AgEcon SEARCH
RESEARCH IN AGRICULTURAL & APPLIED ECONOMICS

The World's Largest Open Access Agricultural & Applied Economics Digital Library

This document is discoverable and free to researchers across the globe due to the work of AgEcon Search.

Help ensure our sustainability.

Give to AgEcon Search

AgEcon Search
<http://ageconsearch.umn.edu>
aesearch@umn.edu

*Papers downloaded from **AgEcon Search** may be used for non-commercial purposes and personal study only. No other use, including posting to another Internet site, is permitted without permission from the copyright owner (not AgEcon Search), or as allowed under the provisions of Fair Use, U.S. Copyright Act, Title 17 U.S.C.*

Queen's Economics Department Working Paper No. 863

An Empirical Model of Strategic Choice with an Application to Coordination Games

Debra J. Holt

Department of Economics
Queen's University
94 University Avenue
Kingston, Ontario, Canada
K7L 3N6

11-1992

Discussion Paper #863

**An Empirical Model of Strategic Choice With
an Application to Coordination Games**

by

Debra J. Holt
Queen's University
Kingston, Ontario
K7L 3N6

November 1992

**An Empirical Model of Strategic Choice
With an Application to Coordination Games**

Debra J. Holt

Department of Economics
Queen's University
Kingston, Ontario K7L 3N6

email: holt@econ01.econ.queensu.ca

Queen's University Discussion Paper #863

February 1991
Current Draft
November 1992

I would like to thank Robert Forsythe for the use of the coordination game data. Valuable comments were offered by James Bergin, Dan Bernhardt, Chris Ferrall, Elizabeth Hoffman, Bart Lipman, James MacKinnon, Bob Miller, Hal Varian, and seminar participants at Queen's University, Washington University, and The International Conference on Game Theory. Rob Porter and two anonymous referees provided very useful comments and suggestions. Special thanks to Michael Shostak. Some parts of this paper circulated with the title, "An Empirical Model of Strategy Choices in Coordination Games."

ABSTRACT

This paper develops structural estimation techniques which can be applied to experimental game data to consistently estimate and test models of strategic choice. I assume that the true game is a Bayesian game of incomplete information and that the observed actions are supported by an equilibrium of this game. The implied structure permits estimation of the distribution of player types. This method provides a unified framework for rigorously testing hypotheses about behavior in games within the theory.

Behavior observed in coordination game experiments is inconsistent with the joint hypotheses of rational play and complete information of the game. One possible explanation is that some players are altruistic; another is that subjects are learning as they play the game. Neither can be clearly rejected or supported with usual empirical approaches. The structural approach adopted in this paper does give clear evidence regarding the importance and validity of the alternative explanations.

I. Introduction

This paper develops structural estimation techniques which can be applied to experimental game data in order to consistently estimate and test models of strategic choice. I assume that the true game, which includes all repetitions of the announced game in the experimental session, is a Bayesian game of incomplete information and that the observed actions are supported by an equilibrium of this game. The structure implied by this assumption allows one to estimate the distribution of player characteristics which determines the equilibrium outcomes of the game. Questions about behavior in games can be formally stated in terms of the players' characteristics; hence a consistent statistical test of the theory's predictions can be constructed as a restriction on the set of possible relevant characteristics, or types.

In experiments on games, the observed outcomes typically include violations of the equilibrium predictions of the complete information game. One explanation for the deviations is that they are random mistakes or errors. But a subject who takes account of the errors that we observe can often do better by choosing an action other than that predicted by the Nash equilibrium of the complete information game. So play of the predicted Nash equilibrium strategy is not synonymous with rational choice nor is deviation from that prediction necessarily an indication of irrationality or ineptness. In many experimental games the deviations from the equilibrium prediction are not distributed among other options as a "mistakes" story would indicate; instead the observed actions differ systematically from the prediction and are consistent with an equilibrium of a game in which some assumption of the intended game has been relaxed. Also the deviations from the predicted equilibrium path typically decline with experience in the experimental session, indicating that learning may be an impor-

tant component of a model of strategic behavior. A statistical approach which does not incorporate alternative explanations for the observations into the model can be expected to generate misleading estimates and test statistics. The expectation of an opponent's errors or uncertainty about an opponent's preferences or uncertainty about the extent of an opponent's knowledge about one's own rationality affect the best response — in that sense actions labelled “deviations” are informative about the equilibria of the true game.

Unlike the Nash equilibrium of the assumed complete information game, the Bayesian equilibrium of the true incomplete information game does not provide directly refutable predictions. Anything which is individually rational, or non-dominated, can be supported as an equilibrium outcome as the players' priors and risk attitudes are varied. In the experimental setting, where preferences may not even be a monotonic function of the monetary payoffs, one cannot identify the dominated actions; any outcome at all can be supported by a Bayesian equilibrium. Since there are no zero probability events, there is support within the theory for any observed path.

Thus, statistical analysis of the data does not require specific, typically unjustifiable, assumptions which *a priori* eliminate certain paths. Given the Bayesian game framework one can maintain minimal assumptions and let the data indicate the extent to which other assumptions are justified. Modelling the experimental sessions as Bayesian games provides a rich theoretical structure in which to analyze the data and implies a statistical framework for making inferences and testing hypotheses. Under the assumption that the observed path is supported by a Bayesian equilibrium, one can determine the subjects' types from that path.¹ An additional assumption that the experimental subjects' types have been drawn indepen-

¹ Identification problems may arise; some types may play the same strategy. These issues are addressed in later sections.

dently from the set of possible types leads directly to a completely specified and tractable empirical model.² The parameters I estimate are the proportions of different types of players in the population and players' beliefs on the distribution of types; most hypotheses are stated as restricting one or more of these parameters to zero.

This analysis follows in the tradition of work by Palfrey and Rosenthal [1988] and Camerer and Weigelt [1988] in viewing the true game as one of incomplete information and making inferences about players' beliefs. Papers by Blume, Holt, and Salant [1987] and McKelvey and Palfrey [1992] also explicitly model the true game as one different from the game the experimenters intended and obtain parameter estimates to describe the true game under the assumption that subjects are playing an equilibrium of that game. In Blume, Holt, and Salant players are assumed to be playing an epsilon perfect or proper Nash equilibrium of a voting game; they estimate the probability of trembling and learning parameters which describe the decline of the rate of trembling with experience. In McKelvey and Palfrey the players are assumed to be playing a sequential equilibrium (with trembles) of a centipede game with two types of players, selfish and altruistic. Each player can have different beliefs about the proportion of altruistic types but is assumed to believe that all other players share his beliefs. The players' probability of trembling may decline over time. They estimate the proportion of players who are altruistic, average beliefs about the proportion of altruists and the trembling and learning parameters.

The econometric approach I use is closely related to structural estimation procedures such as those developed in Miller [1984], Wolpin [1984], Pakes [1986], and Rust [1987]. Brown and Rosenthal [1990] use an approach

² In principle, one may prefer to assume that the distribution of the subjects' types differ from the population distribution; some types may be more likely to respond to a call for subjects. For example, the subjects may be the least risk averse of the pool from which they are recruited.

similar to the one in this paper to analyze a complete information game with a unique mixed strategy equilibrium.

The innovation of this paper and Holt [1988] is to apply structural estimation techniques to the model of the session as a Bayesian game of incomplete information. This approach provides a unified and explicit method for determining which forms of incomplete information are important to the outcomes of a game. It differs from previous work incorporating considerations of incomplete information in several ways. First, I assume that the observed session path is exact equilibrium play and from that determine the informational structure and distribution of types which are consistent with the observed actions. In particular, unlike McKelvey and Palfrey [1992], the type as defined in the theoretical model is the same type that I identify in the empirical analysis. For example, an equilibrium prescribes a play path for an altruistic type; in this paper the subjects who are identified as altruistic types are observed to play that precise path, rather than that path with an error rate of 15 or 20 percent. I analyze apparent mistakes and learning behavior by allowing for various learning types that have imprecise information about the game. Second, I require that the players be correct in their assessment of others' beliefs; common beliefs are indeed commonly held. Third, there are no *a priori* limitations on the types of players or the forms of incomplete information which can be incorporated in the analysis. Finally, as in the Brown and Rosenthal analysis of a complete information game, I fully develop the statistical model which is implied by the theory.

The model is developed in Section II and the empirical model is presented in Section III. Section IV describes the coordination game experiments and Section V develops a specific model for analyzing that data. There are competing hypotheses which are loosely consistent with the observed behavior in coordination games. One is that there are some players who are altruistic in the sense that they care about fair outcomes or maximizing joint payoffs — or there are people who believe that there are some altruists in the population. The other is that players are learning over

the course of the experiment. Less structured analyses provide no way of testing the validity of the alternative explanations. However, modelling the behavior explicitly in the Bayesian game framework allows a rigorous test of these hypotheses. We obtain very strong evidence that there are no players who care about fair outcomes or maximizing joint payoffs and there are no players who believe others have these preferences. On the other hand, learning behavior obtained strong support. Section VI presents and discusses these results. Section VII concludes.

II. The Model

The model is motivated by the observation that explanations of observed behavior in experimental games are informally presented as, and can be formally formulated as, questions about the players' types. Equilibrium strategies and the game form implicitly define a mapping from the type space into the action space. Therefore, given the observations of actions in a particular game, one can use that mapping to infer the players' types.³

Ledyard [1986] points out that “[any] observed behavior can be rationalized as the outcome of a Bayesian equilibrium of some game.” Further, without resorting to strong assumptions on the form of beliefs and preferences, one can only rule out dominated strategies with Bayes-Nash equilibria. In the experimental setting, where true preferences may not be represented by the monetary payoffs, dominated strategies are not unambiguously identified.

Assumption: The observed session path is a history of a Bayesian equilibrium of the session-game.

This assumption assures that all observations are explained within the theory while more restrictive models of behavior can still be identified. The observed path identifies the players' types or the subsets to which their types belong. The data reveals which, if any, restrictions or additional as-

³ The observed action path in an experimental session does not always identify a unique type for a player. This is less of a problem than one might expect, primarily because experiments are typically designed so that the types of interest in the analysis will take different actions at some point in the experiment. Secondly, other information about a player, such as his play in a different experimental session, can be used to make a determination of type. In this paper, I select the type combination that maximizes the likelihood function in the instances of non-unique identification.

sumptions on the game are justified. So, for example, if everyone is playing the Nash equilibrium of the complete information game, this is revealed by the data. More generally, we obtain evidence on whether “there are natural restrictions on utility [and beliefs] which prevent the uninformative explanation that all experimental and market generated observations are Bayesian equilibrium outcomes.”⁴

The session-game.

The experimental session is modelled as a multi-stage game with incomplete information in which both the subjects and experimenters are players. The experimenter-players may or may not have (perceived) payoff-relevant roles. The first move is by nature — certain subjects show up, subjects are randomly assigned roles, and so on. The next stage may be one in which information is revealed, perhaps by an experimenter-player giving instructions for the experimental game. The specificity of the instructions vary widely — some might be considered correlation devices while others will just provide enough information to make some aspects of the subsequent stage-games known to the subjects. The instructions include the information that subjects’ payoffs are a function of their own and other players’ actions and how the payoffs are jointly determined. This stage is sometimes followed by practice games — an unpaid training period. This stage can also, in some cases, serve as a signalling or correlation device. Next, the subject takes the part of a player in the announced game for a set of repetitions. Typically some feedback regarding performance and contribution to expected total payoff is given at the end of each repetition. For example, a player may be informed of the points earned in that repetition where the points will be converted to dollars at the end of the session by some known, perhaps stochastic, process. Finally, subjects are informed of their actual earnings and the session is terminated.

⁴ Ledyard [1986].

Describing the session-game formally requires considerable specific information about the procedure used in running the experiment as well as a clearly delineated procedure. Supposing the least amount of information and for simplicity in notation, we write all repetitions of the announced game in the session as the true game of incomplete information. There are n players, or subjects. There are R repetitions of the game in the session. Let $h_i^r \in H_i^r$ give the history of the game for player i at repetition r . This will include all moves that player i has made through repetition $(r - 1)$. It also includes all moves of the other players which player i has been able to observe through repetitions $(r - 1)$. Let S_i be the set of pure strategies available to player i . So S_i is a function from histories into the set of available actions.

The move by nature at the beginning of the game determines an n -vector of types $t = (t_1, \dots, t_n)$. The type, t_i , describes i 's private information about the way he will play the game, or equivalently, summarizes those characteristics which affect his play of the game. Let T_i be the set of all possible types for player i ; $\forall i$; $t \in T = \times T_i$. In finite games, one can only distinguish among a finite number of subsets of T_i ; so without loss of generality, $\#T_i < \infty$. I assume that the true distribution over T_i places non-zero probability on each point in T_i . Note that no possible beliefs or payoffs are being excluded *a priori*. Therefore it follows from Ledyard's [1986] result that the players do not assign zero probability to any information set.

The equilibrium of the Bayesian game of incomplete information is a function from the players' types to strategies, such that each player's strategy choice maximizes his expected utility conditional on his own type and his beliefs about the other players.

A player's strategy is $s_i \in S_i^{T_i}$. A Bayesian equilibrium of the game in pure strategies is $s = (s_1, \dots, s_n)$ such that, for all i , for all $t_i \in T_i$,

$$s_i(t_i) \in \arg \max_{r_i \in S_i} \sum_{t_{\sim i}} p(t_{\sim i} | t_i) u_i((s_{\sim i}(t_{\sim i}), r_i), t)$$

where p is the players' beliefs about the distribution from which the types are drawn. The equilibrium can also be written as a mapping from the players' types into histories; it therefore defines a sequence of histories for each type, i.e. for each subject.

Recovering types from observed histories.

While it is common to include, in the type description, beliefs and beliefs about beliefs and so on, the tractability of the method of analysis used in this paper relies on maintaining a clean separation between beliefs and other characteristics which define a type. For the remainder of this paper, I use 'type' to describe the set of relevant characteristics minus beliefs. Both beliefs and the distribution of types are estimated. In the application to coordination games, I assume common priors over the set of types because of the small amount of data. The algorithm is easily adapted to the case in which different types have different beliefs.

The types of the experimental subjects are inferred from the observed session path by the process of checking (numerically) for beliefs p and type combinations t that are consistent with observing $\{h_1^R, \dots, h_n^R\}$ in equilibrium. Let $\theta = (\theta_1, \dots, \theta_M)$, where θ_m gives the proportion of the population who are of type m , fully characterize the distribution of types. Then this process yields estimates of θ and the players' common beliefs, p .

The algorithm begins with a preliminary analysis of each players' observed action path and a tentative assessment of the players' possible type(s). That is, there are some assignments of types that can be ruled out *a priori* and doing so when possible increases the efficiency of the algorithm.

The second step is to enumerate the possible equilibrium strategies for each type.

The third step is to put together all theoretically feasible type-strategy

combinations for each session. For example, in the coordination games illustrated in Table I, we might observe some players choosing action 1 at every repetition and others choosing action 2 at every repetition. Consider just two of the possible types: type 1 has $u_i(s, t) = y_i$ and type 2 has $u_i(s, t) = y_i + y_j$ where y_i represents point payoffs to player i and y_j represents the payoffs to i 's opponent. Then a feasible type-strategy combination has all type 1 players using strategy 'always play 1' and all type 2 players using strategy 'always play 2'. The converse is also a feasible type-strategy combination; however, some type 1 players with strategy 'always play 1' and other type 1 players with strategy 'always play 2' is not feasible. There is no attempt to check whether the type-strategy combinations are sensible given the payoffs at this step.

The fourth step is to generate a value for p and then, given that value and one of the feasible type-strategy combinations, compute the implied distribution over actions.

Finally, for each p , each session, and each feasible type-strategy combination, check that no player wants to deviate from the proposed strategy. (This is a check that a set of inequalities hold.)

If more than one type-strategy combination survives this check for one or more of the sessions then another step is required. The algorithm has now generated at least one combination for each session, for the given p . Construct all possible overall type-strategy combinations across sessions; select that one which maximizes the likelihood function conditional on p .⁵ In the coordination game data, for representative values of p , there were from one to three type-strategy combinations per session.

Repeat from step four for (almost) all values of p .

⁵ The next section describes the estimation procedure and the appendix contains a detailed algorithm for both the type detection and the estimation procedures.

This algorithm is implementable for any game since it is essentially an exhaustive set of consistency checks, but it may not be the most efficient algorithm for all applications. There are simplifications that can be made in some cases; for example, optimal bidding strategies are simple functions of (some aspects of) the player's type.

This process generates, for each p , an assignment of a type to each player in the set of experimental sessions.

III. Estimation

Each of Z experimental session has n players, each taking an action at R repetitions.⁶ The sequence of actions for one player is one observation. The observed session path, the $n \times R$ actions, is used to construct $h = \{h_i^R\}$, the set of histories realized in the session. Then, the process described in the preceding section finds a type combination t and beliefs p which are consistent with the observation of h in equilibrium.

I assume that the n players' types are drawn independently from the population of types. The distribution of types is given by $\theta = (\theta_1, \dots, \theta_M)$ where θ_m is the proportion of type m players in the population. Summarize the draw of n types by $g = (g_1, \dots, g_M)$ where g_m is the number of type m players and $\sum_m g_m = n$. Then the probability of obtaining a particular draw of types is given by

$$\frac{n!}{g_1! \dots g_M!} \cdot \theta_1^{g_1} \dots \theta_M^{g_M}.$$

For the multinomial, $E(g) = n\theta$; $\text{var } g_m = n\theta_m(1 - \theta_m)$; and $\text{cov}(g_l, g_m) = -n\theta_l\theta_m$ for all $l \neq m$. Also note the marginal distribution of any g_m is binomial with parameters n and θ_m . Finally, since $g^1, \dots, g^Z$, the summaries of the type draws for the Z sessions, are independent k -dimensional random vectors each with a multinomial distribution with parameters n and θ then $\sum_z g^z = G$ has multinomial distribution with parameters $nZ = N$ and θ . The assignment of types consistent with h in equilibrium does depend on beliefs; therefore, g is a function of p . The likelihood function, conditional on p , is the joint probability function:

$$L(s|\theta)|_p = \frac{N!}{G_1! \dots G_M!} \cdot \theta_1^{G_1} \dots \theta_M^{G_M}.$$

⁶ The number of players, and the number of repetitions may in fact vary across sessions in some experimental data. The extension is straightforward.

Note that this parameter sub-space is an $(M - 1)$ -dimensional simplex. There are M first-order conditions characterizing the maximum of the constrained log likelihood function:

$$G_m/\theta_m = \lambda$$

where λ is the Lagrange multiplier on the constraint that the θ_m 's sum to one. Sample frequencies are unbiased estimators and $\lambda = N$ in this case. It is straightforward to show that

$$N^{1/2}(\hat{\theta} - \theta) \sim \mathcal{N}(0, \Sigma)$$

where the diagonal terms in Σ are $\theta_m(1 - \theta_m)$ and the off-diagonal terms are $-\theta_l\theta_m$.

The true likelihood function, $L(h|\theta, p)$, is continuous with respect to θ , but not with respect to p . That is, the mapping from session paths into types is a discontinuous function of p . The (unconditional) likelihood function is

$$L(h|\theta, p) = \frac{N!}{G_1(p)! \cdots G_M(p)!} \cdot \theta_1^{G_1(p)} \cdots \theta_M^{G_M(p)}.$$

Maximum likelihood estimates for θ and p are obtained in a two step process. First I compute the set of θ 's that maximize the conditional likelihood functions — that is, for each p , obtain $\theta|p$ which maximizes the likelihood function conditional on that p . In this first step, the estimator for θ_m conditional on p is $G_m(p)/N$.

In the second step, I find the maximizing value(s) for p . This is a search over a surface in an $M - 1$ dimensional simplex made up of flat spots and discontinuities. Recall there are convex subsets of p which are consistent with the same type combination in equilibrium, which generate the same G and the same estimate of θ . The method I use to find the set of p 's that maximize the likelihood function makes use of the fact that the identification of types required checking p exhaustively.

The estimation procedure embeds the type detection algorithm. For each p , carry out the stages described in the preceding section and obtain the assignment of types. Then compute G and obtain $\hat{\theta}|p$ and evaluate the likelihood function. Finally, keep an ordered list of likelihood function evaluations with pointers to the associated $\hat{\theta}|p$ and p or set of p 's. Note that one cannot sort only on the value of the likelihood function since there may be more than one flat with the same log likelihood value (each with a different conditional estimate of θ and set of p 's). So the pointer list should have both $\hat{\theta}|p$ and the value of the likelihood function. The appendix contains a detailed description of the combined estimation and type detection algorithm.

Obtaining standard errors for conditional estimates of θ is straightforward. The method that can be used to obtain actual standard errors will depend both on the experimental game and the experimental design. In some games, the strategy choices are quite sensitive to changes in p and in others they are very insensitive. Regardless of the class of games, when there is sufficient variation in game payoffs in the set of experimental sessions, $G(p)$ approaches a continuous function. In the coordination game analysis presented in this paper, the flats in p were very large. I do not obtain standard errors; instead, I numerically determine a confidence region on a dimension-by-dimension basis. From a point in the set of p 's that maximize the log likelihood function, I search outward in each dimension until the new value of the log likelihood is significantly different from that at the maximizing value according to the likelihood ratio test.

IV. The Coordination Game Data

The data analyzed are from experiments run by Cooper, DeJong, Forsythe, and Ross. The focus of the experiment,⁷ is a symmetric, simultaneous move, complete information game with multiple, Pareto-ranked Nash equilibria. The game has two players, each of whom has three choices available. Several versions of the game are used but all share these basic characteristics. The purpose of the experiments was to obtain information about the selection of equilibria in this type of game. In these games even an assumption that everyone always plays a Nash equilibrium strategy is insufficient to predict the outcome. Table I gives payoff matrices for the coordination games analyzed. In all these games both $(1, 1)$ and $(2, 2)$ are trembling-hand perfect and proper Nash equilibria. The equilibrium $(1, 1)$ is Pareto dominated by $(2, 2)$. The symmetric outcome $(3, 3)$ is a cooperative outcome which pays an amount greater than the Pareto dominant Nash equilibrium in some of the games and an amount between the two Nash equilibrium payoffs in others. In one game, the cooperative outcome is also a Nash equilibrium and Pareto dominates $(2, 2)$. However, it is not trembling-hand perfect.

The coordination games differ only in payoffs to dominated strategies. Under the assumption of complete information the theory predicts these payoffs should have no effect on behavior in the games. However, as we describe below, there are marked and systematic differences in the outcomes of the games — evidently related to the payoffs to the dominated strategies.

There are eleven subjects in each experimental session. The sessions are conducted as follows. There is a random, anonymous assignment of

⁷ The experimental results are discussed in Cooper, DeJong, Forsythe, and Ross [1990].

See also Cooper et al, [1989] for a more detailed discussion of the results.

players to groups of two,⁸ and each pair plays a symmetric, simultaneous move game with a unique dominant strategy equilibrium. This random assignment to play the dominant strategy game is repeated ten times altogether. Following these first ten stages of the game, the players are randomly assigned (again in groups of two) to play one of the coordination games. So $n = 11$, $Z = 7$, $N = 77$, and $R = 30$. This random assignment to anonymous opponents in the (identical) coordination games is repeated twenty times altogether. Payoffs are assigned according to the Roth-Malouf [1979] procedure. The anonymity and random assignment of opponents rule out certain kinds of signalling behavior and path dependence. However, a player can clearly learn about the game and about the distribution of types of the other players.

⁸ One player is idle at each stage.

Table I

	1	2	3
1	320, 320	440, 420	500, 180
2	420, 440	600, 600	660, 360
3	180, 500	360, 660	420, 420

Dominant Strategy Game

	1	2	3
1	350, 350	350, 250	1000, 0
2	250, 350	550, 550	0, 0
3	0, 1000	0, 0	600, 600

Game 3

	1	2	3
1	350, 350	350, 250	700, 0
2	250, 350	550, 550	0, 0
3	0, 700	0, 0	600, 600

Game 4

	1	2	3
1	350, 350	350, 250	700, 0
2	250, 350	550, 550	1000, 0
3	0, 700	0, 1000	600, 600

Game 5

Table I (con't)

	1	2	3
1	350, 350	350, 250	700, 0
2	250, 350	550, 550	650, 0
3	0, 700	0, 650	600, 600

Game 6

	1	2	3
1	350, 350	350, 250	700, 0
2	250, 350	550, 550	0, 0
3	0, 700	0, 0	500, 500

Game 7

	1	2	3
1	350, 350	350, 250	1000, 0
2	250, 350	550, 550	0, 0
3	0, 1000	0, 0	500, 500

Game 8

	1	2	3
1	350, 350	350, 250	1000, 0
2	250, 350	550, 550	0, 0
3	0, 1000	0, 0	1000, 1000

Game 9

Table II
Nash Equilibrium and “Cooperative” Outcomes

Game	Nash 1	Nash 2	Coop.
3	0.51	0.04	0.01
4	0.51	0.01	0.03
5	0.00	0.85	0.02
6	0.00	0.64	0.04
7	0.02	0.75	0.00
8	0.09	0.46	0.00
9	0.15	0.00	0.35
overall	0.18	0.39	0.06

Table II provides a summary of outcomes observed in the sessions. Overall, a Nash equilibrium outcome is observed 57 percent of the time.⁹ The first Nash equilibrium outcome is observed 18 percent of the time, and the second (Pareto dominant) Nash equilibrium outcome is observed 39 percent of the time. The cooperative outcome is observed only in only 6 percent of the plays. However, the pattern differs markedly between the games. In the two games (three and four) for which the payoff to action 2 when the opponent chooses the dominated strategy is zero, nearly 51 percent of the outcomes are the first Nash equilibrium and only about 2 percent are the second. In other games (five and six), one may actually be better off choosing action two when the opponent chooses the dominated strategy. In those games, the second Nash equilibrium outcome occurs 75 percent of the plays on average.

There are two more games (seven and eight) in which the payoffs are

⁹ Cooper et al find support for Nash equilibrium behavior despite this proportion of Nash equilibrium outcomes because they consider equilibrium choices instead of outcomes. Overall, 85 percent of the choices were either action 1 or action 2 — both of which are Nash equilibrium strategies.

identical to the first set of games (in that the payoff to action 2 against the dominated strategy is zero) except that the payoffs to the cooperative outcome are now lower than the payoffs in the second Nash equilibrium outcome. Yet the frequency with which the second Nash equilibrium is obtained is quite high — about 60 percent of the outcomes are the second Nash equilibrium, while only about 5 percent are the first Nash equilibrium. An obvious interpretation for this result is that the players' assessment of the probability the opponent will play a dominated strategy is lower in these games.

Table III
Empirical Distributions of Actions and Best Replies

Game	Act 1	Act 2	Act 3	BR	pct BR
3	0.71	0.20	0.09	1	0.71
4	0.71	0.14	0.15	1	0.71
5	0.02	0.92	0.06	2	0.92
6	0.07	0.79	0.15	2	0.79
7	0.13	0.86	0.01	2	0.86
8	0.30	0.68	0.02	2	0.68
9	0.38	0.03	0.59	1	0.38
overall	0.33	0.52	0.15		0.56

Table III indicates that, except for the game with three Nash equilibria, the percentage of actions which were best replies to the empirical distribution for that particular version of the game is quite high — ranging from 68 to 92 percent. The empirical distributions and best replies (and proportion of players choosing the best reply to that empirical distribution) change across games in a way which suggests that players are responding to the information conveyed by the payoffs to the dominated strategy.

The outcomes and payoffs in the introductory dominant strategy game repetitions are very similar across games and across players. In all sessions,

the dominant strategy equilibrium is realized over 85 percent of the time. There is no correlation between the outcomes in the dominant strategy game repetitions and the outcomes in the subsequent repetitions of the coordination games. Therefore we have strong indications that the differences observed in the coordination games are in fact related to the different payoffs and induced differences in beliefs about actions instead of being an artifact of drawing very different samples of players to play the various games. Finally, there are noticeable changes in the actions chosen over time and these patterns of change are different in different games.

The descriptive statistics reported in Tables II and III are consistent with the interpretation of the true game (the experimental session) as one of incomplete information. If the session game were actually a complete information game comprised of 20 repetitions of one of the games shown in Table I, then we would expect no systematic differences in outcomes across games 3 through 8. The differences we do observe suggest that players may be uncertain about their opponents' preferences or unsure about how to play the game, or both. The changes in actions over time and the differences in temporal patterns across games indicate that the behavior could be consistent with Bayesian updating of priors as information is obtained during play of the session-game.

V. A Simple Model for Coordination Games.

The behavior observed in coordination game experiments is inconsistent with the joint hypotheses of rational play and complete information of the game. One explanation for the observed play is that some players have preferences for fair outcomes or beliefs that others prefer fair outcomes. Another explanation is that the subjects begin the session with some uncertainty about the game and learn over repetitions of the game in the session. Both explanations are consistent with the observed patterns of play in the experiments and neither can be clearly rejected or supported with the usual empirical approaches. We develop a model which is capable of giving clear evidence regarding the importance and validity of the alternative explanations.

The questions I consider in specifying the model are: Are there players with preferences for fair outcomes — that is, are there players whose utility functions are increasing in the opponent's experimental payoffs as well as their own? Do players' expectations (perhaps erroneous) that there are altruistic types explain some of the behavior we observe? Are there players who are uncertain about the game and who learn as they play? Can we disentangle the effects of altruistic players from learning players and beliefs about altruistic vs. learning players?

Consider a model in which we can restrict our attention to six relevant subsets of T_i . The first subset contains only one type: type 1. Type 1 players are perfectly rational and fully informed about the game and have beliefs, p , about the distribution of types. The utility associated with a given outcome is assumed to be represented by the points earned in that outcome.

The second subset also contains one type: type 2. Type 2 players have a preference for fair allocations. These types differ from type 1 players only with respect to their preferences. A type 2 player has utility associated with an outcome given by $u_i = \sum_j y_j - |y_i - y_j|$ where y_i and y_j indicate the points awarded players i and j at the outcome. Several other specifications of the utility function for type 2 players were also used. The results were insensitive to these variations. All variations included either the sum of the payoffs (positively) or the absolute difference between payoffs (negatively), but not necessarily both.

There are several equivalent interpretations of the less informed types. In the underlying model these types simply have more uncertainty about the possible types of other players in the game. So types 1 and 2 place positive probability only on those types that are actually in the game while the less informed types place positive probability on those plus some other types. (This could be a difference in the priors or in the information.) As a consequence, the less informed types have beliefs about the distribution of actions taken in equilibrium which does not match that induced by the true distribution of types.

The direct way to characterize the uninformed types is according to their beliefs about types. (I have assumed that the preferences of types 3, 4, and 5 are represented by their experimentally generated payoffs.) For any given beliefs about types, there is an equivalent characterization in terms of some learning algorithm. That is, for any learning algorithm there exists a type with preferences and beliefs such that that type's equilibrium behavior yields the same action path as the learning algorithm and vice versa. I choose to describe the less informed types in terms of learning algorithms instead of their observationally equivalent beliefs over types (and

implied equilibrium play). For the informed types, the situation is the same. Their equilibrium strategies depend, not directly on the less informed types' beliefs about the true game, but on the induced distribution over actions taken by those types in equilibrium.

Players whose types are in subset three are rational Bayesians who are incompletely informed about the game and update their priors over actions during the course of the game.¹⁰ Players in subset three are assumed to have utility represented by points earned. These types have uniform beliefs over actions, q , at the beginning of the session by the principle of insufficient reason. Their updated probability of observing an action at repetition r is a function of h , an initial weight on priors (described below) plus the number of repetitions, and is given by:

$$q_r(a) = [q_{r-1}(a) + \frac{1}{h}] [\frac{h}{(h+1)}]$$

for a observed in $r - 1$ and

$$q_r(b) = q_{r-1}(b) [\frac{h}{(h+1)}]$$

for actions b not observed at $r - 1$. This is an implementation of fictitious play. If beliefs over the mixture of actions the player faces at each period is from the Dirichlet family, then the updating rule above is the correct Bayesian updating rule. For players who do not know what action the opponent will take in each repetition and who believe that their own action will have no effect on future play, a myopic best response is the proper behavior rule. In these experiments, since the players are assigned randomly and anonymously to opponents, a player might reasonably believe his action will not effect the play of his future opponents.

¹⁰ These priors over actions are perhaps induced by priors over payoffs or over some other aspect of the game.

The players begin as though they have a fixed weight on the priors over actions — a player with very low weight on the priors might have $h = 3$ while a player with high weight on the priors might have $h = 30$. Suppose they each begin the session with uniform beliefs over three possible actions. Then upon observing action 2, the updated beliefs will be $(1/4, 1/2, 1/4)$ for $h = 3$ and $(10/31, 11/31, 10/31)$ for $h = 30$. The weight on the priors is an indication of the amount of confidence the player has in the priors. For example, a player who comes to one of these sessions after previously participating in fifty similar sessions would place a higher weight on the beliefs with which he begins this session. On the other hand, an inexperienced player would place a lower weight on his beliefs entering the session. The types in this subset differ only with respect to the initial value of h they use. From this point on, I will refer to these types collectively as type 3 players. I also tested an alternative model in which all but ϵ weight was on actions one and two. The results changed very little and netted no change in the number of players judged to be type 3.

Types in subset four are like those in three except that there are minimal restrictions on beliefs. The only behavioral constraint for these types is: if the action taken at repetition r is different from the one taken at $r - 1$ then the action at r must be a best response to the opponent's play at $r - 1$. Types in subset 3 do satisfy this restriction; the behavior of types in subset 4 is a generalization of types in 3. Therefore in the assignment to types, I assign a player to type 3 if his behavior is consistent with both and to type 4 if it is only consistent with 4. The types in this subset differ with respect to their initial priors and with respect to their updating rule. I will refer to these types collectively as type 4 players.

The fifth subset contains one type; that type uses a very simple learn-

ing algorithm.¹¹ The basic idea of this learning algorithm is: if things are going as well as possible for the player then he will not change his action; if he gets a payoff worse than expected then he will change; in ambiguous circumstances, he may experiment but is not required to do so. So, for example, in game 4, a type 5 player might believe that everyone will coordinate at (3,3); if he gets a payoff of zero, he will change his action. The detailed algorithm is provided in the appendix.

The sixth subset contains all types not included in subsets 1 through 5. While in terms of a partition of the type space this element can be thought of as a residual type, it plays an integral role in the formal model and statistical analysis. In the statistical analysis, including these types excludes the possibility of encountering an observation out of the support. Furthermore, correctly computing the equilibrium mapping requires consideration of these types. That is, equilibrium strategies depend on beliefs about all types, including those in this subset.

The learning types described above are myopic; this specification seems natural given the particular experimental design which produced the data. Recall these subjects played ten other players in random order and anonymously; while they played each other player twice, the order of meeting was completely unpredictable. So though a learning type's choices early in the session could have an impact on later play, it is clear that overall uncertainty about the game overwhelms attempts to predict the effect of early choice on the distribution of actions one would face later in the session. The belief revision process assumes that the players are facing a stationary distribution; this is inconsistent with the learning types being perfectly

¹¹ The learning algorithm is related to one developed in Kelly and Glymour [1989].

rational but a tractable alternative is not obvious.

VI. Results

Results are reported in Tables IV and V. An important feature of the results is that the mapping from session paths into types varied little with changes in p . The unconditional estimates of θ are given in Table IV. The range of p which support those estimates of θ has $p_1 \in (0.95, 1.0)$; p_2 through p_6 can take any value consistent with such a high p_1 . They indicate that approximately 29 percent of the population play a Bayesian equilibrium strategy of the repeated game and care only about their own experimentally generated payoffs in the game. About 24 percent of the players are learning about the game; 19 percent are type 4 players and 5 percent are type 5.

The estimates of θ with the second highest value of the likelihood function are identical except for the estimates of the proportions of types 2 and 4. These estimates are given in Table V. About 3 percent are playing a Bayesian equilibrium strategy of the repeated game but have preferences over the sum of their own and their opponent's experimentally generated point payoffs. Now only 16 percent of the players are estimated to be type 4. The range of p consistent with these maximum likelihood estimates of is quite large and includes p equal to θ 's reported in Table V. I can reject the hypothesis that $p_1 < 0.95$ at the 0.005 level of significance using the $-2\ln\lambda$ test.

Table IV
Unrestricted Estimates of θ and p

index	$\hat{\theta}$	std. errors
1	0.29	0.052
2	0.00	0.000
3	0.00	0.000
4	0.19	0.045
5	0.05	0.025
6	0.48	0.057

$$\ln L = -30.58$$

Table V
Estimates with p equal to θ

index	$\hat{\theta} = \hat{p}$	std. errors
1	0.29	0.052
2	0.03	0.019
3	0.00	0.000
4	0.16	0.042
5	0.05	0.025
6	0.48	0.057

$$\ln L = -35.22$$

I checked the robustness of these results by testing other assignment rules, estimating with other learning specifications, and analysis of simulated data.

The assignment rule I use gives precedence to types 1 and 2. Between types 1 and 2, the algorithm chooses the type-strategy combination across sessions that maximizes the value of the conditional likelihood function. Type 3 gets precedence over type 4 because type 3 behavior is a subset of type 4 behavior. Type 3 also gets precedence over type 5 because in my detection algorithm, a type three would pass the checks for type 5 behavior. This is because type 5 is allowed to experiment but is not required to do so. The appendix gives the necessary details. Type 4 is given precedence over type 5, primarily to enable clear comparisons with type 3.

The preliminary assessment of types found 18 percent whose behavior was consistent with being type 3; but they were also all playing in a way consistent with being a type 1. The preliminary assessment also found that 25 percent of the players were playing in a way consistent with being type 4; approximately 9 percent of those players are also consistent with either type 1 or type 2.

When I changed the order of precedence between types 4 and 5; the estimates of θ_1 , θ_2 , θ_3 and θ_6 were unchanged. The new estimate of θ_4 is 0.17 and the new estimate of θ_5 is 0.04. At uniform p , the log likelihood value was -34.641 .

I report two alternative checks on types one and two. Instead of choosing the type-strategy combinations that maximize the likelihood function I consider arbitrary assignment under two conditions: in the first, the algorithm looks for type 1's first; in the second, it looks for type 2's first.

In both cases, the first combination that passes the equilibrium check is selected. Both are run with $p_1 < 0.95$ since that is a more robust region for comparing alternate specifications.

The first case, with type 1's getting arbitrary precedence, is quite similar to the results when the maximizing combination is chosen. We have: $\theta_1 = 0.26$, $\theta_2 = 0.03$, $\theta_3 = 0.05$, $\theta_4 = 0.17$, $\theta_5 = 0.04$ and $\theta_6 = 0.45$. The log likelihood value is -46.169 . The second case, in which type 2's get arbitrary precedence gives quite different results. We have: $\theta_1 = 0.19$, $\theta_2 = 0.21$, $\theta_3 = 0.0$, $\theta_4 = 0.10$, $\theta_5 = 0.04$ and $\theta_6 = 0.45$. The log likelihood value is -43.132 .

These results illustrate some advantages to this method of analyzing experimental games. First, one has consistent estimators of parameters which have a clear interpretation within the theory. Second, there is a best test for hypotheses which are clearly stated within the theoretical framework. Thus, one can distinguish between different plausible explanations and rigorously test whether they are supported by the data.

As discussed in previous sections, there are competing hypotheses which are loosely consistent with some observed behavior in coordination games. One is that there are some players who are altruistic in the sense that they care about fair outcomes or maximizing joint payoffs — or there are people who believe that there are some altruists in the population. The other is that players are learning over the course of the experiment. Consider a perfectly informed, selfish player's choice if he believes that some proportion of the population is altruistic. A naive altruistic player will choose action 3 in games 3 through 6. The best response is to select action 1 in games 3 and 4 and action 2 in games 5 and 6. This is in fact consistent with summary statistics reported in Table III. At the same time, a model

of players who view this as a game against nature, start the session with positive probability on facing an opponent's choice of action 3, and learn over the course of the session can also explain exactly the same switching between actions 1 and 2.

However, using the Bayesian game framework, one can test the alternative explanations rigorously. This analysis yields very strong evidence that there are no players who care about fair outcomes or maximizing joint payoffs and there are no players who believe others have these preferences. On the other hand, learning behavior obtained strong support. Recall that we always assigned a person to type 2 if their behavior was consistent with both that and a learning type. Thus these results against altruistic preferences obtain even when the altruism argument is given an advantage.

The estimated proportion of type 1 players, 29 percent, indicates that a relatively small group of the population play games according to the standard game theoretic model — even when we account for incomplete information.

VII. Conclusion

Modelling experimental games as Bayesian games of incomplete information and then applying structural estimation techniques provides a promising method for analyzing strategic behavior. This method allows an interpretation of all actions within the theory — we are never required to simply categorize a set of actions as “irrational” or “non-Nash” or “deviations”. Rather all observed actions can be tied to the sets of player characteristics which support them in equilibrium. As a result, one can characterize systematic patterns of behavior which are unexplained by the equilibrium predictions of the complete information game.

The empirical model is directly implied by the theoretical model. Thus it is straightforward to adapt the statistical analysis for any specialization of the game theoretic model. That is, there is no restriction on the sort of strategic behaviors that can be examined in this framework.

The application to coordination games illustrates some of the advantages to following this approach. Competing explanations which are consistent with the observations cannot be directly refuted or cleanly supported by naive analyses of the experimental results. However, modelling these explanations explicitly in terms of the players’ types and estimating the distribution of types under the assumption that the actions are supported by a Bayesian equilibrium of the game yields clear evidence regarding the importance or validity of the alternative explanations. In addition, this method promises to be extremely useful in designing experiments which fully exploit the power of the statistical analysis.

Appendix I

This appendix contains a description of the algorithm which determines the players' types given the observed action paths, finds the maximum likelihood estimates of the parameters θ and p , and constructs an approximate confidence region for the estimates. There are two separable components of the algorithm: the preliminary analysis makes tentative, feasible type assignments and summarizes play paths; in the second component, the estimation procedure is inextricably linked with the process of checking for types and beliefs that support the action path in equilibrium.

Preliminary Analysis

Check for possible assignment to types 3, 4, and 5. These alternative type assignments are held throughout the process of checking for equilibria. If a player is not one of the rational, informed types then this remains his type assignment. If the player is possibly one of the rational, informed types, then the assignment of type is made according to which assignment maximizes the likelihood function (or in an arbitrary ordering to test robustness).

Checking type 3. For each player, start with uniform priors over actions, check whether the player is choosing a best response to that distribution at the first repetition. If so, then update beliefs given observation at the first repetition and check whether the player is choosing a best response to that distribution or the original distribution. Continue for as long as the player is playing a best response to beliefs in the interval between the initial priors and the beliefs updated at each repetition based on the observed actions. The updating rule is

$$q_r(a) = [q_{r-1}(a) + \frac{1}{h}] \left[\frac{h}{(h+1)} \right]$$

for a observed in $r - 1$ and

$$q_r(b) = q_{r-1}(b) \left[\frac{h}{(h+1)} \right]$$

for the actions, b , not observed. The value of h is incremented by one at each repetition. The starting value of h indicates the measure of the strength of the initial priors relative to the weight given new information. For each person, the algorithm begins with $h = 10$ and, so long as the person's actions are not consistent with the implied updating rule, decrements h until $h = 1$. At $h = 1$, the person puts little weight on the initial beliefs and responds strongly to the new information. The maximum value of $h = 10$ was chosen because experimentation and simulation revealed that at that value, for most sequences of opponents' actions, the beliefs change little enough that a type 3 player would be choosing a constant action path over repetitions. The reason for having the person choose a best response to beliefs in a cone defined by the initial beliefs and the updated ones is that the idea that the person is really uncertain about the correct weight to attach to the initial priors and that h is a parameter describing the upper bound to the weight attached to new information. Then the person may proceed playing best response to the initial uniform beliefs over actions for some number of repetitions until a sufficient number of "unlikely" observations indicates that a stronger weight should be given to the new observations.

Checking type 4. A player is given an alternative assignment to type 4 if, at each repetition, he takes the same action as in the previous repetition or takes an action which is a best response to what he observed in the previous repetition.

Checking type 5. Checking type 5 is a matter of checking for prohibited transitions. Transitions which are prohibited are: a move from action 1 to action 3 except in game 9; a move to action 1 after matching on 2; and remaining with action 3 after failing to match on 3. The prohibited transitions vary across payoff structures as follows: in games 3, 4, and 9, a move to 2 from a matched 3 cannot yield an improvement and so is prohibited; and in games 5 and 6, action 2 strictly dominates 3 so a move to 3 from 2 is prohibited.

A player whose action path is not consistent with either of types 3, 4, or 5 is conditionally assigned to the sixth subset of types. It is possible for a player to be assigned to more than one of the learning types. In the primary model, players were assigned to learning types in order by the restrictiveness of the learning algorithm — first to 3, then 4, and finally 5. Alternative specifications were also tested and the results from those are presented in Appendix II. A simple procedure at the end of the preliminary analysis assigns one of the learning types for each player whose behavior is consistent with at least one of them. Recall that in choosing between types 1 and 2 or between one of the informed rational players and a learning type, the type (combination) is chosen which maximizes the conditional likelihood function; when the types are related as types 3, 4, and 5 are, that approach does not seem reasonable. One interpretation of the types 3, 4, and 5, is that they are each proper subsets of a set of uninformed types — that those in subset 3 are more informed than those in either subset 4 or 5 and that those in subset 4 are more informed than those in subset 5 along some dimensions and less informed in others.

The preliminary analysis includes an assessment of each player's action path. The action paths are summarized as: all one's, all two's, all three's, one's and two's, one's and three's, two's and three's, or all actions. Given the player's action path, a potential type assignment of 1 or 2 or both is given to each player. This initial assignment is simple — an action path of all one's, all two's, or one's and two's is consistent with a player being type 1. Every action path is consistent with a possible assignment to type 2. For game 9 and for some beliefs a type 1 might also take action 3 in equilibrium but it is not the case for this data and ignoring that possibility (after a preliminary determination that no mis-specification results) makes the algorithm significantly more efficient.

Checking for Equilibrium Support of the Action Paths

The central task of the algorithm is to ascertain, for each set of histories observed in a session, the beliefs and assignment of types that support that set of histories in equilibrium and to concurrently obtain estimates of θ and p .

To summarize this second part of the algorithm: For each p in a grid, for each session,

- 1) construct all possible type-strategy combinations,
- 2) compute a distribution over actions from p , the empirical distribution of actions taken by types 3, 4, 5, and 6, and the proposed strategies of any types 1 or 2 in this possible combination,
- 3) check that a type 1 and 2 players have no incentive to deviate from the proposed equilibrium strategy,
- 4) put surviving type-strategy combinations in a matrix which keeps all possible type-strategy combinations for each session, for one p .

For each p , after all passing all sessions through procedures 1 through 4,

- 5) find all permutations of type-strategy combinations across sessions,
- 6) find $\hat{\theta}|p$ and evaluate the conditional log likelihood for each of the type-strategy combinations,
- 7) select the type combination that maximizes the log likelihood, conditional on the current p .

Parts 5 through 7 are primarily an equilibrium selection mechanism. The question of how to assign types when there are multiple equilibria is an important question. The approach taken here — that of selecting the one which implies the maximizing type combination — is only one that could be applied.

Run procedures 1 through 7 for each p in a grid. For each p , the log likelihood is also evaluated at $\theta = p$ for testing the hypothesis that the

players' beliefs are correct (in the sense that they are the same as $\hat{\theta}$).

The grid search over p evaluates the likelihood function at 64 million (?) points. Conditional on p , the standard errors for $\hat{\theta}$ are obtained analytically. True standard errors for $\hat{\theta}$ and for $\hat{p}$ are not computed or derived. Instead, approximate confidence regions for the maximum likelihood estimates of θ and p are numerically constructed. The sample likelihood function is searched in each direction for the point at which the value is significantly different from the maximizing value. (This process also yields the approximate shape of the flat which maximizes p and θ .)

Details of Procedures 1 through 7.

Detail 1.

This procedure constructs all permissible type-strategy combinations and also calls the procedures which implement parts 2 through 4. Permissible type-strategy combinations are those that are allowed by the theory. For example, a permissible combination is: 2 type 1 players who play all one's and 5 type 2 players who play all two's while an impermissible combination is: 2 type 1 players who play all one's and 5 other type 1 players who play all two's. The procedure loops through all possible type 2 strategies looking for a player who matches. If it finds a player match it looks for all who match — those are type 2. Within that loop it then loops through all possible type 1 strategies looking for a player match. When it finds a type 1 match, all players are given their type assignment: type 2 with the current loop strategy or type 1 with the current loop strategy or the default type assignment from the preliminary analysis. The procedure also checks the “null” strategy for each type — that is, it checks for type-strategy combinations in which there are no type 1's and for those in which there are no type 2's. Each possible type-strategy combination is then passed through parts 2 through 5.

Detail 2.

This procedure computes a distribution over actions, q , from the beliefs over types, p . The beliefs over actions induced by the beliefs over types depends on the session so $q(s)$ is computed for each session s . The procedure takes one proposed assignment of types to players in a session and the beliefs p along with some estimate of the distribution of actions taken by types 3, 4, 5, and 6.

These estimates can be obtained in two different ways: one method is to use the empirical distribution of actions from this data and the other

is to simulate the learning behavior against the real data and use that empirical distribution. The problem with the former approach is that there are proposed type combinations in sessions for which there may be no type 3 players, for example, while p_3 is strictly positive. Then there is no way to estimate the distribution of actions that the type 1 players expect to observe when their opponent is a type 3. An alternative way of using empirical distributions from the data is to just compute them using the assignment to types 3, 4, 5, and 6 from the preliminary analysis. Then the problem is that, for some sessions, one is using an estimate based on a subject who is now in the proposed type combination as a type 1 instead of as its preliminary type, 3. Using simulations avoids these complications. It is straightforward to apply in the case of type 3 players but involves some arbitrary elements when simulating types 4 and 5. In particular, with type 4, one has to specify when the player will decide to respond to the accumulating evidence and with type 5, one has to specify when the player will decide to experiment (given that it is allowed). Both are determined by a random process in which the frequency is matched to the apparent aggregate frequency in the data. The second method was used for the results reported in this paper.

Given estimates of the distribution of actions taken by each of types 3, 4, 5, and 6, denoted by $e(s, m)$, compute $q(s)$ by the following rule:

$$q(s, a) = 0$$

$$q(s, a) = q(s, a) + e(s, m) \cdot p_m$$

for $m = 3 \dots 6$ and for all actions a . Then, if type 1 players are playing all one's in this proposed type-strategy combination:

$$q(s, 1) = q(s, 1) + p_1$$

or if they are playing all two's in this proposed combination:

$$q(s, 2) = q(s, 2) + p_1.$$

Similarly for type 2 players: if they are playing all three's in this proposed type-strategy combination:

$$q(s, 3) = q(s, 3) + p_3.$$

If a mixed strategy is used, then use the distribution of actions observed by the relevant type, assume the strategy is stationary, and apply the weighted probability. For example, suppose type 1 will play a mix of actions 1 and 2 in this proposed strategy. Count actions 1 and actions 2 taken by the type 1 players in this possible assignment — call them c_1 and c_2 . Then

$$q(s, 1) = q(s, 1) + (c_1/(c_1 + c_2)) \cdot p_1$$

$$q(s, 2) = q(s, 2) + (c_2/(c_1 + c_2)) \cdot p_1$$

Detail 3.

This procedure checks that no player will want to deviate from the proposed strategy by verifying that a set of inequalities are satisfied. The procedure takes the q 's computed as above and the payoffs for types 1 and 2. For each possible type combination, all except types 1 and 2 are ignored, and the types 1 and 2 are checked against the appropriate inequalities by use of nested case statements. Processing of players continues as long as the inequalities are satisfied — if one is not satisfied, then that type-strategy combination fails.

Detail 4.

This procedure writes passing type-strategy combinations to a matrix that holds all passing combinations for each session, for one p . It also collects summary information about the data.

Detail 5.

This is a combinatorics routine. Within the routine, for each cross-session combination, it implements 6 and keeps a running tab on the max-

imum value of the log likelihood function. For this data the number of type-strategy combinations per session typically ranged from one to three with the number of cross-session combinations averaging in the mid-fifties. Therefore there is no problem with having an inordinately large number of combinations to evaluate for each p . However, this point in the algorithm seems the most likely for hitting computational constraints — especially as the number of sessions increase.

Detail 6.

Find $\hat{\theta}$ conditional on p by computing sample frequencies. The conditional log likelihood is the log of the multinomial evaluated at $\hat{\theta}$.

Detail 7.

Select the type combination that maximizes the conditional log likelihood. Begin with a comparison value set to a *very negative* number. Then compare the new log likelihood value with the comparison value — keep the larger of the two as the new comparison value, along with the associated $\hat{\theta}$ and the log likelihood function evaluated at p .

References

- Aho, Alfred V., John E. Hopcroft, and Jeffrey D. Ullman, *Data Structures and Algorithms*, Addison-Wesley Publishing Company, 1983.
- Aumann, Robert J., "Correlated Equilibrium as an Expression of Bayesian Rationality," *Econometrica* Vol. 55, No. 1, 1-18, January 1987.
- Binmore, Ken and Adam Brandenburger, "Common Knowledge and Game Theory," London School of Economics and University of Michigan, Churchill College and Harvard Business School, 1988.
- Blume, Larry, Stephen Salant, and Debra Holt, "Coping with Rejection: The Interpretation of Experimental Evidence from Multistage Games," University of Michigan, 1987.
- Brandenburger, Adam and Eddie Dekel, "Rationalizability and Correlated Equilibria," *Econometrica*, Vol. 55, No. 6, November 1987.
- Brown, James N., and Robert W. Rosenthal, "Testing the Minimax Hypothesis: A Re-examination of O'Neill's Game Experiment," *Econometrica* Vol 58, No. 5, September 1990.
- Camerer, Colin and Keith Wiegelt, "Experimental Tests of a Sequential Equilibrium Reputation Model," *Econometrica* Vol. 56, No. 1, January 1988.
- Cooper, Russell, Douglas V. DeJong, Robert Forsythe and Thomas W Ross, "Selection Criteria in Coordination Games: Some Experimental Results," University of Iowa Working Paper, 1989.
- Cooper, Russell, Douglas V. DeJong, Robert Forsythe and Thomas W. Ross, "Selection Criteria in Coordination Games: Some Experimental Results," *American Economic Review* Vol. 80, No. 1, March 1990.
- Crawford, Vincent P. and Hans Haller, "Learning How to Cooperate: Optimal Play in Repeated Coordination Games," *Econometrica*, Vol 58, May 1990.
- Fudenberg, Drew and David M. Kreps, "A Theory of Learning, Experimentation, and Equilibrium in Games," 1988.

- Gul, Faruk, "Rational Strategic Behavior and the Notion of Equilibrium", Stanford University Working Paper, September 1989.
- Hansen, Lars Peter, "Large Sample Properties of Generalized Method of Moments Estimators," *Econometrica* Vol. 50, No. 4, July 1982.
- Harsanyi, J.C., "Games with Incomplete Information Played by 'Bayesian' Players," *Management Science* 14: 159-189, 320-334, 486-502, 1967-68.
- Holt, Debra J., *Understanding Strategic Choice: The Statistical Analysis of Experimental Games*, University of Michigan, 1989.
- Holt, Debra J., "Detecting Learning in Experimental Games," Queen's University, November 1992.
- Jordan, J.S., "The Exponential Convergence of Bayesian Learning in Normal Form Games," *Games and Economic Behavior*, April 1992.
- Kalai, Ehud and Ehud Lehrer, "Rational Learning Leads to Nash Equilibrium," Northwestern University, Discussion Paper No. 858, March 1990.
- Kelly, Kevin T., and Clark Glymour, "Inductive Inference from Theory Laden Data", Report No. CMU-LCL-89-5, Laboratory for Computational Linguistics, Carnegie Mellon University, October 1989.
- Kreps, David, Paul Milgrom, John Roberts, and Robert Wilson, "Rational Cooperation in the Finitely Repeated Prisoners' Dilemma," *Journal of Economic Theory* 27, 245-252, August 1982.
- Ledyard, John O., "The Scope of the Hypothesis of Bayesian Equilibrium," *Journal of Economic Theory* 39, 59-82, June 1986.
- Lipman, Barton L., "How to Decide How to Decide How to ...: Modeling Limited Rationality," *Econometrica*, Vol 59, July 1991.
- McKelvey, Richard D., and Thomas Palfrey, "An Experimental Study of the Centipede Game," *Econometrica* 60, July 1992.
- Mertens, J.-F, and S. Zamir, "Formalization of Bayesian Analysis for Games with Incomplete Information," *International Journal of Game The-*

ory, 14, 1985.

Milgrom, Paul and John Roberts, "Rationalizability, Learning, and Equilibrium in Games with Strategic Complementarities," *Econometrica* 58, November 1990.

Miller, Robert A. "Job Matching and Occupational Choice", *Journal of Political Economy* 92, December 1984.

Myerson, Roger B., "Bayesian Equilibrium and Incentive-Compatibility: An Introduction," in *Social Goals and Social Organization*, 1985.

Ordeshook, Peter C., and Thomas Palfrey, "Agendas, Strategic Voting, and Signaling with Incomplete Information," California Institute of Technology, 1986.

Pakes, Ariel, "Patents as Options: Some Estimates of the Value of Holding European Patent Stocks," *Econometrica* 54, July 1986.

Pakes, Ariel and David Pollard, "Simulation and the Asymptotics of Optimization Estimators," *Econometrica*, Vol 57, September 1989.

Palfrey, Thomas R., and Howard Rosenthal, "Private Incentives in Social Dilemmas," *Journal of Public Economics* 35, 1988.

Press, William H., Brian P. Flannery, Saul A. Teukolsky, and William T. Vetterling, *Numerical Recipes: The Art of Scientific Computing*, Cambridge University Press.

Roth, Alvin E. and M.K. Malouf, "Game Theoretic Models and the Role of Information in Bargaining," *Psychological Review* 86, 1979.

Roth, Alvin E. and J. Keith Murnighan, "The Role of Information in Bargaining: An Experimental Study," *Econometrica* Vol. 50, No. 5, September 1982.

Roth, Alvin E., ed., *Laboratory Experimentation in Economics: Six Points of View*, Cambridge University Press, 1987.

Rosenthal, R., "Games of Perfect Information, Predatory Pricing, and the Chain-Store Paradox," *Journal of Economic Theory* 25, 1981.

Rubinstein, Ariel, "The Electronic Mail Game: Strategic Behavior Un-

- der 'Almost Common Knowledge'," *The American Economic Review*, June 1989.
- Rust, John, "Optimal Replacement of GMC Bus Engines: An Empirical Model of Harold Zurcher," *Econometrica* 55, September, 1987.
- Salant, Stephen and Eban Goodstein, "Committee Voting Under Alternative Procedures and Preferences: An Experimental Analysis," *Rand Journal of Economics*, forthcoming.
- Samuelson, Larry, "A Note on Uncertainty and Cooperation in a Finitely Repeated Prisoner's Dilemma," *International Journal of Game Theory* Vol. 16, Issue 3, 1987.
- Selten, R., "Reexamination of the Perfectness Concept for Equilibrium Points in Extensive Games," *International Journal of Game Theory* Vol. 4, Issue 1, 1975.
- Wolpin, Kenneth I., "An Estimable Dynamic Stochastic Model of Fertility and Child Mortality," *Journal of Political Economy* 92, 1984.