



Duopoly Strategies Programmed by Experienced Players

Reinhard Selten; Michael Mitzkewitz; Gerald R. Uhlich

Econometrica, Vol. 65, No. 3. (May, 1997), pp. 517-555.

Stable URL:

<http://links.jstor.org/sici?sici=0012-9682%28199705%2965%3A3%3C517%3ADSPBEP%3E2.0.CO%3B2-C>

Econometrica is currently published by The Econometric Society.

Your use of the JSTOR archive indicates your acceptance of JSTOR's Terms and Conditions of Use, available at <http://www.jstor.org/about/terms.html>. JSTOR's Terms and Conditions of Use provides, in part, that unless you have obtained prior permission, you may not download an entire issue of a journal or multiple copies of articles, and you may use content in the JSTOR archive only for your personal, non-commercial use.

Please contact the publisher regarding any further use of this work. Publisher contact information may be obtained at <http://www.jstor.org/journals/econosoc.html>.

Each copy of any part of a JSTOR transmission must contain the same copyright notice that appears on the screen or printed page of such transmission.

JSTOR is an independent not-for-profit organization dedicated to and preserving a digital archive of scholarly journals. For more information regarding JSTOR, please contact support@jstor.org.

DUOPOLY STRATEGIES PROGRAMMED BY EXPERIENCED PLAYERS

BY REINHARD SELTEN, MICHAEL MITZKEWITZ, AND GERALD R. UHLICH

The strategy method asks experienced subjects to program strategies for a game. This paper reports on an application to a 20-period supergame of an asymmetric Cournot duopoly. The final strategies after three programming rounds show a typical structure. Typically, no expectations are formed and nothing is optimized. Instead of this, fairness criteria are used to determine cooperative goals, called “ideal points.” The subjects try to achieve cooperation by a “measure-for-measure policy,” which reciprocates movements towards and away from the ideal point by similar movements. A strategy tends to be more successful the more typical it is.

KEYWORDS: Duopoly, strategy method, computer tournament.

1. INTRODUCTION

AFTER 150 YEARS SINCE COURNOT (1838) the duopoly problem is still open. An empirically well supported duopoly theory has not yet emerged. Field studies meet the difficulty that cost functions, demand functions, and prices are often insufficiently observable. Game playing experiments permit the control of these basic data. However, only plays are observed and strategies remain hidden. Usually, any given play of a duopoly supergame can be the result of a great multitude of strategy pairs.

More than 20 years ago, one of the authors described a method of experimentation which makes strategies observable (Selten (1967)). This procedure, called the “strategy method,” first exposes a group of subjects to the repeated play of a game, and then asks them to design strategies on the basis of their experiences. The strategy method was applied to an oligopoly situation with investment and price variation (Selten (1967)). In view of the special character of the dynamic oligopoly game investigated there, the issue of cooperation which will be important in the paper did not arise in this earlier study. Here we are concerned with a much more basic duopoly situation, namely a finite supergame of an asymmetric Cournot duopoly. Asymmetry is essential for this study, because we are interested in whether and how cooperation can evolve in such situations.

Cournot’s quantity variation model is the most popular one in the oligopoly literature. Many theories have been developed in this framework. Supergames of the Cournot model have also been explored in the newer game-theoretical literature (e.g., Friedman (1977), Radner (1980), Abreu (1986), Segerstrom (1988)). Therefore, it seems to be interesting to apply the strategy method to a supergame of an asymmetric Cournot duopoly.

Infinite supergames cannot be played in the laboratory. Attempts to approximate the strategic situation of an infinite game by the device of a supposedly fixed stopping probability are unsatisfactory since a play cannot be continued beyond the maximum time available. The stopping probability cannot remain

fixed but must become one eventually. Therefore, we decided to base our study on a finite supergame. The experimental literature shows that apart from the end effect there seems to be no significant behavioral difference between infinite and sufficiently long finite supergames (Stoecker (1983), Selten and Stoecker (1986)).

Our subjects were participants of a seminar who first gained experience in playing a 20-period supergame in the Bonn laboratory of experimental economics which is equipped with a network of personal computers. After having gained experience with the game, the participants had to program strategies. These strategies were played against each other in computer tournaments. The participants had the opportunity to improve their strategies in the light of their experience in such tournaments.

Our evaluation will mainly concern the strategies programmed for the final computer tournament. We shall only shortly report on some interesting phenomena observed in the initial game playing rounds and the intermediate tournaments.

The first step in the evaluation of the final tournament strategies was a classification according to structural properties. These properties, called "characteristics," were suggested by a close look at the strategies. We found 13 characteristics, all of which are present in the majority of cases to which they can be applied. A typical structure of a strategy emerges from these characteristics. The programs usually distinguish among an initial phase, a main phase, and an end phase. The initial phase consists of the first one to four periods with outputs depending only on the number of the period. In the main phase, outputs were made dependent on the opponent's previous outputs. By the initial phase the strategies try to prepare cooperation with the opponent to be reached in the main phase. In an end phase of the last one to four periods cooperation is replaced by noncooperative behavior.

Typically, the participants tried to approach the strategic problem in a way which is very different from that suggested by most oligopoly theories. These theories almost always involve the maximization of profits on the basis of expectations on the opponent's behavior. It is typical that the final tournament strategies make no attempt to predict the opponent's reactions and nothing is optimized. Instead of this, a cooperative goal is chosen by fairness considerations and then pursued by an appropriate design of the strategy. Cooperative goals take the form of "ideal points." An ideal point is a pair of outputs at which a player wants to achieve cooperation with his opponent. Such ideal points guide the behavior in the main phase. A move of the opponent towards the player's ideal point usually leads to responses which move the player's output in the direction of his ideal point. Similarly, a move of the opponent away from the ideal point is usually followed by a response which shifts the output away from the ideal point. We refer to this kind of behavior as a "measure-for-measure policy."

The fairness criteria underlying the selection of ideal points are different for different participants, but in most cases not completely arbitrary. Measure-for-

measure policies for the effectuation of ideal points may be quite different in detail, but they are all based on the same general idea.

On the basis of the 13 characteristics which express structural properties common to most of the strategies we have constructed a measure of typicality which is applied both to characteristics and strategies. The typicality of a strategy is proportional to the sum of the typicalities of its characteristics and the typicality of a characteristic is proportional to the sum of the typicalities of the final tournament strategies with this characteristic. It was an unexpected result of our investigation that there is a highly significant positive correlation between the typicality of a final tournament strategy and its success in the final tournament. Moreover, it turned out that for each of the 13 characteristics separately those strategies which have it are more successful than those which do not have it.

In order to get a better insight into the implications of the typical structure of final tournament strategies, we constructed a family of "simple typical strategies." In these strategies the details left open by the 13 characteristics are filled in the simplest possible way. The behavior in the main phase is described by a piecewise-linear continuous reaction function.

Two game-theoretical requirements on simple typical strategies are discussed: "conjectural equilibrium conditions" and "stability against short-run exploitation." These requirements impose restrictions on the ideal points. The first requirement is rarely satisfied but the second one is fulfilled by the vast majority of the ideal points used in final strategies. This condition also turns out to be of descriptive value for the profit combinations reached in the last tournament.

We do not claim that our results are transferable to real duopolies. First of all, it is doubtful whether a supergame of the Cournot duopoly is a realistic description of duopolistic markets. Nevertheless, the structure of behavior in such supergames is of great theoretical interest. Our results throw a new light on the duopoly problem posed in this framework. The choice of an ideal point by fairness consideration combined with the pursuit of this cooperative goal by a measure-for-measure policy constitutes a surprisingly simple approach which avoids optimization and the prediction of the opponent's behavior. The connection between typicality and success in the final tournament shows that this approach is not only simple and practicable but also advisable in the pursuit of high profits.

The participants of our seminar did not develop their strategy programs independently of each other. Interaction in the game playing rounds and the preliminary tournaments was unavoidable. It cannot be completely excluded that our results are due to a cultural evolution which might have a different outcome in a different experimental group. One application of the strategy method alone is not sufficient to establish a firm basis for far-reaching behavioral conclusions.

The tit-for-tat strategy which did so well in Axelrod's tournaments (Axelrod (1984)) is the natural consequence of the transfer of the strategic approach emerging from this study to the prisoner's dilemma supergame. There one finds only one reasonable ideal point, namely the cooperative choice taken by both

players, and only one measure-for-measure policy fitting this cooperative goal, namely tit-for-tat.

More recently, a paper by Fader and Hauser (1988) reports on programs written for two symmetric price triopolies. The players had no opportunity to play the games before writing their strategies and submitted a program only once for each of both models. Fader and Hauser classified strategies according to "features," but it cannot be said that a typical structure emerges from this classification. Perhaps the lack of a typical structure is due to the fact that in comparison to our students the participants of the tournaments were much less experienced with their task. Maybe it is necessary to provide the opportunity to gain extensive game-playing experience and to permit repeated program revisions after preliminary tournaments in order to obtain strategies which show a typical structure.

Nevertheless, this study shows that strategies based on the measure-for-measure principle are very successful against the strategies submitted. The agreement of our findings with those of Axelrod and of Fader and Hauser confirms our impression that the pursuit of ideal points by measure-for-measure policies is more than the accidental result of an isolated study.

The model and the experimental procedure are described in Sections 2 and 3. Then the results of the game playing rounds and the results of the tournaments are discussed in Sections 4 and 5. The evaluation of the strategies programmed for the final tournament begins with Section 6. There the 13 characteristics are introduced and explained in detail. The strategic approach underlying typical strategies is discussed. Section 7 is devoted to the connection between typicality and success. A family of simple typical strategies is introduced in Section 8 as an idealization of the general pattern observed in the programmed strategies. Theoretical properties of these strategies are discussed and game-theoretic stability conditions are compared with the data of the final tournament. Section 9 looks at the implications of our results for duopoly theory. A summary of our findings is given in Section 10.

2. THE MODEL

The experiment is based on a fixed nonsymmetric Cournot duopoly with linear cost and demand functions. Strategies had to be programmed for the 20-period supergame of this Cournot duopoly. The duopolists were fully informed about cost and demand functions, the length of the supergame, and the opponent's decisions in past periods. The decision variable of duopolist i in period t is the quantity $x_i(t)$ for $i = 1, 2$ and $t = 1, \dots, 20$. Quantities must be chosen from nonnegative real numbers. The costs $C_1(t)$ and $C_2(t)$ of duopolists 1 and 2 and the price $p(t)$ in period t are given as follows:

$$\begin{aligned} C_1(t) &= 9820 + 9x_1(t), & x_1(t) &\geq 0, \\ C_2(t) &= 1260 + 51x_2(t), & x_2(t) &\geq 0, \\ p(t) &= \max(0; 300 - x_1(t) - x_2(t)). \end{aligned}$$

TABLE I
SOME THEORETICAL POINTS IN THE SOURCE GAME

Concept	Player 1's Output	Player 2's Output	Price	Player 1's Profit	Player 2's Profit
Cournot	111.0	69.0	120.0	2501.0	3501.0
Monopoly of player 1	145.5	0.0	154.5	11350.3	-1260.0
Monopoly of player 2	0.0	124.5	175.5	-9820.0	14240.3
Stackelberg with player 1 as leader	166.5	41.3	92.3	4041.1	441.6
Stackelberg with player 2 as leader	93.8	103.5	102.7	-1030.9	4096.1
Nash product maximum	86.8	49.5	163.7	3615.0	4313.5
Pareto optimum A of Figure 1	79.1	56.1	164.8	2503.8	5124.2
Pareto optimum B of Figure 1	94.8	42.7	162.5	4731.8	3501.0

The supergame payoff of each duopolist is the sum of his profits over all twenty periods.

Table I and Figure 1 show some theoretical features of the Cournot duopoly described above. The row "Nash product maximum" presents the output combination which maximizes the Nash product with the Cournot solution as fixed threat point. Point A in Figure 1 is the Pareto optimum which yields Cournot equilibrium profits for player 1. Analogously, B is the Pareto optimum which yields Cournot profits for player 2. Figure 1 shows that the model is quite asymmetric. Even point A is below the 45-degree line.

3. EXPERIMENTAL PROCEDURE

The experiment was performed in a seminar lasting over the whole summer semester 1987 at Bonn University, Federal Republic Germany. The subjects were 24 students of economics in the third or fourth year with some knowledge of micro- and macroeconomics and some experience with computer programming, but without special training in price theory and game theory. No introduction in these fields was given in the seminar and no references to the relevant literature was supplied. The seminar was organized in five plenary sessions, three rounds of game playing, and three computer tournaments of programmed strategies.

Plenary sessions: In the first plenary session the participants were informed about the organization of the seminar and the model presented in Section 2,

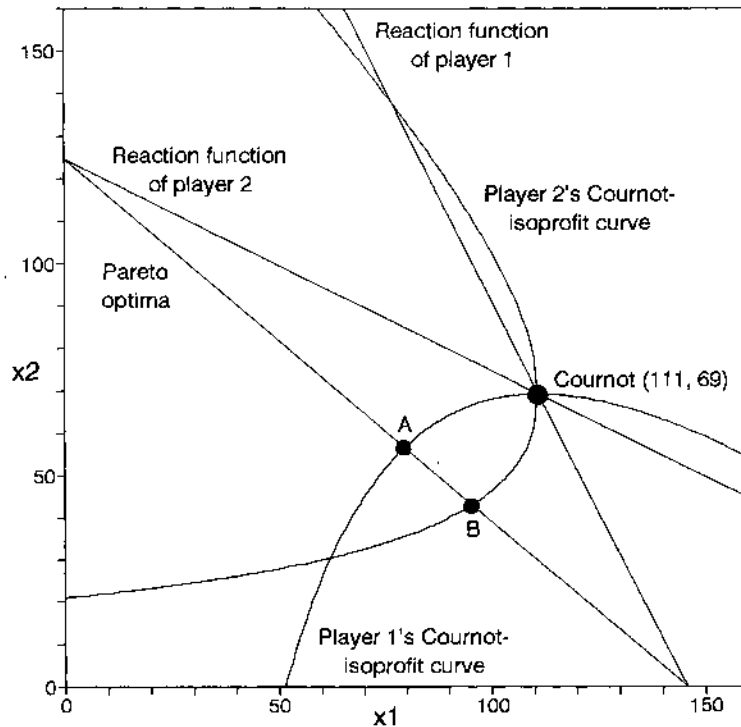


FIGURE 1.—Graphical representation of theoretical features of the one-period Cournot duopoly.

but, of course, without the theoretical features. Moreover, an introduction to the programming of strategies in PASCAL was given. It was not necessary to explain more than an excerpt of PASCAL, since strategies were conceived as subroutines in a game program.¹

The participants had the task first to gain experience by three rounds of playing the 20-period supergame and then to program strategies for both players in the 20-period supergame. They were told that their objective should be to attain a sum of profits as high as possible in a final tournament in which the strategies of all participants compete against each other. Final strategies had to be documented and reasons had to be given for the decisions embodied in the strategies.

The second plenary session took place after two rounds of game playing. The results of these games were presented, but in a way which left players anonymous. The participants were asked to comment on their experiences.

¹The Pascal source code of the students' strategies is available on request.

Each of the three tournaments was followed by one plenary session. Results were presented and students received printouts of the games in which their own strategies were involved. Opponents remained unidentified. The participants were encouraged to discuss strategic problems.

In the last of the five plenary sessions, the most successful participant explained his strategy. Anonymity was not completely preserved in this final plenary session at the end of the seminar.

Game playing rounds: Twenty-two subjects played three 20-period supergames against changing anonymous opponents, two subjects played only two supergames. The subjects were visually isolated from each other in cubicles containing computer terminals. The players interacted only by their decisions via the computer network. The decision time for each period was limited to three minutes. One week passed between one supergame and the next one. In this time the participants had the opportunity to reflect on their experiences. Each subject played with each of both cost functions at least once.

Strategy programming: After the game playing rounds the students had to program strategies in PASCAL for the 20-period supergame. Every student had to write a pair of strategies, one for each player of the supergame. We shall refer to this pair as the student's strategy. PC-owners could program at home, but all participants had the opportunity to develop their strategies at the Bonn laboratory of experimental economics with our technical assistance. A special program called "trainer" could be used by the students to play against their own programmed strategies. The "trainer" was a valuable tool for the development of strategies. No restrictions were imposed on strategies. Decisions could depend on the whole previous history of the play.

Computer tournaments: At three fixed dates the students had to hand in a programmed strategy. In the first two tournaments all workable strategies submitted at this date competed with each other. In the third tournament the last workable strategy of each participant was used. Each of the 24 students succeeded in writing at least one workable strategy.

The tournament program proceeded as follows: Let n be the number of workable strategies. Each of the n strategies played against all others in the role of both players. Payoff sums for player 1 and player 2 were computed on the basis of the $n - 1$ games played in the concerning role.

The procedure has the consequence that for each pair of strategies and each assignment of player roles, two supergames are simulated even if the payoff summation for one strategy makes use only of one of these games. Since sometimes random decisions are used in strategy programs, both games may be different. Altogether, $2n(n - 1)$ supergames were simulated in a tournament. The success of a strategy can be measured for the roles of both players separately by the corresponding payoff sums. The sum of these two measures is a measure for the overall success of a strategy in a tournament. This measure of overall success was the goal variable in the tournament. Strategies were ranked according to the measure of overall success, but also for the success of both player roles separately. Each participant received period-by-period printouts of

all $2(n - 1)$ games underlying the computation of his success measure. Moreover, all participants received lists of success measures, but without identification of the other writers of strategy programs. On the basis of this information the students could try to improve their strategy programs from one tournament to the next one.

Motivation: In view of the length of the experiment, it was not possible to provide an appropriate financial incentive. Presumably, money payoffs in the framework of a student seminar are not legal anyhow. The students were told that their grades would strongly depend on their success in the last tournament. It was emphasized that the absolute payoff sum rather than the rank was important in this respect. We had the impression that for almost all participants the task itself provided a high intrinsic motivation.

4. RESULTS OF THE GAME PLAYING ROUNDS

In this section we give a brief summary of the results of the game playing rounds. The games served the purpose to provide experiences which could be used in the development of strategy programs. Of course, it is plausible to assume that the subjects were intrinsically motivated by the game payoffs, but it is also possible that some of the behavior in these games was exploratory rather than directly payoff-oriented. Nevertheless, it is interesting to look at the results of the game playing rounds. However, our discussion will not be very detailed because our main interest is in the investigation of the final strategy programs.

First game playing round: Although the participants had been informed one week in advance about the structure of the game, their behavior seemed to be confused. Figure 2 shows the supergame payoffs of the 11 groups (two partici-

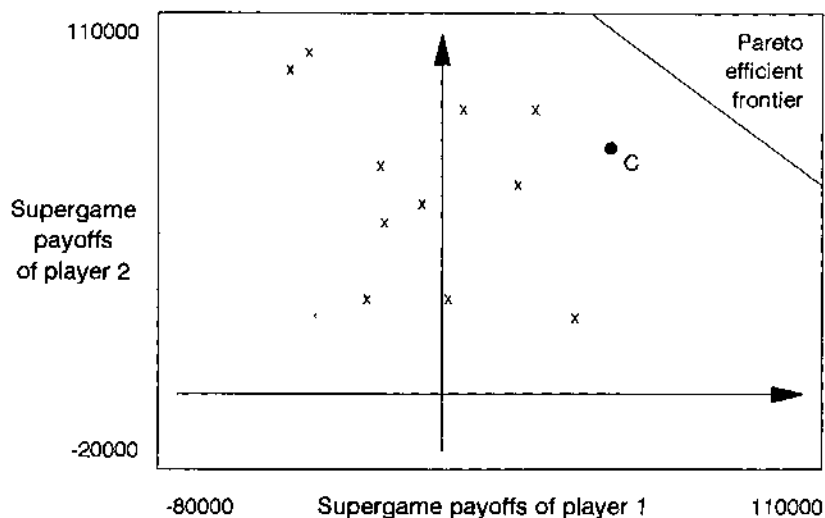


FIGURE 2.—Supergame payoff pairs in the first game playing round.

pants were absent). The repeated Cournot solution (point C in the diagram) yields 50020 for player 1 and 70020 for player 2. It never happened that *both* players achieved at least their Cournot payoffs. Furthermore, in all 11 cases the sum of both supergame payoffs was below the sum of the Cournot payoffs. In seven cases both players earned less than the Cournot payoffs. The role of player 1 (low variable and high fixed costs) was relatively less successful than the role of player 2. In the mean, subjects in the role of player 1 earned 79% of the Cournot gross profit (gross profit is profit plus fixed costs), whereas the corresponding figure for player 2 is 91%. The correlation coefficient between the payoffs of the two players within the groups is $-.36$. This suggests that some players succeeded to exploit their opponents. Figure 2 also shows part of the Pareto efficient frontier.

Second game playing round: The results of the second game playing round are shown in Figure 3. Here, two groups reached a Pareto improvement over the Cournot payoffs. In one group both players supplied the Cournot outputs in almost all periods. "It's the best thing you can do," they commented afterwards. In the remaining nine groups, both players sustained a loss in comparison with the Cournot solution. The mean gross profits of subjects in the role of player 1 was higher than in the first game playing round (87% of the Cournot gross profit), but the mean gross profit of subjects in the role of player 2 was lower than in the first game playing round (77% of the Cournot gross profit). The mean joint profit of both players was only slightly improved compared with the first game playing round.




FIGURE 3.—Supergame payoff pairs in the second game playing round.

There is one striking difference to the first game playing round. In the second game playing round, the correlation coefficient between both players' payoffs is now $+ .91$. This suggests that in the second game playing round the aggressiveness of both players shows a stronger coordination than in the first one. Even if most of the subjects did not yet succeed to play the game well, they seemed to have learned something about the power relationship in the game.

Third game playing round: This round shows an enormous improvement in mean payoffs (Figure 4). Now, subjects in the role of player 1 achieved 101% of the Cournot gross profit and the corresponding figure for those in the role of player 2 is 107%. Eight of the twelve groups succeeded to obtain Pareto improvements over the Cournot payoffs. One group reached a result almost at the Pareto efficient frontier. This group was the only one among those with Pareto improvements over the Cournot payoffs which did not show an end effect. The end effect consists in the breakdown of cooperation in the last periods of the supergame. It is clear that payoffs at the Pareto efficient frontier cannot be achieved if an end effect occurs.

The correlation coefficient between the payoffs of both players is $+ .72$ in the third game playing round. In this respect, the third game playing round is similar to the second one.

It is clear that most of the subjects had learned to cooperate in the supergame in the third game playing round. The results of the three game playing rounds are not dissimilar to those obtained in other experimental studies where finite supergames were repeatedly played against changing anonymous opponents (Stoecker (1983), Selten and Stoecker (1986)). Subjects tend to learn to cooperate but they also learn to exhibit end effect behavior.

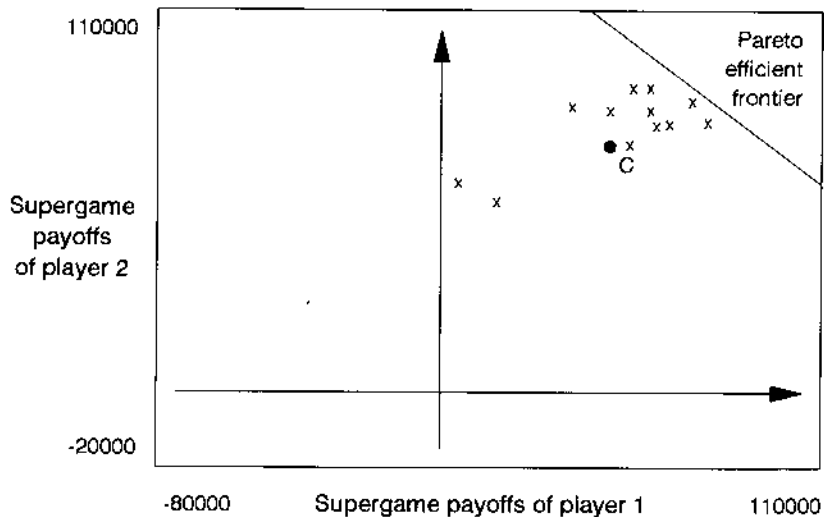


FIGURE 4.—Supergame payoff pairs in the third game playing round.

5. RESULTS OF THE TOURNAMENTS

In the following section we shall discuss the results of the tournaments without giving a detailed account of the strategies used. The typical structure of the strategies of the final tournament will be examined in the next section.

First tournament: Two weeks after the third game playing round the participants had to hand in a programmed strategy for the supergame. Unfortunately, 4 of the 24 strategies had to be excluded from the first tournament since programming errors like dividing by zero or taking the root of a negative number prevented the execution of these programs. The outcome of the first tournament is presented in Figure 5. The significance of the points in Figure 5 is not the same as in Figures 2, 3, and 4. A point now shows the combination of mean payoffs achieved by one participant's strategy in both player roles. Moreover, a larger scale has been chosen. One of the 20 strategies competing in tournament 1 is not shown in Figure 5 since it achieved a very bad result, namely $(-6484, +58178)$, which is outside the scope of the drawing. We omitted this point in order to be able to present the results of all three tournaments with the same scale without losing the distinguishability of different points.

The participant with the omitted bad result programmed a strategy which supplied the respective Stackelberg leader output each period regardless of the behavior of the other player. Only a few times he succeeded in forcing his opponent to the Stackelberg follower position. In most cases his "aggressive" behavior was punished by high opponent's outputs.

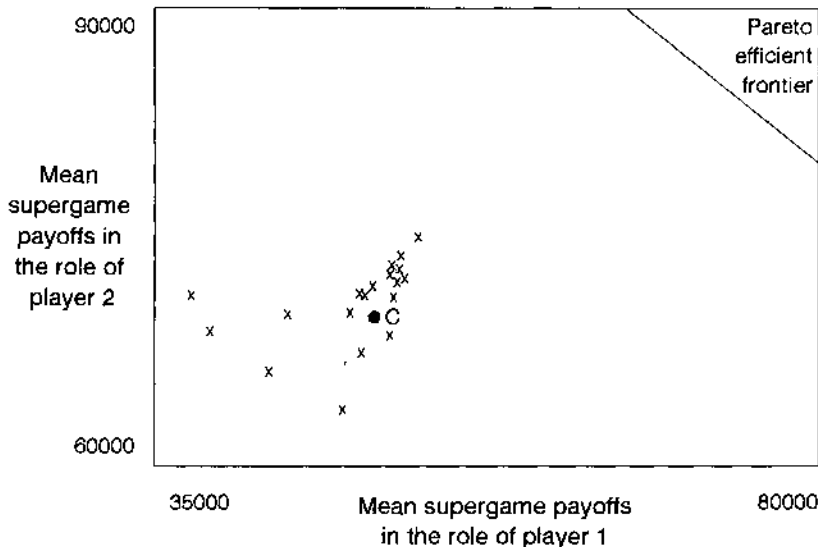


FIGURE 5.—Mean supergame payoffs for both player roles in the first tournament. Each "X" refers to one participant.

The mean gross profit over the whole simulation was 98% of the Cournot gross profit for the role of player 1 and 99.8% of the Cournot gross profit for the role of player 2. The mean performance is inferior to the third game playing round. Maybe the subjects did not yet succeed sufficiently well to mold their game playing intuition into computer programs.

Second tournament: Within three weeks after the first tournament the participants had the opportunity to improve their strategies. Unfortunately, this time only 16 participants presented workable strategies. In the same way as in Figure 5, the results are shown in Figure 6. One point, namely (23860, 63691) is omitted in Figure 6. Each of the other 15 subjects achieved results higher than Cournot payoffs in both player roles. The mean gross profit was now 104% of Cournot gross profit for Player 1 and 109% of Cournot gross profit for player 2. This is a considerable improvement in comparison with the first tournament. It must be admitted, however, that the comparison with the first tournament is difficult in view of the smaller number of workable strategies. Moreover, the result of the second tournament is also influenced by a "conspiracy" of two subjects represented by the two points nearest the right border of Figure 6. In the first period both participants used special outputs specified up to many decimal places in an unusual way. With the help of this code they recognized each other when they played together in the tournament. They then played in the remaining periods the output combination that maximizes joint profits. In order to prevent this type of behavior in the final tournament, we replaced the 8th digit behind the decimal point of each output decision by a random number.




FIGURE 6.—Mean supergame payoffs for both player roles in the second tournament. Each "X" refers to one participant.

This has a negligible influence on the computation of profits. In the plenary session after the second tournament, we announced that similar conspiracies will be prevented in the future. We did not observe any attempt to conspire in the final tournament.

Third tournament: After two more weeks the final strategies had to be turned in. Again four participants did not succeed to program a workable strategy. Fortunately, each of these participants had completed at least one workable strategy in the two preceding tournaments. The last workable strategy entered the final tournament.

A superficial examination of the programs revealed that one strategy consisted of two sequences of fixed outputs for every period, one sequence for each player. The numbers varied unsystematically from period to period. The seminar paper of this student loosely described a completely different strategy which was much more reasonable. Obviously, this student wanted to avoid investing time and effort into the programming of the strategy described in his paper. The irregularity of the output sequences served the purpose of hiding the discrepancy between the program and its description in the seminar paper. Obviously, the programmed strategy cannot be taken seriously and therefore has been excluded from the third tournament for the purposes of this paper.

The mean gross profit was 105% of the Cournot gross profit for player 1 and 111% of the Cournot gross profit for player 2. These figures are only slightly higher than those of the second tournament. Figure 7 shows the results of the third tournament. Computations of standard deviations of mean payoffs confirm the visual impression that the points in Figure 7 are more strongly concentrated than those in Figure 6.

In 983 of the 1012 supergames simulated in the third tournament, the payoffs of *both* players were greater than their Cournot payoffs. In this sense, we can speak of successful cooperation in 97.1% of all cases. It is also worth mentioning that in none of the remaining 29 supergames did *both* players obtain smaller payoffs than their Cournot payoffs.

In the third game playing round only eight out of twelve supergames resulted in payoffs which were greater than the corresponding Cournot payoffs for both players. The comparison with the results of the third tournament shows that the final programmed strategies tend to be much more cooperative than the behavior in the third game playing round. This suggests that the learning process which began with the three game playing rounds was continued in the three tournaments. The results of the third tournament do not seem to be very different from that which could be expected as the outcome of spontaneous game playing after a comparable amount of experience.

6. THE STRUCTURE OF PROGRAMMED STRATEGIES

In this section we shall concentrate our attention on the structure of the final strategies. We shall not be concerned with the success of the strategies. For the reasons which have been discussed in Section 5 (third tournament), one of

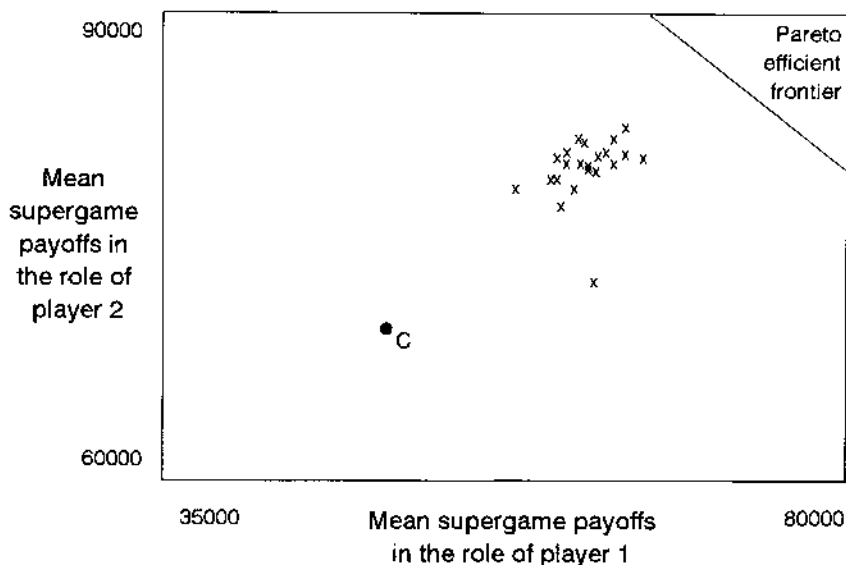


FIGURE 7.—Mean supergame payoffs for both player roles in the third tournament. Each “X” refers to one participant.

the programs will not be considered here. The remaining 23 programs and the underlying ideas expressed in the seminar papers are the basis of the evaluation of structural properties.

A preliminary examination of the strategies and the seminar papers conveyed the impression of a typical structure which is more or less present in almost all programs. Most programs deviate from this typical structure in some respects but the degree of conformity is remarkable.

Usually a program distinguishes three phases of the supergame: an initial phase, a main phase, and an end phase. The initial phase consists of one to four periods and the end phase is formed by the one to four last periods. The main phase covers the periods between the initial phase and the end phase. Different methods of output determination are used in the three phases. The initial phase is characterized by fixed outputs which do not depend on the behavior of the opponent. In the main phase the decisions are responsive to previous developments with the purpose to establish cooperation. In the end phase decisions are guided by the attempt to maximize short-run payoffs.

Different strategies approach the decision problems of the three phases in different ways, but nevertheless a typical structure emerges in this respect, too. In order to describe similarities and differences among the 23 strategies, we introduce 13 *characteristics*. A characteristic is a property of a strategy whose presence or absence can be objectively determined by the examination of a program and its description in the seminar paper. In some cases our characteristics are indicators of strategic ideas underlying the program; in other cases the

characteristics directly refer to the structure of decision rules. We shall distinguish characteristics concerning general principles and the three phases of the supergame.

All characteristics are typical in the sense that they are present in the majority of all strategies to which they can be meaningfully applied. Characteristic 7 is meaningful only if Characteristic 6 holds, too, and Characteristics 12 and 13 presuppose that the strategy has an end phase. These three characteristics are present only in the majority of all relevant cases. All other characteristics hold for the majority of all final strategies.

6.1. *General Principles*

The first three characteristics are indicators of general principles underlying the typical approach to the problem of writing a strategy program.

CHARACTERISTIC 1: *No prediction.*

Many oligopoly theories proceed from the assumption that a player has a method to predict his opponent's behavior and tries to optimize against his predictions. The predictions may involve reactions to own output changes and the payoff maximization may be long-term rather than short-term. In the final tournament, only 5 of 23 strategies involved any predictions of the opponent's behavior.

In the first two tournaments, predictions were more widespread. Subjects tried to obtain a satisfactory payoff against the predicted output of the opponent in the next period. Several subjects who initially wrote programs involving predictions later expressed the opinion that it is useless to try to predict the opponent's behavior. It seems to be more important to react in a way which indicates willingness to cooperate and resistance to exploitation.

The fact that the absence of any predictions is a typical feature of final strategies seems to be of great significance, precisely because it is in contrast with most oligopoly theories.

CHARACTERISTIC 2: *No random decisions.*

At the beginning of the seminar we observed that several students preferred to build random decisions into their strategies. They motivated this by the belief that a deterministic strategy could possibly be outguessed and exploited by the opponent. In the course of the seminar, most of them learned that in an attempt to achieve cooperation, it is important to signal one's intentions. It may be preferable to be outguessed by the opponent. Cooperation requires reliability and random decisions may be counterproductive in this respect. Twenty-two of the 23 final strategies never make a random decision.

CHARACTERISTIC 3: *Non-integer outputs.*

It is natural that real persons playing at computer terminals use mostly integer outputs. This was actually the case in the game playing rounds. Usually, a programmed strategy employs functions which make the output dependent on previous quantities. In general, the values obtained are not integers. However, four of the final strategies did not specify such functions but rather made case distinctions; for each case a different integer output or integer output change was prescribed. Since only relatively few cases are distinguished, this way of programming output decisions is less flexible than the specification of a function. In this light, Characteristic 3 is an indicator of smoothness and flexibility of the response pattern.

6.2. *The Initial Phase*

Two characteristics describe the typical behavior in the initial phase.

CHARACTERISTIC 4: *Fixed outputs for at least the first two periods.*

If no randomization takes place the first period is always a fixed output. Therefore, Characteristic 4 is almost equivalent to a nontrivial initial phase where fixed outputs are chosen. Ten strategies make their decision for the second period dependent on their opponent's choice in the first period, but 13 strategies have fixed amounts for more than one period. The length of the initial phase with fixed outputs is two periods for seven strategies, three periods for four strategies, and four periods for two strategies. Twelve of the 13 strategies with nontrivial initial phases play successively reduced outputs. The participants explained this behavior as a signal of their willingness to cooperate. If one's own output is a response to that of the opponent too early, an unsatisfying decision of the opponent in the first period could lead to an aggressive reaction of oneself in the second period that again could annoy the opponent and so forth, so that no cooperation might evolve over the 20 periods. Some subjects observed such unfavorable oscillations in the printouts of the first two tournaments.

CHARACTERISTIC 5: *The last fixed output decision is at least 8% below the Cournot quantity of the concerning player.*

The percentage by which the last fixed output in the initial phase is below the Cournot output can be regarded as a rough measure of a strategy's initial cooperativeness. A Pareto optimum is reached if both players' outputs are about 24.5% below the Cournot output. The criterion of the 8% limit of Characteristic 5 goes roughly a third of the way towards this Pareto optimum. Admittedly, it is arbitrary to measure cooperativeness by percentages of the Cournot output and to fix the limit at exactly 8%. Characteristic 5 is present in 13 of the 23 strategies. If the limit were increased to 10%, only a minority of 10 strategies would meet the criterion.

6.3. *The Main Phase*

The decision rules for the main phase are the most important part of a strategy program. Characteristics 6 to 11 concern the main phase. The rules given there do not apply to the initial phase and the end phase. This will not be mentioned explicitly in the text of the characteristics.

Typically participants approached the problem of decision making in the main phase by first looking at the question of where cooperation should be achieved. They tried to find an output combination which gives higher profits than Cournot profits to both players and can be considered as a reasonable compromise between the interests of both players. An output combination of this kind which guides the decisions in the main phase will be called an "ideal point." Ideal points are usually not far away from Pareto optimality. They are often based on considerations of equity which will be described below. Some participants used different ideal points for the roles of both players.

In Characteristic 6 we shall speak of "decisions guided by ideal points." With these words we want to express that the strategy program makes explicit use of an ideal point in order to determine the next output as a function of the past history. This can be done in many ways. One possible method connects the ideal point and the Cournot point by a straight line segment in the quantity or profit space. The next output then matches the opponent's last output on the line segment as long as the opponent's last output is in the range where this is possible.

CHARACTERISTIC 6: *Decisions are guided by one or two ideal points.*

The property expressed by Characteristic 6 holds for 18 of the 23 final strategies. Twelve strategies use only one ideal point for both players, whereas 6 strategies specify different ideal points for the two player roles.

Table II gives an overview over the equity concepts underlying the construction of ideal points as far as such concepts could be identified on the basis of the seminar papers. The reasons for the choice of 10 of the 24 ideal points are at least partially unclear. To some extent ideal points were adapted to the learning experience of the first two tournaments and thereby shifted away from equity concepts.

The participants who based their ideal points on equity considerations often did not correctly compute the intended ideal points. They rarely used analytical methods but rather relied on more or less systematic numerical search. The values used instead of the correct ones are given in the footnotes below Table II.

The concept described by the first row of Table II looks at equal profit increases in comparison to Cournot profits as a fair compromise. The Pareto optimum corresponding to this idea is the intended ideal point. The concept of the second row requires profit increases proportional to Cournot profits at a Pareto optimal point.

TABLE II
CONCEPTS UNDERLYING IDEAL POINTS

Concept	Quantities		Number of Strategies
	Player 1	Player 2	
Maximal equal absolute additional profits compared to Cournot profits	85.61	50.50	3 ^a
Maximal profits proportional to Cournot profits	84.37	51.56	2 ^b
Profit monotonic quantity reduction along the straight line through the intersections of both Cournot-isoprofit curves	86.53	49.71	2
Profit monotonic quantity reduction proportionally to Cournot quantities	89.73	55.77	1 ^c
Maximal equal profits	89.70	47.01	2 ^d
Prominent numbers	85.00	50.00	2
	90.00	50.00	2
Unclear	—	—	10

^aApproximated by (85.50) in all three cases.

^bApproximated by (86.53) and (84.33, 51.55).

^cApproximated by (89.76, 55.80).

^dIn one case approximated by (89.0, 46.5).

The third and fourth rows of Table II involve a procedure referred to as *profit monotonic quantity reduction*. Along a prespecified positively inclined straight line through the Cournot point in the quantity diagram, quantities are gradually reduced as long as both profits are increased in this way. The output combination reached by the procedure is the ideal point. In the case of row 3 of Table II, the prespecified straight line connects the intersections of both Cournot isoprofit curves. In the case of row 4 the prespecified straight line connects the Cournot point and the origin.

The concept of row 3 yields a Pareto optimum even if Pareto optimality is not a part of the underlying idea. Contrary to this, profit monotonic quantity reduction proportional to Cournot quantities yields an ideal point which is not even approximately Pareto optimal.

The concept of maximal equal profits determines the Pareto optimum where both profits are equal. Obviously, this ideal point does not only depend on variable costs but also on fixed costs. The same is true for maximal profits proportional to Cournot profits. Two of the ideal points classified as unclear also were based on equal profits but without an attempt towards maximization.

Some participants chose pairs of prominent quantities as ideal points. Roundness in the sense of divisibility by 5 seems to be the prominence criterion. More detailed discussions of prominence in the decimal system can be found in the literature (Schelling (1960), Albers and Albers (1983), Selten (1987)).

Figure 8 shows the ideal points used by the final strategies. The ideal points are given as quantity combinations. In the quantity diagram the Cournot-isoprofit curves of the two players enclose a lens-shaped area. The ideal points used by final strategies are in a relatively small area in the middle of this lens. The mean of all ideal points is located at (87.02, 49.43). This combination is almost Pareto optimal.

Characteristics 7 to 11 are described as rules to be followed by a programmer of a strategy.

CHARACTERISTIC 7: If your opponent has chosen an output below his output specified by your ideal point, then choose your ideal point quantity in the next period.

If a strategy is based on two ideal points then the words "your ideal point" refer to the ideal points for the concerning player role. The interpretation of Characteristic 7 is simple. If your opponent is even more cooperative than required by your ideal point, then there is no reason to deviate from your own ideal point quantity. Ten of the 18 final strategies based on ideal points have this characteristic. However, some other strategies increase the output in the situation of Characteristic 7 in order to test the opponent's willingness to cooperate at a point more favorable for oneself.

The remaining characteristics will be applicable to strategies without ideal points, too. Even if a strategy is not based on an ideal point, it may still involve a

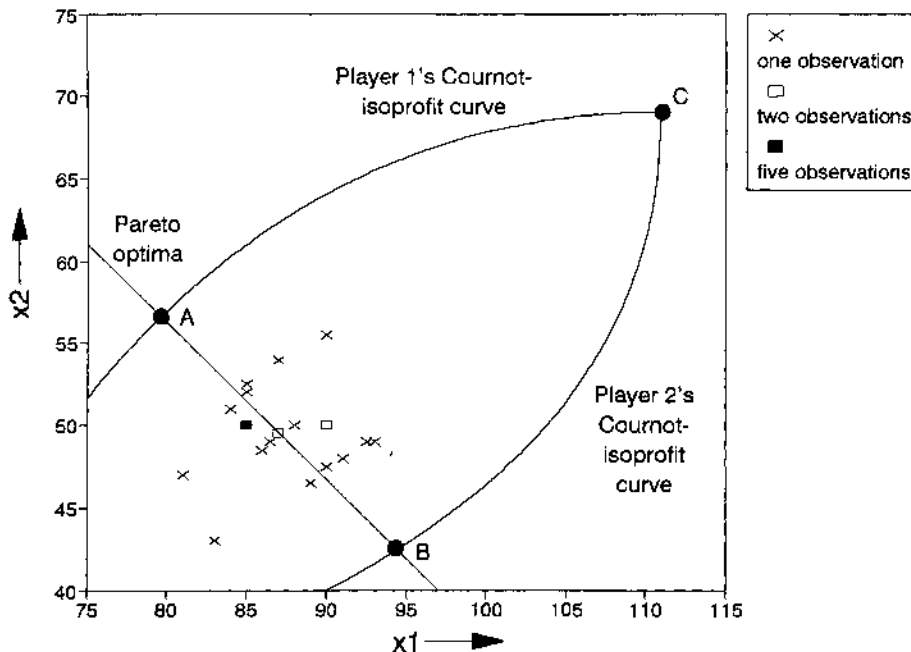


FIGURE 8.—Ideal points used by final strategies.

measure which permits an interpersonal comparison of cooperativeness. Thus a strategy may look at the profit difference achieved at the beginning of the main phase as a standard reference. Both players are judged to be equally cooperative if this profit difference is attained.

CHARACTERISTIC 8: If your opponent has chosen an output above his Cournot quantity, then in the next period choose your Cournot quantity.

Twelve of the 23 strategies obey this rule. The other strategies do not use more rigorous methods of punishment; instead, if they realize that their opponent plays permanently above his Cournot quantity, they abandon the idea of punishment after some periods and reduce their own output below their Cournot quantity to increase their profits. Such strategies run the danger of becoming exploitable by attempts to establish Stackelberg leadership. Characteristic 8 on the one hand avoids excessive aggressiveness and on the other hand provides protection against exploitative opponents.

CHARACTERISTIC 9: If your opponent has chosen his Cournot quantity, then in the next period choose a quantity not higher than your Cournot quantity and 5% at most below your Cournot quantity.

It can be seen with the help of Figure 8 that Characteristic 9 limits the response to the opponent's choice of his Cournot quantity to a relatively small interval. Sixteen of the 23 final strategies satisfy the requirement of Characteristic 9. Among these 16 strategies there are 10 which respond to Cournot quantities by Cournot quantities. The remaining 6 strategies want to indicate their willingness to cooperate by a slightly smaller output. Of course, the number of 5% in Characteristic 9 is to some extent arbitrary.

The following two Characteristics 10 and 11 apply to situations in which the following four conditions hold.

- (i) The last period was a period of your main phase.
- (ii) Up to now you always followed your strategy.
- (iii) In the last period your opponent's output was below his Cournot output.
- (iv) If you have an ideal point (for the relevant player role), then your opponent's output was above his output in your ideal point.

CHARACTERISTIC 10: Suppose that conditions (i), (ii), (iii) and (iv) hold. If in the last period your opponent has raised his output, then your decision raises your output to a quantity below your Cournot output.

CHARACTERISTIC 11: Suppose that conditions (i), (ii), (iii), and (iv) hold. If in the last period your opponent has lowered his output, then your decision lowers your output. If you have an ideal point, then your new output remains above your ideal point output.

To illustrate Characteristics 10 and 11, let us give an example: Consider a strategy which in the main phase matches the opponent's last output on a straight line between the Cournot point and an ideal point in the quantity space, of course, only as long as the opponent's last output was in the relevant range. A strategy of this kind satisfies Characteristics 10 and 11. However, it is necessary to impose condition (i) since in the first period of the main phase matching on the line may require an increase of output even if the opponent has lowered his output.

As long as condition (ii) is satisfied matching on the line in later periods of the main phase will move in the right direction. Conditions (iii) and (iv) make sure that Characteristics 10 and 11 apply only in the relevant range.

Both characteristics can be satisfied for strategies not based on a line between the Cournot point and an ideal point in any space. They may even be satisfied for strategies without ideal points. Thus a strategy's response may be guided by the criterion of a profit difference equal to that at the Cournot point without any regard to Pareto optimality. Two of the final strategies were of this kind.

Fourteen final strategies have Characteristic 10. The number of final strategies with Characteristic 11 is also 14, but only 11 final strategies have both characteristics.

6.4. *The End Phase*

A strategy with an end phase has a special method of output determination for the last one to four periods. Attempts towards cooperation which are typical for the main phase are not continued in the end phase. Instead of this, short-run profit goals are pursued.

Only 2 of the 23 final strategies do not have an end phase. One of these 2 strategies was typical in many other respects but the other was the most atypical. This atypical strategy tries to estimate response functions of the opponent and then computes the output decision by an elaborate approximative method for the solution of the dynamic program of maximizing expected profits for the remainder of the game. Even if something like an end effect is automatically produced by the dynamic programming approach, no end phase is present here since the method of output determination is always the same.

CHARACTERISTIC 12: *The strategy has an end phase of at least two periods.*

Characteristic 12 is shared by 11 of the 21 final strategies with end phases. Ten of these strategies planned an end effect only for the last period.

CHARACTERISTIC 13: *The strategy has an end phase and specifies the Cournot output of the relevant player as the output for all periods of the end phase.*

This characteristic is present in 12 final strategies. Other strategies sometimes optimized short-run profits against the opponent's last output or approached the Cournot output in several fixed steps.

6.5. *The Strategic Approach Underlying Typical Strategies*

A typical strategy does not try to optimize against expectations on the opponent's behavior (Characteristic 1). The strategic problem is not viewed as an optimization problem but rather as a bargaining problem. The first question to be answered concerns the point where cooperation should be achieved. Of course, cooperation should be favorable for oneself but it also must be acceptable for the opponent. A failure to reach cooperation is expected to lead to Cournot behavior. Therefore, cooperation requires that both players obtain more than their Cournot profits. Ideal points are constructed as reasonable offers of cooperation within these constraints. Various kinds of fairness considerations but also prominence (divisibility by five) and prior experience may influence the selection of ideal points.

After the choice of an ideal point the question arises as to how cooperation at this point or in its neighborhood can be achieved. It is necessary to indicate one's willingness to cooperate there and to show that one is not going to accept less favorable terms.

A decreasing sequence of outputs in the initial phase is a natural signal indicating cooperativeness. In the main phase a typical strategy evaluates the cooperativeness of the opponent's last output and responds by an output of a similar degree of cooperativeness according to some criterion. The response may depend on whether the opponent decreased or increased his output. If there is such a difference, it is natural to respond more aggressively to the same output after an increase.

One may say that main-phase behavior is guided by a principle of "measure for measure." Small changes of the opponent's output lead to small reactions and big changes cause big reactions.

Many oligopoly theories are based on the idea that a player anticipates the reaction of his opponent in order to maximize his profits. Contrary to this, a strategy based on an ideal point and a response rule guided by the principle "measure for measure" does not involve any anticipation of the opponent's reactions. The aim is to exert influence on the opponent rather than to adapt to his behavior. In order to achieve this aim one's own behavior has to provide a clear indication of one's own intentions. If the implied offer of cooperation is reasonable, one can hope that the aim will be reached. A response guided by the principle "measure for measure" protects against attempts to exploit one's own cooperativeness and rewards cooperative moves of the other player.

Of course, cooperation breaks down in the end phase. The strategies have been written for the 20-period supergame. This game permits only one subgame perfect equilibrium path, namely Cournot outputs in every period. The participants were aware of the backward induction argument which came up in the discussions of the plenary sessions. They accepted the idea that cooperation must break down in the last periods but as the strategies show they did not accept the full force of the backward induction argument. An explanation of this phenomenon is given elsewhere (Selten (1978a)).

7. TYPICITY AND SUCCESS

All characteristics are typical for the final strategies in the sense that they are present in the majority of the cases to which they are applicable. Of course, they are not all equally typical. Some appear in more of the final strategies than others. Moreover, the extent to which a characteristic is typical should not only be judged by the number of strategies with this characteristic, but also by the extent to which these strategies are typical. In the following, we shall construct a measure of typicity applicable to both characteristics and strategies which tries to do justice to these considerations.

The measure of typicity assigns a real number to each characteristic and to each strategy. The sum of the typicities of all 13 characteristics is normed to 1. The measure of typicity can be thought of as the outcome of an iterative procedure. At the beginning, all characteristics have the same typicity $1/13$. Then, in each step, first a new typicity is computed for each strategy as the sum of the typicities of its characteristics. Afterwards, a new typicity for each characteristic is computed as proportional to the sum of the typicities of the strategies with this characteristic. The sum of the typicities of all characteristics is again normed to 1.

In order to give a more precise mathematical definition of our measure, it is necessary to introduce some notation. The typicity of characteristic i is denoted by c_i and s_j stands for the typicity of strategy j . The symbol c is used for the column vector with the components c_1, \dots, c_{13} and s denotes the column vector with the components $s_1 \dots s_{23}$. Let A be the 13×23 -matrix with entries a_{ij} as follows: $a_{ij} = 1$ if strategy j has characteristic i , and $a_{ij} = 0$ otherwise. In our case c and s are uniquely determined by the following equations.

$$\begin{aligned} c &= \alpha As, \\ s &= A^T c, \\ \sum_{i=1}^{13} c_i &= 1, \end{aligned}$$

where A^T is the transpose of A and $1/\alpha$ is the greatest eigenvalue of AA^T . It is a consequence of elementary facts of linear algebra that the iterative process described above converges to vectors c and s which can be described as the solution of this system of equations.

Table III shows which strategy has which characteristics. The rows correspond to the 13 characteristics and the columns to the 23 final strategies. The strategies have been numbered according to the success in the final tournament. Strategy 1 is the most successful one, strategy 2 the second most successful one, etc. A black mark indicates that the strategy corresponding to the column has the characteristic corresponding to the row.

Obviously, the black marks in Table III describe the matrix A . A black mark corresponds to an entry 1 and the absence of a black mark corresponds to an entry 0. The typicities of the characteristics are given at the right margin and the typicities of the strategies can be found at the bottom of Table III.

opposite relationship would also seem to be possible. It is not inconceivable that typical characteristics reflect nothing else than typical mistakes. However, in our case the characteristics seem to embody advisable strategic principles. Maybe the positive correlation between typicity and success is the result of the learning process which produced the final strategies.

For each characteristic the mean success rank of those strategies which have it is smaller than the mean success rank of those which do not have it. This shows that *each of the characteristics separately is positively connected to the success in the final tournament*. In this sense all 13 characteristics are favorable structural properties of a strategy.

Our judgments of the advisability of the characteristics must be understood relative to the strategies developed by the participants of our experiment. We cannot exclude the possibility that a very atypical strategy can be found which turns out to be very good in a tournament against the 23 final strategies. In fact, the participant who wrote a strategy with success rank 20 firmly believes that this approximative dynamic programming approach based on an estimated response function of the opponent can be improved to a degree which will make it superior to all final strategies in a tournament against them. We doubt that this is the case. The difficulty with the dynamic programming approach is the problem of forming a correct estimate of the opponent's behavior. A best response to a wrong prediction can have disastrous consequences.

Admittedly, our experiment does not really justify strong conclusions since the final strategies have not been developed independently of each other. Perhaps a different picture of a typical strategy would emerge in a repetition of the experiment. Nevertheless, the results reported in this section seem to be of considerable significance for the further development of oligopoly theory.

8. A FAMILY OF SIMPLE TYPICAL STRATEGIES

The 13 characteristics do not completely determine a strategy. Many details are left open. In this section we shall construct a family of strategies which are typical in the sense that they have all 13 characteristics and the missing details are furnished in a particularly simple way. The members of the family differ only by the pair of ideal points used for both player roles. The special case of only one ideal point is not excluded.

For our family of simple typical strategies we shall discuss the question of what happens if two strategies with different ideal points play against each other. This exercise conveys some insight into the strategic properties implied by the 13 characteristics. We shall also look at the question of what is a reasonable choice of ideal points. For this purpose we have determined that member of the family which did best in a tournament against 22 of the final strategies. (The only strategy which involved random decisions was eliminated in order to avoid time consuming Monte-Carlo simulation.)

8.1. *Description of the Simple Typical Strategies*

The ideal points are described by output pairs u and v , one for each player role:

Ideal point for the role of player 1: $u = (u_1, u_2)$.

Ideal point for the role of player 2: $v = (v_1, v_2)$.

The first components of the vectors u and v denote player 1's output and the second stands for player 2's output. As mentioned above, the special case $u = v$ is not excluded. We also introduce the following notation for the output combination in the Cournot equilibrium of the underlying duopoly.

Cournot equilibrium: $c = (c_1, c_2)$.

We now can describe the decision $x_i(t)$ specified by the *simple typical strategy* with ideal points u and v . The following conditions (i) and (ii) have been imposed on the ideal points:

(i) The ideal points u and v are Pareto superior to the Cournot equilibrium.

(ii) $u_1 \leq .92c_1$ and $v_2 \leq .92c_2$.

Condition (ii) is necessary to make the specification of the initial phase compatible with Characteristic 5.

Initial phase:

$$x_1(t) = \frac{t}{3}u_1 + \frac{3-t}{3}c_1,$$

$$x_2(t) = \frac{t}{3}v_2 + \frac{3-t}{3}c_2 \quad \text{for } t = 1, 2, 3.$$

Main phase:

$$x_1(t) = \begin{cases} u_1 & \text{for } x_2(t-1) \leq u_2, \\ c_1 & \text{for } x_2(t-1) \geq c_2, \\ u_1 + \frac{c_1 - u_1}{c_2 - u_2}(x_2(t-1) - u_2) & \text{otherwise;} \end{cases}$$

$$x_2(t) = \begin{cases} v_2 & \text{for } x_1(t-1) \leq v_1, \\ c_2 & \text{for } x_1(t-1) \geq c_1, \\ v_2 + \frac{c_2 - v_2}{c_1 - v_1}(x_1(t-1) - v_1) & \text{otherwise.} \end{cases}$$

End phase:

$$x_i(t) = c_i \quad \text{for } i = 1, 2 \text{ and } t = 19, 20.$$

The initial phase can be thought of as a sequence of three equal "concessions" moving from the Cournot output c_i to the ideal point output u_1 or v_2 respectively. The first period already makes the first concession. Obviously, the initial phase satisfies Characteristic 4 which requires at least two periods. Characteristic 5 is satisfied since u_1 and v_2 are not greater than $.92c_1$ and $.92c_2$, respectively.

Characteristic 6 requires that the strategy make use of ideal points. Obviously, this is the case for our family of simple typical strategies.

We now turn our attention to the equation for the main phase. The upper line on the right-hand side secures Characteristic 7. The middle line is in agreement with Characteristics 8 and 9. Characteristic 9 concerns the special case $x_j(t-1) = c_j$ and permits a response $x_i(t)$ up to 5% lower than c_i . As has been pointed out before, the majority of these final strategies which conformed to Characteristic 9 specified a response of exactly c_i . Therefore, this response can be considered as typical.

The lower line on the right-hand side of the equation for the main phase is a very simple version of the principle "measure for measure." The last output of the opponent is matched by the corresponding output on the straight line which connects the ideal point and the Cournot point in the quantity space. Obviously, this has the consequence that Characteristics 10 and 11 are present in the strategies of our family.

The end phase has two periods and, therefore, conforms to Characteristic 12. The output in the end phase is always c_i , as required by Characteristic 13.

The strategies of our family also have the Characteristics 1, 2, and 3. In accordance with Characteristic 1, no attempt is made to predict the opponent's behavior and to optimize against this prediction. As required by Characteristic 2, the strategies are completely deterministic. In the main phase the strategies permit a continuum of possible responses and therefore have Characteristic 3.

8.2. Simple Typical Strategies Playing Against Each Other

Consider a play of the 20-period supergame where each of both players uses a member of the family described above as his strategy. Let u and v be the ideal points of the strategy of player 1. Similarly, let u^* and v^* be the ideal points of the strategy of player 2. Actually, only u and v^* are of interest here since we have fixed the player roles.

The behavior in the main phase can be described by two "reaction functions," r and r^* :

$$r(x_2) = \begin{cases} u_1 & \text{for } x_2 \leq u_2, \\ c_1 & \text{for } x_2 \geq c_2, \\ u_1 + \frac{c_1 - u_1}{c_2 - u_2}(x_2 - u_2) & \text{otherwise;} \end{cases}$$

$$r^*(x_1) = \begin{cases} v_2^* & \text{for } x_1 \leq u_1^*, \\ c_2 & \text{for } x_1 \geq c_1, \\ v_2^* + \frac{c_2 - v_2^*}{c_1 - u_1^*} (x_1 - u_1^*) & \text{otherwise.} \end{cases}$$

The development of the play in the main phase is given by the following equations:

$$\begin{aligned} x_1(3) &= u_1, \\ x_2(3) &= v_2^*, \\ x_1(t) &= r(x_2(t-1)) \quad \text{for } t = 4, \dots, 18, \\ x_2(t) &= r^*(x_1(t-1)) \quad \text{for } t = 4, \dots, 18. \end{aligned}$$

Figure 9 shows four examples for the development of this system of difference equations. In Figures 9a and 9b the path of output combinations moves towards the Cournot equilibrium. In Figure 9c the path stays at (u_1, v_2^*) for $t = 3, \dots, 18$.

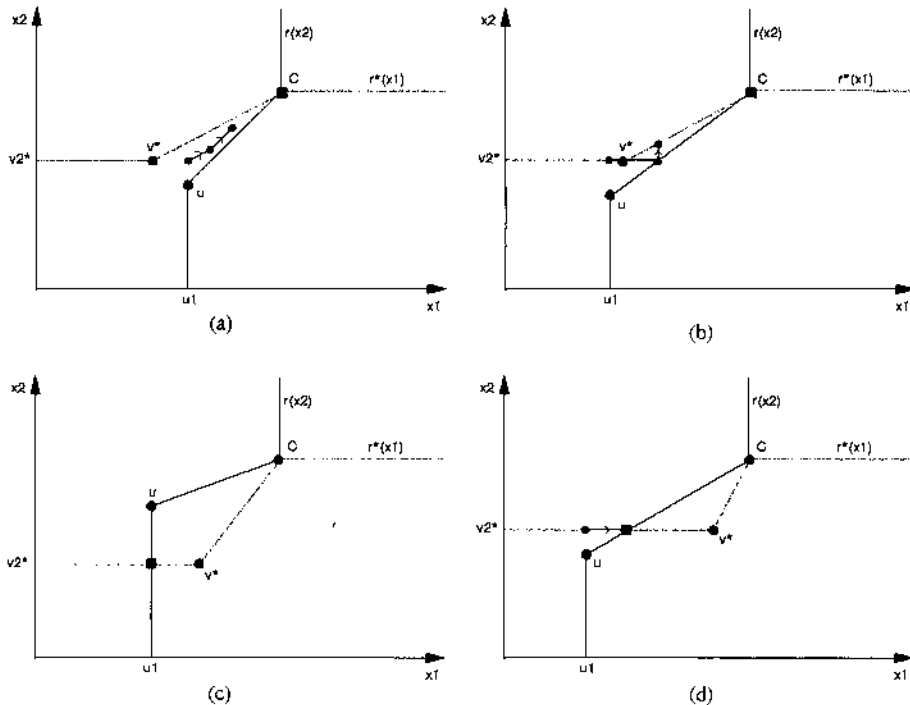


FIGURE 9.—Simple typical strategies playing against each other. Four examples with different ideal point pairs u and v^* .

In Figure 9d we have

$$\begin{aligned}x_1(t) &= r(v_2^*) & \text{for } t = 4, \dots, 18, \\x_2(t) &= v_2^* & \text{for } t = 4, \dots, 18.\end{aligned}$$

We shall speak of a *conflict case* if the output combination path moves towards the Cournot equilibrium, and of an *agreement case* if the output path becomes stationary in periods 4 to 18.

It can be seen without difficulty that a conflict case is obtained whenever the Cournot output combination is the only common point of r and r^* . All other cases are agreement cases. An agreement case is also characterized by the condition that player 2's ideal point is not above the straight line through the Cournot point and player 1's ideal point. This is the case if and only if player 1's ideal point is not below the straight line through the Cournot point and player 2's ideal point. From what has been said it follows that an agreement case is obtained if and only if the following inequality holds:

$$\frac{c_2 - v_2^*}{c_1 - v_1^*} \geq \frac{c_1 - u_2}{c_1 - u_1}.$$

Figure 9 will illustrate the consequences of this condition: In the special case in which both ideal points are Pareto optimal, an agreement case is reached if each player does not ask for more than the other player will grant him. The ideal points are like bargaining offers. The less one asks for oneself and the more one grants to the other player, the better are the chances for agreement.

In view of the condition for an agreement case it seems to be quite reasonable to specify two different ideal points for the two player roles in such a way that player 1's ideal point is more favorable for player 2 and vice versa. However, those 6 participants who specified two different ideal points did this in a way which leads to a conflict case if the strategy plays against itself. In each player role these subjects wanted more for themselves than they would grant to the other player if he were in this role.

It can be seen without difficulty that the condition which distinguishes agreement cases from conflict cases does not depend crucially on the special way in which our simple typical strategies specify the initial phase. As long as at the end of the initial phase both outputs are below the respective Cournot outputs, the output combination path moves towards the Cournot point in a conflict case and towards stationary cooperation in an agreement case.

8.3. The Best Ideal-Point Selection Against the Final Strategies

It is interesting to ask the question of what is the best selection of ideal points within the family of simple typical strategies defined above in a tournament against the final strategies. Actually, we simulated tournaments only against 22 of the final strategies since we omitted the only strategy which uses random

choices. The best choice of ideal points turned out to be as follows:

$$u = (89.4, 55.6),$$

$$v = (86.6, 50.4).$$

Both components of u are greater than the corresponding components of v , but if this strategy plays against itself an agreement case is obtained; the quantity combination (89.4, 52.6) is played in periods 4 to 18.

The ideal point (86.6, 50.4) is nearly Pareto optimal whereas $u = (89.4, 55.6)$ is relatively far from the Pareto optimal line. However, $u = (89.4, 55.6)$ has the advantage that it yields agreement cases against all ideal points which have been specified for the role of player 2 by those of the 22 participants who used ideal points. This is due to the fact that $u_2 = 55.6$ is rather large.

The ideal point $v = (86.6, 50.4)$ does yield conflict cases against some of the ideal points specified for the role of player 1 by participants. These ideal points for player 1 are too aggressive to make it worthwhile to reach agreement with them by a more generous ideal-point choice which, of course, would diminish payoffs against other strategies.

The simple typical strategy with $u = (89.4, 55.6)$ and $v = (86.6, 50.4)$ is not only the best among its family but is also the winner of the tournament against the 22 final strategies. This seems to indicate that the way in which the simple typical strategies fill in the details left open by the 13 characteristics is not an unreasonable one. One may say that the structure of these strategies provides an appropriate idealized image of typical behavior of experienced strategy programmers, at least as far as our experiment is concerned.

8.4. *Game-Theoretic Properties of Simple Typical Strategies*

The 20-period supergame has only one subgame perfect equilibrium point. In this equilibrium point both players always choose their Cournot quantities regardless of the previous history. If both players use simple typical strategies of the family described above the resulting strategy pair is always a disequilibrium, simply because it would be advantageous to deviate in the fourth last period.

Game theoretically there is a fundamental difference between finite and infinite supergames. It is known from the experimental literature that this difference seems to have little behavioral relevance. In sufficiently long finite experimental supergames cooperation is possible until shortly before the end, even if the source game has only one equilibrium point (Stoecker (1983), Selten and Stoecker (1986)). If one wants to connect finite supergame behavior with game-theoretical equilibrium notions, one has to take the point of view that the players behave as if they were in an infinite supergame.

It is shown in another paper of one of the authors that it is possible to construct equilibrium points for the infinite supergame of our duopoly model based on the main phase of our simple typical strategies (Mitzkewitz (1988)). In these equilibrium points both players have the same ideal point. This ideal point

is chosen in the first period of the game; later the strategies respond to the previous period as specified by the reaction functions r and r^* . Under certain conditions which have to be imposed on the ideal points, equilibrium points are obtained in this way. However, these equilibrium points are not subgame perfect. This is a consequence of a result in the literature which shows that equilibria where output continuously depends on the opponent's last-period output only cannot be subgame perfect unless the Cournot output is specified regardless of the previous history (Stanford (1986), Robson (1986)). Mitzkewitz (1988) shows that an appropriate modification of the main phase of the simple typical strategies yields subgame perfect equilibrium points for a wide range of ideal points.

Among the newer game-theoretical literature on the duopoly problem we have only found one paper which shows some similarities with the approach taken here (Friedman and Samuelson (1988)).

8.5. Reasonable Conditions for Ideal Points

One may ask the question whether it is possible to impose reasonable restrictions on the choice of ideal points in our simple typical strategies. A strategy programmer who considers an ideal point for one of the player roles will probably explore what happens if his opponent uses the same ideal point for the opposite player role. Therefore, it is natural to focus on the case in which both opponents use the same ideal point $u = (u_1, u_2)$ for both player roles.

Suppose player 1 knows that player 2 plays a simple typical strategy as defined above with the ideal point $u = (u_1, u_2)$. Suppose that for some output x_1 the profit $G_1(x_1, r^*(x_1))$ is greater than $G_1(u_1, u_2)$. Then player 1 has a better alternative than to agree to player 2's ideal point (u_1, u_2) . This consideration and an analogous one for player 2 lead to the following conditions:

$$G_1(u_1, u_2) = \max_{x_1} G_1(x_1, r^*(x_1)),$$

$$G_2(u_1, u_2) = \max_{x_2} G_2(r(x_2), x_2).$$

We refer to these two equations as "conjectural equilibrium conditions" since there is an obvious relationship to conjectural oligopoly theories (see Selten (1980)).

Another reasonable condition on ideal points is connected to the possibility of attempts of short-run exploitation. Suppose that a player deviates just once from the ideal point and then returns to cooperation at the ideal point. It should not be possible to improve profits in this way. This leads to the following conditions:

$$2G_1(u_1, u_2) = \max_{x_1} [G_1(x_1, u_2) + G_1(u_1, r^*(x_1))],$$

$$2G_2(u_1, u_2) = \max_{x_2} [G_2(u_1, x_2) + G_2(r(x_2), u_2)].$$

We refer to these equations as “stability against short-run exploitation.” In our numerical case the conjectural equilibrium conditions imply stability against short-run exploitation, but this is not the case for all possible parameter values.

As has been explained in subsection 8.4 it will be shown elsewhere (Mitzkewitz (1988)) that subgame perfect equilibrium points for the infinite supergame can be constructed on the basis of the reaction functions (but with memory also of the own behavior) embodied in the main phase of simple typical strategies if certain conditions on the ideal point are satisfied. These conditions are nothing else than the conjectural equilibrium conditions and the stability against short-run exploitation.

Perhaps it is also of interest that only one Pareto optimal point satisfies the conjectural equilibrium conditions, namely the point described in the third row of Table II: profit monotonic quantity reduction along the straight line through the intersections of both Cournot-isoprofit curves (see Mitzkewitz (1988)). It is tempting to look at this ideal point as distinguished among others by its special theoretical properties. In the final strategies it has been employed twice. However, as can be seen in Table II, other ideal points based on different principles have proved to be at least as attractive to the participants.

Figure 10 shows the ideal points used in final strategies of the participants and the restrictions imposed by the conjectural equilibrium conditions (the smaller lens-shaped area) and by stability against short-run exploitation (the greater lens-shaped area). The equations for these curves will be discussed elsewhere




FIGURE 10.—The set of ideal points satisfying the conjectural equilibrium conditions (smaller lens), the set of ideal points stable against short-run exploitations (greater lens), and the ideal points used in final strategies.

(Mitzkewitz (1988)). Only 4 of the 24 ideal points satisfy the conjectural equilibrium conditions, but 21 of the ideal points are stable against short-run exploitation.

Obviously, the participants were not concerned about the conjectural equilibrium conditions. Maybe a violation of these conditions is not perceived as a serious danger since in the case of an optimization of the other player along one's own reaction function, cooperation will still be reached, even if the resulting output levels are higher than in the ideal point.

Some strategies which were not yet the final ones contained attempts at short-run exploitation. Most participants seemed to be aware of this possibility since the "trainer"-program enables them to play against their own strategy. They were able to check short-run exploitability without analytical computations. Of course, such numerical checks will sometimes fail to reveal the right answer. Maybe it is of interest in this connection that two of the three ideal points without stability against short-run exploitation are very near to the corresponding area in Figure 10.

8.6. *Stability against Short-Run Exploitation and Outcomes of Plays in the Final Tournament*


In the tournament among 23 final strategies (including the strategy with random choices) 1012 plays were simulated (two plays for each strategy pair). Table IV shows the distribution of the pairs of total profits in the 1012 plays. The inner cells of the table correspond to profit intervals of four thousand for both players.

The curve superimposed on this table is connected to stability against short-run exploitations. The curve encloses all profit pairs which can be reached by plays in which the same ideal point with the property of stability against short-run exploitation is played in all 20 periods. We call the region enclosed by this curve the "exploitation stability region."

Consider two simple typical strategies whose ideal points are stable against short-run exploitation. Whenever such strategies are played against each other, the resulting profit combination of the 20-period supergame must be in the exploitation stability region, regardless of whether the ideal points of both players are equal or not. However, the set of all profit combinations which can be reached in this way is a proper subset of the exploitation stability region. This is due to the behavior in the initial phase and the end phase. The exploitation stability region can be obtained by pairs of modified simple typical strategies, strategies in which the initial phase and the end phase are of different length, but the behavior in the main phase remains the same.

In the final tournament 983 (97.1%) of the 1012 plays resulted in total profit combinations in the exploitation stability region. In those few total profit combinations outside the exploitation stability region, one of both profits is below the corresponding Cournot profit.

TABLE IV
 SUPERGAME PROFIT PAIRS IN THE FINAL TOURNAMENT
 AND THE EXPLOITATION STABILITY REGION.



The evidence of Figure 10 and Table IV strongly suggests that stability against short-run exploitation has some relevance for the prediction of outcomes of plays between strategies written by experienced players.

9. IMPLICATIONS FOR DUOPOLY THEORY

The results presented in this paper suggest a new view of the duopoly problem. Traditional duopoly theories and game-theoretical approaches rely heavily on optimization ideas. Usually, a duopolist is assumed to optimize against expectations on his opponent's behavior. Contrary to this, it is typical for the strategies programmed by the experienced players in our experiment that no expectations are formed and nothing is optimized.

The approach to the duopoly problem suggested by our results can be described as the "active pursuit of a cooperative goal." First, one has to answer the question of where one wants to cooperate. The goal of cooperation is made precise by the concept of an ideal point. The ideal point should be a reasonable

compromise between both players' interests; otherwise, one cannot hope to achieve cooperation. Concepts of fairness such as those listed in Table II are the basis for judgments on the reasonableness of compromises.

It is well known in the experimental literature that considerations of fairness have a strong influence on observed behavior. Many of the empirical and experimental phenomena can be subsumed under an equity principle (Selten (1978b)). Further literature can be found there and in a newer paper which contains many illustrative examples (Kahneman, Knetsch, and Thaler (1986)). Fairness considerations also have been proved to be useful in the explanation of behavior in duopoly experiments (Friedmann (1970), Selten and Berg (1970)).

Once an ideal point has been chosen one has to determine a policy for its effectuation. Formally, an effectuation policy may be described by a reaction function as in the simple typical strategies of Section 8. However, contrary to conjectural oligopoly theory, such reaction functions are not to be interpreted as hypotheses on the opponent's behavior. Effectuation policies are more like reinforcement schedules which serve the purpose to guide the opponent's behavior rather than to optimize against it.

The typical structure of an effectuation policy is based on the principle of measure for measure. This principle requires an interpersonal comparison of the degree of cooperativeness of the players' actions. The degree of cooperativeness measures the nearness to the ideal point. The response matches the opponent's last action according to this measure.

A player who plays the dynamic game may try to learn how to do best against his opponent's behavior. A player who does this takes a "learning approach." It is also possible to take a "teaching approach," which means that one behaves in a way which induces the other player to conform to one's own goals.

It seems to be very difficult to design a reasonable strategy which takes the learning approach. One participant tried to do this in a sophisticated way. His strategy involved an approximate intertemporal optimization against statistical estimates of his opponent's strategy. His success rank was 20. As Table III shows, his strategy has only one of the thirteen characteristics, namely the absence of random decisions. Obviously, the optimization attempt, of this participant failed badly. The reason for this lies in the difficulty of forming an accurate estimate of the opponent's behavior on the basis of relatively few observations.

The difficulties connected to the learning approach point in the direction of a teaching approach. Of course, somebody who takes the teaching approach does not necessarily expect that the other player takes a learning approach. The other player may very well take a teaching approach, too. This will not lead to difficulties if both players pursue compatible cooperative goals. However, if the opponent tries to adapt to my strategy, this should not endanger my cooperative goal.

Maybe in a very long supergame of thousands of periods, a good strategy would involve both, teaching and learning, but within 20 periods not much can

be learned which still can be used within this time. Real duopoly situations rarely are analogous to very long supergames. Maybe a relatively short supergame more adequately captures the decision problem of managers who want to be successful within a foreseeable time.

The new view of the duopoly problem emerging from our results may be described by the slogan "fairness and firmness." One must first choose a fair goal of cooperation and then devise an effectuation policy which shows one's willingness to cooperate and firmly communicates resistance to unfair behavior.

As we have seen, the requirement of stability against short-run exploitation seems to be a restriction obeyed by the participants' choices of ideal points, even if their effectuation policies were not exactly the same as those of the simple typical strategies. It is clear that one should not give rise to the possibility of being exploited. Moreover, in the case in which the other player selects one's own ideal point, he should not be exploitable. This criterion of stability against short-run exploitation is in good agreement with our data.

It is clear that the theory of fairness and firmness can be easily transferred to different contexts, e.g. price-variation duopoly supergames. The tit-for-tat strategy which was the winner of Axelrod's contests (1984) is in harmony with the fairness-and-firmness theory. In the prisoner's dilemma the choice of an ideal point is not an issue. In view of the symmetry of the situation there is only one natural cooperative goal. Since there are only two choices available, measure for measure cannot mean anything else than tit-for-tat.

It must be admitted that no strong conclusions can be drawn from our data since the final strategies cannot be regarded as statistically independent observations. The participants interacted in game playing rounds and tournaments. Moreover, there was some verbal communication, even if the participants seemed to be reluctant to reveal the principles underlying their strategies.

More studies similar to the investigation presented here are necessary to establish the empirical relevance of the fairness-and-firmness theory. It should also be kept in mind that the final strategies of our participants are the result of a long experience with the game situation. It is quite possible that real duopolists have much less experience with their strategic situation and therefore do not achieve the same extent of cooperation. The experimental literature shows that only after a considerable amount of experience, subjects learn to cooperate (Stoecker (1980), Friedman and Hogatt (1980), Alger (1984, 1986), Benson and Faminow (1988)).

It would be wrong to assert that there is no difference between a programmed strategy and spontaneous behavior. The strategy method cannot completely reveal the structure of spontaneous behavior. However, it seems to be plausible that somebody who writes a strategy program is guided by the same motivational forces which would influence his spontaneous behavior. Of course, a strategy program is likely to be more systematic. Obviously this is an advantage from the point of view of theory construction.

10. SUMMARY OF RESULTS

1. Mean profits increased from one game playing round to the next.
2. The correlation between both player profits was negative in the first game playing round and became positive in the second and the third game playing round. This can be interpreted as a growth of understanding of the strategic situation.
3. Mean profits increased from one computer tournament to the next. In the final tournament 97.1% of all plays had profits above Cournot profits for both players.
4. Typically, a strategy program for the final tournament distinguishes among an initial phase, a main phase, and an end phase. Outputs independent of the opponent's previous behavior are specified for the initial phase of one to four periods. In the main phase the strategies aim at a cooperation with the opponent. Noncooperative behavior characterizes an end phase of one to four periods.
5. Typical structural features of strategies programmed for the final tournament can be described by 13 characteristics. These characteristics imply a strategic approach which begins with the selection of a cooperative goal described by an "ideal point." (A different ideal point may be chosen for each player role.) Cooperation at the ideal point is then pursued by a "measure-for-measure policy." If the opponent moves towards the ideal point or away from it, the response of a measure-for-measure policy is of similar force in the same direction. In the end phase a typical strategy always chooses Cournot outputs.
6. Typically, no predictions about the opponent's behavior are made and nothing is optimized.
7. The extent to which a strategy or a characteristic is typical can be measured by an index of typicity. There is a highly significant positive rank correlation between the index of typicity and the success of a strategy in the final tournament.
8. For each of the 13 characteristics separately those final strategies which have this characteristic have a higher average success rank than those which do not have it.
9. Ideal points are often based on various fairness considerations (see Table II).
10. A family of "simple typical strategies" has been introduced as an idealized description of the structure implied by the 13 characteristics. The simple typical strategy which performed best against the final tournament strategies was determined by a computer simulation. This "best" simple typical strategy is also the winner in the tournament against the final strategies.
11. Two game-theoretical requirements for simple typical strategies impose restrictions on ideal points. One of these restrictions, the "conjectural equilibrium conditions," is rarely satisfied by the ideal points in the final strategies.

However, most of these ideal points satisfy the weaker restriction of "stability against short-run exploitation."

12. An "exploitation stability region" for profit combinations reached in the supergame can be derived from the requirement of stability against short-run exploitation. The profit combinations of all plays in the final tournament in which both players received more than their Cournot profits are in the exploitation stability region. These are 97.1% of all plays in the final tournament.

Universität Bonn, Wirtschaftstheoretische Abteilung I, Adenauerallee 24-42, D-53113 Bonn, Germany.

Manuscript received October, 1990; final revision received April, 1996.

REFERENCES

- ABREU, D. (1986): "Extremal Equilibria of Oligopolistic Supergames," *Journal of Economic Theory*, 39, 191-225.
- ALBERS, W., AND G. ALBERS (1983): "On the Prominence Structure of the Decimal System," in *Decision Making Under Uncertainty*, ed. by R. W. Scholz. Amsterdam: Elsevier, pp. 271-287.
- ALGER, D. (1984): "Equilibria in the Laboratory: Experiments with Oligopoly Markets where Goods Are Made to Order," Working Paper No. 121, Bureau of Economics, Federal Trade Commission, Washington.
- (1986): *Investigating Oligopolies within the Laboratory*. Washington: Bureau of Economics, Federal Trade Commission.
- AXELROD, R. (1984): *The Evolution of Cooperation*. New York: Basic.
- BENSON, B. L., AND M. D. FAMINOW (1988): "The Impact of Experience on Prices and Profits in Experimental Duopoly Markets," *Journal of Economic Behavior and Organization*, 9, 345-365.
- COURNOT, A. (1938): *Recherches sur les Principes Mathématiques de la Théorie des Richesses*. Paris: Hachette.
- FADER, P. S., AND J. R. HAUSER (1988): "Implicit Coalitions in a Generalized Prisoner's Dilemma," *Journal of Conflict Resolution*, 32, 553-582.
- FRIEDMANN, J. W. (1970): "Equal Profits as a Fair Division," in *Beiträge zur Experimentellen Wirtschaftsforschung*, Band 2, ed. by H. Sauermann. Tübingen: J. C. B. Mohr, pp. 19-32.
- (1977): *Oligopoly and the Theory of Games*. Amsterdam: North Holland.
- FRIEDMANN, J. W., AND A. C. HOGATT (1980): *An Experiment in Noncooperative Oligopoly Research in Experimental Economics*, Vol. 1, Supplement 1. Greenwich: JAI Press.
- FRIEDMANN, J. W., AND L. SAMUELSON (1988): "Subgame Perfect Equilibrium with Continuous Reaction Functions," Mimeo, University of Bielefeld, Center of Interdisciplinary Research.
- KAHNEMAN, D., J. L. KNETSCH, AND R. H. THALER (1986): "Fairness and the Assumptions of Economics," *Journal of Business*, 59, 285-300.
- MITZKEWITZ, M. (1988): "Equilibrium Properties of Experimentally Based Duopoly Strategies," Working Paper No. B-107, University of Bonn.
- RADNER, R. (1980): "Collusive Behavior in Noncooperative Epsilon-Equilibria of Oligopolies with Long but Finite Lives," *Journal of Economic Theory*, 22, 136-154.
- ROBSON, A. (1986): "The Existence of Nash Equilibria in Reaction Functions for Dynamic Models of Oligopoly," *International Economic Review*, 27, 539-544.
- SCHELLING, T. C. (1960): *The Strategy of Conflict*. Cambridge: Harvard University Press.
- SEGERSTROM, P. S. (1988): "Demons and Repentance," *Journal of Economic Theory*, 45, 32-52.
- SELTEN, R. (1967): "Die Strategiemethode zur Erforschung des eingeschränkt rationalen Verhaltens im Rahmen eines Oligopol-experiments," in *Beiträge zur Experimentellen Wirtschaftsforschung*, ed. by H. Sauermann. Tübingen: J. C. B. Mohr, pp. 136-168.

- (1970a): "The Chain Store Paradox," *Theory and Decision*, 9, 127–159.
- (1978b): "The Equity Principle in Economic Behavior," in *Decision Theory and Social Ethics: Issues in Social Choice*, ed. by H. W. Gottinger and W. Leinfellner. Dordrecht: Reidel, pp. 289–301.
- (1980): "Oligopoltheorie," in *Handwörterbuch der Wirtschaftswissenschaft*, ed. by W. Albers et al. Stuttgart-New York: Fischer, pp. 667–678.
- (1987): "Equity and Coalition Bargaining in Experimental Three-person Games," in *Laboratory Experimentation in Economics*, ed. by A. E. Roth. Cambridge: Cambridge University Press, pp. 42–98.
- SELTEN, R., AND C. C. BERG (1970): "Drei experimentelle Oligopolspielserien mit kontinuierlichem Zeitablauf," in *Beiträge zur Experimentellen Wirtschaftsforschung*, Band 2, ed. by H. Sauermann. Tübingen: J. C. B. Mohr, pp. 162–221.
- SELTEN, R., AND R. STOECKER (1986): "End Behavior in Finite Prisoner's Dilemma Supergames," *Journal of Economic Behavior and Organization*, 7, 47–70.
- STANFORD, W. G. (1986): "Subgame Perfect Reaction Function Equilibria in Discounted Duopoly Supergames are Trivial," *Journal of Economic Theory*, 39, 226–232.
- STOECKER, R. (1980): *Experimentelle Untersuchung des Entscheidungsverhaltens im Bertrand-Oligopol*. Bielefeld: Pfeffer.
- (1983): "Das Erlernte Schlußverhalten—eine Experimentelle Untersuchung," *Zeitschrift für die Gesamte Staatswissenschaft*, 139, 100–121.