variables are prices or quantities.[3] The predicted "competitive" result in all cases other than monopoly in this context is the basis of the design of the experiment discussed in section III.

II. EVIDENCE FROM PREVIOUS EXPERIMENTS

The first question that should be addressed is whether the popular static Nash equilibrium approach can explain behavior in multiperiod-market experiments. F. Trenery Dolbear, Lester Lave, et al. (1968) reported data showing that behavior in multiperiod duopoly experiments deviates systematically from a static Nash equilibrium. Their subjects were students who chose prices simultaneously at the beginning of each "period." I will only discuss the "complete information" experiments in which subjects were given a payoff table that relates price choices to payoffs in pennies.[4] The subject's price choice determined a row in the table, and the average of the prices of the subject's competitors determined a column. The payoff entries in the table were computed with a quadratic profit function that resulted from a demand function with some product differentiation. Payoffs were rounded off to the nearest penny, and as a result, there were two symmetric Nash equilibria in prices at common prices of 17 or 18. If subjects had been able to collude, they could have maximized their joint profit by raising prices to 23. However, subjects were not able to communicate.

In each market experiment, the subjects made simultaneous price decisions 15 times, but they were not told the number of repetitions in advance. The average price for each experiment was obtained by averaging all prices for periods 8 through 12. There were 12 duopoly experiments with complete information, and the average price in each experiment is represented by a dot on the horizontal price scale in figure 1. The average price across all 12 experiments was 19.5, and the authors concluded that these data indicate some tacit collusion in the sense that average prices and profits exceed the levels determined by a static Nash equilibrium in prices. Using the parameter values for the Dolbear, Lave, et al. profit function, Holt (1980) calculated the consistent-conjectures equilibrium price to be 19.2 in this context, and this is quite close to the observed price average.

Of course, these experiments were not designed to test the consistent-conjectures equilibrium, and there are several obvious ways in which the experiments do not provide a satisfactory test of this equilibrium concept. First, the subjects were required to make integer-valued price choices, but the CCE price was not an integer. Second, there is not much difference between the static Nash and the CCE prices. (This problem was even more severe for the oligopoly experiments with four subjects.) Finally, the word "competitors" in the subjects' payoff table may have suggested a particular type of behavior.[5]

Next, consider the previous section's quantity-choice model with a linear market-demand function and a common, constant average cost. The symmetric, static Nash (Cournot) equilibrium when strategies are output quantities will result in a price that is greater than average cost and less than the price resulting from perfect collusion. In contrast, the consistent-conjectures equilibrium in this context will result in competitive outputs that drive price down to average cost and profits to zero. Therefore, homogeneous-product oligopoly experiments with quantity-setting subjects and constant average costs may provide a good opportunity to discriminate between the static Nash and the consistent-conjectures theories.

Lawrence Fouraker and Sidney Siegel (1963) reported the results of some complete-information duopoly and triopoly experiments with these characteristics. The columns in their payoff table corresponded to a subject's own output choices, which were integers between 8 and 32. The row was determined by the "Quantity produced by my competition," and this quantity could be between 8 and 64. Outputs between 33 and 64 were actually possible in the triopoly experiments because the "competition" consisted of two other subjects. For a duopoly, the collusive industry output was 30 (15 per subject), the theoretical Nash/ Cournot industry output was 40 (20 per subject), and the competitive industry output was 60 (30 per subject). As indicated in the previous section, 60 is the output predicted by the CCE in this context. Fouraker and Siegel do not seem to have noticed that the rounding off of payoffs to the nearest half-penny resulted in two symmetric Nash equilibria: one at an industry output of 40 and another at an industry output of 44.[6]

There were 16 complete-information duopoly experiments in this series (Experiment 10). Instead of averaging, Fouraker and Siegel used the subjects' decisions in the 21st period as an indicator of equilibrium behavior. The period-21 industry outputs were scattered fairly uniformly over the range from the collusive industry output (30) to the competitive (and CCE) industry output (60).[7]

The failure of outputs to rise to the CCE level in many markets may have been due to the fact that the profit was zero because price equaled average cost at this level. Subjects were told in the instructions that if they follow instructions and make "appropriate decisions," they "may earn an appreciable amount of money . . . but poor choices will result in small or no profit to you." Thus there is a possibility that the wording of the instructions made it less likely that the CCE result with zero profits would be observed. In my own experiments, subjects often appear to be frustrated after periods of very low profits, and such periods are usually followed by large output reductions that raise profits considerably.

There is, for me, a more compelling reason to expect that the outputs of 30 per duopolist would not be frequently observed in the Fouraker and Siegel duopoly experiments. Note that each subject is restricted to choose an output that is less than or equal to 32. The payoff table used by Fouraker and Siegel shows profits for values of the output of the "competition" between 8 and 64. In my opinion, each subject in the duopoly experiments was likely to realize that the outputs from 33 to 64 were irrelevant, and of course, no outputs above 32 were ever observed. If the output of the competition is less than 33, then it is a property of their table that any output below 28 will guarantee the subject a positive profit, regardless of what the competitor does. This truncation of the relevant payoff table caused by exogenous limits on output choices implies that the CCE profit of zero can be strictly dominated.[8] In particular, if both subjects were choosing outputs of 30 and earning no profit, then either one could cut output to 15 and earn at least 7.5 cents per period because the other seller's output cannot exceed 32.

This truncation argument does not apply in the triopoly experiments because the "competition" consists of two subjects, and there is no output decision a subject can make that will ensure a positive profit when each of the other two sellers chooses an output of 32. In fact, behavior in the triopoly experiments did seem to be much more competitive. The static Nash-equilibrium industry output for the triopoly was either 45 or 48.[9] The "competitive" and CCE output was 60, and the actual outputs in period 21 for the 11 triopoly markets were: 40, 44, 46, 47, 51, 58, 59, 59, 62, 63, and 70.[10] The median industry output of 58 is quite close to the CCE prediction of 60. An industry output of 58 with an approximately symmetric output configuration would result in earnings of only $.02 per subject per period in 1960 dollars.

III. AN EXPERIMENTAL TEST OF THE CONSISTENT-CONJECTURES HYPOTHESIS

It follows from the discussion in the previous section that an experiment designed to test the consistent-conjectures hypothesis should have the following characteristics: (a) potentially suggestive words such as "competitors" and "oligopoly" should not appear in the instructions and payoff tables, (b) the CCE decisions should be integers, (c) the profit per hour per subject at the CCE should not be too different from the payment that subjects expected to earn after reading the instructions or the announcement that solicits subjects, (d) there should be no decision a subject can make that ensures a profit that will always exceed the CCE profit level, and (e) the CCE decisions should not be "close" to the decisions implied by either static Nash or collusive behavior.

A. THE PAYOFF STRUCTURE

The instructions for the experiment reported in this section are reproduced in the appendix, and the "Profit Table" is reproduced as table 1. This Profit Table was computed from equation (1) with A=12, B=1/2, and $.45 was added to each of the resulting profit entries. A simple calculus argument can be used to show that the outputs in a symmetric, collusive equilibrium are six per subject and the static Nash/Cournot outputs in a symmetric, collusive equilibrium are eight per subject. Outputs are constrained to be integer-valued in the experiment, but this discreteness does not affect the collusive and Nash equilibria. For example, if both subjects choose outputs of eight, then a unilateral, integer-valued deviation will not increase a subject's profit, given the Cournot conjecture. Because of the rounding off of profits to the nearest penny, there are also two asymmetric Nash equilibrium configurations: one with outputs of 7 and 9 and another with outputs of 6 and 10. In all cases, however, the industry output is 16 in a Nash equilibrium.

It follows from the calculations in section I that the consistent conjecture is -1 in this context, and any combination of outputs that sum to 24 constitutes a CCE. These output combinations lie on the diagonal with $.45 profits in the Profit Table. Starting on the diagonal, if one subject increases or decreases output by an integer amount, the other subject is conjectured to make an equal output change in the opposite direction. Thus the new output pair would again be on the $.45 profit diagonal, so the deviation would not increase the subject's profit, given the consistent conjecture.

The collusive industry output of 12 yields earnings of $.81 per subject, the static Nash/Cournot industry output of 16 yields earnings of $.77 per subject in the symmetric case, and the CCE industry output of 24 yields earnings of $.45 per subject. The experiment was not designed to distinguish noncooperative and collusive behavior, but neither of these modes of behavior yields outputs and profits that are close to those implied by the consistent-conjectures hypothesis in this context.[11] The high output levels (13 to 22) were included so that no output decision would guarantee a profit that exceeds the CCE level of $.45 per period.

The $.45 can be thought of as a normal rate of return when price equals average cost and economic profits are zero. Subjects were also given an initial stake of $.50 to cover any early losses. The announcement used to solicit subjects stated: "Although earnings cannot be predicted precisely, they will average about $6 per hour." The experiments were run at a pace of about 13 periods per hour, so the $.50 stake and the CCE profit of $.45 per period would result in earnings of about $6 per hour.

B. SUBJECTS AND PROCEDURES

The subjects were students in introductory and intermediate economics classes at the University of Minnesota. The instructors in these classes had not discussed experimental economics or formal oligopoly theory. The subjects had no previous experience with economics experiments.

Subjects were given about 10 minutes to read the instructions in the appendix. An additional paragraph (also in the appendix) was read aloud by one of the people conducting the experiments. The purpose of this additional paragraph was to convince the subjects that the "other seller" was a real person (not a computer).

The subjects were also given a "Decision Sheet" that revealed the "position number" of the "other seller" in that subject's market. The "other sellers" were seated in a separate room. First there was a "trial period," in which subjects marked their "output choices" on their Decision Sheets. Then they were told the output choice of the other seller, and they were asked to use the payoff table to compute both their own and the other seller's profit. This allowed us to check the subjects' understanding of the payoff table without suggesting anything by the use of hypothetical outputs to illustrate the computation of profits. In each subsequent period, we collected the Decision Sheets, computed profits, and paid the profits earned before the beginning of the next period. Subjects in the same room were spaced so that they would not be able to see exactly how much money others were earning. Subjects were also invited to write brief "explanations" of their decisions on their "Explanation Sheet."

Subjects will naturally be curious about when the experiment will end, and I think the best way to deal with this is to be explicit about the stopping rule. A random stopping rule was used to avoid end effects. Subjects were told that there would be at least seven periods and that there was a probability of 1/6 that period seven and each following period would be the final period. The final period was determined by a six on the throw of a die, but we used the same sequence of die throws for all subjects. The throw of the die was recorded on the Decision Sheet.

There were 24 subjects that will be labeled S1, S2, etc. There were 12 initial pairings of subjects, and all subjects participated in a "first market" that was terminated by a throw of the die after 13 periods for all pairs. In order to check for experience effects, 16 of these subjects were rematched and given a new Decision Sheet with the new position number of the other seller. A different sequence of throws of the die was used, and this "second market" was terminated after nine periods.

C. THE DATA

The output choices for the 24 subjects who participated in the first market are shown in table 2, and choices for the 16 experienced subjects who participated in the second market are shown in table 3. There was some collusive behavior resulting in outputs of six per subject, and there was some rivalistic behavior resulting in industry outputs greater than the static Nash/Cournot industry output of 16. Regardless of whether the first-market and second-market data are considered separately or together, the mean and median (or medians) of the final period industry outputs are between 14 and 16. Earnings averaged about $8.50 per subject per hour.

The data are clearly inconsistent with the CCE prediction of an industry output of 24, in my opinion. None of the final-period industry outputs exceed 21. There was only one pair of subjects (subjects S7 and S2 in the second market) with combined outputs that were often closer to the CCE level of 24 than to the static Nash/Cournot level of 16. The occasional high outputs of other subjects usually appear to be attempts to punish a rival for not reducing output. For example, subject S3 had been in a collusive duopoly in the first market, but S3 was not able to induce S6 to collude in the second market. Apparently frustrated, S3 increased output from 6 to 19 in period 4 and then returned to 6 in period 5.

A statistical analysis should begin with a consideration of why some duopoly pairs are more collusive than others. Variations in market outcomes may be due to variations in variables not included in the oligopoly models discussed above, variables such as individuals' willingness to experiment with output changes. Suppose that individuals' characteristics are independent drawings from some population of possible characteristics. Then it is natural to think of final-period industry outputs for either the first or second market (not both together) as being independent realizations of a random variable. In the following discussion, the eight final-period industry outputs in the second market will be denoted by Q_1, $Q_2, \ldots Q_8$, and the vector of these outputs will be denoted by $\overline{Q}$. Consider a family of hypotheses of the form: $\Pr\{Q_i < y\} \leq 1/2$ for some $y > 21$; $i=1,\ldots 8$. This family includes a hypothesis that the median of the industry outputs is 24, the theoretical prediction of the consistent conjectures equilibrium. Let H_y denote a particular hypothesis in this family that corresponds to a particular value of y. It can be seen from a binomial probability table that $\Pr\{\overline{Q}|H_y\} \leq .0039$ because all eight industry outputs are less than 21. However, a rejection of H_y using a classical hypothesis test would be misleading if there were no other hypothesis that is reasonable given the data observed. But there are many reasonable alternatives in this case.

For example, consider a hypothesis H_{16}: $Pr\{Q_i < 16\} = 1/2$, i=1,... 8. This hypothesis implies that a median of the distribution is 16, the theoretical prediction of the static Nash equilibrium. It follows from simple binomial probability calculations that $Pr\{\overline{Q}|H_{16}\} = .2734$, so the likelihood ratio is greater than .2734/.0039. If the posterior probabilities for H_{16} and H_y are denoted by $Pr\{H_{16}|\overline{Q}\}$ and $Pr\{H_y|\overline{Q}\}$ respectively, then the ratio $Pr\{H_{16}|\overline{Q}\}/Pr\{H_y|\overline{Q}\}$ is more than 70 times as great as the corresponding ratio of prior probabilities. A Bayesian analysis of the final-period outputs for the first-market experiments yields even stronger conclusions.

IV . A SINGLE-PERIOD DUOPOLY EXPERIMENT

The experimental design discussed in the previous section induces an infinite horizon in which the probability of termination determines the tradeoff between profit in the current period and profit in the future. In other words, the probability of termination determines the rate of which profits are discounted. If the probability of termination is low enough, subjects may be willing to make unprofitable output reductions in the hope of inducing the other seller to cut output in the future.

Roughly speaking, the behavior in the experiments discussed in section III can be categorized as either collusive or noncooperative. I expected that an increase in the termination probability from 1/6 to 1 would result in no collusion. From a game-theoretic perspective, the static Nash equilibrium is appropriate for single-period games in which subjects are not able to use strategies that are contingent on decisions made in previous periods. Thus, single-period experimental markets would give the static Nash equilibrium its best chance. These markets may also yield even more rivalistic behavior.

I conducted one set of experiments with 12 subjects who participated in a series of 11 single-period duopoly markets with the same payoff table that was used in the multiperiod experiments. The subjects were drawn from a pool of people who had previous experience with a different series of duopoly experiments with different payoff tables. Six subjects were seated in each of two large rooms, and subjects were spaced so that they were unable to determine the "position number" of any other subject in their own room. A research assistant was present in each room at all times. The instructions for these single-period experiments are also reproduced in the appendix.

The experiment began with a trial period in which profits were computed but not paid. This was followed by 10 single-period markets. The aggregate data on individual choices for these markets are graphed in figure 2, and data for particular subjects and their rivals are given in table 4. The output choices are initially quite diverse, but by period 7 two-thirds of the subjects are choosing outputs of 9. This is followed by a trend toward the symmetric Nash/Cournot outputs of 8, and 7 of the 12 subjects choose 8 in the final period. As expected, there was no successful collusion in the later periods of this experiment.

The frequency of rivalistic outputs of 9 in the intermediate periods is interesting. First, note that 9 is not very far from a Nash equilibrium in terms of profits. For the range of sellers' outputs in the final periods, any seller with an output of 9 could only increase profit by $.01 by switching from 9 to 8. If the outputs are 9 for one seller and 8 for the other, the profit $.76 for the high-output seller and $.73 for the other. At outputs 8 and 8, they each make $.77. To see why some individuals were willing to give up a penny of profit per period, I looked at the explanation sheets. There were several rivalistic comments about relative profits. For example, one person remarked: "Only a $.01 loss occurs producing at 9 instead of 8. This keeps the other firm's profits down." This subject did switch to 8 in the final period. Another subject, the only one to have an output of 10 in the final period, remarked in period 4 that when paired ". . . against a firm with lower output then mine, I make the larger profit, 9 is an interesting number to produce . . . ." However, it is clear that no subject's objective was to maximize the difference between profits; if the other seller produces either 8 or 9, then an output of 12 will maximize the difference between a subject's own profit and that of the other seller. In retrospect, there probably would have been less variability in the data if subjects in these experiments had not been given the complete information necessary to compute the other sellers' profit.

V. CONCLUSION

In this paper, I compare the theoretical predictions of the consistent-conjectures hypothesis with data for individuals' behavior in several laboratory experiments. In all experiments discussed, subjects simultaneously choose either price or quantity in a sequence of market periods, and subjects are given payoff tables that provide "complete information" about the relationship between decisions and profits for all participants.

My interpretation of the previously published experimental results is: The consistent-conjectures hypothesis provides a good explanation of the price choices made by subjects in the Dolbear, Lave, et al. experiments, but the predictions of the consistent conjectures and static Nash equilibria are quite close. The predictions of these two equilibria are not close for the Fouraker and Siegel experiments with quantity-setting subjects. The CCE does not provide a good explanation of the output choices in the Fouraker and Siegel duopoly experiments, but its predictions look more reasonable in the triopoly experiments. The poor performance of the CCE in the duopoly case may have been because subjects' profits were zero at the CCE and there were other output choices a duopolist could make that would ensure a strictly positive profit.

This paper reports the results of a new set of duopoly experiments with complete information in which payoffs are positive at the CCE, and there is no decision that can guarantee a profit that exceeds the CCE profit. The consistent-conjectures equilibrium does not provide good predictions in these experiments. The data are more consistent with the Cournot equilibrium, although several duopoly pairs managed to achieve perfect collusion tacitly. Thus, there is at least one simple payoff structure (with homogeneous products, linear demand, and constant average variable costs) in which the CCE predictions are clearly inaccurate.

There are, however, several questions a skeptical reader may wish to consider. First, can laboratory experiments with individual decision makers be used to evaluate theories of the behavior of business firms? Many economists will give a negative answer, but I see nothing in the computation of a consistent-conjectures equilibrium that suggests that the arguments apply to business organizations but not to individuals. One obvious difference between businessmen and the student subjects is that businessmen have more experience with the markets in which they operate. But when experience has been shown to have a significant impact on behavior in experiments, the effect has been to increase the frequency of collusion.[12] Increased collusion in the experiments reported here would further skew the data away from the "competitive" CCE output prediction.

A second issue is whether the inaccuracy of the CCE prediction derived in section I is due to something other than the inconsistency of conjectures. In particular, could it be the case that conjectures are consistent but that subjects are maximizing something other than profit? There was a slight tendency toward rivalistic behavior in the single-period experiment, so one may wish to consider an objective function R_i for the ith subject of the form: $R_i = \pi_i + w_i \pi_j$; $(i = 1, 2; j \neq i)$; where $\pi_i = x_i (A - Bx_1 - Bx_2)$, $-1 < w_i < 1$. If the w_i parameter is zero the subject is a profit maximizer, and as the w_i parameter approaches -1 the subject becomes very rivalistic and seeks to maximize the difference in profits. The first-order condition analogous to (1) is:

$$(4) \quad A - Bx_j - 2Bx_i - Bx_i\lambda_j + w_i [(A - Bx_i - 2Bx_j) \lambda_j - Bx_j] = 0.$$

The consistency condition analogous to (3) is:

$$(5) \quad -B - 2B\lambda_i - B\lambda_j\lambda_i + w_i (-B - 2B\lambda_j - B\lambda_j\lambda_i) = 0,$$

$(i = 1, 2; j \neq i)$. The two equations in (5) imply that $\lambda_i = \lambda_j = -1$, so the consistent conjectures are not affected by the possible rivalistic nature of objectives. These conjectures and (4) imply that $x_i + x_j = A/B$, so the CCE industry output is unchanged. Thus the inaccuracy of the CCE predictions in this context cannot be attributed to the possibility of non-zero values of the w_i parameters.

Finally, there is the question of the choice of the rule for ending the experiments. In experiments reported in this paper, the stopping rule was explicit, and a termination probability of 1/6 was used in the multiperiod experiment. The choice of this particular termination probability was arbitrary because there is no parameter in the theoretical analysis of consistent-conjectures equilibria that corresponds to a termination probability nor is there a discount rate. The CCE concept is not explicitly dynamic; the timing of output deviations, initial reactions, and subsequent reactions by the deviant is not clear. As Perry (1982, p. 200) points out:

> The conjectural variation model is a simple static representation of the potentially complex dynamics of an oligopoly, and consistency as defined [in a CCE] . . . is the simplest adequate static condition for rational behavior in such a model.

The CCE did not provide a satisfactory representation of the dynamics in experimental markets with a termination probability of 1/6. I would expect to observe more collusion and less rivalistic behavior if the termination probability were even less than 1/6. For termination probabilities that exceed 1/6, I would expect behavior to conform more closely to the predictions of the static Cournot model. In the single-period market experiments with a termination probability of 1, the Nash/Cournot equilibrium provided accurate predictions, and there was no tendency to collude.

APPENDIX

1. Multiperiod experiment: instructions read by the subjects

You are about to take part in a decisionmaking experiment. You will be able to make choices which, together with the choices of other participants, determine the payoff that you will receive. Whatever payoffs you accumulate will be yours to keep as your payment for participating in the experiment.

There are two sellers in this experiment. Sellers produce a hypothetical product, and each seller must decide how much of the product to offer for sale. This decision will be called an "output choice." Your monetary earnings in this experiment will depend on your own output choice and on the output choice of the other seller.

Before you is a profit table. The numbers across the top represent your own output choice. The numbers down the left side of the table represent the output choice of the other seller. The output chosen by the other seller identifies a row in the table, and your output identifies a column. The cell where that column and that row intersect reveals the profit you will receive for that specific combination of outputs. Profit is in cents. The other seller has a profit table that is exactly like yours, so the profit opportunities are symmetrical. The other seller is a student, and both of you are in separate rooms.

Before you is a plate containing 50 cents. This is yours to keep, along with any profits you accumulate during the experiment.

However, if you sustain losses in excess of your profits, the amount of your losses will be taken out of the original 50-cent stake. You cannot lose any of your own money. If your losses should reach 50 cents, you will be excused from the experiment.

Profits and losses will be determined by both your own and the other seller's output choices in each "decision period." During each decision period, you and the other seller will choose outputs from the choices available in the profit table. You will record your decisions on a Decision Sheet found in front of you. Each period, we will collect your Decision Sheet, record the other seller's choice, determine payoffs, and return the sheet. While we have your Decision Sheet, please note reasons for your output choice on the Explanation Sheet.

Each experiment will begin with a single "trial period" in which you and the other seller make a decision. Then we will record the other seller's output choice on your Decision Sheet, and we will let you use the payoff table to compute the profit or loss for each of you. Someone will check your calculations to be sure that you understand how to read the payoff tables. Profits will not be paid and losses will not be collected for the trial period. After each subsequent decision period, we will collect output choices, compute profits, and pay your profit or take away your loss.

The number of decision periods in each experiment will be determined by a random device. In particular, there will be at least seven decision periods. After the seventh period, a single die will be thrown, and there will be no more decision periods if the throw of the die yields a six. If the throw results in any number one through five, there will be an eighth period, and the number obtained by the throw of the die will be recorded in the right column on the Decision Sheet. Then the die will be thrown again after the eighth period, and a six will end the experiment. The die is thrown after each subsequent period to determine whether the experiment continues or not, and the probability that it will terminate is 1/6 in each case.

When the random device determines that an experiment is terminated, you will start a new experiment with a different person as the other seller. At this time, note your position number, which is written on your money plate. At the beginning of each new experiment you will be told the position number of the other seller in your market.

As you participate in the experiment, it is very important that you not communicate with other subjects who may be in the same room. This means that you will have to suppress elation, disgust, or other emotions, the expression of which may reveal how you feel about outcomes during the experiment. It will do you no good to try to influence the behavior of another person in the room or to try to observe another person's output choices, because the other sellers in your market are seated in other rooms.

However, we still ask that no communication occur between subjects, since the experiment becomes useless for our purposes if communication occurs.

We request also that you not talk to other persons about any details of the experiment after you leave. They might participate in later experiments and be influenced to play differently. Since the experiments are all different, this could work to their disadvantage, and it will bias our results as well.

Are there any questions?

2. Multiperiod experiment: instructions read to the subject

Have you finished reading the instructions? Are there any questions? As you can see, there are ___ of you in this room. There is also another room nearby with ___ people who are students like yourself. In the first market experiment, each of you is matched with one of the people in the other room, and each of them is matched with one of you. Thus there are ___ pairs of people in this market experiment. If the throw of the die causes this market experiment to end early, there will be another market experiment in which each of you is paired with a different person. In total, the session will last about 2 hours. Are there any questions? If not, go ahead and mark your output choice for the trial period that begins the experiment.

Instructions to be read to the subjects before the second market experiment

For the second experiment, the identity of the other seller has changed, as you can see on your Decision Sheet. Thus, each of you is now matched with a different person in the other room, and each person in the other room is matched with a different person in this room. The procedure for the second experiment will be the same as that of the first, and as before, we will begin throwing dice after period seven to determine when the experiment terminates. There will be no trial period this time, so you may now mark your output decision for period 1, which begins the second market experiment.

3. Single-period experiments: instructions read by the subjects (substitute the following three paragraphs for the fifth through ninth paragraphs in the multiperiod instructions read by the subjects)

Profits and losses will be determined by both your own and the other seller's output choices in each "decision period." The identity of the other seller changes after each decision period; see the "position number of other seller" column on the attached Decision Sheet. During each decision period, you and the other seller who is matched with you for that period will choose outputs from the choices available in the profit table. You will record your output decision on the Decision Sheet. Each period, we will collect your Decision Sheet, record the other seller's choice, determine payoffs, and return the sheet. While we have your

Decision Sheet, please note reasons for your output choice on the Explanation Sheet.

The experiment will begin with a single "trial period," in which you and the other seller matched with you for the trial period will make a decision. Then we will record the other seller's output choice on your Decision Sheet, and we will let you use the payoff table to compute the profit or loss for each of you. Someone will check your calculations to be sure that you understand how to read the payoff tables. Profits will not be paid and losses will not be collected for the trial period. After each subsequent decision period, we will collect output choices, compute profits, and pay your profit or take away your loss. Again, note that the position number of the other seller changes after each decision period. The experiment will end after you have been paired once with each of the other sellers.

As you participate in the experiment, it is very important that you not communicate with other subjects who may be in the same room. This means that you will have to suppress elation, disgust, or other emotions, the expression of which may reveal how you feel about outcomes during the experiment. We ask that no communication occur between subjects, since the experiment becomes useless for our purposes if communication occurs.

4. Single-period experiments: instructions read to the subjects

The participants in this experiment are students like you. Participants are located in this room and in another room nearby. In the trial period, each of you is matched with one of the people in the other room, and each of them is matched with one of you. In period 1, which follows the trial period, each of you will be matched with a different person. This switching continues so that for each of you, the identity of the other seller changes each period. The experiment will end after you have been paired once with each of the other sellers.

TABLE 1

PROFIT TABLE (in pennies)

Your Output (columns); Other Seller's Choice (rows)

Other Seller's Choice	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22
22	41	37	33	27	21	13	5	-5	-15	-27	-39	-53	-67	-83	-99	-117	-135	-155	-175
21	43	40	36	31	25	18	10	1	-9	-20	-32	-45	-59	-74	-90	-107	-125	-144	-164
20	45	42	39	34	29	22	15	6	-3	-14	-25	-38	-51	-66	-81	-98	-115	-134	-153
19	47	45	42	38	33	27	20	12	3	-7	-18	-30	-43	-57	-72	-88	-105	-123	-142
18	49	47	45	41	37	31	25	17	9	-1	-11	-23	-35	-49	-63	-79	-95	-113	-131
17	51	50	48	45	41	36	30	23	15	6	-4	-15	-27	-40	-54	-69	-85	-102	-120
16	53	52	51	48	45	40	35	28	21	12	3	-8	-19	-32	-45	-60	-75	-92	-109
15	55	55	54	52	49	45	40	34	27	19	10	0	-11	-23	-36	-50	-65	-81	-98
14	57	57	57	55	53	49	45	39	33	25	17	7	-3	-15	-27	-41	-55	-71	-87
13	59	60	60	59	57	54	50	45	39	32	24	15	5	-6	-18	-31	-45	-60	-76
12	61	62	63	62	61	58	55	50	45	38	31	22	13	2	-9	-22	-35	-50	-65
11	63	65	66	66	65	63	60	56	51	45	38	30	21	11	0	-12	-25	-39	-54
10	65	67	69	69	69	67	65	61	57	51	45	37	29	19	9	-3	-15	-29	-43
9	67	70	72	73	73	72	70	67	63	58	52	45	37	28	18	7	-5	-18	-32
8	69	72	75	76	77	76	75	72	69	64	59	52	45	36	27	16	5	-8	-21
7	71	75	78	80	81	81	80	78	75	71	66	60	53	45	36	26	15	3	-10
6	73	77	81	83	85	85	85	83	81	77	73	67	61	53	45	35	25	13	1
5	75	80	84	87	89	90	90	89	87	84	80	75	69	62	54	45	35	24	12
4	77	82	87	90	93	94	95	94	93	90	87	82	77	70	63	54	45	34	23

Table 2. First-Market Output Choices for Subjects S1-S24

Subject S1 was paired with S2, S3 with S4, etc.

Period	S1	S2	S3	S4	S5	S6	S7	S8	S9	S10	S11	S12
1	10	6	10	8	8	10	12	8	10	8	10	10
2	9	10	8	10	8	8	14	9	9	10	4	8
3	10	11	8	6	7	7	13	6	11	9	10	10
4	8	4	6	7	7	9	13	7	9	8	4	8
5	8	10	6	6	8	7	11	8	8	7	10	10
6	10	9	6	6	8	10	11	10	7	7	7	10
7	8	10	6	6	10	8	9	10	7	7	7	8
8	9	8	6	6	10	8	11	9	8	7	7	8
9	10	10	6	6	10	8	10	10	7	10	7	8
10	9	9	6	6	9	9	10	9	8	10	22	8
11	8	8	6	6	9	13	10	9	9	9	7	10
12	7	7	6	6	9	6	10	9	8	8	7	8
13	6	6	6	6	9	8	10	8	7	7	7	8

Table 2. First-Market Output Choices for Subjects S1-S24
(continued)

Period	S13	S14	S15	S16	S17	S18	S19	S20	S21	S22	S23	S24
1	8	9	7	10	8	7	9	10	6	9	8	4
2	7	8	9	13	6	8	5	9	6	8	9	5
3	6	9	6	7	8	6	9	9	6	8	7	9
4	9	8	7	9	8	9	8	8	6	8	8	13
5	8	8	7	8	8	6	6	9	8	8	11	9
6	7	8	14	6	7	9	8	9	8	8	10	8
7	8	8	8	11	8	6	10	9	8	6	9	7
8	8	8	7	5	8	8	9	8	7	6	8	7
9	8	8	6	6	8	9	7	8	6	7	7	8
10	8	8	6	5	6	10	9	8	6	6	7	8
11	8	8	6	6	8	7	8	8	6	6	8	7
12	8	8	6	6	8	6	8	8	6	6	7	8
13	8	8	6	6	8	8	8	8	6	6	8	6

Table 3. Second-Market Output Choices for Subjects S1-S16

Period	S1	S4	S3	S6	S5	S8	S7	S2	S9	S12	S11	S14	S13	S16	S15	S10
1	6	6	6	9	9	12	11	8	8	10	7	7	6	10	10	8
2	6	6	6	9	9	12	11	10	9	9	7	7	8	10	8	9
3	6	6	6	9	10	9	11	10	8	9	7	7	8	6	7	8
4	6	6	19	9	9	9	11	10	8	8	7	7	6	8	6	7
5	6	6	6	8	9	8	11	9	7	8	7	7	8	9	6	7
6	6	6	7	8	8	8	11	11	7	8	7	7	8	8	7	7
7	6	6	6	8	8	8	11	10	8	8	7	7	8	8	7	7
8	6	6	8	8	7	8	11	10	8	8	7	7	8	8	7	7
9	6	6	8	10	8	7	11	10	8	8	7	7	8	8	7	7

Table 4. Single-Period Experiments: Subjects' Output Choices With Rivals' Choices Shown in Parentheses

	Subjects											
Period	S25	S26	S27	S28	S29	S30	S31	S32	S33	S34	S35	S36
Trial	7(8)	8(7)	22(10)	10(22)	5(7)	7(5)	13(7)	7(13)	6(8)	8(6)	10(5)	5(10)
1	5(10)	8(7)	9(6)	10(5)	8(6)	6(9)	11(8)	6(8)	9(7)	8(11)	7(8)	7(9)
2	6(10)	8(9)	9(7)	10(7)	8(8)	10(6)	10(8)	7(9)	9(8)	8(8)	7(10)	8(10)
3	6(7)	8(10)	9(8)	9(9)	8(9)	8(8)	10(8)	7(6)	9(9)	8(9)	8(8)	9(8)
4	6(8)	9(8)	9(9)	8(9)	8(9)	8(9)	9(8)	6(8)	9(8)	8(6)	8(6)	9(9)
5	6(9)	9(9)	9(9)	9(9)	9(9)	8(9)	9(8)	8(9)	9(8)	8(8)	8(8)	9(6)
6	7(9)	8(9)	9(7)	9(8)	9(8)	9(8)	8(9)	8(9)	9(8)	8(9)	8(9)	9(8)
7	7(8)	9(9)	9(9)	9(9)	9(9)	9(9)	9(9)	8(8)	9(9)	8(8)	8(7)	9(9)
8	7(9)	9(8)	9(9)	9(8)	9(7)	8(9)	8(8)	8(9)	9(9)	8(9)	8(8)	9(8)
9	7(8)	9(8)	9(8)	9(9)	9(8)	8(8)	8(7)	8(9)	8(9)	8(8)	8(9)	9(9)
10	7(8)	9(8)	9(8)	8(8)	9(8)	8(10)	8(9)	8(8)	8(7)	8(9)	8(9)	10(8)

FOOTNOTES

* This research is partly funded by the National Science Foundation and the Sloan Foundation. Laura Cohen, Brad Hauck, and Anne Villamil assisted in setting up and administering the experiments. I am grateful to Dan Alger, Alfonso Novales, Robert Porter, and Joel Slemrod for comments and criticisms of an earlier draft.

[1] James Friedman (1977) discusses the existence of Nash equilibria in a general class of reaction function strategies, but one cannot actually compute nondegenerate equilibrium reaction functions for even the simplest quadratic payoff structures. More severe restrictions on the strategy spaces can produce results. For example, Richard Cyert and Morris DeGroot (1970) use backward induction to compute Nash equilibrium sequences of outputs for a finite horizon duopoly model in which firms make output decisions in alternate periods. Friedman's (1977) "balanced temptation equilibrium" is a Nash equilibrium for a supergame in which firms choose contingent strategies that specify an equilibrium output level and a commitment to a permanent switch to the firm's static Cournot output if another firm increases its output above its equilibrium level. Ed Green and Robert Porter (1981) have analyzed a stochastic generalization of this "balanced temptation equilibrium."

[2] Bresnahan (1981) shows that the consistent conjectural variations will be constants when the profit function is quadratic.

[3] See Morton Kamien and Nancy Schwartz (1981). If marginal costs are increasing or there is product differentiation, Bresnahan (1981) and Perry (1982) have shown that price can exceed average cost in a consistent-conjectures equilibrium.

[4] Dolbear, Lave, et al. (1968) also considered an "incomplete information" condition. The average level of price choices was approximately the same under each information condition, but there was less dispersion in the incomplete-information experiments. Their paper provides an interesting analysis of the effects of information and the number of sellers on the degree of tacit collusion.

[5] Roger Sherman warned me about using suggestive words, but I made the same mistake myself. In one of my pilot experiments, the term "oligopoly game" appeared on the receipt form to be completed by subjects at the end of the experimental session. This form was passed out at the beginning of the experiment, and one of the subjects who noticed the oligopoly phrase later remarked that the phrase "gave it away." He remembered seeing an assertion in a textbook that oligopolists would collude to maximize joint profit. This subject was in the only duopoly pair (out of four pairs) that was able to reach the collusive output combination in the first market experiment. All data from this pilot experiment were disregarded, and the wording of the receipt form was changed.

[6] See the profit table in their appendix IV.

[7] The industry outputs in period 21 were: 25, 30, 30, 32, 33, 38, 39, 40, 40, 44, 45, 49, 50, 55, 59, and 60.

[8] This is a serious limitation of the Fouraker and Siegel experiments because the main objective of these experiments seemed to have been to determine the proportions of duopoly pairs which could be best classified as either collusive, Cournot, or competitive. The competitive or "rivalistic" outputs of 30 probably did not have a chance. In a different context, Murphy (1966) has shown that truncation of the payoff table can have a major effect on experimental results.

[9] The output of 45 was implied by the profit-function parameters, but outputs of 16 for each subject constituted a Nash equilibrium for the payoff table that was used.

[10] Fouraker and Siegel also conducted duopoly and tripoly experiments with "incomplete information." The results of all of their experiments are summarized and discussed in Vernon Smith et al. (1982).

[11] An increase in the A parameter will increase the spread between the Cournot and collusive output decisions, but this will increase profits and make the experiments more expensive to run. The use of a fixed cost to lower all profit entries is not possible because the profit at the consistent-conjectures equilibrium should be sufficiently positive. A reduction in the B parameter will also increase the spread between the Cournot and collusive outputs, but the resulting flatness in the payoff structure results in multiple Cournot equilibria when profits are rounded off to the nearest penny.

FOOTNOTES (continued)

[12] See Plott (1981) for a discussion of the relationship between experience and collusion in laboratory experiments. Plott also has an excellent summary of the arguments for and against using laboratory experiments to test industrial-organization theories.

REFERENCES

Bresnahan, T. F. "Duopoly Models with Consistent Conjectures," American Economic Review 71 (December 1981):934-45.

Cyert, R. M. and DeGroot, M. H. "Multi-period Decision Models with Alternating Choice as a Solution to the Duopoly Problem," Quarterly Journal of Economics 84 (August 1970): 410-29.

Dolbear, F. T.; Lave, L. B.; Bowman, G.; Lieberman, A.; Prescott, E. C.; Rueter, F.; and Sherman, R. "Collusion in Oligopoly: An Experiment on the Effect of Numbers and Information," Quarterly Journal of Economics 82 (May 1968):240-59. Reprinted in The Journal of Reprints for Antitrust Law and Economics 10(1) (1980):415-36.

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

Friedman, J. W. Oligopoly and the Theory of Games. Amsterdam and New York: North-Holland, 1977.

Green, E. J., and Porter, R. H. "Noncooperative Collusion Under Imperfect Price Information," Discussion Paper No. 81-142. Center for Economic Research, University of Minnesota, February 1981.

Holt, C. A. "Equilibrium Models of Tacit Collusion in Oligopoly Experiments with Price-Setting Firms," Discussion Paper No. 80-138. Center for Economic Research, University of Minnesota, October 1980.

Kamien, M. I., and Schwartz, N. L. "Conjectural Variations," Managerial Economics and Decision Sciences, No. 466S, Northwestern University, March 1981.

Murphy, J. L. "Effects of the Threat of Losses on Duopoly Bargaining," Quarterly Journal of Economics 80 (May 1966): 296-313.

Perry, M. K. "Oligopoly and Consistent Conjectural Variations," Bell Journal of Economics 13 (Spring 1982):197-205.

REFERENCES (continued)

Plott, C. R. "Industrial Organization Theory and Experimental Economics." Social Science Working Paper 405. California Institute of Technology, September 1981.

Smith, V.L.; Williams, A. W.; Bratton, W. K; and Vannoni, M.G. "Competitive Market Institutions: Double Auctions vs. Sealed Bid-Offer Auctions," American Economic Review 72 (March 1982):58-77.

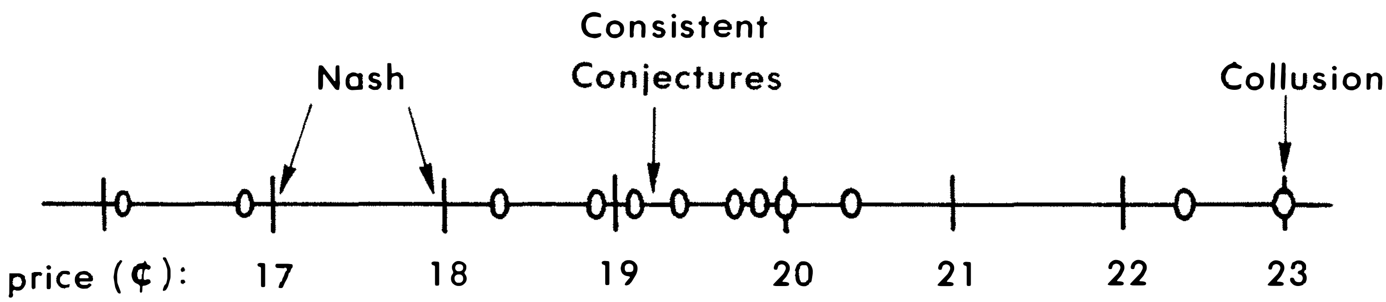

FIGURE 1. PRICE AVERAGES FOR THE TWELVE DUOPOLY EXPERIMENTS WITH COMPLETE INFORMATION REPORTED IN DOLBEAR, ET AL. (1968)

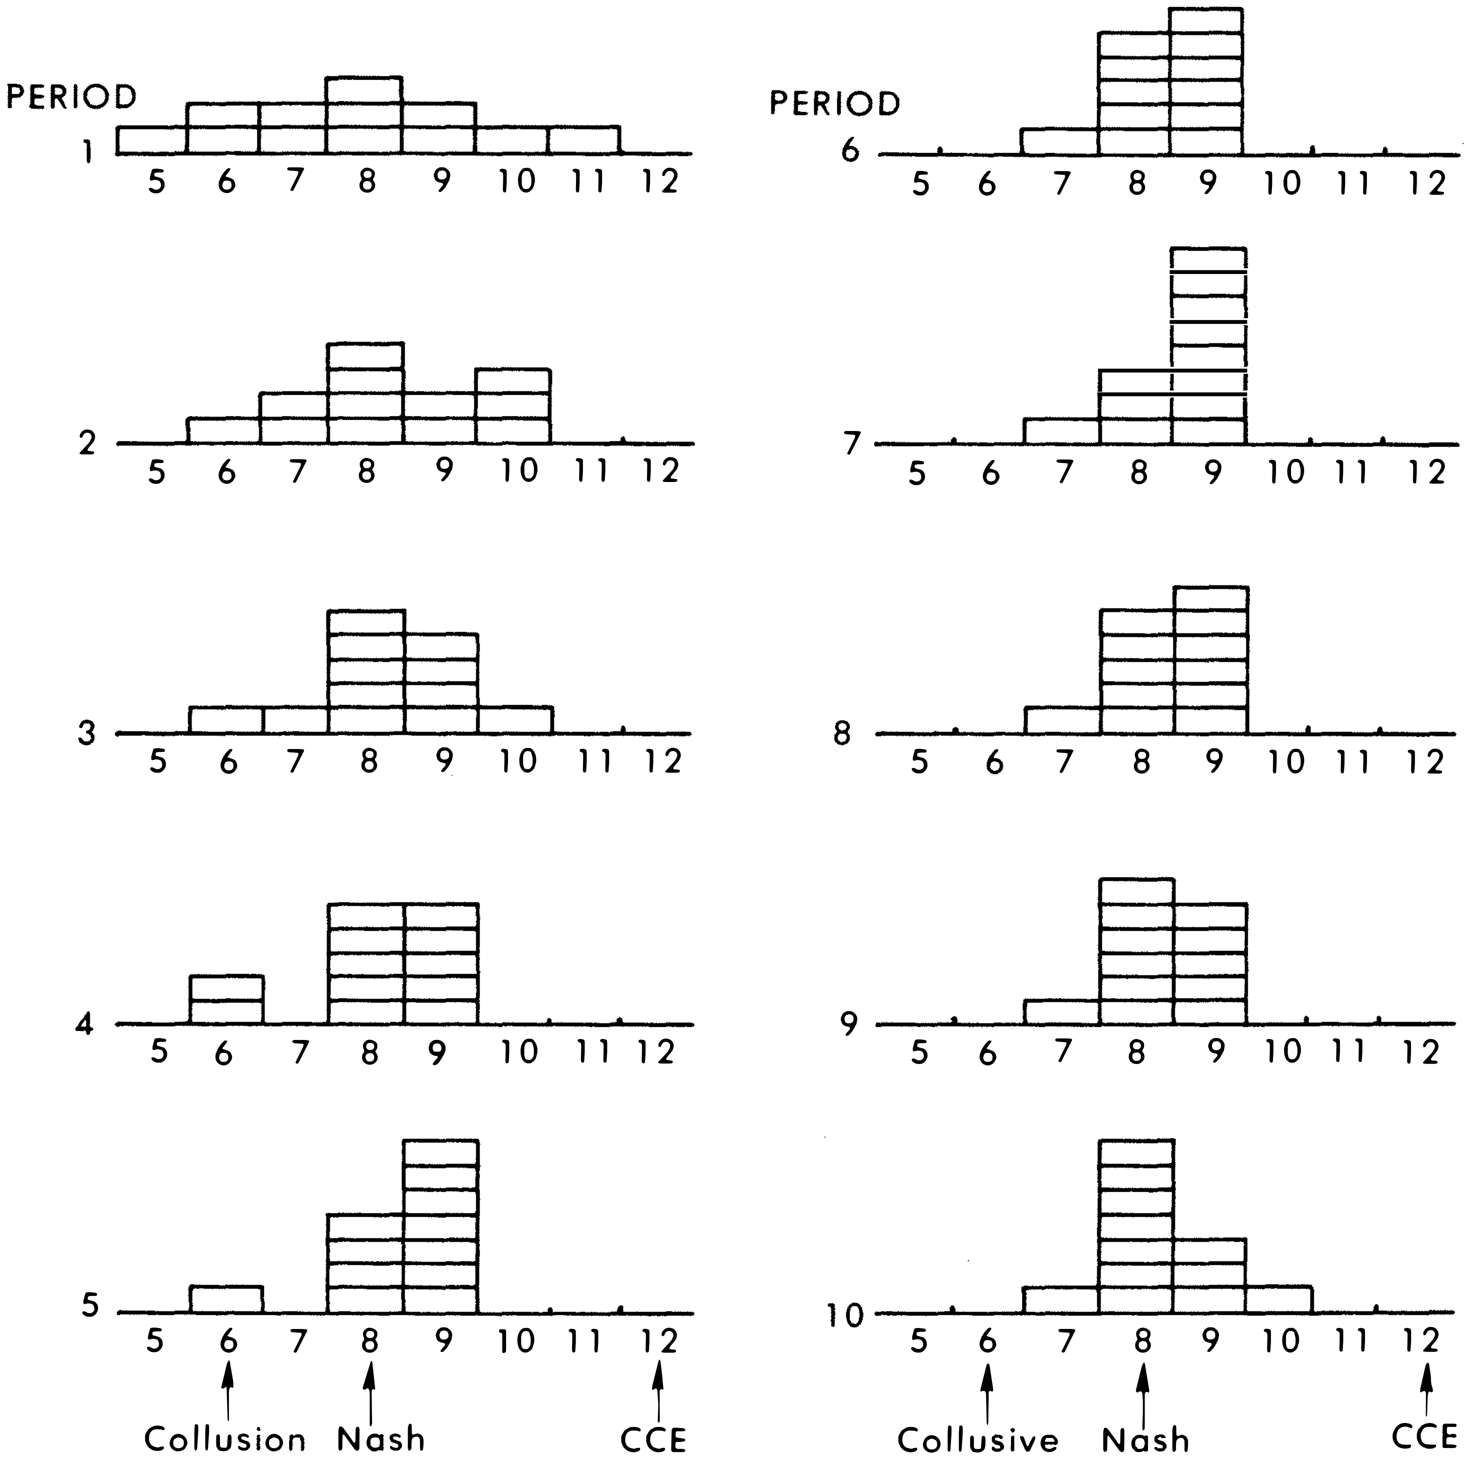

FIGURE 2. FREQUENCY OF INDIVIDUAL OUTPUT CHOICES
IN THE SINGLE-PERIOD MARKET EXPERIMENT