

DIVISION OF THE HUMANITIES AND SOCIAL SCIENCES
CALIFORNIA INSTITUTE OF TECHNOLOGY
PASADENA, CALIFORNIA 91125

ECONOMICAL EXPERIMENTS: BAYESIAN EFFICIENT EXPERIMENTAL DESIGN

Mahmoud A. El-Gamal

Thomas R. Palfrey



SOCIAL SCIENCE WORKING PAPER 884

March 1994

Economical Experiments: Bayesian Efficient Experimental Design

Mahmoud A. El-Gamal and Thomas R. Palfrey

Abstract

We propose and implement a Bayesian optimal design procedure. Our procedure takes as its primitives a class of models, a class of experimental designs, and priors on the nuisance parameters of those models. We select the experimental design that maximizes the information (in the sense of Kullback-Liebler) from the experiment. We sequentially sample with the given design and models until all but one of the models has viable posterior odds. A model which has low posterior odds in a small collection of models will have an even lower posterior odds when compared to a larger class, and hence we can dismiss it. The procedure can be used sequentially by introducing new models and comparing them to the models that survived earlier rounds of experiments. The emphasis is not on running as many experiments as possible, but rather on choosing experimental designs to distinguish between models in the shortest possible time period. The first stage of optimal design is illustrated with a simple experimental game with one-sided incomplete information.

JEL classification numbers: 026,211,215

Economical Experiments: Bayesian Efficient Experimental Design

Mahmoud A. El-Gamal and Thomas R. Palfrey *

1 Prologue

DeGroot: Let's talk a little bit about the current state of statistics. What areas do you think are particularly important these days? Where do you see the field going?

Blackwell: I can tell you what I'd like to see happen. First, of course, I would like to see more emphasis on Bayesian statistics. Within that area it seems to me that one promising direction which hasn't been explored at all is Bayesian experimental design. In a way, Bayesian statistics is much simpler than classical statistics in that once you're given a sample, all you have to do are calculations based on the sample. Now, of course, I say "all you have to do" – sometimes those calculations can be horrible. But if you are trying to design an experiment, that's not all you have to do. In that case, you have to look at all the different samples you might get and evaluate every one of them in order to calculate an overall risk, to decide if the experiment is worth doing and to choose among the experiments. Except in very special situations, such as when to stop sampling, I don't think a lot of work has been done in this area.

DeGroot: I think the reason there hasn't been very much done is because the problems are so hard. It's really hard to do explicitly the calculations that are required to find *the* optimal experiment. Do you think that perhaps the computing power that is now available would be helpful in this kind of problem?

Blackwell: That's certainly going to make a difference. ...

From: M. DeGroot (1986) "A Conversation with David Blackwell", *Statistical Science*, vol.1, no1., p.47.

* We acknowledge the financial support from NSF grant #SES-9223701 to the California Institute of Technology. We also acknowledge the research assistance of Eugene Grayver who wrote the software for the experiments. This paper was presented at the Stony Brook workshop of Experimental Game Theory, summer 1993.

2 Introduction

One of the main goals of experimental economics is to distinguish models of choice that are in accordance with actual human behavior from those that are not. The method experimental economists have traditionally used to distinguish “good” (empirically justified) from “bad” (empirically rejected) models is classical hypothesis testing. Following in the tradition of the majority of econometric specification tests, one typically uses experimental data to estimate the nuisance parameters under a given theory, and uses an asymptotic approximation to log likelihood to test that the specified theory generated the observed data. This is a specification test with unspecified alternatives (e.g. McKelvey and Palfrey (1992), and others). Alternatively, one can specify a collection of potential models of behavior, estimate the nuisance parameters under each of the models, and then conducting a likelihood ratio test, with the asymptotic distribution of that test determining the significance level. This is a test with specified alternatives (e.g. Harless and Camerer (1992)).

The emphasis of experimental studies of the type discussed above is orthogonal to the main theme of this paper. In traditional experimental economics, one decides *ex ante* (without any statistical analysis) what is an economically interesting experiment to run. One then attempts to collect as much data as financially feasible, non-statistically taking into consideration the importance of this experiment as a piece of a larger research agenda, how high the payoffs need to be to make the subjects interested in making optimal decisions, the length of time one can keep subjects in the laboratory, etc. After the data is collected, one approaches it in the same fashion an econometrician will apply any data set; hypothesizing the form of the data generating process (either structural or reduced form), and proceeding with the estimation and hypothesis testing.

In El-Gamal et al. (1993), we took a different approach to the issue of model selection in experimental Economics. We started with a pre-specified pair of models of behavior in repeated games of incomplete information. One of our models was inspired that of McKelvey and Palfrey (1992) for a centipede game where it is common knowledge that some of the subjects are altruists, and both altruists and rationals can make errors. The McKelvey and Palfrey (1992) model assumed that the proportion of altruists in the population is common knowledge, and hence subject beliefs about it were degenerate at the true value. In El-Gamal et al. (1993), we allowed subjects to have different beliefs about the proportion of altruists in the population, and hence, they can potentially update those beliefs. One of our models allowed subjects to update their beliefs about the proportion of altruists when they were done playing one subjects, and waited to be matched with the next; and the other model assumed that they did not update those beliefs. Fixing a design (a three-move centipede game with given payoffs, and where each subject plays two other subjects), we employed methods of optimal sequential sampling to decide on a stopping rule. Using an asymptotic approximation of a Bayesian Sequential Probability Ratio Test (BSPRT) as the cost per experiment was very small compared to the cost of choosing the wrong model, we concluded that the fully Bayesian model where subjects updated their beliefs about the proportion of altruists is significantly more likely

than the myopic model where they do not.

In the BSPRT procedure used in El-Gamal et al. (1993), the primitives are the cost of collecting an extra data point, the loss for making the wrong decision, and the prior odds on the competing models. The family of Wald SPRT's to which this procedure belongs has the remarkable property that it minimizes the number of data points needed to reach a decision, for a given type I and type II error. The emphasis in this sequential sampling approach to experimental economics is thus shifted from ex post data analysis to ex ante assessment of the informativeness of the experiment, and the relative losses to obtaining data and making a decision. This helps deal with the problem: when do we know when we are done running this experiment? It does not deal with the problem: which experiment should we run. In the prologue to this paper, Blackwell implied that the problem of “look[ing] at all the different samples you might get and evaluat[ing] every one of them in order to calculate an overall risk, to decide if the experiment is worth doing and to choose among the experiments” is more important than the problem of “when to stop sampling”. In this paper, we propose a procedure for “choosing among experiments” and apply it to the same class of problems addressed in El-Gamal et al. (1993); namely, whether subjects in experimental games with incomplete information are Bayesians.

The remainder of this paper will proceed as follows. Section 3 will define what we mean by optimal statistical design of an experiment, and discuss its relation to classical hypothesis testing. In section 4, we introduce a simple game of one-sided incomplete information, and derive the equilibrium for that game under our rival models. In section 5, we report the optimal design calculations for that game, and report the results that we obtained using the optimal design. In section 6 we discuss the general framework of sequential model selection which we are advocating, and section 7 concludes the paper.

3 The problem of optimal statistical design

The main primitives of our procedure will be a class of experiments and a class of models of subject behavior in those experiments. The class of experiments is indexed by the “design parameters” $\theta \in \Theta$. For example, in a market experiment, θ will be a vector determining the number of subjects, the number of markets, the length of each trading period, the redemption values of the tokens endowed to the subjects, etc. In a game experiment, the design parameters will be the number of subjects, the number of moves available to each subject, the information structure, randomization devices, payoff tables, etc. The class of models of behavior may contain models implied by economic theory, psychological models, reduced form functional relationships, etc.

To make the problem tractable, we start by selecting a finite dimensional vector of design parameters $\theta \in \Theta$, and a finite number of classes of models n . Typically, a class of models will have a number of nuisance parameters (e.g. utility function parameters, subject-specific beliefs, error rates, etc.), and hence each of those classes actually encom-

passes an infinite number of models. Our optimal statistical design problem now becomes choosing the “best” experimental design from the selected class of designs to discriminate among the selected classes of models. Let X be the space of all possible data sets under all of our designs. Denote a typical data set by x . Let the likelihoods of a given data set $x \in X$ under design θ for each of our n competing models be $l_1(x; \theta), \dots, l_n(x; \theta)$.¹ Given a collection of priors on models $1, \dots, n$, say p_1, \dots, p_n , we can define for each model the Kullback-Liebler information number measuring how informative a given design is expected to be if that model were correct. For example, the information number for model 1 under design θ is simply:

$$I(1; \theta) = \sum_{x \in X} l_1(x; \theta) \log \left(\frac{(1 - p_1)l_1(x; \theta)}{\sum_{i=2}^n p_i l_i(x; \theta)} \right).$$

The design which maximizes our expected separation between model 1 and the other $n - 1$ models, if model 1 were indeed the correct model, is $\theta_1^* = \max_{\theta \in \Theta} I(1; \theta)$. In standard hypothesis testing procedures, we would pick one of the models to be our null hypothesis, and perform a test of a particular size under that null. An equivalent bias of giving one model as good a chance as possible to win if it is indeed the correct model, we choose in this paper to treat one model as the primary candidate by choosing the design that maximizes the informativeness of the experiment if it were indeed the correct model. In all other aspects, however, the models can start on the same footing by letting $p_i = 1/n; \forall i$.²

One point must be made clear. After selecting a design, obtaining the data, and “selecting a model” from the chosen class, we are not actually “accepting” this model as the correct one in the classical sense of the term. As is the case with classical statistics, one cannot “accept” a model. We can reject models if the probability that the data we observe comes from those models is too small (tradition after Fisher (1950) dictates treating probabilities less than 1% or 5% as sufficiently close to zero; hence using those magic numbers as sizes for hypothesis test). Since we employ a Bayesian posterior that provides posterior odds on the selected (finite) class of models, those posteriors have to add up to unity. Hence, if there are say two models, and one of them has a very low posterior probability, then the other must have a very high posterior probability. That is not to say that we accept the model with the high posterior as the true model. We know that the posterior probability of the less likely model within the given finite class is an upper bound on its posterior probability if we considered all potential models. Hence, if a model is discarded by our procedure due to having too low a posterior probability, that result can be accepted regardless of the class of models we started with. However, the fact that one of our models performs well is not an indication that it will continue to

¹In our Bayesian design framework, each of the models will also have nuisance parameters. As shown below, the likelihood functions will then be computed by assessing our priors on those nuisance parameters, and integrating them out.

²One could alternatively make the problem completely symmetric by maximizing $\sum_{i=1}^n p_i I(i; \theta)$ instead of any one of the given $I(i, \theta)$'s. We chose in this paper to follow the procedure of highlighting one model as the benchmark to beat since that is closer in spirit to the more familiar hypothesis testing framework.

do so when compared to other models. The process for selecting “the best” model can never terminate since there is an infinite number of eligible classes of models, and we can only analyze a finite number of such classes at a time.

An alternative criterion for optimizing the design can be intuited from the above mentioned classical hypothesis testing point of view. In the classical hypothesis testing framework as culminated in Lehmann (1959), the emphasis is on choosing a point on the indifference curves of “size” (probability of rejecting a model when it is true) and “1-power” (probability of accepting a model when it is false). This would suggest choosing a design to minimize our expected loss (as a function of accepting a false model or rejecting a true one) subject to some budget constraint. We do not take this approach due to the well known shortcomings of the classical hypothesis testing approach, most notable of which is that “classical statistics tends to divert attention from (the likelihood ratio) to the two conditional probabilities of making errors” Edwards et al. (1963), resulting in the well known Jeffrey’s / Lindley’s paradox Lindley (1957) that given any significance level, and any posterior odds in favor of the null hypothesis, there exists a datum which produces that posterior odds, and is rejected at the given significant level. By defining our “utility” function in terms of likelihood ratios, and therefore defining expected utility (our objective function, the Kullback-Liebler information number) as the expectation of that utility function under the null hypothesis, we avoid those problems. For any given dollar value of the experiment, we simply maximize its informativeness under then null hypothesis.

In this paper, we perform the first step of this infinite sequential procedure of selecting a class of models, deciding on the design, sampling until one model gets very high posterior odds within the class, then add more models, and repeat the whole process. This paper takes that first step by choosing a class of games of one-sided incomplete information that are parametrized by a two-dimensional design vector. We start with three models of behavior that differ in their assumptions about how agents learn. The first model (which we choose as the benchmark) is that subjects play the Bayes-Nash equilibrium of the game. The second model postulates that subjects play a version of the Bayes-Nash equilibrium without Bayesian updating (the uninformed players do not use the information implicit in their opponents’ actions to update about the state of nature). The third model is that subjects myopically choose moves to maximize expected payoff given beliefs on their opponents’ moves that are updated using the fictitious play heuristic. In all three models, subjects will be allowed to have errors in beliefs (their perception of probabilities of moves by nature may be erroneous), and errors in actions (they may tremble and choose an action that does not agree with their strategy). In the following section, we shall introduce the game and the three models of behavior that we wish to compare.

4 The Game of Stop-Go:

We use for our experimental study a game which, for lack of a better name, we call “stop-go”. There are two possible states of nature “a”, and “b”. Nature chooses state “a” with probability π , and state “b” with probability $(1 - \pi)$. There are two players “1” and “2”. Player 1 is informed of the state of nature, but player 2 is not. In both states of nature, player 1 has an outside option *stop* which guarantees both players a payoff of 1. If player 1 does not choose *stop* and decides to choose *go*, then player 2 gets to choose *right* or *left*, and the payoffs are shown in the game tree in Figure 1, where the payoff A is assumed to be strictly greater than 2.

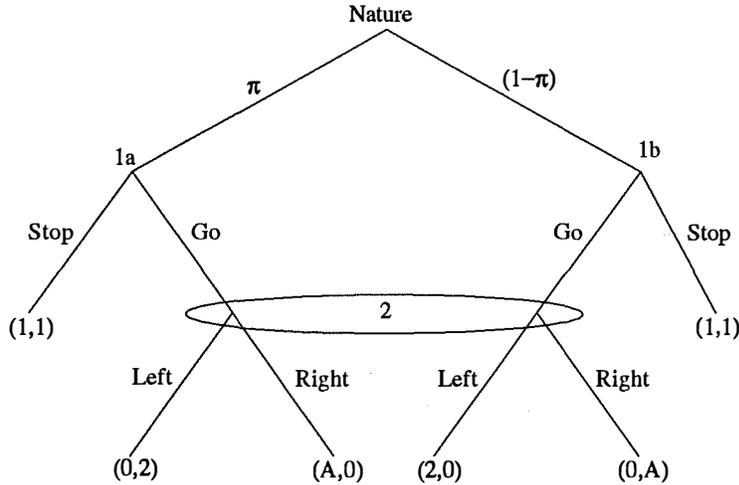


Figure 1: The Game of Stop-Go ($A > 2$)

This game is as simple a game of incomplete information as we could construct while still maintaining the main features that we need. The first and most important of those features is that the Bayesian updating needed to play the Bayes-Nash equilibrium of the game is quite simple (given 1 goes, what is the probability that the state of nature is “a”). Another useful characteristic is that except for a very small set of designs (which we avoid), the game has a *unique* Bayes-Nash equilibrium, which is easy to find as a function of π and A . This in turn makes the calculation of the likelihood of various data sets under the model (after introducing the errors in actions and errors in beliefs) rather straightforward, which is needed for calculating the above described information numbers for various design parameters $\theta = (\pi, A)$. The game also has the advantage of having a very simple structure which is very easy to explain to subjects, and it mimics a variety of confrontation games with an outside option that makes it easy for subjects to understand.

4.1 Perturbing the Game: Errors in actions and errors in beliefs

A generic problem in using standard economic models to study experimental data is the so called *zero likelihood problem*. We have characterized this problem in previous work (e.g. Boylan and El-Gamal (1993), El-Gamal et al. (1993), El-Gamal and Palfrey (1994)) and proposed a number of methods of introducing errors in actions to make every observable data set a positive-likelihood set under all models under consideration. In this paper, we follow the same approach by letting every agent at every decision node have a probability ϵ of trembling and making their choice by flipping a fair coin (i.e. if player 1 trembles, he “stops” with probability 1/2 and “goes with probability one-half, and if player 2 trembles, he chooses *right* with probability 1/2 and *left* with probability 1/2). For reasons of parsimony we assume that the probability of tremble ϵ is the same for all decisions and all players (see McKelvey and Palfrey (1993) for an alternative specification where the probability of errors depends on potential payoff loss), but that there is learning by doing which forces ϵ_t to decline exponentially in t , where t indexes the number of times the subject has played that particular game before. Therefore, if model i states that a player should choose some action with probability p_i , then we append to it this error probability and say that he chooses that action with probability $(1 - \epsilon_t/2)p_i + \epsilon_t/2$ at the t^{th} round (where $\epsilon_t = \epsilon_0 e^{-\alpha t}$, ϵ_0 is the *baseline error rate*, and α is the *learning by doing rate*). The nuisance parameters ϵ_0 and α , will therefore be parameters that all three of our models described below will contain.

We add a third nuisance parameter to all of our models by recognizing the fact that due to perceptive limitations, subjects tend not to fully understand probabilities, and hence it is very difficult to induce the correct priors in an experimental setting. At best, there will be a distribution of priors that different subjects will actually hold when they are given the same set of experimental instructions. We model this by making the perceived prior π_{per} uniformly distributed over $(\pi - \delta, \pi + \delta)$, where π is the induced prior. In other words, we assume that when we announce to the subjects that state of nature “a” will be drawn with probability π , they draw a $\pi_{per} \sim U(\pi - \delta, \pi + \delta)$. We can again make the distribution of perceived priors more sophisticated, but we choose this parsimonious specification which still allows us to estimate a *perception error* nuisance parameter δ . The randomly drawn π_{per} is assumed to be common knowledge to the subjects, but unknown to us.

4.2 Models of Behavior

We specify three models of behavior that could potentially describe the behavior of the subjects in this game. In this section we shall state the predicted behavior under each of the models as a function of π_{per} , A , and ϵ . When finally calculating the likelihoods below, we shall integrate out the nuisance parameters δ (implicitly defining the distribution of π_{per}), ϵ_0 , and α with respect to our priors.

Model I: Each individual i plays the Bayes Nash Equilibrium of the game, defined by (π_{per}, A, ϵ) .

Model II: Player 2 does not update π_{per} following go , and this is common knowledge.

Model III: Individuals use fictitious play to construct beliefs about opponents' play.

The most difficult model to solve is the Bayes-Nash equilibrium, where ϵ , the tremble rate is common knowledge, and all players take into consideration that they as well as their opponents will tremble with probability ϵ , and that that is common knowledge. The equilibrium strategies will be calculated by enumerating five different cases. There are two main features of this (unique) equilibrium:

1. For ϵ sufficiently large, there is a pure strategy equilibrium.
2. For all cases, at most one type of player 1 mixes.

4.3 Predicted behavior under Model I:

For model I, p_a is the probability that player 1 chooses go if the state of nature is “a”, p_b is the probability that player 1 chooses go if the state of nature is “b”, and q is the probability that player 2 will choose $left$ (if given the chance to move), provided that they do not tremble. In other words, the actual probability of a player 1 choosing go in state “a” is $(1 - \epsilon/2)p_a + \epsilon/2$, the actual probability of choosing go in state “b” is $(1 - \epsilon/2)p_b + \epsilon/2$, and the actual probability of player 2 choosing $left$ is $(1 - \epsilon/2)q + \epsilon/2$.

$$\hat{\pi}_{per} = \frac{A}{2 + A}$$

Case A: $\pi_{per} > \hat{\pi}_{per}$, $\epsilon \leq \frac{2\hat{\pi}_{per}(1-\pi_{per})}{\hat{\pi}_{per}+\pi_{per}-2\hat{\pi}_{per}\pi_{per}}$ and $\epsilon \leq \frac{2}{A}$:

$$p_a^I = \frac{A(1 - \pi_{per})}{2\pi_{per}} - \frac{[(A + 2)\pi_{per} - A]\epsilon/2}{(1 - \epsilon)2\pi_{per}}$$

$$p_b^I = 1$$

$$q^I = \frac{1}{1 - \epsilon} \left[\frac{A - 1}{A} - \epsilon/2 \right]$$

Case B: $\pi_{per} > \hat{\pi}_{per}$, $\epsilon > \frac{2\hat{\pi}_{per}(1-\pi_{per})}{\hat{\pi}_{per}+\pi_{per}-2\hat{\pi}_{per}\pi_{per}}$ and $\epsilon \leq \frac{2}{A}$:

$$p_a^I = 1 - p_b^I = 1 - q^I = 0$$

Case C: $\pi_{per} \leq \hat{\pi}_{per}$, $\epsilon \leq \frac{2\pi_{per}(1-\hat{\pi}_{per})}{\hat{\pi}_{per}+\pi_{per}-2\hat{\pi}_{per}\pi_{per}}$:

$$\begin{aligned} p_a^I &= 1 \\ p_b^I &= \frac{2\pi_{per}}{A(1-\pi_{per})} + \frac{[(A+2)\pi_{per}-A]\epsilon/2}{(1-\epsilon)A(1-\pi_{per})} \\ q^I &= \frac{1}{2} \end{aligned}$$

Case D: $\pi_{per} \leq \hat{\pi}_{per}$, $\epsilon > \frac{2\pi_{per}(1-\hat{\pi}_{per})}{\hat{\pi}_{per}+\pi_{per}-2\hat{\pi}_{per}\pi_{per}}$:

$$p_a^I = 1 - p_b^I = 1 - q^I = 1$$

Case E: $\pi_{per} > \hat{\pi}_{per}$:

$$p_a^I = p_b^I = q^I = 1$$

4.4 Behavior Under Model II:

Case AA: $2\pi_{per} > A(1-\pi_{per})$

$$\begin{aligned} p_a^{II} &= \begin{cases} 1 & \text{if } \epsilon A/2 > 1 \\ 0.5 & \text{if } \epsilon A/2 = 1 \\ 0 & \text{if } \epsilon A/2 < 1 \end{cases} \\ p_b^{II} &= \begin{cases} 1 & \text{if } 2(1-\epsilon/2) > 1 \\ 0.5 & \text{if } 2(1-\epsilon/2) = 1 \\ 0 & \text{if } 2(1-\epsilon/2) < 1 \end{cases} \\ q^{II} &= 1 \end{aligned}$$

Case BB: $2\pi_{per} < A(1-\pi_{per})$:

$$\begin{aligned} p_a^{II} &= \begin{cases} 1 & \text{if } (1-\epsilon/2)A > 1 \\ 0.5 & \text{if } (1-\epsilon/2)A = 1 \\ 0 & \text{if } (1-\epsilon/2)A < 1 \end{cases} \\ p_b^{II} &= q^{II} = 0 \end{aligned}$$

Case CC: $2\pi_{per} = A(1-\pi_{per})$:

$$\begin{aligned} p_a^{II} &= 1 \\ p_b^{II} &= 0.5 \\ q^{II} &= 0.5 \end{aligned}$$

4.5 Behavior Under Model III:

Player 1 starts with a belief that probability of player 2 choosing left is 1/2, and player 2 starts with beliefs that probability of player 1 going when it is game a is 1/2, and when it is game b is 1/2. Both players play optimal responses to their beliefs, and update them using fictitious play. In other words, initializing $emp_1 = 0.5$, $empa_1 = 0.5$, $empb_1 = 0.5$, the agents are assumed to use fictitious play to obtain player 1's belief emp_t about the proportion of the time player 2's choose *left*, and player 2's beliefs ($empa_t, empb_t$) about the proportion of time player 1's choose *go* when game *a* or *b* is drawn (respectively), such that in round t ,

$$\begin{aligned} emp_t &= \frac{emp + \#\{left|go\}}{\#\{go\} + 1} \\ empa_t &= \frac{empa + \#\{go|gamea\}}{\#\{gamea\} + 1} \\ empb_t &= \frac{empb + \#\{go|gameb\}}{\#\{gameb\} + 1} \end{aligned}$$

where $\{left|go\}$ is the event that an opponent (player 2) in some previous round chose *left* (they have to be given the chance to move first), $\{go|gamex\}$ is the event that an opponent (player 1) in some previous round chose *go* when the state of nature was $x \in \{a, b\}$, $\{go\}$ is the event that player 1 chooses *go*, $\{gamex\}$ is the event that the state of nature in a previous round was x , and $\#\{\cdot\}$ is the number of times the event has occurred so far.

Using those definitions, we can now define player 2's updated empirical belief at round t (via fictitious play) that game *a* is being played conditional on *go*, which we call $emp\pi_t$:

$$emp\pi_t = \frac{empa_t \pi_{per}}{empa_t \pi_{per} + empb_t \pi_{per}}.$$

With these definition, we can define the predicted strategy in round t when a player does not tremble as follows (again the actual likelihoods of moves under the model will take into consideration the probability of trembles, replacing 0's by $\epsilon/2$'s and 1's by $1 - \epsilon/2$'s).

$$p_{a,t}^{III} = \begin{cases} 1 & \text{if } (1 - emp_t)A > 1 \\ 0.5 & \text{if } (1 - emp_t)A = 1 \\ 0 & \text{if } (1 - emp_t)A < 1 \end{cases}$$

$$p_{b,t}^{III} = \begin{cases} 1 & \text{if } 2 emp_t > 1 \\ 0.5 & \text{if } 2 emp_t = 1 \\ 0 & \text{if } 2 emp_t < 1 \end{cases}$$

$$q_t^{III} = \begin{cases} 1 & \text{if } 2emp\pi_t > (1 - emp\pi_t) A \\ 0.5 & \text{if } 2emp\pi_t = (1 - emp\pi_t) A \\ 0 & \text{if } 2emp\pi_t < (1 - emp\pi_t) A \end{cases}$$

4.6 Likelihood Functions

Now that we have defined the strategy profiles (p_a, p_b, q) for players who do not tremble under each of our three models, we can define the likelihoods of a particular move by defining (for $M \in \{I, II, III\}$):

$$\begin{aligned} \text{obs}p_{a,t}^M &= (1 - \epsilon_t)p_{a,t}^M + \epsilon_t/2, \\ \text{obs}p_{b,t}^M &= (1 - \epsilon_t)p_{b,t}^M + \epsilon_t/2, \\ \text{obs}q_t^M &= (1 - \epsilon_t)q_t^M + \epsilon_t/2. \end{aligned}$$

The likelihood function under model M is now:

$$\text{like}_T^M = \int_{\epsilon_0} \int_{\alpha} \int_{\delta} \int_{\pi-\delta}^{\pi+\delta} \left(\prod_{t=1}^T \prod_{i=1}^n \text{like}(\text{action}_i^i | M; \epsilon_0, \alpha, \pi_{per}) \right) d\pi_{per} \text{prior}(d\epsilon_0, d\alpha, d\delta)$$

Where $\text{like}(\text{action}_i^i | M; \epsilon_0, \alpha, \pi_{per})$ is the likelihood under model M , with agents' perceived prior being π_{per} , that agent i (with his observed history) in round t , with tremble rate $\epsilon_0 e^{-\alpha t}$ chooses action_i^i . If agent i is a player 1, then this contribution to the overall likelihood is $\text{obs}p_{a,t}^M$ or $\text{obs}p_{b,t}^M$ if the action was *go* and the draw of nature was a or b respectively, and it is $\text{obs}q_t^M$ if the agent is a player 2 and chose *left* (all computed with the appropriate $\epsilon_t = \epsilon_0 e^{-\alpha t}$, and π_{per}). Now, for a given number of rounds and number of players, we can calculate the Kullback-Liebler information numbers defined in section 2 by calculating the likelihoods of the models under all possible data sets (i.e. all possible collections of actions_i^i ; $i = 1, \dots, n$; $t = 1, \dots, T$), and choose the optimal design.

5 The Experimental Design:

In this section, we shall discuss a number of practical design issues that we have so far ignored. It must be apparent to the reader that the problem of choosing *the* optimal design is not well posed. There are too many parameters over which to optimize (the class of all possible games, the number of rounds, the number of individuals, etc.). In the general discussion of section 2, we assumed that we limited our attention to a class of experimental designs parametrized by a finite dimensional vector $\theta \in \Theta$. The optimization of the design was then reduced to a problem of choosing the value of θ that maximizes a Kullback-Liebler information number. By limiting attention to the class of games defined above by the game of stop-go, and parametrized by $(\pi \in (0, 1), A > 2)$, we made the information number of the experiment a function of (π, A) , i.e. we generate a two dimensional information surface whose peak we seek. Now, we have to take into consideration relevant issues of perception of payoff differentials (e.g. an $A = 2.01$ cannot be assumed clearly distinguishable from 2.0, and a $\pi = 0.401$ cannot be assumed distinguishable from 0.4). We also need to take into consideration issues of not having too few a number of people (to avoid reputation effects and the possibility of guessing who one is paired with) or too many (collecting too much data all at once without getting a chance to analyze

it and see if we have all the information we need). We cannot have too few rounds (not allowing fictitious play types to update), or too many (inducing subject boredom and/or making the payoff per decision insignificant). Moreover, there are technical constraints to the calculation of the information, since each information number calculation requires looping over all possible data sets (whose number grows exponentially in the number of rounds, and geometrically in the number of subjects).

We decided on running experimental sessions with 10 subjects (5 player 1’s and 5 player 2’s), which lasts for three rounds, yielding “batches” of data, 15 games at a time.³ The experiment was conducted at the social science experimental laboratory at Caltech, and as subjects entered the lab, they were randomly assigned a role as player 1, or player 2, and their role as player 1 or 2 remained the same throughout the experiment. Then, the instructions (reproduced in Appendix B) were read aloud to the subjects. Subjects were told that they will each play three rounds of the game of stop-go described above, and that they were never going to be paired with the same opponent twice. This specification (the number of subjects $n = 10$, and the number of rounds $T = 3$) was at the limit of our computational abilities, and as will be seen below, they were enough to very strongly distinguish among our three models. The generation of this Bayesian information surface required (as discussed in section 2) the integration out of the nuisance parameters $(\epsilon_0, \alpha, \delta)$ parametrizing the errors in actions and errors in beliefs in all three models. We used independent priors on the three parameters with the prior on ϵ_0 and α being $U[0, 1]$, and the prior on δ being $U[0, 0.2]$. For that design with $n = 10$, $T = 3$ and the priors specified above, we computed the information surface under the assumption that model 1 is the true model. This information surface is shown in Figure 2, and its contours are shown in Figure 3.

³As it turned out, we only needed one batch of data.

Information surface:

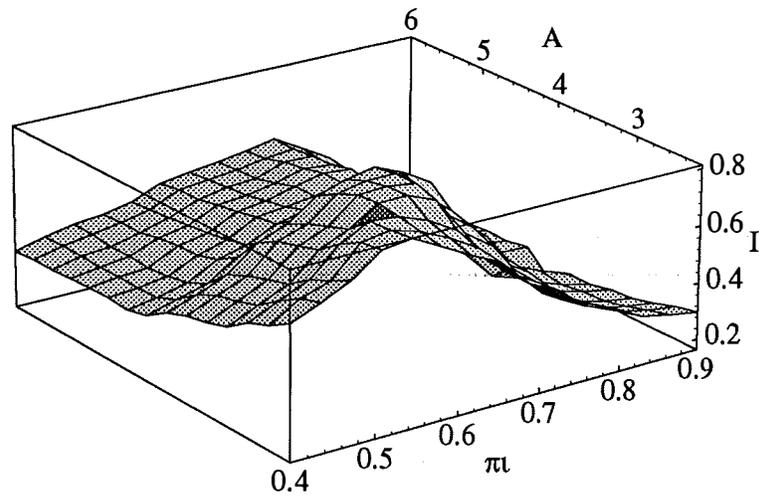


Figure 2: Information as a function of π and A

Contours:

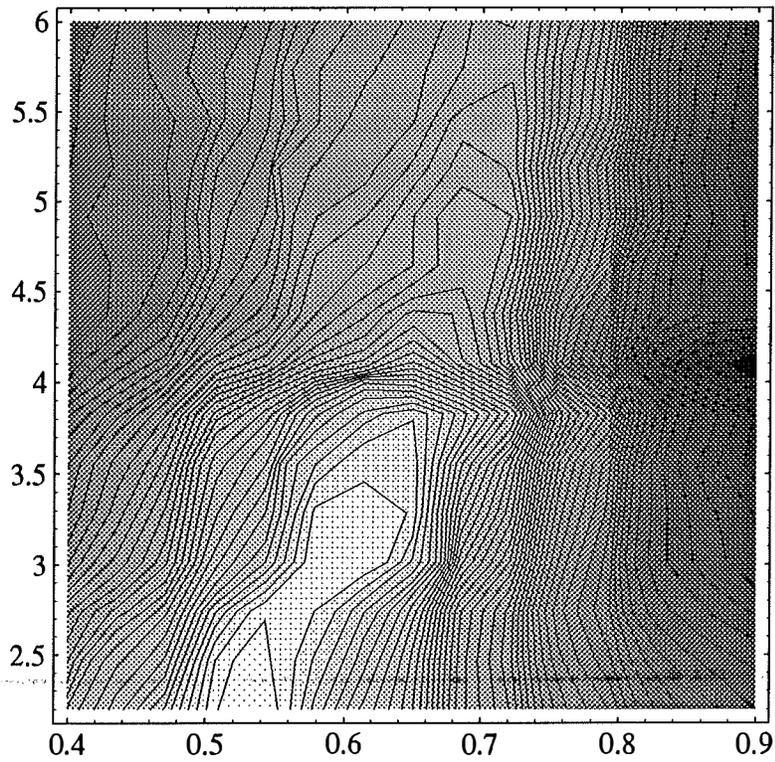


Figure 3: Contours of the information surface

A blind maximization of the information surface would drive us all the way to $A = 2$, which is not allowed. Hence, a true maximum with $A > 2$ does not exist. Moreover, we wish to choose a point that is clearly distinct from 2.0 to minimize the possibility that subjects do not find A and 2.0 distinguishable. If we look for an interior point very close to the optimum, we get ($A = 3.33, \pi = 0.6$) (inside the “bulb” in Figure 3). Since probability 0.6 is relatively hard to explain to subjects, and since the surface drops rather slowly in the direction of lower π 's, we selected the nearby design

$$\pi = 0.5 \quad A = 3.33$$

We chose payoffs of \$6.00 and \$10.00 (for 2 and A , respectively).

6 Results

We ran an experimental session with 10 subjects playing three rounds each as described above. The raw data collected from that experimental session is shown in Table 1.

R-ID	B-ID	R-Mv	B-Mv	R-Pay	B-Pay	Game
1	1	G	L	0	6	a
2	2	G	L	6	0	b
3	3	G	R	0	10	b
4	4	G	R	10	0	a
5	5	G	L	6	0	b
1	2	G	R	0	10	b
2	3	G	R	0	10	b
3	4	G	R	0	10	b
4	5	G	L	6	0	b
5	1	G	L	0	6	a
1	3	G	L	0	6	a
2	4	G	L	6	0	b
3	5	G	R	10	0	a
4	1	G	R	0	10	b
5	2	G	L	6	0	b

Table 1: Raw Data

Analysis of this data results in the likelihoods of the three models shown on the first line of Table 2. Since we start with priors 1/3 on each of the models, the posterior odds on the three models shown on the second line of Table 2 are easily calculated by just dividing the likelihood of each of the models by the sum of all three likelihoods. The result is very strong, with the posterior odds on model I (the Bayes-Nash equilibrium prediction) being higher than 99.9%.

	Model I	Model II	Model III
likelihoods	3.527×10^{-7}	9.655×10^{-11}	1.042×10^{-10}
post. odds	0.99943	0.00027	0.00030

Table 2: Likelihoods and Posterior Odds

The story behind the performance of the three models can be seen by looking at the posteriors on the nuisance parameters. Since it is not easy to portray four-dimensional figures, we show in Appendix A the marginals of our posteriors under the three models. We first note that the posterior under model I (the strong winner) puts a lot of mass on high levels of α (learning by doing), and low levels of δ (errors in beliefs), whereas the two losing models do the exact opposite. Models I and II both predict high levels of ϵ_0 (initial tremble rates), but model I, with a significantly higher α predicts much lower values of ϵ_T . The interpretation one would give is that the two losing models II and III are trying to get a better fit to the data by increasing the persistence in the error probabilities, as well as increasing the probability of large errors in beliefs. A safe conclusion seems to be that the Bayes-Nash model very significantly outperforms the other two, and in future investigations, when we choose future designs, we need not take those other two models into consideration. Of course, after getting the data from the next round of experiments, we can still evaluate the likelihoods of those abandoned models, and they may get resurrected if their performance puts back into contention. At this point, however, our prior belief that these models can be resurrected is very low. Without any need for explicitly introducing a penalty function for selecting the wrong model (as in El-Gamal et al. (1993)), we note that any sequential sampling procedure with a reasonable loss function will stop with posterior odds higher than 1000:1.

We can (with perfect hindsight) see in the data why the performance of models II and III was so much inferior to the performance of model I. At the aggregate level, the data looks like player 1's always choosing *go*, and player 2's choosing *left* with probability 1/2. Under model II, this can best be explained by high ϵ_t and low α (which always drive observable proportions of *left* towards 1/2), as well as $\pi_{per} = 5/8$, which has probability zero of occurring, this explains why the contribution of player 2's to the likelihood of model II will be low. As for player 1's, if they think that player 2's will always be choosing *left* with probability 1/2, they will always *go* when game *a* is drawn, and be indifferent when game *b* is drawn. Since we model indifference via choosing *go* with probability 1/2, that also makes player 1's contribution to the likelihood of model II rather low. For model III, as player 2's update their beliefs to player 1's choosing *go* all the time, the model reduces to model II, and therefore will run into the same low likelihood of generating the given data. Moreover, in the first round, player 2's should, under model III, choose *right* all the time. We can therefore see why models II and III perform poorly.

7 Concluding Remarks

It is clear from our experimental results how choosing a design to maximize the separation between our models (in terms of a Kullback-Liebler information number) subject to some common sense constraints (making probabilities and payoffs easily recognizable by subjects) can yield very strong results from very small data sets. Since given any class of models, the stage of distinguishing the most viable among them is only one of many stages towards improving our understanding of behavior subjects in similar experimental settings, we wish to minimize the time we spend on that stage and move on to the next stage. In the following stage, we would introduce a large class of models, and use some data mining procedure (e.g. El-Gamal and Grether (1993)) to decide on a small subclass of models with sufficiently high likelihood to carry out the next stage of optimizing the design and running new experiments to choose the best within the current class of models, and so on. We should not worry too much about models that are currently ignored due to their low posterior odds being permanently rejected. In future data mining phases of the sequential procedure, those models could be resurrected and taken into consideration in future designs. At the point of optimal design, however, the expected utility of taking those models into consideration is too low to justify the computational and modeling effort.

Appendix A
Marginals of the posterior under model I

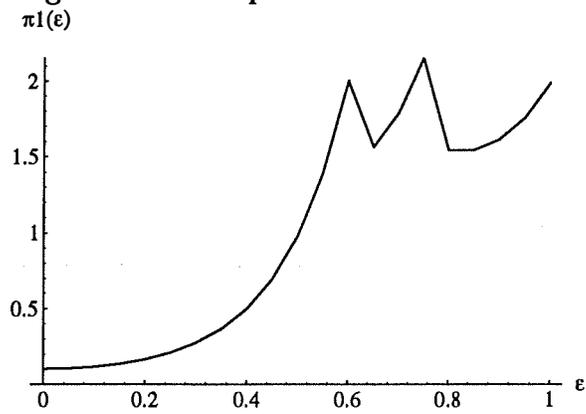


Figure 4: Marginal posterior on ϵ_0 under model I

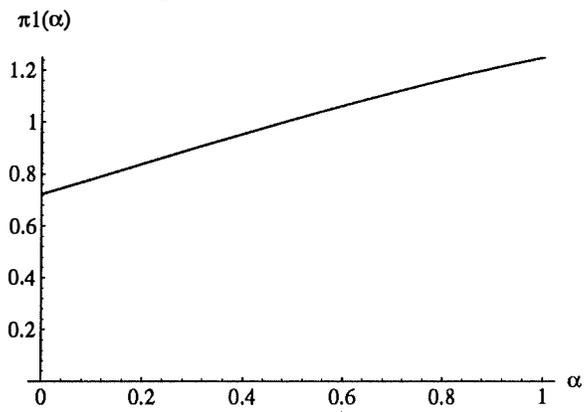


Figure 5: Marginal posterior on α under model I

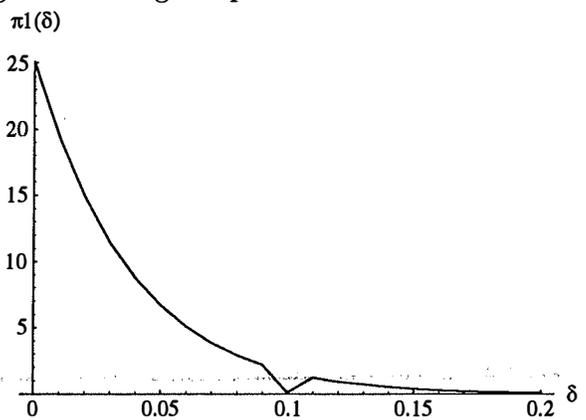


Figure 6: Marginal posterior on δ under model I

Marginals of the posterior under model II

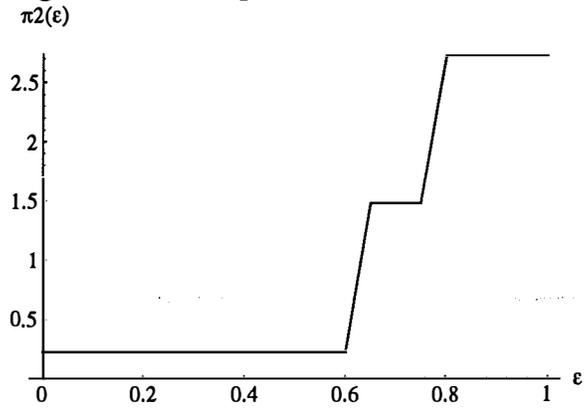


Figure 7: Marginal posterior on ϵ_0 under model II

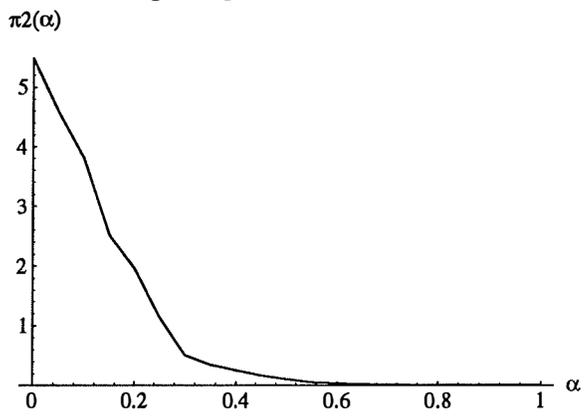


Figure 8: Marginal posterior on α under model II

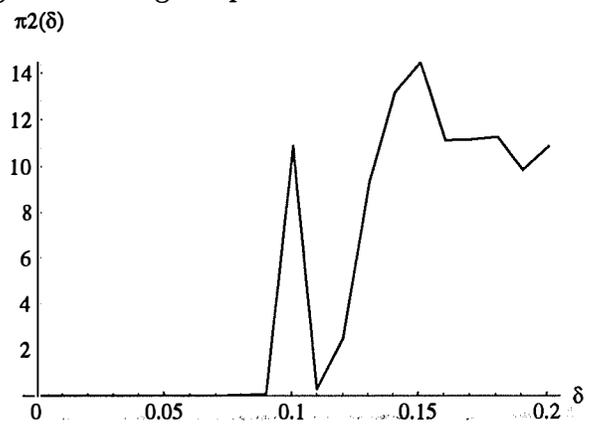


Figure 9: Marginal posterior on δ under model II

Marginals of the posterior under model III

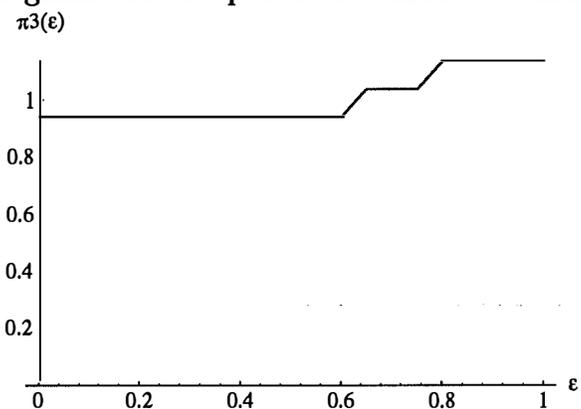


Figure 10: Marginal posterior on ϵ_0 under model III

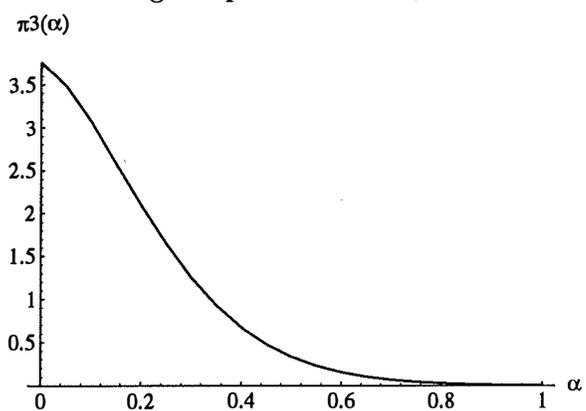


Figure 11: Marginal posterior on α under model III

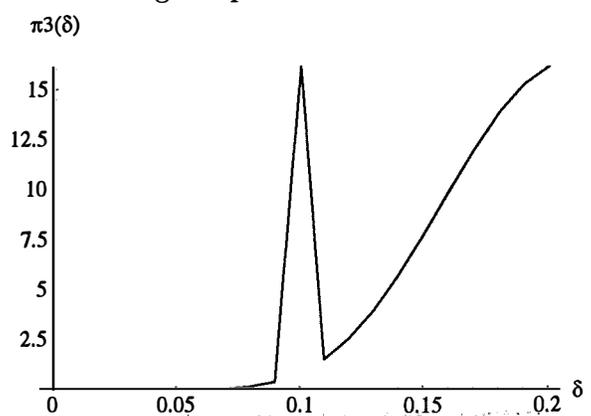


Figure 12: Marginal posterior on δ under model III

Appendix B

June 18, 1993 (STOP-GO)

Decision-Making Experiment

This session is part of a laboratory study in decision making. You will be paid IN CASH at the end of the session. The amount of money you earn will depend upon the decisions you make and on the decisions other people make. We ask that you do not talk or otherwise attempt to communicate with the other subjects except according to the specific rules below. If you have a question, feel free to raise your hand. One of us will come over to where you are sitting and answer your question in private.

This session you are participating in is broken down into a sequence of 3 separate rounds called MATCHES. At the end of the last match, you will be paid the total amount you have earned during the course of all 3 matches. Everyone will be paid in private and you do not have to tell anyone how much you earned. Half of you have been assigned to be "red" players and half of you have been assigned to be "blue" players. This assignment will stay the same in all 3 matches.

[Make sure red players are the ones sitting closest to the wall.]

Each of the 3 matches goes as follows. Each red player is randomly and anonymously matched with a blue player. This divides the N? of you in this room into N/2? pairs. You will never find out whom you were matched with in any of the matches. You will always be rematched with a completely different player at the beginning of the next match. After you have been matched with a player, you will never again be matched with them for the rest of the session. We next describe how each match proceeds for a given pair. Every pair follows the same instructions. The computer first chooses for that pair either payoff table I or payoff table II, in a random way, so that for any given pair, there is a 50/50 chance that table I is chosen and a 50/50 chance that table II is chosen. The easiest way to think about this random device is that it is exactly like tossing a fair coin several times, once for each pair. A HEAD means that pair will use table I and a TAIL means that pair will use table II. For each pair in each match the computer randomly selects I or II as if it is tossing another fair coin. Therefore, no matter how many times you ended up with table I in previous matches, there is always a 50/50 chance of ending up with table II in the current match. Furthermore, different pairs in the same match might have different payoff tables. After the computer has determined which payoff table your pair uses in the current match, ~~ONLY THE RED PLAYER IS TOLD WHICH IT IS.~~ ~~Blue players will not be told which payoff table is being used until that match is over.~~ The red player moves first in the match and has two choices, called STOP and GO. If the red player chooses STOP, the match is over for that pair. When this happens, red and blue each receives a STOP payoff of \$3.00. When this happens, both player's screens will highlight the STOP move. If the red player chooses GO, it is the blue player's turn to choose LEFT or RIGHT. After Blue chooses, Red is told Blue's choice and the match ends. No one is told what happens to the other pairs. The payoffs in this case depend on whether that pair uses payoff table I or payoff table II. The payoffs for this pair are not affected by what happens to the other pairs in this match.

Payoff Table I:

Left	Right
Red gets \$0.00 Blue gets \$6.00	Red gets \$10.00 Blue gets \$0.00

Payoff Table II:

Left	Right
Red gets \$6.00 Blue gets \$0.00	Red gets \$0.00 Blue gets \$10.00

Notice that which payoff table was selected matters only if the red player has chosen GO. If the red player chooses STOP, the match ends and payoffs are \$3.00 for each player in the pair. We also ask blue players to do one additional thing. If the red player you are paired with chooses STOP, please enter a decision (either LEFT or RIGHT) anyway. It makes absolutely no difference at all which decision you enter after red has chosen STOP, since the match is automatically over already. Let me emphasize this fact. This choice of yours after red has already ended the match by choosing STOP will not affect your payoffs in any way whatsoever. Since it has no effect on payoffs, the red player you are matched with is not told which decision you entered in this case. At the end of a match, each BLUE player is then told which payoff table their pair was using (recall that the RED players already had been told this at the beginning of the match). Your payoffs for that match are then displayed and highlighted on your screen. For those pairs in which red selected "GO", Blue's choice of LEFT or RIGHT is then highlighted on the screen of the red player in that pair to plainly see. After all of this has happened you will be prompted to record on your record sheet the match number, Red's move, Blue's move, which payoff table was used, and what your payoff was for that match. Please do not record anything on your record sheet until you are prompted by the computer to do so. When everyone in the room has finished recording and is ready to begin the next match, we will then assign new pairings for everyone, and the computer will randomly select a payoff table for each pair. Remember that table I is always selected with 50/50 chance and II is selected with 50/50 chance, and which tables you have used in previous matches have no bearing at all on which table the computer happens to select for you in this match. Also remember that as soon as a payoff table has been selected for your pair, the red player is told which one it is, but the blue player is not told until the end of the match.

This procedure will continue for a total of 3 matches. After the last match, add up your earnings for all 3 matches, and record this amount at the bottom of your record sheet. As a double check, the computer will also calculate your total earnings and display it on the screen. The experimenter will pay you this amount in cash, one at a time, in

the next room. Please bring all forms and personal effects with you when it is your turn to be paid. You will exit by the door in the other room. Are there any questions?

We will now explain how to enter your decisions into the computer, and how to read the information on the screen. Please do not press any keys on your keyboard unless prompted to enter specific information.

[Start Practice Rounds]

Please enter your name when the screen prompts you to do so. Then type in your color assignment (r or b) when prompted. Remember that you will have this color for the entire session of 3 matches. Please do not talk and do not press any keys. Is everyone ready to begin a practice match?? You will not be paid for the practice match

[Hand out "practice record sheet", but remind them that no payments will be made for the practice match.]

[Practice match begins. The experimenter then explains what is on the screen. Blue players walk over to see what the red players' screens look like, and vice versa, to emphasize that red players are informed which payoff table it is, and blue players are not so informed.]

[Practice match 1 begins. All red players are asked to play Go and all Blue players are asked to play Left. The experimenter then explains where the opponent's move is listed on the screen, where the payoff is listed, and makes sure everyone understands this. Everyone records this information.]

[Experimenter then marches through what would have happened if the other table had been chosen and the same move had been made.]

[Practice match 2 starts. All red players are asked to play Stop and blue players are asked to play whatever they wish. Everyone records the information.]

[Experimenter then marches through what would have happened if the other table had been chosen and the same move had been made. No difference in this case.]

[In practice match 3, all red players are asked to play Go and blue players are asked to play Right. Experimenter then marches through what would have happened if the other table had been chosen and the same move had been made.]

[The practice experiment is over and subjects are reminded to record everything.]

Are there any questions? [Answer questions]

We will now hand out a short quiz to make sure that everyone understands the details of what we will be doing in these 3 matches. Please raise your hand when you have completed the quiz, so one of us can come by and check your answers.

[Restart experiment and begin real payoffs]

References

- Boylan, R. and M. El-Gamal. 1993. Fictitious play: A statistical study of multiple economic experiments. *Games and Economic Behavior* 5:205–222.
- Edwards, W., H. Lindman, and L. Savage. 1963. Bayesian statistical inference for psychological research. *Psychological Review* 70:193–242.
- El-Gamal, M. and D. Grether. 1993. Uncovering behavioral strategies: Likelihood based experimental data mining. SSWP #850. California Institute of Technology.
- El-Gamal, M., R. McKelvey, and T. Palfrey. 1993. A Bayesian sequential experimental study of learning in games. *Journal of the American Statistical Association* 88:428–435.
- El-Gamal, M. and T. Palfrey. 1994. Vertigo: Comparing structural models of imperfect behavior in experimental games. *Games and Economic Behavior* page Forthcoming.
- Fisher, R. 1950. *Contributions to mathematical statistics*. New York: Wiley.
- Harless, D. and C. Camerer. 1992. The predictive utility of generalized expected utility theories. mimeo. University of Chicago.
- Lehmann, L. 1959. *Testing statistical hypotheses*. New York: Wiley.
- Lindley, D. 1957. A statistical paradox. *Biometrika* 44:187–192.
- McKelvey, R. D. and T. Palfrey. 1992. An experimental study of the centipede game. *Econometrica* 60:803–836.