COMPLETE INFORMATION OUTCOMES WITHOUT COMMON KNOWLEDGE

Andrew F. Daughety
Departments of Economics
and Management Sciences
The University of Iowa
Iowa City, IA 52242

Robert Forsythe
Department of Economics
The University of Iowa
Iowa City, IA 52242

ABSTRACT

In this paper we examine the robustness of existing models of agent decision making which rely on the assumption of common knowledge. The research we report on involves testing equilibria predicted by traditional, common-knowledge intensive models in laboratory settings wherein the underlying common knowledge assumptions are violated in a specified manner. This approach provides insight into when we may employ narrow but tractable models in an "as if" mode: when common knowledge is lacking, are models that assume common knowledge still likely to perform well? In general we find that with an opportunity to coordinate plans, either through communication or via histories, cooperative outcomes emerge. Furthermore, these collusive outcomes often extensively and efficiently exploit the multi-period nature of the interaction, achieving the complete information collusive allocation of profits. This strongly suggests that predictions of collusive behavior in repeated play settings are quite robust with respect to relaxing common knowledge assumptions: a small degree of information exchange or coordination yields outcomes that are generally indistinguishable from the complete information outcomes.

INTRODUCTION

In this paper we examine the robustness of existing models of agent decision making which rely on the assumption of common knowledge. Common knowledge refers to the condition wherein an attribute (e.g. strategies, utility functions, relevant information on environments, etc.) of an agent is known to all agents, and this fact is known to all agents, and all agents know that all agents know that this fact is known to all agents, and so on (see Aumann [1976]). Common knowledge of the characteristics of an attribute means that while the precise value of the attribute itself may not be common knowledge (it might be privately held), the possible values it could take on (e.g. the possible utility functions that a player might have) and the probability distribution of those possible values is common knowledge. We refer to both uses of common knowledge as the common knowledge assumptions. The research we report on involves testing equilibria predicted by traditional, common-knowledge intensive models in laboratory settings wherein the underlying common knowledge assumptions are violated in a specified manner. This approach provides insight into when we may employ narrow but tractable models in an "as if" mode: when common knowledge is lacking, are models that assume common knowledge still likely to perform well?

More precisely, our focus is on voluntary collusion by duopolists. At least since Adam Smith economists have harbored the deep suspicion that, given the opportunity, firms in an industry will collude. The problem for the firms, of course, was how to make the collusion stick. For many years collusion was viewed as requiring individually irrational behavior on the part of the participants: short of enforcement via government intervention, non-cooperative behavior on the part of the agents should lead to defections from the collusive solution. Specifically, at least one firm would find it advantageous to expand output and cut price.

The recognition that firms engage in repeated interactions with each other, and that retaliation for defection was therefore a possibility (e.g., Stigler [1964] and Green and Porter [1984]), rekindled the old concern: market forces and individual incentives may not be able to drive out collusive behavior. More recently, papers have appeared showing that even finite-horizon repeated games can yield collusive outcomes when the appropriate threats can be employed (e.g., Kreps, et. al. [1982], Benoit and Krishna [1985] and Friedman [1985]). Thus, theory suggests that collusion is not only feasible, but individually rational and optimal. Significantly, these predictions are based on extensive use of the common knowledge assumptions, which conflict with typical circumstances in the "real world". In reality, firms don't know each other's profit functions (types) or even the distribution of types. In such naturally occurring settings firms make errors as they learn about their environment and their opponents, thereby making it more difficult to use past behavior as a guide to future behavior (previous moves may be errors, which need not ever occur again under the otherwise same circumstances, or they may be experiments designed to extract information which, once discovered, need not be extracted again). Thus, in contrast to the central assumptions of game theoretic models of strategic conflict under incomplete information, real world firms face decision settings wherein there may be little common knowledge: in such settings it is difficult to infer types or intent. Can't we expect such wholesale violations of the common knowledge assumptions inherent in the repeated game models to work in our favor--once again limiting the value of cooperation, and increasing the value of defection?

In this paper we report on a laboratory analysis of repeated play by duopolists. In one series, duopolists are endowed with asymmetric payoffs and each player has little or no information about the other player in the game. In some of the games players are allowed to exchange messages for the early periods of play, while in other games they are given no opportunity to coordinate. In another series, players (in this case in symmetric duopolies) are provided with histories of play for previous players (i.e. players other than themselves) and are provided no opportunity to communicate. In all cases players had no previous experience with each other, or with the specific market at hand. Moreover, all players were kept separated, knew they would never learn each other's identity, and stood to make a lot of money. In general we find that with an opportunity to coordinate plans, either through communication or via the histories, cooperative outcomes emerge. Furthermore, these collusive outcomes often extensively and efficiently exploit the multi-period nature of the interaction, achieving the complete information collusive allocation of profits. This strongly suggests that predictions of collusive behavior in repeated play settings are quite robust with respect to relaxing common knowledge assumptions: a small degree of information exchange or coordination yields outcomes that are generally indistinguishable from the complete information outcomes.

This paper is organized as follows. The next section summarizes a simple model of multi-period collusive behavior and the standard one period non-cooperative model; these are used to provide the equilibrium predictions for the experiments to be discussed. The third section will briefly describe the experiments that were performed, followed by the presentation of the results in section four.

MODELS

There are a number of criteria that a model of collusive behavior should satisfy if we are to employ it as a predictive tool and if we are to be able to draw conclusions about the relevance of game-theoretic predictions for analyses of industrial organization. First, the model should reflect non-cooperative play on the part of the agents. Second, outcomes should not depend upon direct side payments. This means that some traditional solutions, such as operating only the lowest cost firm and shutting down the others, are not possible. This last observation yields a third consideration: any proposed collusive solution (and associated experimental test) must explicitly allow for asymmetries in cost structure. Such asymmetry presumably makes collusion more difficult to achieve and maintain, especially if agents lack complete information on rivals.

Since in general no such solution exists in the one-shot case, it should be clear that our approach will be to examine a multi-period solution, which when imbedded in the infinitely repeated setting, yields collusion as a Nash equilibrium. To accomplish this we will first motivate consideration of "alternating strategies" (see Friedman [1986] and Shubik [1982]) in a repeated game setting and then use a specific alternating strategy to define the set of collusive equilibria to be used for predictive purposes.

ALTERNATING STRATEGIES

In what follows we will consider a homogeneous product industry comprised of n sellers. Let x_{it} denote seller i's output choice in period t and let $\mathbf{X}$ be the n×T matrix of outputs, where T is a finite number of periods. Denote the t-th column of $\mathbf{X}$ as x^t and the i-th row of $\mathbf{X}$ as s^i; x^t is the vector of seller outputs during period t, while s^i is the vector of outputs for seller i for the T-period horizon (i's T-period strategy). Demand is a function of current aggregate output and is represented by the inverse demand function $p(\Sigma_i x_{it})$; we assume that $p(\cdot)$ is twice continuously differentiable with $p'<0$ and $p''\leq 0$. Seller i faces cost function $c_i(x_{it})$, which is also twice continuously differentiable, increasing and convex (i.e. $c'>0$, $c''\geq 0$). Finally, period t profits for seller i are $\pi^i(x^t)=p(\Sigma_i x_{it})x_{it}-c_i(x_{it})$.

While a strategy for a seller in the infinitely repeated game is an element of $\Re^{\infty}_{+}$, our interest will center on finite length vectors of outputs that sellers repeat over the infinite horizon; the most interesting of these will involve sellers adjusting output levels so as to make implicit inter-seller transfers. Repeating such a finite sequence means that players are alternating their output levels so as to rotate profit opportunities amongst the game's participants. We will call a strategy s^i a T-period alternating strategy if there are at least two elements of the vector s^i which are unequal. For example, in a duopoly, a possible 2-period alternating strategy is one wherein each seller produces its monopoly output for one period and produces zero the other period with outputs sequenced so that neither seller produces in the same period. For expositional purposes, we will consider the following duopoly example (this is the first parameter set discussed in the next section): $p(x_1+x_2)=33-x_1-x_2$, $c_1(x_1)=3x_1+4$ and $c_2(x_2)=x_2+40$. A pair of alternating monopoly strategies for this case is $s^1=(15,0)$ and $s^2=(0,16)$.

An alternating strategy provides a "point of reference" for a collusive solution: it provides a means for coordinating output levels so as to distribute the benefits of collusion to the participants without engaging in overt acts of making actual side payments. Moreover, if its repeated use yields strategies that constitute a Nash equilibrium in the repeated game, then it is a candidate collusive solution itself. Note also that we are not restricted to alternating monopoly solutions, though these are clearly easy to construct.

We can formalize this by posing the following optimization problem, denoted (CS_T), the solution of which (the matrix of outputs $\mathbf{X}$) is referred to as a <u>T-period collusive solution</u>:

$$(CS_T): \quad \max_{X\geq 0} \sum_{t=1}^{T} \sum_{i=1}^{n} \alpha_i \pi^i(x^t)$$

$$\text{s.t.} \quad T^{-1} \sum_{t=1}^{T} \pi^i(x^t) \geq \tilde{\pi}^i, \quad i=1,\ldots,n,$$

where the α's are varied over all possible values such that $\Sigma_i \alpha_i = 1$ and $\alpha_i \geq 0$. The solution to CS_T: 1) provides the T-period strategies for all n sellers that maximizes joint T-period profits; 2) guarantees that each seller i receives, on average during the T periods, at least its reference profit $\tilde{\pi}^i$ per period. For the duopoly example above, when T=2, use of the alternating monopoly strategies as reference values means that $\tilde{\pi}^1=108.5$ and $\tilde{\pi}^2=88$. Note that use of the alternating monopoly strategy as the reference point guarantees that the constraints in (CS_T) are always feasible.

THE PROFIT FRONTIER

The foregoing discussion suggests that the employment of alternating monopoly strategies in a repeated game provides a "decentralized" means for asymmetric sellers to collude without direct side payments. Moreover, even when such a solution can be dominated by some other strategy (i.e. even when the average profits from employing an alternating monopoly solution are not on the Pareto frontier), such a strategy provides a reference outcome to which other, more efficient, collusive strategies can be compared. To see this, let

$$P = \{\Pi \in \Re^n \mid \Pi_i = \pi^i(x),\ i=1,\ldots,n,\ \forall x \in \Re^n_+\}$$

be the set of one-period profits for the n players and assume that: 1) P is compact and 2) that the origin is contained in the interior of P. Let conv(P) denote the convex hull of P. If $(\text{conv}(P)-P)\cap\Re^n_+$ is nonempty then there will be a value of T, and an $n\times T$ matrix of outputs **X** involving alternating strategies, that yields undominated average profits for all sellers. This is simply a restatement of the well-known result that the convex hull of P reflects the use of mixed strategies; alternating strategies are simply a finite period implementation of a mixed strategy. (More precisely, some mixed strategies can be achieved via alternation of pure strategies over a finite number of periods.)

The importance of alternating monopoly strategies becomes most obvious when one examines the standard two seller, linear demand $(a-b(x_i+x_j))$, linear cost $(c_i x_i + F_i)$ case so often manipulated in the industrial organization literature. Let $\pi^i(x_i,x_j)=(a-b(x_i+x_j))x_i - c_i x_i - F_i$ and let $\Pi_2(\Pi_1)$ be the frontier profit function (extended to the real line):

$$\Pi_2(\Pi_1) = \begin{cases} \{\max \pi^2(x_1,x_2) \mid \pi^1(x_1,x_2) \geq \Pi_1\}, & \pi^i \geq -F_i, \quad i=1,2 \\ -\infty & \text{otherwise.} \end{cases}$$

When cost functions are identical then it is straight-forward to show that the non-infinite portion of $\Pi_2(\Pi_1)$ (i.e. the portion defined over the effective domain of $\Pi_2(\cdot)$) is a straight line with slope of -1 drawn between the two points in (Π_1,Π_2)-space $(\pi^M,-F)$ and $(-F,\pi^M)$, where π^M is the monopoly level of profits. Note that these are also the profit pairs that would result (over two periods) by use of an alternating monopoly strategy.

However, when the sellers are not identical, then $\Pi_2(\Pi_1)$ is not generally concave, and in particular the non-infinite portion lies inside conv(P). This characteristic of the Pareto surface appears to have escaped general recognition in the literature (except in some numerical examples in Shubik [1984]).

This means that in such cases, alternating the role of monopolist between sellers produces an average profit that is on the convex hull of the one-period profit possibility set, but "above" the one period profit frontier described by $\Pi_2(\Pi_1)$. Thus in the case of a convex profit frontier, for any one period collusive solution there is a multi-period alternating solution that strictly dominates it in terms of average profits. Alternatively, if the profit frontier function is concave then the average profits associated with alternating lies in the interior of P. The use of the alternating point as a reference outcome means that the players restrict action to the portion of the frontier that "lies to the northeast" (i.e. involves profits no less than the alternating average profits) of this reference point.

ALTERNATING STRATEGIES AND BARGAINING SOLUTIONS

Another reason for considering alternating strategies is their relationship to some of the axiomatic solutions to two-person bargaining problems. Two solutions, the Raiffa solution [1953] and the Kalai-Smorodinsky solution [1975] involve either paths or outcomes that reflect the alternating monopoly solution. In the Raiffa solution the alternating monopoly point is first found. If it is within the Pareto frontier (i.e. inside the bargaining set) the it is use to construct a translated orthant that restricts attention to the "northeast direction" portion of the set. This procedure is continued in each reduced set until the frontier is reached. In the Kalai-Smorodinsky procedure, a ray from the origin is constructed which always passes through the alternating monopoly point; where it passes through the frontier is identified as the solution. Either solution yields the alternating monopoly outcome whenever the profit frontier is convex.

THE SINGLE-PERIOD SOLUTION

Finally, in what follows we will compare the prediction from the CS_T model with the classical one-period non-cooperative outcome, the Cournot solution (note that we suppress the superscript/subscript t in what follows). The Cournot solution is the output vector x^C such that no seller wishes to adjust its element of the vector if all other sellers employ their respective elements of the vector; in other words it is the one-period Nash equilibrium. More precisely, x^C solves the following system of optimization problems:

$$\left. \begin{aligned} x_i^C = \; & \arg\max_{x_i} \; \pi(x) \\ & \text{s.t. } x_j = x_j^C \quad j=1,\dots,n,\; j\neq i \end{aligned} \right\} \; i=1,\dots,n$$

THE DUOPOLY GAMES

The experiment consisted of three sets of two player, repeated games. Players were asked to make production decisions in each of a sequence of periods. These decisions determined each player's profits. In all instances each player knew his own profit as a function of both players' production decisions. Players were never given any information about their opponent's profit function. The procedures employed were those which have become standard in experimental economics (see Plott [1982], Daughety and Forsythe [1987a,b]). Moreover, all games were conducted in a way to ensure anonymity so as to remove any potential perceived gains from cooperation among players who thought they might encounter each other in the future.

We conducted a control set of games in which players simply made production decisions for a specified number of periods. As treatments, we conducted two additional sets of games: communication games and "history" games. The communication games consisted of two phases. There were 10 periods in Phase I and 15 periods in Phase II. During the first phase of each game, sellers were never given any indication that there would be a second phase. Moreover, sellers were never told the number of periods in a phase (they were only told that the experiment would last three hours). During Phase I players were allowed to communicate prior to making their production decisions, while in Phase II players could no longer communicate and could only make their production decisions. At the end of each period both players were informed of their opponents' decisions.

The players in the history games had previously participated in one of the communication games. Upon their arrival we gave them a "history" of play which we told them had been previously adopted by earlier sellers in their industry for the previous 25 periods. They were also informed that the other player was equally experienced but that they had never previously played against each other in the same game. Players were then asked to make production decisions during a sequence of periods and they were never permitted to communicate with each other. Players were informed of their opponent's decision at the end of each period. Each game lasted for 13 periods although players were never told how many periods there would be. The instructions for the control games appear in Daughety and Forsythe [1987b] and those for the communication and history games appear in Daughety and Forsythe [1987c].

Although subjects were given only a profits table, each table was constructed from demand and cost schedules so as to be consistent with the models in Section 2. We display these schedules in Table 1. The first letter in the game number identifies whether the game is a communication game (C) or a history game (H). For the communication games, the second number identifies the parameter set, while for the history games the second letter indicates the history which was provided to the players: an N indicates that no history was given; an E indicates an efficient history was given (a solution to CS_T); and an I indicates an inefficient history was given to the players (an equal output history with payoffs less than the Cournot solution). Finally, the number to the right of the hyphen gives the number of the game that was run under the corresponding environment.

Table 1. Experimental Parameters and Model Predictions*

Communication Games: (Market Demand is $p = 33 - (x_1+x_2)$ in C1, C2 & C3; it is $p = 200 - 4(x_1+x_2)$ in C4)

Game Number	Player 1's Cost	Player 2's Cost	Model Predictions: Cournot	Model Predictions: CS_T
C1-1-6	$3x_1 + 4$	$x_2 + 40$	(9,11);(10,11);(9,12)	alternate (15,0),(0,16)
C2-1-6	x_1	$5x_2$	(12,8)	alternate (16,0),(0,14)
C3-1-9	x_1	$17x_2$	(16,0)	alternate (16,0),(0,8)
C4-1-9	$x_1^2 + 600$	$x_2^2+20x_2+300$	(15,12)	(9,10);(10,10);(10,11); (10,12);(11,10);(11,11); (11,12);(12,10)

History Games: (Market Demand is $p = 33 - .5(x_1+x_2)$ for all games)

Game Number	Player 1's Cost	Player 2's Cost	Cournot	CS_T
HE-1-10				$\{(x^t,x^{t+1}) \mid x_1^t + x_2^t = 24,$
HI-1-3	$9x_1$	$9x_2$	(16,16)	$x_1^{t+1} + x_1^t = 24,$
HN-1-8				$x_2^{t+1} + x_2^t = 24\}$

*The CS_T solution uses alternating monopoly profits as the reference point.

The parameter sets employed in the communication games are purposely asymmetric. This was done to maintain the uncertainty by each player about the other's payoff. Further, to the extent that the collusive solution best explains the data, the third parameter set was designed to provide a sterner test of that model. In this third set, the Cournot solution is also the monopoly solution. Player 1 has a sufficient cost advantage that he can make any production by player 2 unprofitable.

Since no communication was allowed in the history games, a symmetric parameter set was used. In all cases no information on opponents' payoffs was provided by the experimenter so as to force players to rely on communication and observation (in the communication games) or history and observation (in the history games) for their inference about their opponents' payoffs.

Table 1 also provides the two competing equilibrium predictions: the Cournot outcome and the solution to CS_T for T=2 where the average profits from the alternating monopoly solution is used as the reference profits, $\bar{\pi}$. For the sets which use linear cost functions, the alternating monopoly solution is efficient. For the parameter set with quadratic costs, the alternating monopoly solution is not efficient; rather it provides a subset of the Pareto frontier (a reduced quadrant). Solving CS_T provides the indicated one period elements of the (integer) Pareto frontier.

To obtain substantial differences in the level of each seller's profits, the cost of using dollars directly would have been prohibitive. To overcome this difficulty, we used a currency which we called "francs" in these experiments. This artificial currency has been used in market experiments in Friedman [1967] and Daughety and Forsythe [1987a,b]. The payoffs for a given game are of the form bz, where b is the exchange rate of francs into dollars and z is the number of francs earned by a subject. In the communication games, b was .01 for each player in games C1, C2 and for player 1 in C3; b was .03 for player 2 in C3, and .002 for each player in C4. In the history games, b was always .005 for each player. With these exchange rates, players earned between $20 and $32 in the communication games (which lasted between 2 and 3 hours) and they earned between $5 and $10 in the history games (which lasted well under one hour).

RESULTS

We first present the results of our control games followed by the outcomes from each of our other two classes of games: the communication games and the history games. Two findings emerge: First, while players' decisions do not settle down to a readily characterized equilibrium in the control set of games, convergence to an equilibrium was quickly observed in the communication games and the games with efficient histories. Second, these games provide strong support for the collusive solution; given an opportunity to coordinate plans, either via communication or an efficient history, players choose and adhere to decentralized actions so that efficient cooperative outcomes result. This further suggests that common-knowledge intensive model predictions are very robust with respect to violations of these assumptions.

The control set (involving no communication or history), which used the same demand and cost functions as C1, consisted of 13 games. The outcomes from these games are displayed in Figure 1. The efficient frontier is shown, along with points which indicate the average profits from periods 14 and 15 for each player in the indicated game. All games were run for at least 15 periods and the average profits from periods 14 and 15 are plotted for each of these games. Three of these games were continued for 30 periods and the average profits from periods 29 and 30 for each player are also shown on the figure. In this and subsequent figures, the average

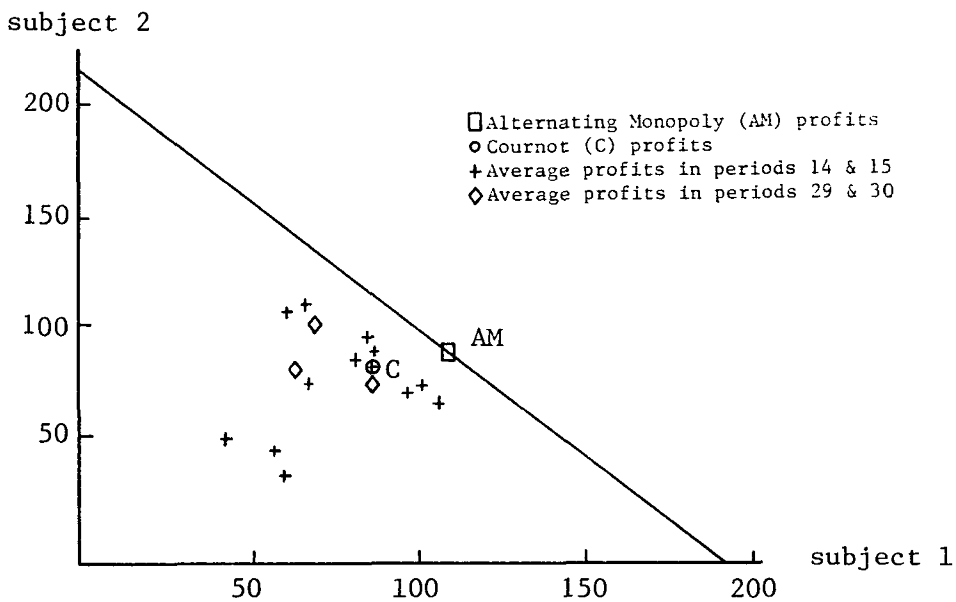

FIGURE 1. Results from Control Games.

of the final two period profits are displayed so that the average profits from playing a 2-period alternating strategy can be correctly shown. The Cournot profits (C) and the alternating monopoly profits (AM) are also displayed.

From this figure it can be seen that the Cournot prediction is uniformly supported against the CS_T prediction. The joint profit of both players in the 15 period games exceeds the Cournot prediction in only 5 of the 13 cases and all of these exceptions are within ten cents of the profits per period that players would jointly earn at the Cournot outcome; alternatively, these exceptions are all twenty cents or more away from the joint profit per period which players would earn at the AM solution. In only one of 30 period games do players jointly earn more per period than they would at the Cournot prediction and even here it is within five cents of the joint profit at the Cournot outcome.

Although it is not apparent from the figure, all players in these games continued to vary their strategies even at the game's conclusion: there is simply no evidence that these players' choice of strategies has settled down. As we report in Daughety and Forsythe [1987b], there is evidence that players use their Cournot strategies more often as play progresses. However, they don't tend to do it jointly; they tend to move in and out of "equilibrium." Players appear to have a problem of properly anticipating what each other is likely to do.

Figure 2 provides a summary of the outcomes of the communication games for each parameter set. Again, the efficient frontier is displayed, along

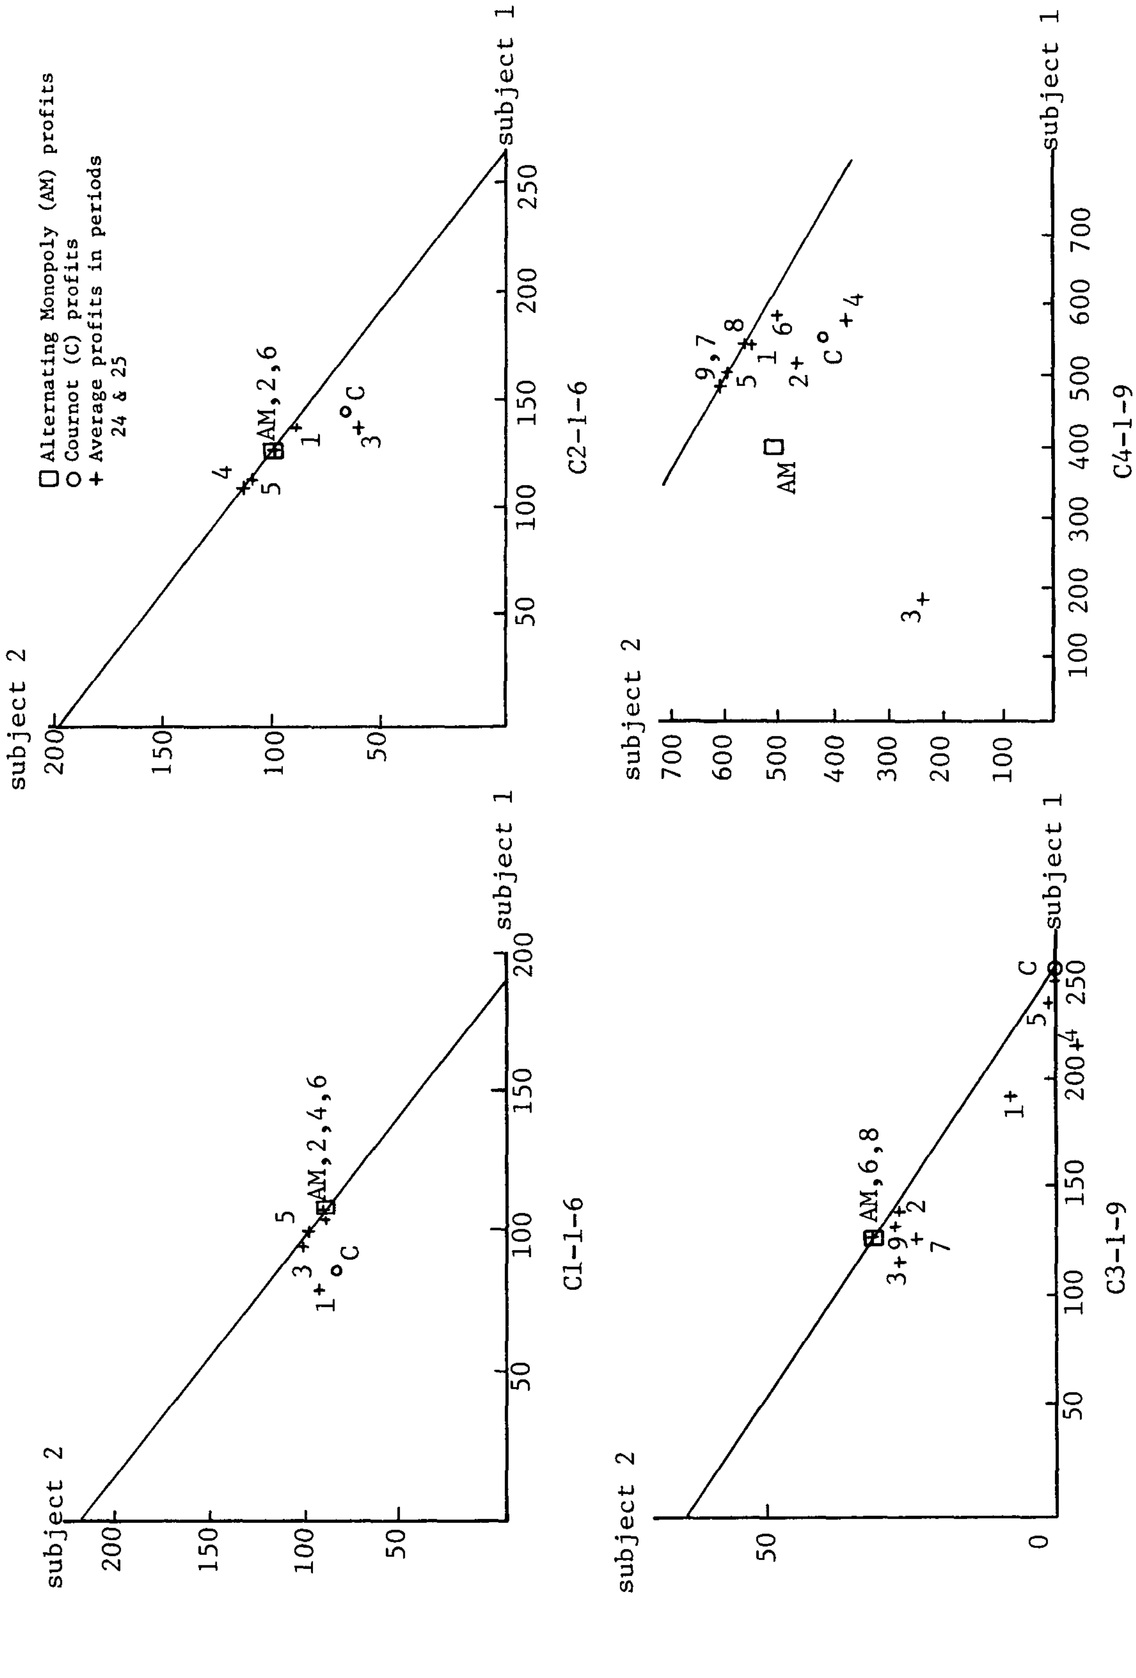

Figure 1: Results from Communication Games

with points which give the average profits for the last two periods for each player in the indicated game. (For C1-3, the average profits from the last four periods are shown since these players persisted in playing a 4-period alternating strategy.) Thus, the profits shown indicate outcomes in <u>Phase II</u>, long after communication has ceased.

In all three linear sets, C1, C2, and C3, the AM solution is also the CS_T solution since it lies on the efficient frontier. In these three sets, 7 of the 21 observations are precisely at the AM solution while 9 others are very near or on the frontier and "close" to the AM prediction. Even in C3, where player 1 has an enormous cost advantage over player 2 (see Table 1), six of the nine games are best predicted by AM. Of the 5 games (C1-1, C2-3, C3-1, C3-4, C3-5) which fail to support the AM prediction, there are only two (C2-3 and C3-5) in which players found their way to the cooperative solution during Phase I. In the other 4 games players were either unable to come to an agreement or were unable to adhere to an agreement even with communication during Phase I. Thus, the failure of these players to establish a history of cooperative play in Phase I carried over into Phase II.

Finally, in the quadratic set, C4, the AM profits serve as a reference point for the CS_T prediction. In 5 of 9 games, this CS_T solution provides very accurate predictions of the observed behavior. Game 6 is just barely outside the quadrant defined by the AM profits. Moreover, it is the only game which lies in the quadrant defined by the Cournot profits as a reference point.

All of the twenty-one games which support the CS_T solution seem to have converged to an equilibrium in the sense that players were continuing to play the same strategy they had adopted by period 10 when communication was terminated. Of the remaining nine games, players show no sign of having settled down; in seven of these games, players were unable to agree on a strategy during the periods of communication, while in the other two games, the players defected from their Phase I play early in Phase II. As in the control games, these players are unable to coordinate their strategies after communication ceases due to their inability to forecast their opponents' behavior.

Thus, the great preponderance of the outcomes are at or near the efficient frontier. This is particularly notable in view of the fact that this reflects play long after communication was halted. Moreover, in no case do subjects have complete information on their opponent's profits; they only know what the other player has communicated to them. Furthermore, the rapid achievement of collusive outcomes is in stark contrast with the earlier results of Friedman [1967] and Alger [1987], both of which required many periods of play before observing collusion.

The observations support the following two-stage model of behavior. In the first stage players choose between finding a collusive outcome and playing the Cournot solution. In the second stage they pick among the collusive outcomes guided by the AM profits. More precisely, the Cournot solution acts as the defection threat from the CS_T solution which, in turn, employs the AM solution as the reference point. In only two of the thirty observations do players reach a collusive agreement during Phase I and adhere to it, only to defect from this arrangement early in Phase II. It is only these two observations which are inconsistent with the two-stage model of the decision process.

The average profits from the final two periods of each history game along with the efficient frontier is shown in Figure 3. Further, the outcome of each of the different histories -- efficient, inefficient and none -- are plotted using a different symbol in the figure. As a casual inspection of this figure indicates, different histories led to much different outcomes. Efficient histories act as focal points: the alternating monopoly solution was attained in 7 of the 10 games with efficient histories. In two additional games with efficient histories the outcome was within a few cents of the efficient frontier. The strength of the efficient history as a focal point is further illustrated by the fact that in two of the games, players were able to return to the efficient solution even after they had deviated from it for several periods. Further, players who choose to adopt the efficient history settled down to this pattern of play very quickly; over the last eight periods of play, all of these players choose this outcome.

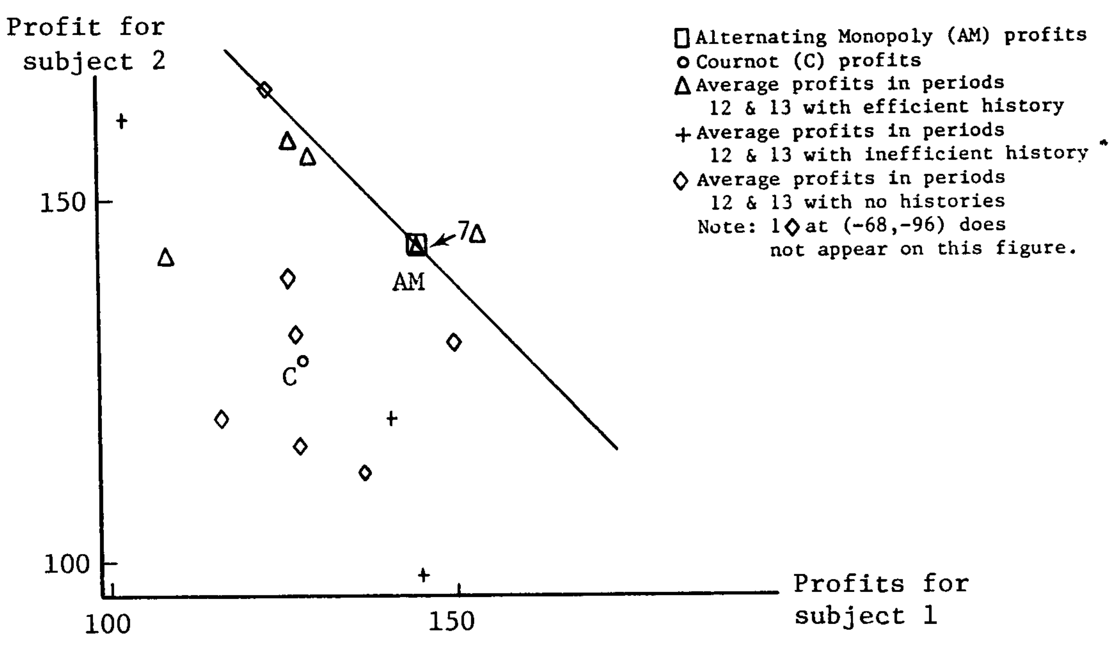

FIGURE 3. Results from History Games.

Players seem to discard inefficient histories. While we currently have only a few observations, the outcomes in games with inefficient histories as well as in those with no history are clustered about the Cournot prediction in 8 of 11 cases (2 out of 3 with inefficient histories and 6 out of 8 with no histories). In 5 of these 8 games, players jointly earn less than they would at the Cournot outcome while in the other 3 games their joint profits are within 10 francs (5 cents) of this outcome. Thus, players seem to recognize that a history is inefficient and simply ignore it.

However, once players discard their history they are in a similar situation as they faced in the control games. Without any coordinating mechanism, players again seem to have difficulty in forecasting their opponents' choices and due to this, there is no apparent tendency towards an equilibrium in these games. In all games with inefficient histories and with no histories, players never repeated their decisions over the final three periods of play, and only once (in the game with no history which achieved the efficient frontier) did players repeat their decisions over the final two periods of play.

SUMMARY

The two sets of games allowed a limited degree of coordination to be effected by the players. This ability to coordinate allows players to better forecast their opponents' decisions and assists them in achieving efficient outcomes. Thus, in both the communication and efficient history games, the availability of such coordination opportunities generally resulted in the complete information collusive outcome in spite of the lack of complete information. Moreover, lack of such opportunities to coordinate, as seen in the control games as well as the no-history and inefficient-history games, led to inefficient outcomes.

Acknowledgments

Support by NSF Grants SES-8218684 and IST-8610360 is gratefully acknowledged. We thank Marc Knez for assistance in computer analysis and Cynthia Carlson, Lisa Armstrong and David Waldron for help in running the experiments.

References

Alger, D. "Laboratory Test of Equilibrium Predictions with Disequilibrium Data," Review of Economic Studies, 54(1987), 105-145.

Aumann, R., "Agreeing to Disagree," Annals of Statistics, 4 (1976), 1236-1239.

Benoit, J-P. and V. Krishna, "Finitely Repeated Games," Econometrica, 53

(1985),905-922.

Daughety, A.F. and R. Forsythe, "Regulation and the Formation of Reputations: A Laboratory Analysis,: in Public Regulation: New Perspectives on Institutions and Policies, E. Bailey, ed., MIT Press, (1987a), 347-398.

__________, "Regulatory-Induced Industrial Organization: A Laboratory Investigation," Journal of Law, Economics, and Organizations, (1987b) forthcoming.

__________, "Collusion without Common Knowledge," Working Paper, Department of Economics, The University of Iowa, (1987c).

Friedman, J.W., "An Experimental Study of Cooperative Duopoly," Econometrica, 35 (1967), 379-397

__________, "Cooperative Equilibria in Finite Horizon Noncooperative Supergames," Journal of Economic Theory, 35 (1985), 390-398.

__________, Game Theory with Applications to Economics, New York: Oxford University Press (1986).

Green, E. and R. Porter, "Noncooperative Collusion Under Imperfect Price Information," Econometrica, 52 (1984), 87-100.

Kalai, E., and M. Smorodinsky, "Other Solutions to Nash's Bargaining Problem," Econometrica, 43 (1975), 513-518.

Kreps, D., P. Milgrom,, D.J. Roberts, and R. Wilson, "Rational Cooperation in the Finitely Repeated Prisoner's Dilemma, Journal of Economic Theory, 27(1982), 245-252.

Plott, C.R., "Industrial Organization and Experimental Economics," Journal of Economic Literature, 20 (1982), 1486-1527.

Raiffa, H., "Arbitration Schemes for Generalized Two-Person Games," in Contributions to the Theory of Games, eds. H. Kuhn and A. W. Tucker, Princeton: Princeton University Press (1953).

Shubik, M., Game Theory in the Social Sciences: Concepts and Solutions, Cambridge: MIT Press (1982).

__________, A Game-Theoretic Approach to Political Economy, Cambridge: MIT Press (1984).

Stigler, G., "A Theory of Oligopoly," Journal of Political Economy, 47 (1964), 44-61.