DIVISION OF THE HUMANITIES AND SOCIAL SCIENCES

CALIFORNIA INSTITUTE OF TECHNOLOGY

PASADENA, CALIFORNIA 91125

VERTIGO: COMPARING STRUCTURAL MODELS OF IMPERFECT BEHAVIOR IN EXPERIMENTAL GAMES

Mahmoud A. El-Gamal

Thomas R. Palfrey

SOCIAL SCIENCE WORKING PAPER 800

July 1992

Vertigo:
Comparing Structural Models of Imperfect Behavior in Experimental Games

Mahmoud A. El-Gamal and Thomas R. Palfrey *
Division of the Humanities and Social Sciences, Caltech, Pasadena, CA 91125

July 1992

Abstract

This paper investigates learning in games with one-sided incomplete information using laboratory data from a game which we call the game of Vertigo. The predicted Bayes-Nash equilibrium behavior of the agents in this type of game generates overly strong restrictions on the data, including the *zero likelihood problem*: certain actions should never be observed. To circumvent statistical problems, and to allow for deviations from perfectly rational behavior, we introduce the possibility of players making errors when choosing their actions. We compare two competing models depending on whether players take the errors in actions into consideration when formulating their strategies. We also investigate possible deviations from Bayes's rule, producing too fast or too slow an updating rule. In total, we get six models of sophisticated and unsophisticated strategy formation on the first dimension, and fast, slow, or no updating on the second. We apply a fully Bayesian structural econometric approach to compare the statistical performance of these six models, and to obtain posterior estimates of several nuisance parameters governing the errors in actions. The two models where players are unsophisticated and either use no updating at all, or use dampened updating, have a much higher likelihood than any of the others.

*We acknowledge the financial support from NSF grant #SES9011828 to the California Institute of Technology and from the JPL-Caltech Supercomputing project. We have benefited from many discussions with Richard McKelvey. Marty Hahm wrote the computer programs for the experiment.

VERTIGO:
COMPARING STRUCTURAL MODELS OF IMPERFECT BEHAVIOR IN EXPERIMENTAL GAMES

MAHMOUD A. EL-GAMAL AND THOMAS R. PALFREY

1 INTRODUCTION

This paper investigates learning in multistage games with one-sided incomplete information using laboratory data. We focus on two issues related to possible deviations from behavior predicted by game theoretic models focused on perfectly rational behavior by Bayesian players. The first issue is imperfect choice behavior, and employs a model that generates an error structure that permits rigorous statistical analysis of the data. The second issue is imperfect learning behavior; agents' updating may be too fast or too slow in the existence of errors in actions.

Recently, it has become increasingly clear that models of game behavior demanding perfect rationality of the players seem to perform poorly under the scrutiny of careful statistical analysis of data from simple finite game experiments. This is true whether the games are strictly competitive (Brown and Rosenthal (1990)), have clear gains to cooperative behavior (Camerer and Weigelt (1988), McKelvey and Palfrey (1992)), or lie in some intermediate territory Banks et al. (1988), Brandts and Holt (1989), and Palfrey and Rosenthal (1991), Palfrey and Rosenthal (1992)). One reason for the existence of these problems is that most equilibrium concepts are essentially static and fail to address basic issues of learning dynamics, information, and the cognitive limitations of the players. A second, related, reason is that these equilibrium concepts typically place overly strong restrictions on the data. In fact, for most of the games that have been experimentally investigated, the equilibrium predicts the use of pure strategies. More to the point, equilibrium models frequently predict that certain strategies will *never* be observed. In these cases, both classical and Bayesian statistical methods lead to outright rejection of strictly rational models.

The predictions that certain strategies should never be observed may arise from a number of different scenarios ranging from the availability of dominant actions under certain circumstances, to very sophisticated updating rules resulting in full revelation of certain states of nature, and requiring particular pure strategies to be played (e.g. Cooper et al. (1991), and El-Gamal et al. (1991)). The end result is that certain actions cannot be observed at certain points in time during the experiment. If those "infinitely unlikely" events are observed in experimental data (e.g. players choosing a dominated action in Cooper et al. (1991)), one is faced with what we call *the zero likelihood problem.*

A common, but statistically unsatisfactory way to deal with this problem is to ignore the zero-likelihood data points, treating them as missing data. Other popular procedures involve descriptive data analyses and the application of ad hoc measures of association or fit to compare the relative performance of several competing models.

Several recent attempts have been made to endow the strict rationality models with an extra feature that allows all observations to occur with positive probability. This paper follows this approach, and evaluates and compares several different specific ways to do this in the context of a multistage game of incomplete information. The players are assumed to follow the theory as long as they do not "tremble"; and they tremble with some probability $0 < \epsilon < 1$. When players tremble, they are assumed to choose their action in an arbitrary way which allows all possible actions to be chosen with some positive probability.[1] The first question explored along this approach was what value of ϵ to use. In Boylan and El-Gamal (1992)'s study of disequilibrium dynamics, the model comparison was done at all possible ϵ's in $[0, 1]$, and in their case, the results held for almost all ϵ (excluding those values arbitrarily close to 1). In McKelvey and Palfrey (1992), and Harless and Camerer (1992), a value of ϵ was estimated by its maximum likelihood estimator under the relevant model(s). In El-Gamal et al. (1991), the parameter ϵ was viewed as a nuisance parameter, and hence, a prior was induced on ϵ together with all other nuisance parameters of the models they studied, and this prior on ϵ was updated as experimental data was collected. In this paper, we follow this fully Bayesian procedure, by specifying a model with nuisance parameters, integrating out these nuisance parameters with respect to our beliefs, which are updated as more data are observed.

The second issue that the introduction of the trembles poses is how players take it into account. In El-Gamal et al. (1991), and McKelvey and Palfrey (1992), the trembles were assumed to come from a common knowledge distribution which the players incorporated in making their decisions. In this sense, the players were assumed to be very sophisticated in dealing with the fact that their opponents, as well as themselves, may tremble at any point in the game. They incorporated these probabilities of trembles in their strategies (mappings from beliefs to actions), as well as updating rules (mappings from beliefs and observed actions to beliefs). This very high degree of sophistication would seem to put an incredible computational burden on the players. In order to solve for the full sequential equilibrium of a three-move centipede game played with two opponents, El-Gamal et al. (1991) required 120 CPU hours of a fast supercomputer (a Cray Y-MP2E/116). It might seem odd to deviate from perfect rationality by introducing the trembles, and then require such a high level of rationality in dealing with the trembles to make it almost impossible to solve for the equilibrium behavior. Lower degrees of sophistication in dealing with the trembles may be more reasonable.

This paper considers two levels of sophistication in the players' strategies. The *so-*

[1]There are alternative ways to introduce imperfect choice behavior. Besides this "trembling" model, we are stochastic utility models (Logit, Probit, etc.), imperfectly controlled preference models (Palfrey and Rosenthal (1992)), evolutionary models with mutation, and imperfect equilibrium models (Beja (1992)). This paper focuses only on a simple revision of the tremble model.

phisticated model lets the agents take into account the fact that they, as well as their opponents, can tremble at any time, and adjust for that fact when deciding on their strategies. The *unsophisticated model* does not have the agents take account of the tremble probabilities when formulating their best responses, i.e., they decide on their actions as if no one ever trembles.

A second source of deviation from behavior predicted by Bayesian equilibrium, besides errors in actions, derives from errors in updating. The players may, to varying degrees, use available information in ways that are inconsistent with Bayes's rule. We consider three alternative updating models. The first model (which we label the undampened or *fast updating* model) lets the agents update from the observed moves of their opponents via the Bayes updating map, not taking into account the fact that their opponents may have trembled. The second model (which we label the *dampened updating* model) lets the agents update from the observed moves taking into account the fact that they may have resulted from a tremble. The last model of updating we consider is the *no updating* model, where the players do not use the observed actions to adjust their beliefs. Combining the two models of sophistication with the three models of updating generates a total of six alternative models.

We compare these models using a simple game of one-sided incomplete information, which we call the game of Vertigo. The informed player may be one of two-types, corresponding two possible payoff tables of a bimatrix game, called game 1 and game 2. The games are equivalent, up to a relabeling of the players, so one may think of the private information being that the informed player knows whether he is the row player or the column player, but the uninformed player does not.[2] The probability distribution of the informed player's type is common knowledge. No player has a dominant strategy in either game 1 or game 2. There is a unique mixed-strategy equilibrium in each component game.

The game of incomplete information also has a unique equilibrium, which depends on the common knowledge prior over the types. One type mixes, and the other type adopts a pure strategy. Which type mixes depends on the prior.

In the experiment, each uninformed player plays a sequence of informed players, all of whom are the same type. The experimental procedures are organized in a manner that eliminates signaling possibilities by the informed (row) players. The next section will analyze the equilibrium of the game and describe the details of the experimental design and procedures. Section 3 lays out the six models embodying varying degrees of sophistication on the part of the players when accounting for the tremble probabilities in their strategies and updating rules. Section 4 presents the experimental results and our methods of data analysis, and discusses our ranking of the models that we consider. Section 5 concludes the paper.

[2]In the experiment, the two players know their labels, and the payoff entries for game two are the same as the payoff entries to game one, except they have been rotated counterclockwise by one cell in the table. Hence the name "Vertigo." This is equivalent to relabeling the players.

2 The Game of Vertigo:

2.1 The Stage Game:

The Game of Vertigo is defined by the following two tables, called game 1 and game 2:

Game I prob= π:

	L	R
U	$(0, a_2)$	$(a_2, 0)$
D	$(a_1, 0)$	$(0, a_1)$

Game II prob= $(1 - \pi)$:

	L	R
U	$(a_2, 0)$	$(0, a_1)$
D	$(0, a_2)$	$(a_1, 0)$

In words, the players get to play Game I with probability π, and Game II with probability $(1 - \pi)$. The values of π and relevant payoffs a_1 and a_2 are chosen by the experimenter.

This class of games was chosen so that each game has the typical structure of a two person game with an unique mixed strategy equilibrium. The equilibrium of each game also has a symmetry property: each player assigns mixing probabilities to their two strategies in ratios of $p : 1 - p$ where $p = \frac{a_1}{a_1 + a_2}$. These features of symmetry, uniqueness, and lack of dominant strategies make the game one which is convenient from an analytical and statistical point of view, and which is sufficiently complex to warrant the introduction of trembles into the individual's choices. The game was also sufficiently abstract so that we hoped to avoid having our control of subjects' financial motivations contaminated by preconceived notions of fairness or altruism, as in prisoner's dilemma or public goods experiments. We also wished to avoid a structure of payoffs that led to multiple equilibria, as this would generate serious difficulties in the statistical comparison of the competing models.

2.2 The experimental design:

We study a $2n$-person ($n \geq 5$) five-stage repeated game based on the vertigo stage game defined above. There are n row players and n column players, who are seated at computer terminals which are separated by partitions. One payoff game is drawn according to the probability π before the beginning of stage 1. All row players in the room are informed of the outcome of this random draw. In stages 1-5 each participant in the experiment plays the stage game against a sequence of five different opponents, always using the same payoff game that was drawn initially.

The matching sequence is constructed using the technique of McKelvey and Palfrey (1992), which guarantees within the five stage game that it is logically impossible for a player's current move to have an effect on how that player's future opponents make choices in the later stages of the game. This is done by first labeling the players row player 1, row player 2,..., row player n; column player 1, column player 2,..., column player n. In stage i, row player i is matched with column player i, and in stage j, row player i is matched with column player $(i + j - 1)$ (mod n). Notice that, to avoid repetition, at least ten players are needed. This matching sequence is announced to all subjects, and they are informed that the labels (e.g. row player i) are assigned randomly so that they cannot identify any of the players by their labels.

The information structure is controlled in a special way. In each stage players make their choices simultaneously. After everyone has moved, the row and column players are told what move their opponents at that stage chose. However, only the row players are told their payoff for that stage (they can infer it anyway from their knowledge of the drawn game). Since the column players were initially uninformed, and do not see their payoff they cannot infer directly which game was played. Play then proceeds to the next stage. After being rematched, each row player is informed of the history of the moves of other row players when they had previously played his current column player opponent. Thus for each pair, in each stage, the column player's past observations of row player moves is common knowledge to both players. The row player is ***not*** informed of the current column player's past moves, nor is the column player informed of the current row player's past moves. This information structure permits the row player to (in principle) infer the current column player opponent's updated belief of whether the actual game being played uses payoff game I or II, under the maintained hypothesis that earlier games followed Bayesian equilibrium play.

After the last stage (5) game is over, the payoff table is revealed to the column players and all subjects record their earnings. Subjects are then randomly reassigned player number labels (but row players remain row players, and column players remain column players). A second 5-stage game is then played in exactly the same manner as the previous one. A new payoff table is drawn according to the probability π, independently from the past draws. The matching sequence is done in the same manner (but anonymously, since labels have been randomly shuffled), and so on. Over the course of each experimental session, we conduct a total of ten 5-stage games in exactly this manner. After the tenth 5-stage game, subjects are paid in private, one at a time, in a separate room.

All of the above information is publicly announced to the subjects, by reading the instructions aloud at the beginning of the session. Following the instruction periods, subjects were led step-by-step through a series of exercises to assure that they understood the information structure, the matching rule, how they would be paid, and the keyboard and record-keeping tasks. See Appendix A for all the details. The sessions were conducted on a network of computer terminals at the Caltech Laboratory of Experimental Economics and Political Science, using Caltech undergraduate students. No subject participated in more than one session. A total of two sessions were conducted using twenty six different subjects.

2.3 Equilibrium Strategies and Beliefs:

The Bayesian equilibrium of the 5-stage repeated game can be calculated by separately computing the equilibrium of an arbitrary stage s, depending on the prior π_s of a column player at stage s. The reason each stage can be treated independently (after accounting for the updating of the beliefs π_s) is that our rotation and information structure isolates the strategic calculations of each stage. There is no role for signalling equilibria and reputational phenomena. Thus, for stage s, we index the strategies by the prior of the column player, say π_s that game I is currently being played. Denote:
$q(\pi_s)$= column player's mixed strategy [prob(L)].
$p1(\pi_s)$= row player's mixed strategy [prob(U)] if game I.
$p2(\pi_s)$= row player's mixed strategy [prob(U)] if game II.
Finally, let $\pi_{s+1} = \mathfrak{B}(\pi_s, a_s)$ define the Bayes updating rule for a column player who started period s with prior s, and observed the matched row opponent make move a_s. Of course, $\pi_1 = \pi$, the common knowledge probability announced in the instructions.

Thus to compute the equilibrium at stage s if the current prior on game I is π_s, we simply compute the (unique) one-shot Bayesian equilibrium given π_s, and then at stage $(s+1)$, we perform a similar derivation, except that we use the posterior $\pi_{s+1} = \mathfrak{B}(\pi_s, D)$ or $\pi_{s+1} = \mathfrak{B}(\pi_s, U)$, depending on the row player's realized move at stage s.

Notice that in the laboratory sessions, several games are played simultaneously, so (except for stage 1) different pairs of players matched in stage s may have different common knowledge prior of the column player that they are playing game I, depending on a column player's observations of moves by earlier row players. Recall that the design calls for row players to be told their current column opponent's past observations of moves by other row players, which (in theory) preserves common knowledge of the current belief.

The actual derivation of equilibrium proceeds by establishing three facts, each of which may be verified by the reader. First denote $b = a1/(a1 + a2)$, and assume that $b \neq \pi$.[3] The three facts are:

1. There is no pure strategy equilibrium.

2. For all values of π, the column player is mixing.

3. For all values of π, exactly one row type mixes.

From these facts, one can easily compute the unique equilibrium of the game, and compute the updating rule, $\pi_{s+1} = \mathfrak{B}(\pi_s, a_s)$. This is summarized below.

[3] If $b = \pi$, the equilibrium is indeterminate.

$$\text{if } \pi_s \in [0,b) \Longrightarrow \begin{cases} q(\pi_s) = b \\ p_1(\pi_s) = 1 \\ p_2(\pi_s) = \frac{1-b}{1-\pi} \\ \pi_{s+1} = \mathfrak{B}(\pi_s, a_s) = \frac{\pi}{1-b+\pi} & \text{if } a_s = U \\ \pi_{s+1} = \mathfrak{B}(\pi_s, a_s) = 0 & \text{if } a_s = D \end{cases}$$

$$\text{if } \pi_s \in (b,1] \Longrightarrow \begin{cases} q(\pi_s) = 1-b \\ p_1(\pi_s) = \frac{b}{\pi} \\ p_2(\pi_s) = 1 \\ \pi_{s+1} = \mathfrak{B}(\pi_s, a_s) = \frac{b}{1+b-\pi} & \text{if } a_s = U \\ \pi_{s+1} = \mathfrak{B}(\pi_s, a_s) = 1 & \text{if } a_s = D \end{cases}$$

There are several noteworthy features of the equilibrium and the updating rule. First, the uninformed player acts "as if" he were informed. That is, if π_s is sufficiently large ($\pi_s > b$), his strategy is exactly the same as if $\pi = 1$ was common knowledge. If π is sufficiently low, the reverse is true. Second, only hybrid equilibria (one type mixes, the other doesn't) can arise. This has two interesting implications. First, it implies a *zero likelihood problem*: some moves are predicted to never be played by some types of row players. Second, beliefs are updated sharply following some moves, since the mixing type is using a strategy that the non-mixing type never uses. When π_s is low, a choice of D reveals that it is game II, and when π_s is high, a choice of D reveals that it is game I. Both of these implications point to the overly strong restrictions implied by Bayesian equilibrium. The first has to do with overly strong restrictions on actions, while the second has to do with overly strong restrictions on beliefs. What we propose below are some alternative models for coming to grips with these problems.

3 Modeling Imperfect Behavior

3.1 Errors in Actions:

The maintained model is that at each time period, an individual may tremble and randomly choose a move from the menu of moves available to him/her. We allow this tremble rate to change over time. Technically, we introduce the sequence ϵ_t = probability of tremble = $\epsilon_0 e^{\alpha t}$. The changing probability of trembling implied by the parameter α is intended to capture the learning by doing component of the agents' behavior. If $\alpha > 0$, then error rates are declining over time. In what follows, we suppress dependence on t, and write the strategies and updating rules as a function of ϵ. It is convenient to define:

$$E = \frac{\epsilon(1-2b)}{2-2\epsilon}$$

Model S, the sophisticated model:

The first model we introduce is the "perfectly rational" extension of the previous model where the agents take into consideration that both their opponents and they themselves may tremble at any move with probability ϵ, and incorporate that into their strategies. Straightforward analysis yields the following equilibrium strategies.

$$\text{if } \epsilon > 2b \Longrightarrow \begin{cases} q(\pi_s) = 1 \\ p1(\pi_s) = 1 \\ p2(\pi_s) = 1 \end{cases}$$

$$\text{if } \epsilon < 2b \text{ and } \pi - s < b - E \Longrightarrow \begin{cases} q(\pi_s) = b - E \\ p1(\pi_s) = 1 \\ p2(\pi_s) = \frac{1+E-b}{1-\pi_s} \end{cases}$$

$$\text{if } \epsilon < 2b \text{ and } \pi - s > b - E \Longrightarrow \begin{cases} q(\pi_s) = 1 - b + E \\ p1(\pi_s) = \frac{b-E}{\pi_s} \\ p2(\pi_s) = 1 \end{cases}$$

Notice that if ϵ is sufficiently large relative to b, there is a pure strategy equilibrium at (U, L) which does not permit the uninformed players to update.

Model U: The unsophisticated model:

This model does not allow the players to take the ϵ into account when constructing their strategies. the resulting strategies are:

$$\text{if } \pi_s \in [0, b) \Longrightarrow \begin{cases} q(\pi_s) = b \\ p_1(\pi_s) = 1 \\ p_2(\pi_s) = \frac{1-b}{1-\pi} \end{cases}$$

$$\text{if } \pi_s \in (b, 1] \Longrightarrow \begin{cases} q(\pi_s) = 1 - b \\ p_1(\pi_s) = \frac{b}{\pi} \\ p_2(\pi_s) = 1 \end{cases}$$

The unsophisticated model has the same equilibrium as $\epsilon = 0$, since errors are not taken into account when formulating strategies. Nevertheless, different patters of play are predicted under the unsophisticated model if $\epsilon > 0$, since *observed* mixing probabilities will be closer to .5.

3.2 Errors in Beliefs

Model F, fast (undampened) updating:

This model uses the Bayes updating map to update from observed row player moves under the assumption that trembles cannot occur. For models SF (i.e. sophisticated strategies, fast updating) and UF (i.e. unsophisticated strategies, fast updating) the updating rule is:

$$\text{if } \pi_s \in [0,b] \Longrightarrow \begin{cases} \pi_{s+1} = \mathfrak{B}(\pi_s, a_s) = \frac{\pi}{1-b+\pi} & \text{if } a_s\text{=U} \\ \pi_{s+1} = \mathfrak{B}(\pi_s, a_s) = 0 & \text{if } a_s\text{=D} \end{cases}$$

$$\text{if } \pi_s \in [b,1] \Longrightarrow \begin{cases} \pi_{s+1} = \mathfrak{B}(\pi_s, a_s) = \frac{b}{1+b-\pi} & \text{if } a_s\text{=U} \\ \pi_{s+1} = \mathfrak{B}(\pi_s, a_s) = 1 & \text{if } a_s\text{=D} \end{cases}$$

Notice that model F is only correct when $\epsilon = 0$.

MODEL D, DAMPENED UPDATING:

This model uses the Bayes updating map taking into consideration the fact that the observed row player moves may have resulted from trembles. This updating rule will be different for models S and U. For model SD (i.e. sophisticated strategies, dampened updating), the updating rule is:

$$\text{if } \epsilon > 2b \Longrightarrow \{ \; \pi_{s+1} = \mathfrak{B}(\pi_s, a_s) = \pi_s$$

$$\text{if } \epsilon \leq 2b \text{ and } \pi - s \leq b - E \Longrightarrow \begin{cases} \pi_{s+1} = \mathfrak{B}(\pi_s, a_s) = \frac{\pi_s(1-\epsilon/2)}{(1-\epsilon)(1+\pi_s-b+E)+\epsilon/2} & \text{if } a_s\text{=U} \\ \pi_{s+1} = \mathfrak{B}(\pi_s, a_s) = \frac{\pi_s \epsilon/2}{(1-\epsilon)(b-\pi_s-E)+\epsilon/2} & \text{if } a_s\text{=D} \end{cases}$$

$$\text{if } \epsilon \leq 2b \text{ and } \pi - s > b - E \Longrightarrow \begin{cases} \pi_{s+1} = \mathfrak{B}(\pi_s, a_s) = \frac{\pi_s((1-\epsilon)(b-E)/\pi_s+\epsilon/2)}{(1-\epsilon)(1-\pi_s+b-E)+\epsilon/2} & \text{if } a_s\text{=U} \\ \pi_{s+1} = \mathfrak{B}(\pi_s, a_s) = \frac{\pi_s((1-\epsilon)(1-(b-E)/\pi_s)+\epsilon/2}{(1-\epsilon)(\pi_s+E-b)+\epsilon/2} & \text{if } a_s\text{=D} \end{cases}$$

For model UD (i.e. unsophisticated strategies, dampened updating), the updating rule is:

$$\text{if } \pi_s \in [0,b] \Longrightarrow \begin{cases} \pi_{s+1} = \mathfrak{B}(\pi_s, a_s) = \frac{\pi_s(1-\epsilon/2)}{(1-\epsilon)(1-b+\pi_s))+\epsilon/2} & \text{if } a_s\text{=U} \\ \pi_{s+1} = \mathfrak{B}(\pi_s, a_s) = \frac{\pi_s \epsilon/2}{(1-\epsilon)(b-\pi_s)+\epsilon/2} & \text{if } a_s\text{=D} \end{cases}$$

$$\text{if } \pi_s \in [b,1] \Longrightarrow \begin{cases} \pi_{s+1} = \mathfrak{B}(\pi_s, a_s) = \frac{b(1-\epsilon)+\pi_x \epsilon/2}{(1-\epsilon)(1+b-\pi_s))+\epsilon/2} & \text{if } a_s\text{=U} \\ \pi_{s+1} = \mathfrak{B}(\pi_s, a_s) = \frac{(\pi_s-b)(1-\epsilon)+\pi_s \epsilon/2}{(1-\epsilon)(\pi_s-b)+\epsilon/2} & \text{if } a_s\text{=D} \end{cases}$$

Notice that model D is "correct" updating for all values of ϵ.

MODEL N, NO UPDATING:

The final model of updating that we consider is the completely myopic one where the players do not learn from observations and maintain their initial belief $\pi_1 = \pi$ throughout the stages $s = 1, \ldots, 5$. The updating rule for the two models SN (i.e. sophisticated strategies, no updating), and UN (i.e. unsophisticated strategies, no updating) is simply:

$$\pi_{s+1} = \mathfrak{B}(\pi_s, U) = \mathfrak{B}(\pi_s, D) = \pi_s = \pi$$

Model N is correct only if $\epsilon = 1$.

4 Experimental results

4.1 Bayesian econometric method

Now, we have a total of six models which we have labeled SN, UN, SD, UD, SF, UF. The first letter (S or U) refers to whether the model assumes that agents incorporate the tremble probabilities in forming their strategies (denoted by S for sophisticated), or ignore them (denoted U for unsophisticated). The second letter (N, D, or F) refers to whether the model assumes that agents do no updating at all (denoted N for no updating), use Bayes updating but dampened by the tremble probabilities (denoted D for dampened updating), or use Bayes updating under the assumption that no trembles can occur (denoted F for fast updating).[4]

As in subsection 2.3 above, the equilibrium strategic behavior under all the models can be indexed by the beliefs of the column (uninformed players) at time s that game I is being played (which we label π_s, being updated over the five stages of the game using the appropriate updating rule N, D, or F). For each column player, we observe that player's five moves, and the moves of his 5 row opponent (a total of 10 data points) for each round. We ran two experimental sessions. The first session had sixteen subjects. Therefore, we had 8 column subjects playing 10 rounds each, with each round generating 10 data points. Hence, we had 800 data points from that first session. The second session had 10 subjects, i.e., 5 column subjects, and hence generated 500 data points. The first experimental session was conducted with $\pi = 1/6$, and payoffs $a_1 = \$.25$, and $a_2 = \$1.00$. The second experimental session was conducted with $\pi = 0.2$, and payoffs $a_1 = \$.50$, and $a_2 = \$.75$. The results shown later in the section show robustness to the values of π, a_1, a_2, resulting in almost identical results for both sessions separately. We, therefore, confine attention below to the aggregate analysis of the pooled data set. Notice that again by our discussion of subsection 2.3 above, once we condition on the belief π_s of each column player at each stage s (which is directly computable since we know $\pi_1 = \pi$, and the sequence of row-opponent moves that he observes), all 1300 data points that we have can be treated as i.i.d. draws[5] from the appropriate strategies $q(\pi_s)$ and either $p_1(\pi_s)$ or $p_2(\pi_s)$ depending on whether game I or II was being played.

For each model IJ (I=S,U; and J=N,D,F), denote the strategies of the players under that model by $q^{IJ}(\pi_s)$, $p_1^{IJ}(\pi_s)$, and $p_2^{IJ}(\pi_s)$, and denote the updating rule by $\mathfrak{B}^{IJ}(\pi_s, \alpha_s)$. Then, given data for n column players, we compute the likelihood under model IJ of the data generated in all $n \times 50$ games (n players times 5 stages times 10 rounds) by:

$$Like^{IJ}(\pi) = \int_0^1 \int_0^{0.3} \prod_{player=1}^{n} \prod_{r=1}^{10} \prod_{s=1}^{5} \left[(1 - \epsilon_0 e^{\alpha r}) Prob^{IJ}\{a_{rs}^{column}; \pi_{rs}\} + \epsilon_0 e^{\alpha r}/2 \right]$$

[4]The SD and SN models correspond to the "sequential" and "non-sequential" models of El-Gamal et al. (1991), respectively.

[5]This also embodies an implicit assumption of homogeneity of error rates (ϵ) and learning rates (α) across subjects. Allowing for heterogeneity was not computationally feasible.

$$\times \Big[(1-\epsilon_0 e^{\alpha r}) Prob^{IJ}\{a_{rs}^{row}\pi_{rs}\} + \epsilon_0 e^{\alpha r}/2\Big] prior(d\epsilon_0, d\alpha)$$

where r indexes the rounds 1,...,10; s indexes the stages 1,...,5; the tremble probability in round r is the same for all players and defined by the learning by doing model $\epsilon_r = \epsilon_0 e^{\alpha r}$; a_{rs}^{row} is (U or D) the appropriate row player's action in round r and stage s; and a_{rs}^{column} is (L or R) the appropriate row player's action in round r and stage s. π_{rs} is the appropriate column player's belief in stage s of round r, which is set at each stage 1 at $\pi_{r1} = \pi$, and then updated using $\pi_{r,s+1} = \mathfrak{B}^{IJ}(\pi_{rs}; a_{rs}^{row})$. All variables are of course indexed by the column players 1,...,n but that dependence has been suppressed in the above equation since no confusion can arise. The probabilities $Prob^{IJ}\{.\}$ are computed using the Bayes-Nash strategies under model IJ. For instance $Prob^{IJ}\{U; \pi_{rs}\} = p_1^{IJ}(\pi_s)$ if game I was played and $p_2^{IJ}(\pi_s)$ if game II was played; $Prob^{IJ}\{R; \pi_{rs}\} = 1 - q^{IJ}(\pi_s)$, and so on.

After evaluating the likelihoods of our six models $Like^{IJ}$, for $I = U, S$, and $J = N, D, F$, we can compare each of the models to the other five by computing the posterior odds ratio. The posterior odds for model IJ are simply $Like^{IJ}/\sum_{i=U,s;j=N,D,F} Like^{ij}$. This measures the relative likelihood of model IJ within the collection of models that we consider.

4.2 Results

The following table lists the likelihoods for all six models that we study using all 1300 data points.

Model Likelihoods

	N	D	F
U	8.8×10^{-355}	2×10^{-355}	1×10^{-358}
S	4.8×10^{-373}	4.1×10^{-373}	1.9×10^{-377}

The following table shows the posterior odds discussed in the previous subsection.

Model Odds

	N	D	F
U	0.815	0.185	9.3×10^{-5}
S	4.4×10^{-19}	3.7×10^{-19}	1.8×10^{-23}

It is clear that only unsophisticated models have credibility within the class of six models that we analyze. Moreover, model UN (which assumes the least sophistication in updating) seems to be about 4.4 times as likely as model UD, suggesting that models can be ranked according to their degree of sophistication. While UF significantly outperforms all of the sophisticated models, it performs poorly within the class of unsophisticated models. Indeed, a quick look at the two tables above shows that model can be monotonically ranked with less sophistication always being better than more (U uniformly beats S), and with slower or no updating always dominates fast updating (N uniformly beats D which in turn uniformly beats F).

Appendix B contains six figures depicting the joint posterior on (ϵ_0, α) under each of the models. Since we started with uniform priors on those nuisance parameters, those posteriors are simply the likelihood function at each (ϵ_0, α) under each of the models normalized to integrate to 1. It is clear from the pictures that all the models concentrate our posterior belief on ϵ_0, the initial rate of trembling, around 0.7, with the exception of model SF which has the posterior concentrated at a value of ϵ_0 above 0.8. The striking difference between the models in terms of posterior beliefs on the nuisance parameters is between the unsophisticated (U) models on the one hand, and the sophisticated (S) models on the other. All the sophisticated models concentrate our beliefs on α, the exponential rate of learning by doing, at zero. This means that the rate of trembling does not decline over time. This will typically happen when the strategic model fails to explain the behavior of the subjects, and it becomes easier to explain their behavior by trembles. For the unsophisticated models, our posterior belief on α seems to be concentrated around 0.1. This means that we believe that the subjects start the session trembling around 70% of the time, and end up trembling around 26% of the time by the 10^{th} round.

5 Concluding Remarks:

We introduced a simple repeated game with one-sided incomplete information to study models which allow for deviations for perfect rationality. We introduced the probability of players making errors in actions, and studied deviations from Bayes-Nash behavior in two dimensions. On the first dimension, we developed two models where players take or do not take into account the errors in action when formulating their optimal responses. On the second dimension, we developed three models where players use Bayes updating ignoring the errors in actions, use Bayes updating taking account of the errors in actions, and do not use any updating. The results from the experimental data show that on both dimensions, less sophisticated models perform better than the more sophisticated ones. This ranking is particularly strong along the dimension of optimal response formulation (where the computational cost of incorporating the errors in action is more pronounced), the unsophisticated model very impressively outperforms the sophisticated.

As usual with all experimental work, the results may be limited to the class of games that we looked at. It would be interesting to investigate if similar results obtain in different environments. In fact some of the findings here are probably true for a very wide class of experimental games. For example, the estimate of $\alpha > 0$, reflecting learning-by-doing, was also recovered in McKelvey and Palfrey (1992) and El-Gamal et al. (1991), and is supported widely in informal data analysis of most game experiments.

Proceeding cautiously, we hope the reader is left with two strong messages from our results. The first is methodological, and states that the manner in which imperfect behavior in experimental games is introduced is quite non-trivial, and can lead to different conclusions. Nevertheless, rigorous statistical analysis of the data is both feasible and worth doing. Clearly more investigation of alternative ways to statistically model these imperfections is an important job for future theoretical and experimental research. The second is theoretical, and states that it is unreasonable to assume too much rationality on the part of players. This is true not only because of the unreasonbly strong statistical restrictions implied by full rationality, but for substantive reasons as well. Further synthesis of experimental and modeling techniques to uncover useful and theoretically sound ways to incorporate limited rationality seems to be a useful direction to proceed.

Appendix A
Instructions for Experimental participants

This is an experiment in human decision making under uncertainty. As you entered the room, you were randomly assigned a seat. If your seat is on the middle isle of the room, you are an "A" participant and will be one for the duration of this experiment. If your seat is on one of the outside isles of the room, you are a "B" participant and will be one for the duration of this experiment. The instructions are slightly different for A participants and B participants. You should make sure that you understand the instructions for both A and B participants since your payoff will depend on your actions, and the actions of participants of the other type. The experiment will have 10 rounds, each round consisting of 5 periods.

FEEL FREE TO RAISE YOUR HAND AND ASK QUESTIONS AT ANY POINT. MAKE SURE THAT YOU UNDERSTAND THE INSTRUCTIONS BEFORE THE EXPERIMENT BEGINS. IF YOU HAVE ANY QUESTIONS DURING THE EXPERIMENT, FEEL FREE TO RAISE YOUR HAND AND ASK THE EXPERIMENTER WHO WILL BE PRESENT THROUGHOUT THE EXPERIMENT, AND WILL ANSWER YOUR QUESTION PRIVATELY.

AT NO POINT DURING THE INSTRUCTIONS PERIOD OR DURING THE EXPERIMENT ARE YOU ALLOWED TO COMMUNICATE IN ANY WAY WITH OTHER PARTICIPANTS.

In each period of each round, each A participant will be matched with a B participant. A participants will be asked to choose either left (L) or right (R). B participants will be asked to choose either up (U) or down (D). The two moves of the A and the B participants determine their payoffs for that period of that round according to one of the following two payoff tables: (show transparency 1)

Two out of ten chance:		PAYOFF TABLE I
	A participant chooses L	A participant chooses R
B participant chooses U	A gets \$0.75 B gets \$0.00	A gets \$0.00 B gets \$0.75
B participant chooses D	A gets \$0.00 B gets \$0.50	A gets \$0.50 B gets \$0.00

eight out of ten chance:		PAYOFF TABLE II
	A participant chooses L	A participant chooses R
B participant chooses U	A gets \$0.00 B gets \$0.75	A gets \$0.50 B gets \$0.00
B participant chooses D	A gets \$0.75 B gets \$0.00	A gets \$0.00 B gets \$0.50

For example, if the relevant payoff table turns out to be Table I, and if you are participant A3, matched with participant B2, and if you choose L, and B2 chooses D, then you get $0.00, and participant B2 gets $0.50. For a second example, if the relevant payoff table turns out to be Table II, and if you are participant B5, matched with participant A1, and if you choose U, and participant A1 chooses L, then you get $0.75, and A1 gets $0.00.

At the beginning of each round, one of the two payoff tables will be chosen at random by rolling a fair 10-sided die. If the outcome is 0 or 1, the first payoff table will be used for that round, otherwise, (if the outcome is 2,3,4,5,6,7,8, or 9) the second payoff table will be used for that round. All B participants will then be informed of the relevant payoff table for the round as soon as the die is rolled. The die will be rolled so one of the B participants can verify the outcome. The A participants will not be informed of the relevant payoff table until the end of the round.

At the beginning of each round, new identity numbers will be randomly assigned to both A and B participants. For the duration of that round, you will be referred to by your type (A or B), and your number. For instance, if you are an A participant, you may be A1, A2, ...etc. If you are a B participant, you may be B1, B2, ...etc. (Remember that your letter (A or B) remains the same all the time, but your number may be different from round to round.) Each A participant is assigned one white sheet, one red sheet, and 5 green slips. The white sheet contains the following table to be filled by each A participant (show transparency 2):

WHITE SHEET FOR *Name* ID:A—				
ROUND 1 PERIODS	MY MATCH	MY MOVE L or R	MATCH's U or D	PAYOFF in $
1				
2				
3				
4				
5				
TOTAL PAYOFF THIS ROUND IS:				

The red sheet contains the following table to be filled by the B participants with whom our A participant is matched at each period:

RED SHEET FOR A—		
PERIODS	My MATCH	B's Moves
1		
2		
3		
4		
5		

The green slips each contains one cell, and they are to be filled by the A participant.

GREEN SLIP FOR PARTICIPANT A—	
Period	MY MOVE(L or R)

To each B player is assigned a white sheet to be filled out by the B participants:

WHITE SHEET FOR *NAME* ID: B—				
ROUND 1 PERIODS	MY MATCH	MY MOVE U or D	MATCH's L or R	PAYOFF in $
1				
2				
3				
4				
5				
TOTAL PAYOFF THIS ROUND IS:				

Since participants are randomly assigned new identity numbers in each round, you will never know the participant with whom you are matched.

In each round, the following sequence of events occurs:

1. The round begins.

2. Participant identity numbers are randomly drawn. Everyone is informed of their own identity number but no one else's.

3. The relevant payoff matrix is randomly determined by rolling a 6-sided die, the B participants are informed of the relevant payoff table (it is the table with the highlighted background) for that round.

4. The first period begins. During that period:

(a) The A participants enter their move (L or R) at the flashing sign on their green slip (the green slip disappears from the A participant's).

(b) The relevant A participant's red sheet appears on each B participant's screen. The B participants enter their move (U or D) at the flashing sign on the red sheet of the relevant A participants.

(c) After all participants have made their moves, all moves are automatically recorded on the white sheets of A and B participants.

(d) The B participants will see their payoffs after each period whereas the A participants will only be informed of their payoffs at the very end of the round when the relevant payoff table is revealed.

(e) The first period ends.

5. The second period begins. The same procedure is followed.
6. The third period begins. The same procedure is followed.
7. The fourth period begins. The same procedure is followed.
8. The fifth period begins. The same procedure is followed.
9. The round ends.
10. The relevant payoff table for this round is revealed. The white sheets now show your total payoff for this round as well as your payoff for each period. Copy your total payoff for this round onto the record sheet.

There will be 10 such rounds.

Now, instruct participants to turn on their computers, and follow instructions. First we shall walk you through a practice round, you will not be paid for this practice round. The practice round will be conducted both on the computers and using actual paper sheets and slips. Follow the instructions of the experimenter and do not write anything on the sheets or the computers until instructed to do so. During the practice round, the experimenter will tell you what choices to make, make exactly those choices.

first period: DR, second period: UL, third period: UR, fourth period: DL, fifth period: DR.

The practice round is over, check payoffs, etc. Get ready for starting the actual experiments. Remember that the experiment will have 10 rounds like the one you just witnessed. You will be paid what you earn for all 10 rounds. In the actual experiment, you will use the computers only (no sheets of paper will be used).

Outcomes of practice round if the first table is drawn (show transparency 3). Outcomes of practice round if the second table is drawn (show transparency 4).

Appendix B
Posteriors on (ϵ_0, α) under the six models

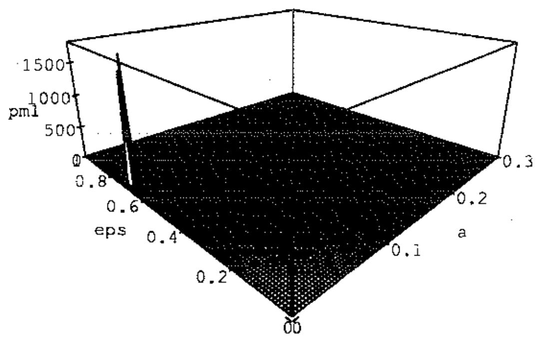

Figure 1: Posterior under Model SN.

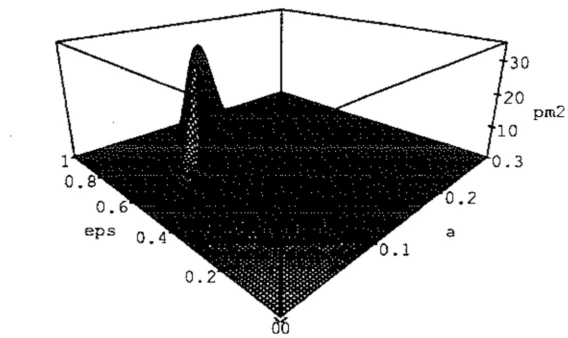

Figure 2: Posterior under Model UN.

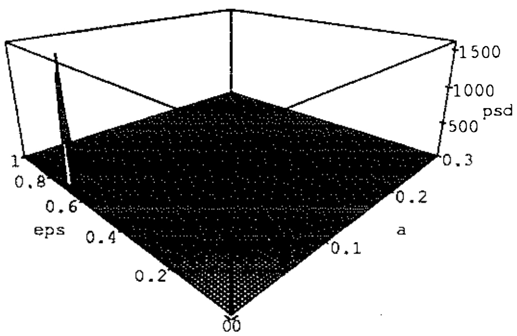

Figure 3: Posterior under Model SD.

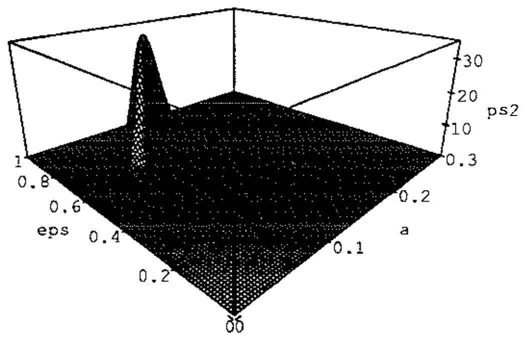

Figure 4: Posterior under Model UD.

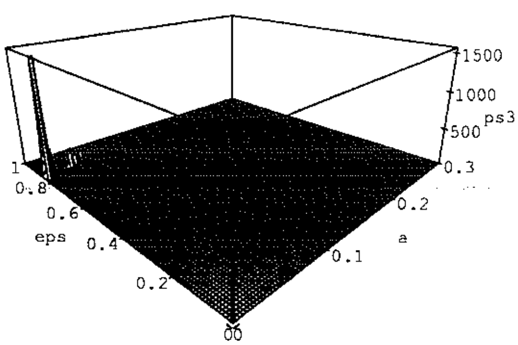

Figure 5: Posterior under Model SF.

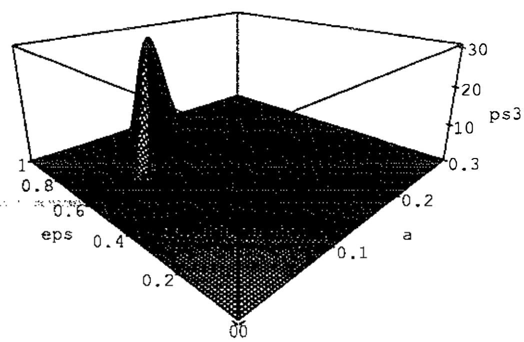

Figure 6: Posterior under Model UF.

REFERENCES

Banks, J., C. Camerer, and D. Porter. 1988. Experimental tests of Nash refinements. Mimeo. California Institute of Technology.

Beja, A. 1992. Imperfect equilibrium. *Games and Economic Behavior* 4:18–36.

Boylan, R. and M. El-Gamal. 1992. Fictitious play: A statistical study of multiple economic experiments. *Games and Economic Behavior.* Forthcoming.

Brandts, J. and C. Holt. 1989. An experimental test of equilibrium dominance in signaling games. Technical report, University of Virginia.

Brown, J. and R. Rosenthal. 1990. Testing the minimax hypothesis: A re-examination of O'Neill's game experiment. *Econometrica* 58:1065–1081.

Camerer, C. and K. Weigelt. 1988. Experimental tests of a sequential equilibrium reputation model. *Econometrica* 56:1–36.

Cooper, R., D. DeJong, R. Forsythe, and T. Ross. 1991. Selection criteria in coordination games. *American Economic Review.*

El-Gamal, M., R. McKelvey, and T. Palfrey. 1991. A bayesian sequential experimental study of learning in games. Soc. Sc. working paper 757, Caltech.

Harless, D. and C. Camerer. 1992. The predictive utility of generalized expected utility theories. mimeo. University of Chicago.

McKelvey, R. D. and T. Palfrey. 1992. An experimental study of the centipede game. *Econometrica* Forthcoming.

Palfrey, T. and R. Rosenthal. 1991. Testing game-theoretic models of free riding: New evidence on probability bias and learning. *Games and Economic Behavior* 3:183–220.

Palfrey, T. and R. Rosenthal. 1992. Repeated play, cooperation and coordination: An experimental study. SSWP 785. California Institute of Technology.