

DIVISION OF THE HUMANITIES AND SOCIAL SCIENCES
CALIFORNIA INSTITUTE OF TECHNOLOGY

PASADENA, CALIFORNIA 91125

FICTITIOUS PLAY: A STATISTICAL STUDY OF MULTIPLE ECONOMIC EXPERIMENTS

Richard T. Boylan and Mahmoud A. El-Gamal



SOCIAL SCIENCE WORKING PAPER 737

May 1990

FICTITIOUS PLAY: A STATISTICAL STUDY OF MULTIPLE ECONOMIC EXPERIMENTS*

Richard T. Boylan

Mahmoud A. El-Gamal

California Institute of Technology

May 31, 1990

Abstract

This paper illustrates the use of a full Bayesian procedure to update an experimenter's belief over various economic behavioral hypotheses using data from a variety of (potentially very different) experiments. Our example uses experimental data to update our belief as to whether individuals select strategies according to fictitious play. We endow the experimenter with priors over the events that players act according to fictitious play and according to the Cournot process. We then numerically compute the likelihood function for each experiment by replicating the experimental design and running the experiment with robots that behave according to each of our hypotheses. Updating experiment by experiment shows that some of the experiments favor Cournot, but most of them favor fictitious play as the more likely hypothesis. This illustrates the limitations of a classical procedure that can take only one experiment into consideration since some of the experiments may be misleading. Indeed, when we did the overall updating using 9 experiments, we found that, for any priors, the overall posterior put probability very close to one on the individuals acting according to fictitious play. Given the heterogeneity in the payoffs and design of the experiments that we combine for that overall posterior, it is clear that there is no classical procedure that would offer the same type of information.

*The authors thank Robert Forsythe and Gary Miller for making their data available. Please Send Communications to DHSS 228-77, Caltech, Pasadena, CA 91125.

1 Introduction

Rational play of a game implies that a player chooses strategies that maximize his payoff given his beliefs over his opponents' strategies. A Nash equilibrium is a choice of strategies such that the expectations are fulfilled. There are two problems with the concept of Nash equilibrium. First, there is no explanation of how the players end up with beliefs that are self fulfilled. Second, the definition of Nash equilibrium does not discriminate between beliefs that are more and those that are less likely to occur. This leads to a great number of Nash equilibria and thus reduces the usefulness of the equilibrium concept ¹.

One way of solving these problems is to examine the situation where individuals play a particular game repeatedly. (Indeed a game of any interest is one that will be played quite often.) Players start with a belief of how the other players will act and update their beliefs after each stage of the game. An equilibrium of this system will be an invariant distribution of strategies among the set of players. Such equilibria will be justified by some learning process and will discriminate between beliefs that are more likely and less likely to prevail. This approach to the problem is not novel. Cournot examined a model where each player assumes that the other players will select the strategy they selected in the previous round and maximizes his payoff accordingly. Pure strategy equilibria of this process will be Nash equilibria. Another process that has been analyzed is the so-called fictitious play (see Brown (1951), Robinson (1951), Shapley (1964), Brock, Marimon, Rust and Sargent (1988)). The process assumes that an individual has Dirichlet priors over his opponents' strategies and that at each round the player updates his beliefs according to Bayes' rules. It turns out that these assumptions are equivalent to assuming that each player has beliefs that are a convex combination of his initial belief and the empirical distribution.

Let us then start with an economist who has beliefs regarding the relative validity of the two hypotheses that humans update according to the fictitious play or the Cournot updating rule. We put full support over smeared versions of these hypotheses that allow errors with some probability. We shall discuss the details of that smearing procedure and its implications in later sections of the paper. The economist observes a sequence of experiments where the

¹One way of reducing the number of equilibria is to restrict the player's set of beliefs. This approach is adopted by Banks and Sobel (1987), Cho and Kreps (1987), and Grossman and Perry (1987). However such restrictions are somewhat arbitrary.

experimental subjects played a repeated game. The experimenter then uses Bayes' rule to update his beliefs about the two hypotheses given the data from all the experiments. Since the experiments are very different in terms of payoffs and designs, combining all the data to compute a single likelihood function is not possible. Moreover, the sequential approach that we use in this paper (computing a likelihood function for each experiment separately, and sequentially using Bayes' rule) cannot be implemented in a classical framework since the distribution of the maximum likelihood statistic will be unknown given the previous experiments.

Two points make our method of combining experiments especially valuable. The first is the fact that the order in which we update the economist's beliefs (and hence the order in which the experiments are observed) does not matter. This is not surprising in the least, but we have a rigorous statement of that result in section 4 for completeness. The other point is that empirically, some of the 9 experiments that we use for updating actually point strongly in favor of the Cournot process which the overall analysis strongly suggests to be less likely than fictitious play.

Figure 1 is a plot of the economist's posterior as a function of his prior and the smearing parameter ϵ to be explained later. For now, just look at the posterior over the smeared fictitious play hypothesis (the height of the graph) at different values of the prior over that hypothesis at small (say less than .5) values of ϵ . At $\epsilon = 0$, the hypothesis is strictly fictitious play (and we cannot compute the posterior), and for positive values, ϵ is the probability that any particular person in any particular stage of any particular experiment gets to choose his action purely randomly. It is clear that the posterior for most reasonable values of epsilon and for all positive priors is very close to unity, and hence we are inclined to think that the smeared fictitious play hypothesis is infinitely more likely than the smeared Cournot hypothesis.

The rest of this paper builds up and justifies the necessary machinery to achieve Figure 1. Section 2 will discuss one class of games three of which we studied, Section 3 will discuss the other class of games, six of which we studied. Section 4 will discuss and justify the econometric procedure that we follow, and the paper will end with a series of 9 appendices for the nine experiments.

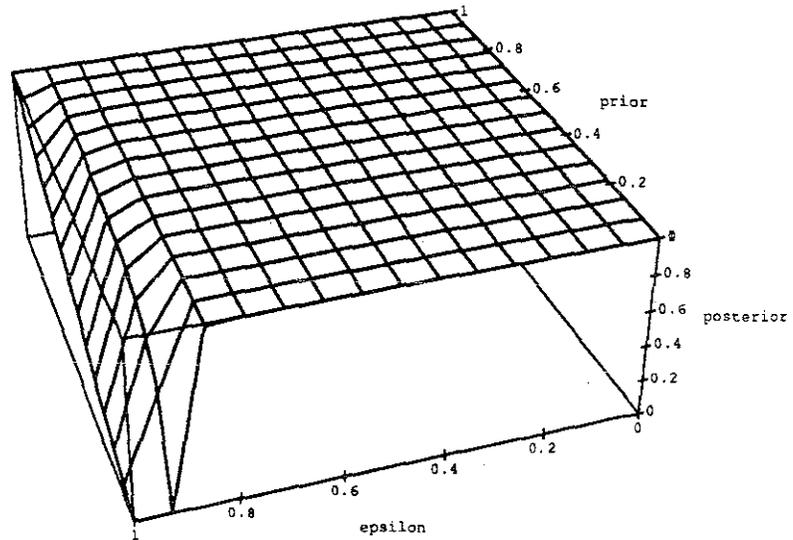


Figure 1: Posterior probability of agents playing according to fictitious play using all 9 experiments.

2 Games with one opponent

This section is based on some experiments that were run by Knott and Miller (1987). In this series of three experiments (labeled A, B, and C), each of which has individuals matched in pairs and play the games reproduced in Figures 2, 3, and 4 respectively ten times. Thus, for instance, if in a given repetition of game A, an individual selects strategy S_2 while his opponent selects strategy S_1 then the individual receives 200 pennies.

For each experiment, all the interesting aspects of the theoretical and observed behavior is depicted in a separate appendix. For instance, Appendix A deals with experiment A, Appendix B with experiment B and Appendix C with experiment C. In each appendix, the first two figures show a Monte-Carlo distribution of actions in each of the ten stages of the game. Specifically, in Appendix A, Figure 11 displays the simulated distribution of those plays under Cournot updating. To obtain this distribution, we simulate 1000 pairs (2000 individuals) and endow them with randomly drawn (uniform over the unit simplex) initial priors over their opponents' possible actions. Figure 11 shows the obtained empirical distribution. We then let each pair go through the 10 stages of the game where they are made to update according

	S1	S2	S3	S4	S5	S6	S7	S8	S9	S10	S11	S12	S13	S14	S15
S1	100	100	75	85	85	90	90	90	90	90	90	90	90	90	90
S2	200	120	100	100	100	100	100	100	100	100	100	100	100	100	100
S3	0	140	125	115	115	110	110	110	110	110	110	110	110	110	110
S4	-100	160	150	130	125	120	120	120	120	120	120	120	120	120	120
S5	-200	180	175	145	135	130	130	130	130	130	130	130	130	130	130
S6	-300	200	200	160	145	140	140	140	140	140	140	140	140	140	140
S7	-400	220	220	175	155	155	150	150	150	150	150	150	150	150	150
S8	-450	0	250	190	165	160	160	160	160	160	160	160	160	160	160
S9	-475	-100	0	205	175	170	170	170	170	170	170	170	170	170	170
S10	-500	-125	-100	0	185	180	175	175	180	180	180	180	180	180	180
S11	-525	-150	-125	-100	0	190	180	180	190	190	190	190	190	190	190
S12	-550	-175	-150	-125	-100	0	195	190	195	195	200	200	200	200	200
S13	-575	-200	-175	-150	-125	-100	0	250	200	200	210	210	210	210	210
S14	-600	-225	-200	-175	-150	-125	-100	0	250	250	250	250	250	220	220
S15	-625	-250	-225	-200	-175	-150	-125	-100	80	100	120	140	160	180	200

Figure 2: Payoff matrix for Knott and Miller's experiment A.

	S1	S2	S3	S4	S5	S6	S7	S8	S9	S10	S11	S12	S13	S14	S15	
S1	100	100	90	90	90	90	80	-100	-125	-150	-175	-200	-225	-250	-625	
S2	800	120	100	100	100	100	100	0	-100	-125	-150	-175	-200	-225	-600	
S3	700	400	125	110	110	110	110	0	-100	-125	-150	-175	-200	-200	-575	
S4	600	450	350	130	120	120	120	120	0	-100	-125	-150	-175	-175	-550	
S5	500	500	375	150	135	130	130	130	130	130	0	-100	-125	-150	-525	
S6	400	600	400	200	200	140	140	140	140	140	50	0	-100	-125	-500	
S7	350	800	425	210	210	150	150	150	150	150	100	50	0	-100	-475	
S8	200	300	450	250	220	160	160	160	160	160	160	120	50	0	-450	
S9	100	200	240	400	230	170	170	170	170	170	170	140	100	20	-400	
S10	90	100	200	200	300	200	180	180	180	180	180	160	150	40	-300	
S11	65	65	65	100	100	280	200	190	190	190	190	180	175	60	-200	
S12	60	60	60	60	80	100	250	200	200	200	200	200	200	200	80	-100
S13	55	55	55	55	55	55	210	220	210	210	210	210	210	100	0	
S14	50	50	50	50	50	50	50	50	220	220	220	220	220	220	220	
S15	40	40	40	40	40	40	40	60	80	100	120	140	160	180	200	

Figure 3: Payoff matrix for Knott and Miller's experiment B.

	S1	S2	S3	S4	S5	S6	S7	S8	S9	S10	S11	S12	S13	S14	S15
S1	100	75	85	85	90	90	90	90	90	90	90	90	90	90	90
S2	350	120	100	100	100	100	100	100	100	100	100	100	100	100	100
S3	300	140	125	115	115	110	110	110	110	110	110	110	110	110	110
S4	275	130	150	130	125	120	120	120	120	120	120	120	120	120	120
S5	250	120	125	145	135	130	130	130	130	130	130	130	130	130	130
S6	200	110	110	125	145	140	140	140	140	140	140	140	140	140	140
S7	100	100	100	100	100	155	150	300	350	250	150	150	150	150	150
S8	-350	0	50	50	75	100	125	250	400	450	450	160	160	160	160
S9	-400	-100	0	25	50	75	100	200	300	425	400	375	350	325	170
S10	-500	-125	-100	0	25	50	75	175	225	325	375	350	325	300	300
S11	-525	-150	-125	-100	0	25	50	125	200	225	350	325	300	275	250
S12	-550	-175	-150	-125	-100	0	25	100	175	200	200	250	250	250	225
S13	-575	-200	-175	-150	-125	-100	0	75	150	175	175	175	200	200	200
S14	-600	-225	-200	-175	-150	-125	-100	0	125	150	150	150	150	175	175
S15	-625	-250	-225	-200	-175	-150	-125	-100	0	100	120	140	140	150	150

Figure 4: Payoff matrix for Knott and Miller's experiment C.

to the Cournot rule. Similarly, Figure 12 shows the distribution obtained by running a similar simulation except that the agents are made to update using the fictitious play rule. The third figure in each appendix shows the actual observed proportions of play of each action in each period. For example Figure 13 shows the proportion of times that each of the 15 actions was played in each of the ten periods. The proportion here is taken over the 8 pairs (16 individuals) that played the game in experiment A. The fourth figure in each appendix is the result of our econometric analysis of each individual experiment which will be explained in section 4.

Notice that each game has a unique pure strategy Nash equilibrium: (S14,S14) for experiments A and experiment B and (S7,S7) for experiment C. Notice also that both the Cournot process and fictitious play converge to the pure strategy Nash equilibrium and that the convergence of fictitious play is much slower. The behavior of the subjects in the three experiment is very similar to fictitious play: subjects seem to converge to the pure strategy Nash equilibrium at a somewhat slower rate than fictitious play.

3 Games with multiple opponents

This section is based on some experiments that were run by Cooper, DeJong Forsythe, and Ross (1990). In each experiment, there are eleven players. Each player plays twice against each of the other players where the matchings in each round are determined at random. Agents do not know the identity of the player they are matched with and after each play find out which strategy the opponent selected. We analyze 6 of the experiments run by the authors. These experiments are labeled experiment 3 through 8, and the payoff matrices are depicted in Figures 5 through 10. The structure of appendices 3 through 8 is identical to that of appendices A through C, with the first three figures in each appendix depicting the simulated proportion under Cournot, the simulated proportion under fictitious play, and the observed proportions respectively. Notice in the third figure of each appendix that there are only three strategies; the fourth strategy was only mandated by the limitations of our graphics package.

In simulating the 1000 experiments for each of the hypotheses and each of the experiments, we actually drew 1000 ensembles of 11 agents and endowed each of them with an initial belief that is drawn uniformly over the unit simplex. We then replicated the exact matching scheme that occurred in each experiment keeping track of the beliefs of all 11 individuals in our simulated ensemble. For each experiment, then, we followed the full evolution of actions and beliefs for the 11 simulated individuals for a total of 1000 ensembles. The depicted simulated distributions for Cournot and fictitious play in appendices 3 through 8 are the empirical distribution over 1000 ensembles where the individuals were made to update according to the Cournot rules or the fictitious play rules as the figure states.

	S1	S2	S3
S1	350,350	350,250	1000,0
S2	250,350	550,550	0,0
S3	0,1000	0,0	600,600

Figure 5: Payoff matrix for Cooper's experiment 3.

All the games analyzed have two pure strategy Nash equilibria: (S1,S1)

	S1	S2	S2
S1	350,350	350,250	700,0
S2	250,350	550,550	0,0
S3	0,700	0,0	600,600

Figure 6: Payoff matrix for Cooper's experiment 4.

	S1	S2	S3
S1	350,350	350,250	700,0
S2	250,350	550,550	1000,0
S3	0,700	0,1000	600,600

Figure 7: Payoff matrix for Cooper's experiment 5.

	S1	S2	S3
S1	350,350	350,250	700,0
S2	250,350	550,550	650,0
S3	0,700	0,650	600,600

Figure 8: Payoff matrix for Cooper's experiment 6.

	S1	S2	S3
S1	350,350	350,250	700,0
S2	250,350	550,550	0,0
S3	0,700	0,0	500,500

Figure 9: Payoff matrix for Cooper's experiment 7.

	S1	S2	S3
S1	350,350	350,250	1000,0
S2	250,350	550,550	0,0
S3	0,1000	0,0	500,500

Figure 10: Payoff matrix for Cooper's experiment 8.

and (S2,S2). Suppose that n_i players select strategy S_i at a given period and suppose that players select strategies according to the Cournot process. Then all the players that were matched with players that selected strategies S1 and S3 will assign probability 1 that their next opponent selects strategy S1 and S3 and thus will select strategy S1. Similarly, all players that are matched with players selecting strategy S2 will select strategy S2. Thus in any period $t \geq 2$, $n_1 + n_3$ individuals select strategy S1 and n_2 individuals select strategy S2. Thus the only difference in the Cournot processes corresponding to different payoff matrices is the proportion of initial beliefs that lead the players to selecting strategy S1 and strategy S2.

In games 3,4,7,8 for almost all initial beliefs strategy S1 is a best response. Suppose an individual acts according to fictitious play. Suppose that an individual has the rare belief for which strategy S2 is a best response and is matched with another individual who adopts strategy S1. Then the second individual increases the probability he assigns to other individuals selecting strategy S2 but not sufficiently to change his strategy choice. On the other hand the first individual increases his belief that individuals select strategy S1 and thus changes his strategy choice to 1. Thus in experiments 3,4,5,7,8, if fictitious play holds, the proportion of individuals that selects strategy S1 converges to 1.

Game 5 is similar to games 3,4,7,8 excepts that for almost all beliefs an individual selects strategy S2 and thus the proportion of individuals that selects strategy S2 converges to 1.

In game 6 the proportion of initial beliefs for which strategy S1 and S2 are best responses are about equal. However strategy S1 is dominated by strategy S2 if strategy S3 is not played. Therefore when individuals behave according to fictitious play the proportion of individuals that selects strategy S2 converges to 1.

The results of experiment 3,4 and 5 are consistent with fictitious play and the Cournot process since in the first two experiments the proportion of the subjects that selects strategy S1 seems to converge to 1 and in the third experiment the proportion of the individuals selecting strategy S2 seems to converge to 1.

In game 6 the two models make very different predictions: the Cournot process predicts that strategy S1 and strategy S2 are equally likely while fictitious play predicts convergence to strategy S2. The experimental data strongly favors fictitious play since the proportion of the subjects selecting strategy S2 seems to converge to 1.

The results of experiment 7,8 are inconsistent with both fictitious play and the Cournot process since both models predict that almost all subjects should select strategy S1 while in fact most of the subjects select strategy S2.

4 Econometric Analysis

4.1 The simulation procedure

In sections 2 and 3, we described the simulation procedure we used to get approximations of the likelihood function. There are two major questions that we expect the reader to ask:

1. In both experiments, why, the reader may ask, do we draw initial beliefs of all of our robots from the uniform distribution over the unit simplex? Surely, the reader may add, the distribution of beliefs should depend on the actual payoff matrix that the players get to observe. We totally agree with that statement. The problem with incorporating this type of theoretical analysis of the payoff matrix, however, should be obvious. For indeed it is the very game-theoretic contemplation of the payoff matrices that would convince us that all players should have point mass beliefs that their opponents will play the Nash strategy, and nothing other than Nash strategies should ever be observed. That clearly takes us back full circle and does not add to our knowledge of the justifiability of the notion of Nash equilibrium on the basis of different learning algorithms.

2. One alternative approach that we chose not to follow is to put our method on its head and start with an assumption about one particular learning algorithm being the true one, and then proceeding to get an estimate of the initial distribution of beliefs. This is the procedure followed by McKelvey and Palfrey (1989). It is clear that for our purposes, we are trying to find a reasonable learning algorithm, and assuming that agents actually follow any particular algorithm will be no more justified than assuming that they all play Nash in the first place. Also, once we use our data to estimate the initial distribution of beliefs (under some parametrization of course), we cannot update our beliefs on the truth of our maintained learning hypothesis.
3. A related question to the previous one is why we do not condition our simulations on the actual observed data. In other words, we could have the robots respond by updating to the actual moves that the human subjects used. The answer is quite simple. We want, for the purposes of our Bayesian updating, to compute the likelihood function under the maintained hypothesis that initial beliefs are drawn at random, and under our two alternative models. We then compare our prior beliefs on the two models with the (theoretically) simulated likelihood function and the observed data, and do our updating. Using the data in computing the likelihood function will interact with our assumption on the distribution of initial priors (and hence with the likelihood function) in ways that we cannot account for and would constitute a form of "data-mining".

4.2 The updating procedure

Now we go back to the issue of Bayesian updating of the economist's belief over the two hypotheses and discuss the derivation of the fourth figure of each appendix and the overall result of our analysis depicted in Figure 1 ². Formally, let p_c^t be the experimenter's subjective probability at time t that individuals act according to the Cournot process. Let p_f^t be the experiment subjective probability at time t that individuals act according to fictitious

²The asymptotic aspects of the evolution of the economist's beliefs following that approach is discussed in more rigor and in more general contexts in El-Gamal and Sundaram (1989,1990).

play. Let q_c^t be the probability of observing the data at time t given that the individuals act according to the Cournot process. Let q_f^t be the probability of observing the data at time t given that the individuals act according to fictitious play. The experimenter updates his beliefs according to Bayes' rule. Then the posteriors, p_c^t, p_f^t ($t > 0$), are determined from the priors, p_c^0, p_f^0 and the observations $\{q_c^t, q_f^t\}$ in the manner described in the following lemma. An obvious corollary to this lemma is that the order with which the experiments are analyzed does not affect the belief of the experimenter.

Lemma 4.1 For all $t \in N$,

$$p_c^t = \frac{p_c^0 q_c^1 q_c^2 \cdots q_c^t}{p_c^0 q_c^1 q_c^2 \cdots q_c^t + p_f^0 q_f^1 q_f^2 \cdots q_f^t}$$

$$p_f^t = \frac{p_f^0 q_f^1 q_f^2 \cdots q_f^t}{p_c^0 q_c^1 q_c^2 \cdots q_c^t + p_f^0 q_f^1 q_f^2 \cdots q_f^t}$$

Proof: We will prove the lemma by induction. By Bayes' rule the equalities hold for $t = 1$. Suppose that the equalities hold for $t = k$; i.e.,

$$p_c^k = \frac{p_c^0 q_c^1 q_c^2 \cdots q_c^k}{p_c^0 q_c^1 q_c^2 \cdots q_c^k + p_f^0 q_f^1 q_f^2 \cdots q_f^k},$$

$$p_f^k = \frac{p_f^0 q_f^1 q_f^2 \cdots q_f^k}{p_c^0 q_c^1 q_c^2 \cdots q_c^k + p_f^0 q_f^1 q_f^2 \cdots q_f^k}.$$

By Bayes' rule,

$$p_c^{k+1} = \frac{p_c^k q_c^{k+1}}{p_c^k q_c^{k+1} + p_f^k q_f^{k+1}}.$$

Substitution the values for p_c^k and p_f^k in the previous expression we get:

$$p_c^{k+1} = \frac{\frac{p_c^0 q_c^1 q_c^2 \cdots q_c^k}{p_c^0 q_c^1 q_c^2 \cdots q_c^k + p_f^0 q_f^1 q_f^2 \cdots q_f^k} q_c^{k+1}}{\frac{p_c^0 q_c^1 q_c^2 \cdots q_c^k}{p_c^0 q_c^1 q_c^2 \cdots q_c^k + p_f^0 q_f^1 q_f^2 \cdots q_f^k} q_c^{k+1} + \frac{p_f^0 q_f^1 q_f^2 \cdots q_f^k}{p_c^0 q_c^1 q_c^2 \cdots q_c^k + p_f^0 q_f^1 q_f^2 \cdots q_f^k} q_f^{k+1}}$$

$$= \frac{p_c^0 q_c^1 q_c^2 \cdots q_c^{k+1}}{p_c^0 q_c^1 q_c^2 \cdots q_c^{k+1} + p_f^0 q_f^1 q_f^2 \cdots q_f^{k+1}}.$$

Thus the equalities hold for $t = k + 1$ and thus by induction for all integers t . ■

4.3 Smearing the hypotheses

If we were to put all of the probabilistic mass in our prior on rational agent theories, then a single observation which is inconsistent with all the theories in the support of the prior will make Bayesian updating impossible (it will produce a zero numerator and denominator in the Bayesian updating formula). In the Cooper et al's experiments there are observations of agents playing strictly dominated actions which cannot be justified under any beliefs. There also observations of actions that, although not strictly dominated, cannot be justified under either of our learning algorithms. To make Bayesian updating applicable, we include in our models a small probability of a completely random action taking place. Specifically, we shall let each model predict actions based on the updating rule in question with probability $(1 - \epsilon)$ and using a uniform (completely arbitrary) decision rule with probability ϵ . We shall then examine the predictive power of the two models for different values of ϵ .

The argument for such a procedure (other than the fact that it allows us to do Bayesian updating) is that people early on in the experiment do not understand the experiment fully, in particular if they do not wish to spend the effort (the payoffs are not particularly high: a payoff of 600, for instance, gives the subject a lottery with a probability .6 of winning one dollar). Cooper et al (1990) show that with time there is a statistically significant decrease in the number of times a dominated strategy is selected. Furthermore, individuals have different mental costs in figuring out the experiment. Thus there is a strong argument for a different error rate for each data point. The problem with this approach is that the error component then explains all the experimental data, and thus does not allow the experimenter to differentiate between the different theories. Thus we opted to select a single ϵ for all out agents at all time periods. As the sample size gets large, the average 'correct' ϵ will get small. Thus we will be interested in the validity of the model as ϵ goes to zero. Alternatively, the value for ϵ can be taken to be the average number of mistakes over the experiment.

The two models we compare are then the following:

$$\text{Model 1: } a' = \begin{cases} a_c & \text{with probability } 1 - \epsilon \\ a & \text{with probability } \frac{\epsilon}{A}, \end{cases}$$

$$\text{Model 2: } a' = \begin{cases} a_f & \text{with probability } 1 - \epsilon \\ a & \text{with probability } \frac{\epsilon}{A}, \end{cases}$$

where a' is the action chosen by the experimental subjects, a_c is the action that maximizes their expected payoff given their beliefs that are updated according to the Cournot procedure, a_f is the same as a_c with fictitious play updating, and $a \in \{1, 2, \dots, A\}$ is any strategy.

The simulation results depicted in the previous sections give us Monte-Carlo approximations to the probability in each period of the experiment that a particular action will be played. Given those probabilities and the actual observed actions of the individuals in each period, we can compute $Pr(\text{data}|\text{Model } i)$ and then proceed with Bayesian updating as described above. The computations we use for $Pr(\text{data}|\text{Model } i)$ will be reduced form in the following sense. We shall treat the actual action of an agent in a period as a data point which is assumed to be a random draw according to the probability distribution obtained from the simulated experiments with probability $(1 - \epsilon)$ and as a uniform random draw with probability ϵ . Let the observed data points be indexed by $(n, t) \in \{1, 2, \dots, N\} \times \{1, 2, \dots, T\}$, let the actions available to each individual n in period t be $a \in \{1, 2, \dots, A\}$ ³, and let the simulated probability under the pure version of *Model* i of action a in period t be $q_{a,t}^i$. Then

$$Pr(\text{data}|\text{Model } i) = \prod_{n=1}^N \prod_{t=1}^T \left((1 - \epsilon)q_{a_{n,t}}^i + \epsilon \frac{1}{A} \right)$$

where $a_{n,t}$ is the actual action chosen by agent n in period t .

$$\text{Bad} = \{(n, t) \in \{1, \dots, N\} \times \{1, \dots, T\} | q_{1,a_{n,t}}^i = 0 \text{ and } q_{2,a_{n,t}}^i = 0\}$$

and let *Good* be the set of points that can be explained by at least one of the theories; i.e.,

$$\text{Good} = (\{1, \dots, N\} \times \{1, \dots, T\}) \setminus \text{Bad}.$$

Lemma 4.2 *For all $\epsilon > 0$ Bayesian updating ignores all points that are incompatible with both models.*

³In general, A may depend on n and t , but since it does not in our experiments, we do not need to include this complication.

Proof: Notice that

$$Pr(data|Model i) = \left(\frac{\epsilon}{A}\right)^{\#(Bad)} \times \prod_{(n,t) \in Good} \left((1 - \epsilon)q_{i,a_n,t}^t + \frac{\epsilon}{A} \right)$$

Then substituting in the Bayesian updating rule above, we can cancel the common factor $\left(\frac{\epsilon}{A}\right)^{\#(Bad)}$ from the numerator and denominator. Thus Bayesian updating ignores these points; i.e.,

$$p_1^{t+1} = \frac{\left(\prod_{Good} (1 - \epsilon)q_{1,a_n,t}^t + \frac{\epsilon}{A} \right) \cdot p_1^t}{\sum_{i=1}^2 \left(\prod_{Good} (1 - \epsilon)q_{i,a_n,t}^t + \frac{\epsilon}{A} \right) \cdot p_i^t}$$

■

The fourth figure in each of the nine appendices shows the posterior using the likelihood function for that experiment alone as a function of the economist's prior and the level of smearing ϵ . As we discussed in the previous section the behavior of the subjects in the experiments ran by Knott and Miller is very similar to the one predicted by fictitious play. Notice that for reasonable values of ϵ (i.e., less than sixty percent of the actions are caused by random error) and as long as the experimenter's priors put positive weight on fictitious play, then after seeing the outcome of any of the Knott and Miller's experiments the experimenter believes with probability very close to 1 that subjects act according to fictitious play.

In Cooper et al's experiments 3 the proportion of subjects that selects strategy S1 converges to 1 but at a much slower rate than fictitious play. Thus for reasonable values of ϵ the experimenter's posteriors give probability very close to 1 to the event that subjects act according to the Cournot process. In experiment 4 subjects converge to strategy S1 at a faster rate; thus for small values of ϵ (such as $\epsilon = 0.1$) the experimenter posteriors put probability very close to 1 on the Cournot process. For high values of ϵ (such as $\epsilon = 0.4$) the experimenter's posteriors will assign almost all the weight to fictitious play. Finally, intermediate values of ϵ will give posteriors that put positive weight on both theories.

In Cooper et al experiment 5 the proportion of the population that adopts strategy S2 converges very quickly to 1. The convergence is not quite as fast as the Cournot process and thus for any reasonable prior and reasonable value of ϵ the posteriors put probability very close to 1 on fictitious play. Cooper et

al's experiment 6 is by far the most interesting example since fictitious play and the Cournot dynamics are totally different. The behavior of the subject is very similar to fictitious play and thus not surprisingly the posteriors of the experimenter put probability very close to 1 on fictitious play.

In Cooper et al's experiments 7 and 8 the Cournot process and fictitious play indicate that almost all subjects should adopt strategy S1. In fact almost all subjects adopt strategy S2. The posteriors after observing these experiments put probability very close to 1 on the Cournot process; however both models are clearly inadequate for this experiment. Overall, these results lead an experimenter to believe that fictitious play describes the data better than the Cournot process and this observation is represented in Figure 1.

References

- Banks, Jeff and J. Sobel (1987), "Equilibrium Selection in Signaling Game." *Econometrica*, 5, 647-663.
- Brock, William, R. Marimon, J. Rust, T. Sargent (1988). "Informationally Decentralized Learning Algorithms for Finite-Player, Finite-Action Games of Incomplete Information." University of Minnesota, mimeo.
- Brown, G. (1951). "Iterative Solutions of Games by Fictitious Play," in Koopmans, T (ed), *Activity Analysis of Production and Allocation*, Cowles Commission Monograph 13, John Wiley & Sons, New York, 377-380.
- Cho, In-Koo and D. Kreps (1987), "Signaling Games and Stable Equilibria." *Quarterly Journal of Economics*, 102, 179-221.
- Cooper, Russell, D. DeJong, R. Forsythe, and T. Ross (1990), "Selection Criteria in Coordination Games: Some Experimental Results," *American Economic Review*, 80, 218-234.
- El-Gamal, Mahmoud and R. Sundaram (1990), " \aleph_0 Experiments in Search of a Theory", work in progress.
- El-Gamal, Mahmoud and R. Sundaram (1989), "Bayesian Economists ... Bayesian Agents I: An Alternative Approach to Optimal Learning", Social Sc. Working Paper # 705, Caltech.
- Grossman, Sandy and M. Perry (1987), "Perfect Sequential Equilibrium." *Journal of Economic Theory*, 39, 97-119.
- Knott, Jack and G. Miller (1987), "The Dynamics of Convergence and Cooperation: Rational Choice and Adaptive Incrementalism." Washington University, mimeo.
- Robinson, Julia (1951), "An Iterative Method for Solving a Game," *Annals of Mathematics*, 54, 296-301.

This page intentionally left blank

Appendix A

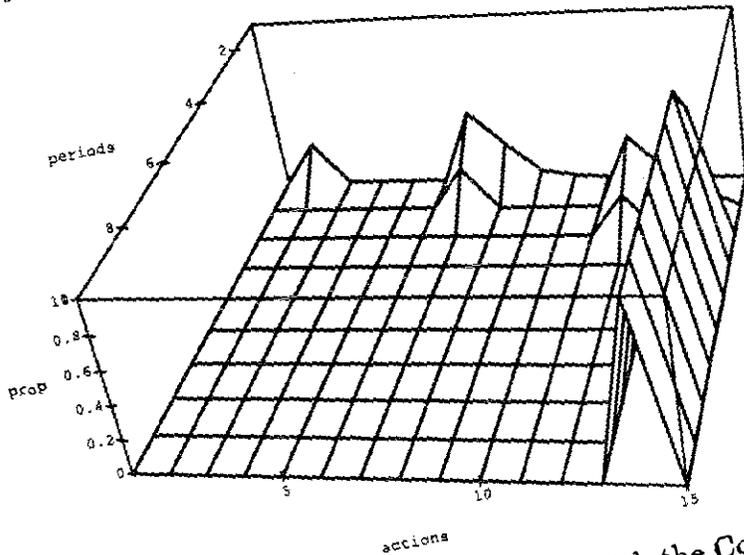


Figure 11: Simulation of Knott and Miller's experiment A with the Cournot process.

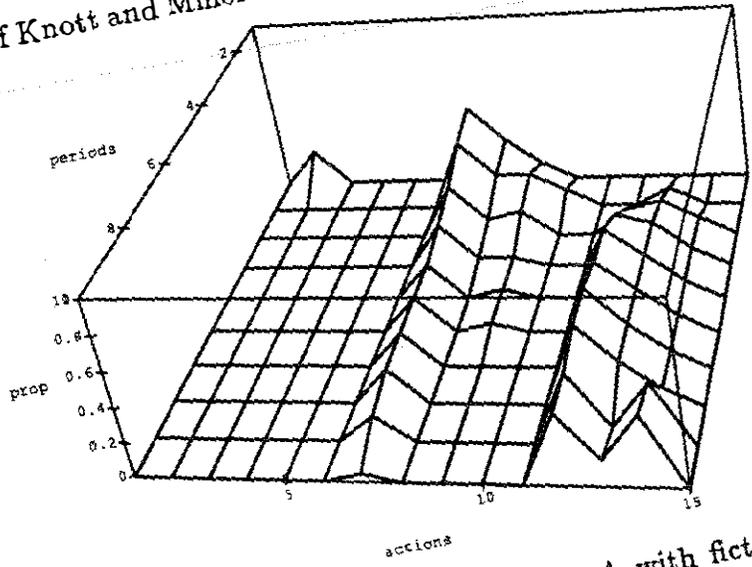


Figure 12: Simulation of Knott and Miller's experiment A with fictitious play.

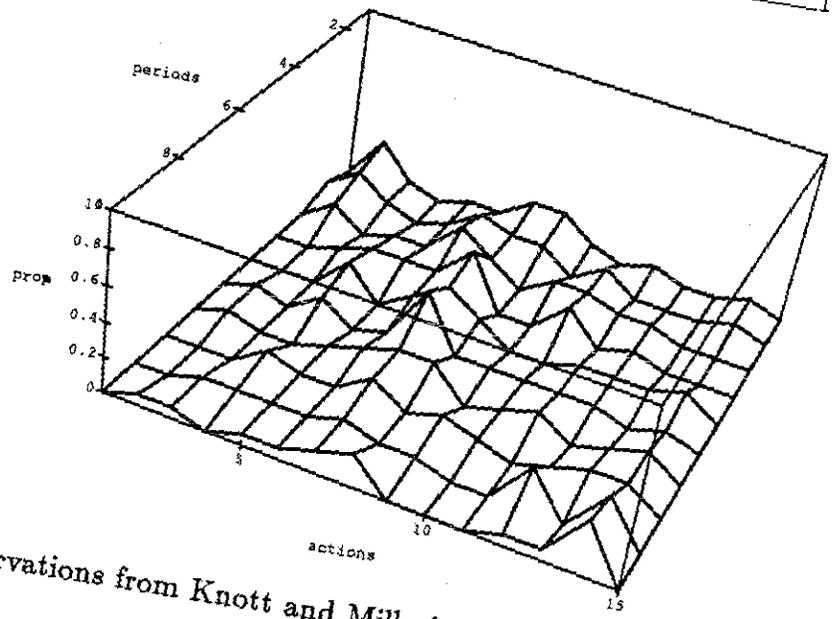


Figure 13: Observations from Knott and Miller's experiment A.

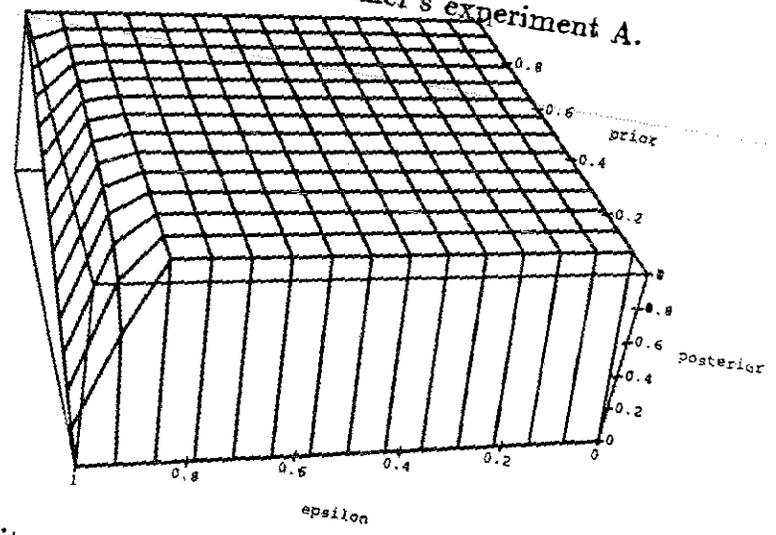


Figure 14: Posterior probability of agents playing according to fictitious play in Knott and Miller's experiment A.

Appendix B

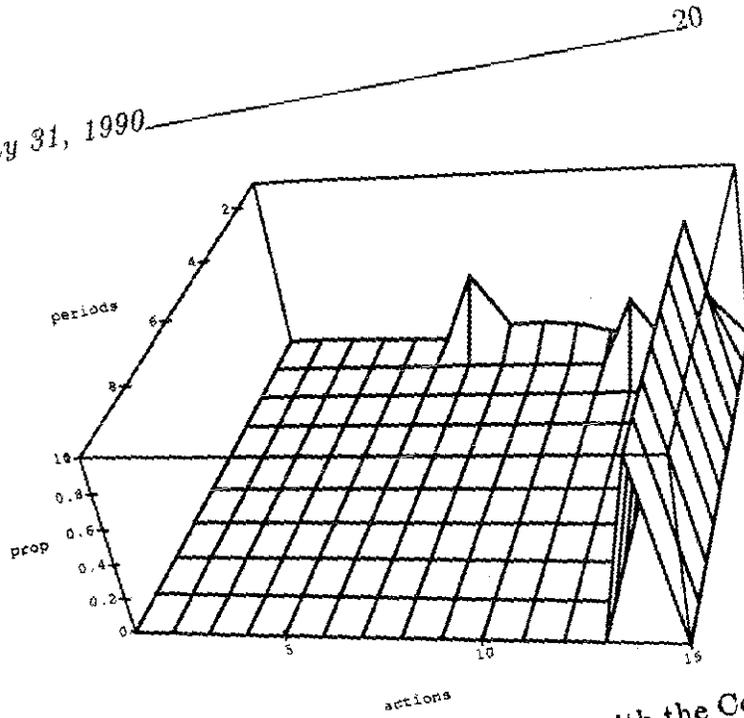


Figure 15: Simulation of Knott and Miller's experiment B with the Cournot process.

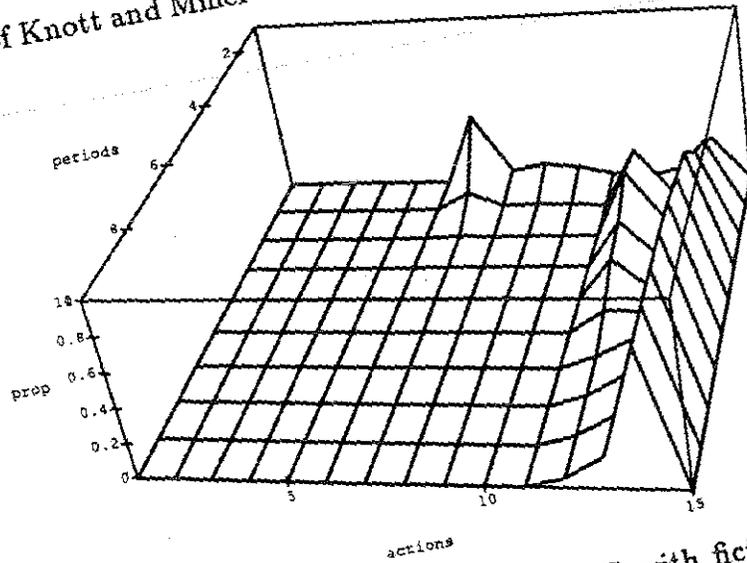


Figure 16: Simulation of Knott and Miller's experiment B with fictitious play.

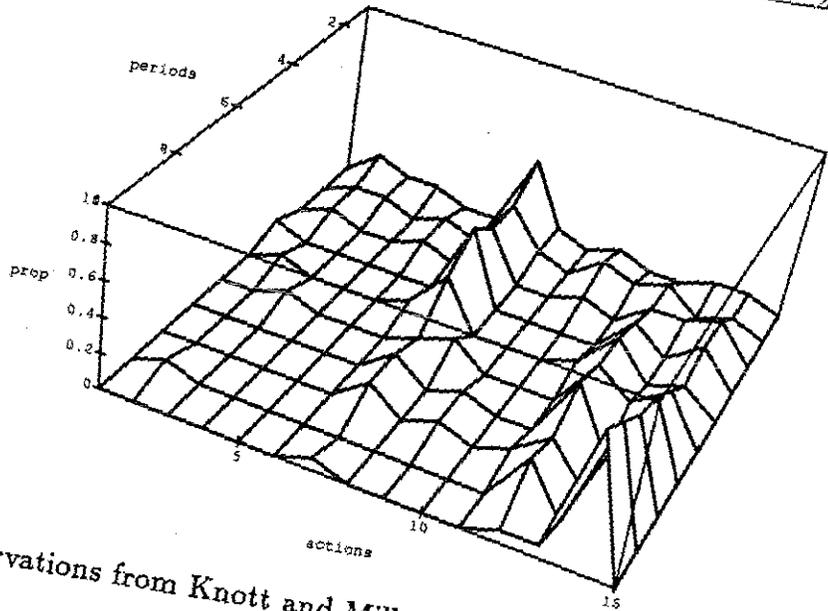


Figure 17: Observations from Knott and Miller's experiment B.

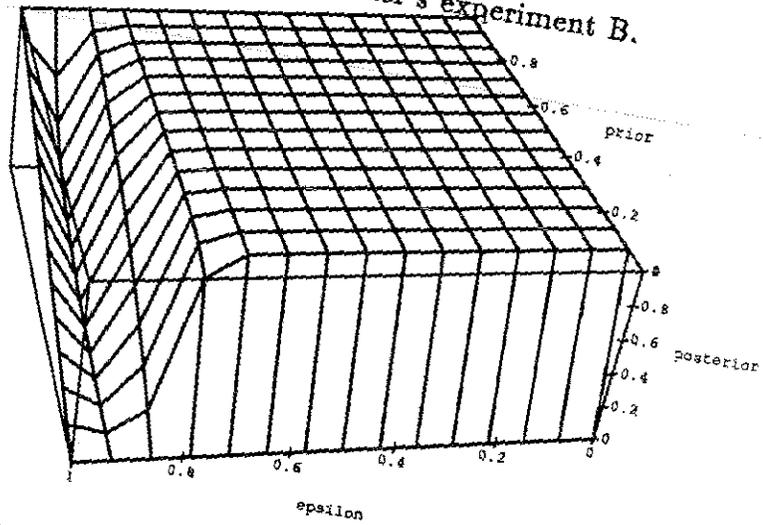


Figure 18: Posterior probability of agents playing according to fictitious play in Knott and Miller's experiment B.

Appendix C

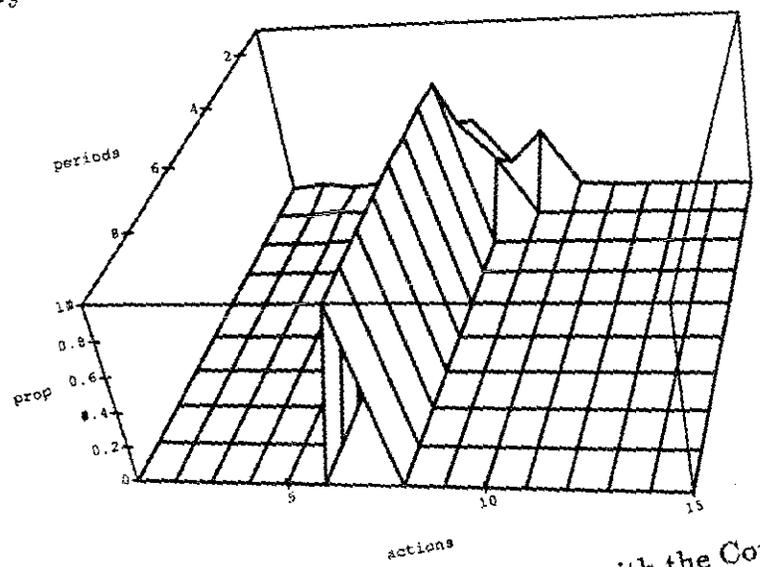


Figure 19: Simulation of Knott and Miller's experiment C with the Cournot process.

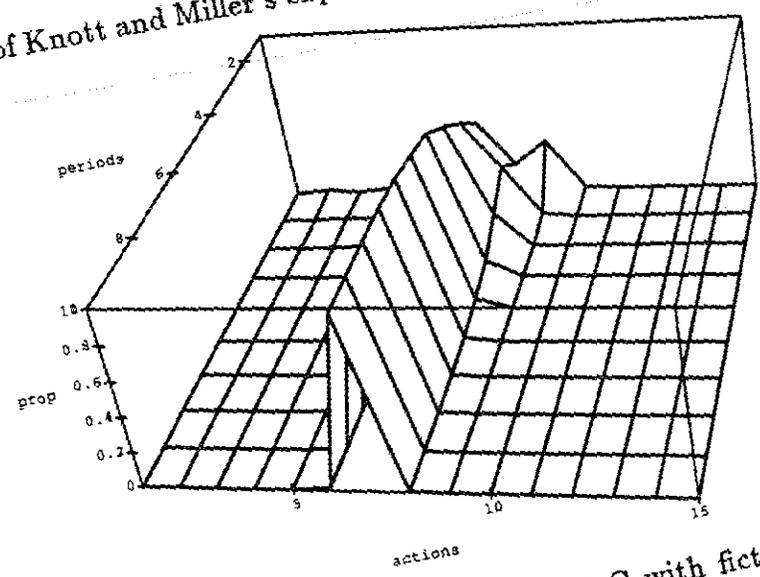


Figure 20: Simulation of Knott and Miller's experiment C with fictitious play.

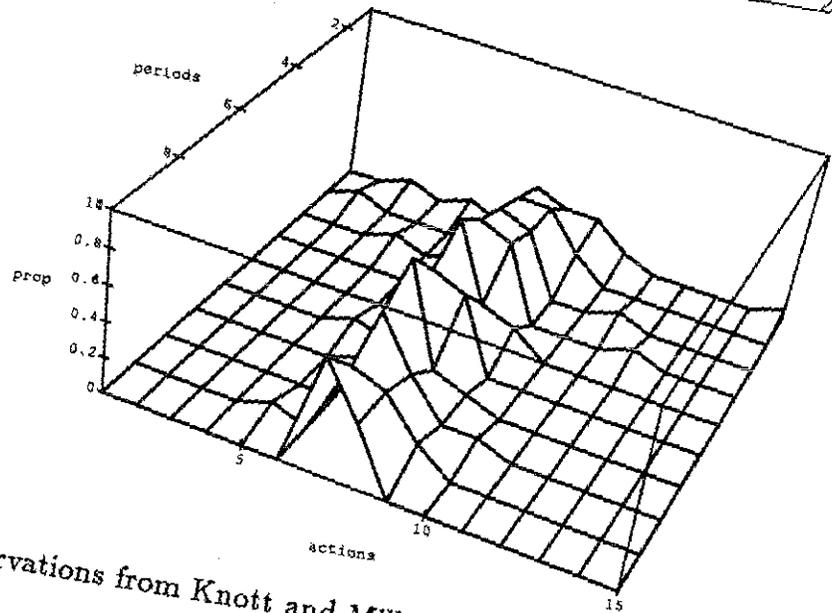


Figure 21: Observations from Knott and Miller's experiment C.

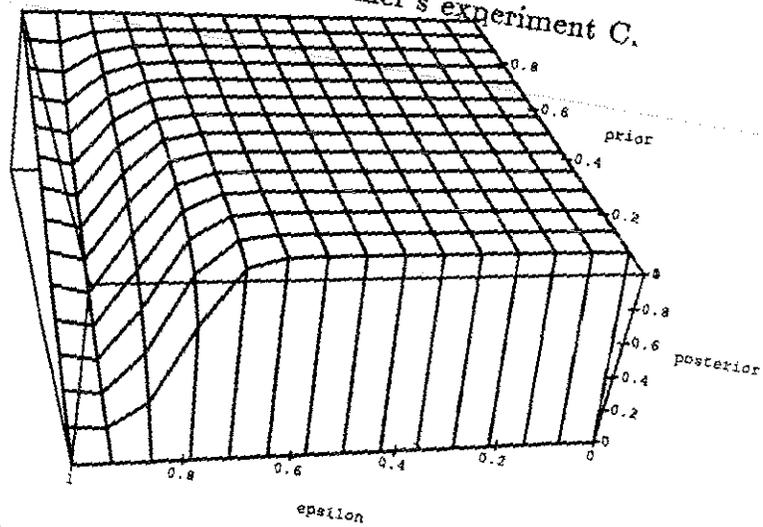


Figure 22: Posterior probability of agents playing according to fictitious play in Knott and Miller's experiment C.

Appendix 3

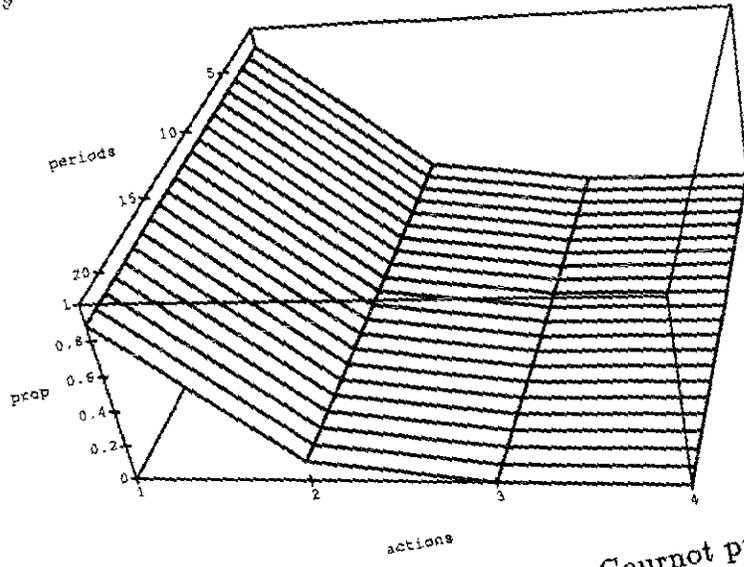


Figure 23: Simulation of Cooper's experiment 3 with the Cournot process.

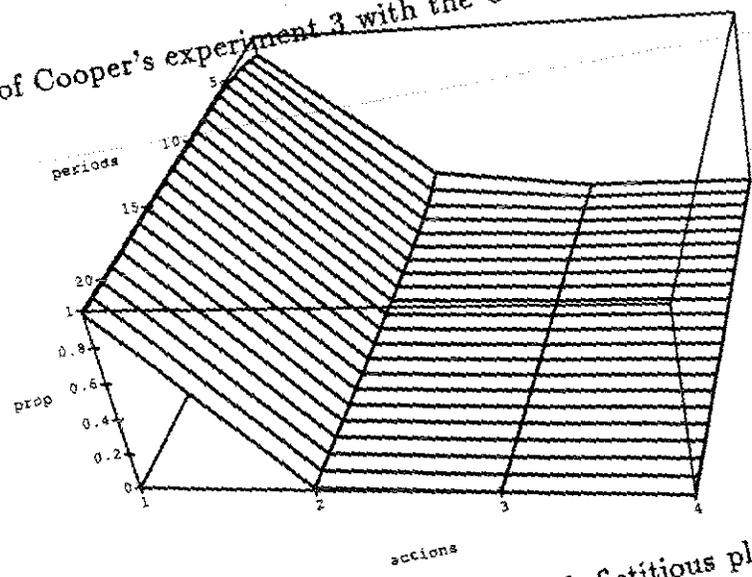


Figure 24: Simulation of Cooper's experiment 3 with fictitious play.

Appendix 3

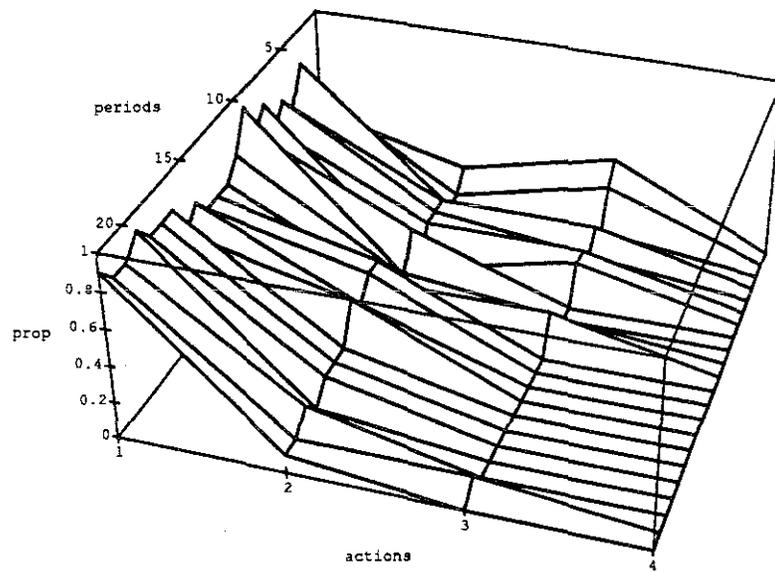


Figure 25: Observations in Cooper's experiment 3.

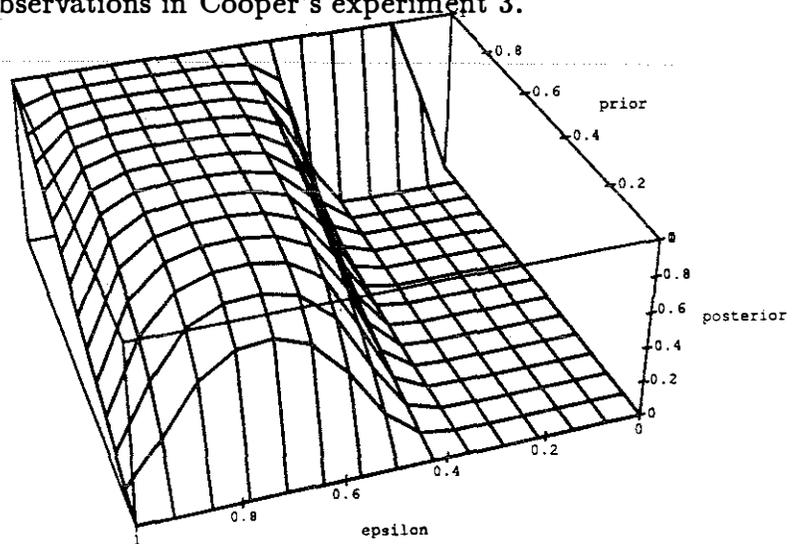


Figure 26: Posterior probability of agents playing according to fictitious play in Cooper's experiment 3.

Appendix 4

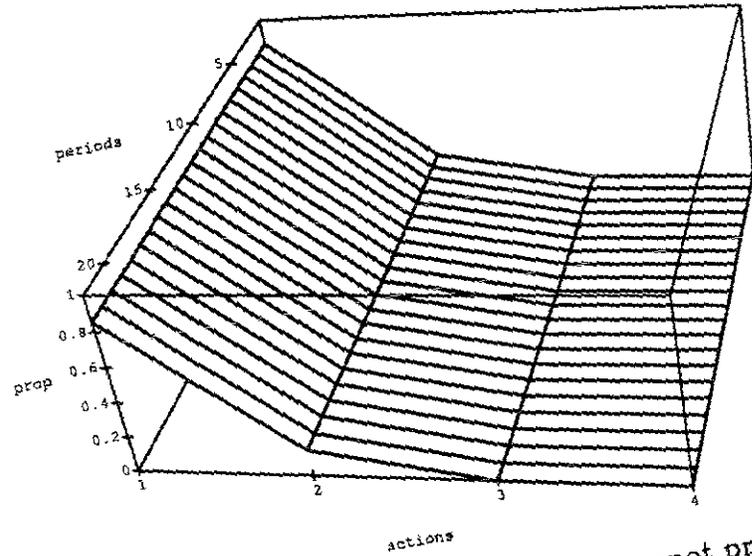


Figure 27: Simulation of Cooper's experiment 4 with the Cournot process.

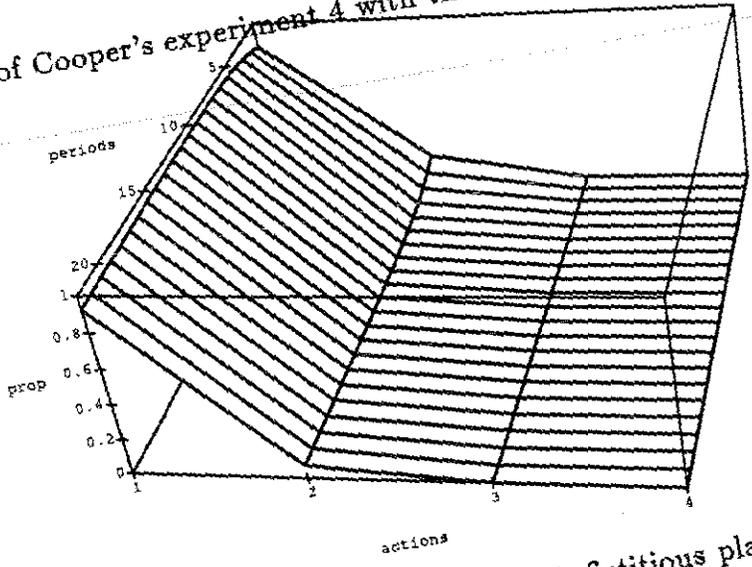


Figure 28: Simulation of Cooper's experiment 4 with fictitious play.

Appendix 4

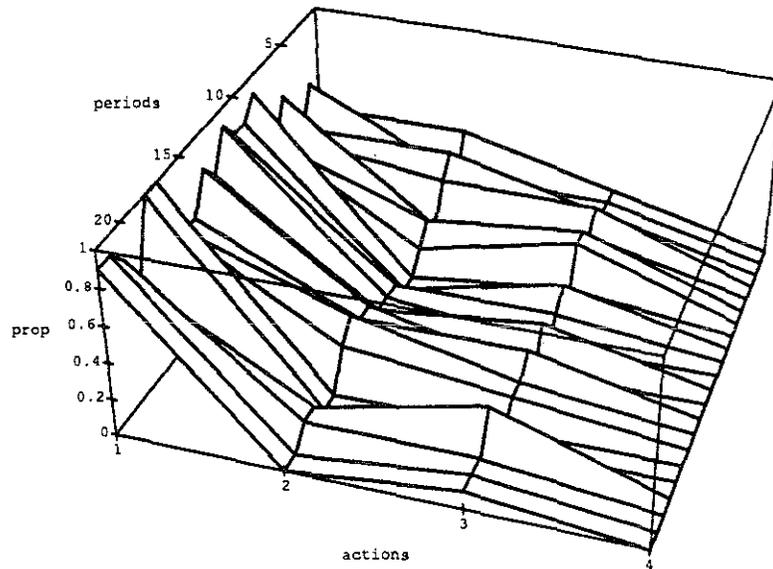


Figure 29: Observations in Cooper's experiment 4.

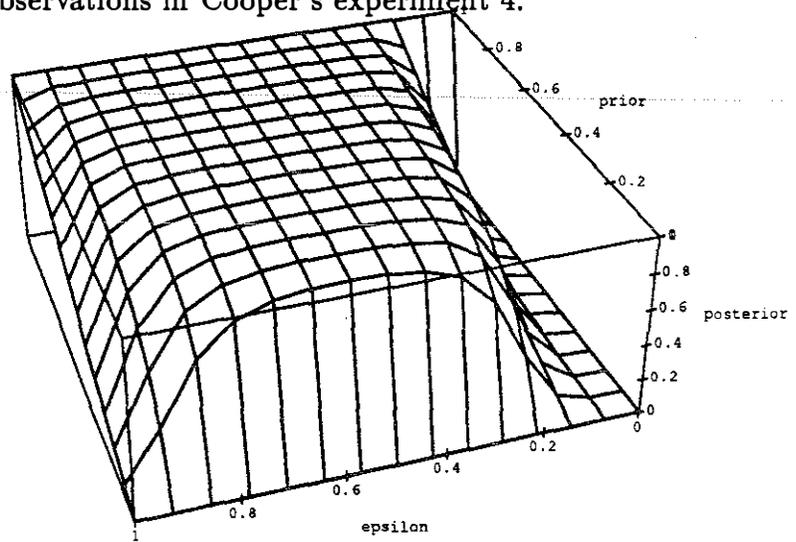


Figure 30: Posterior probability of agents playing according to fictitious play in Cooper's experiment 4.

Boylan & El-Gamal, May 31, 1990

Appendix 5

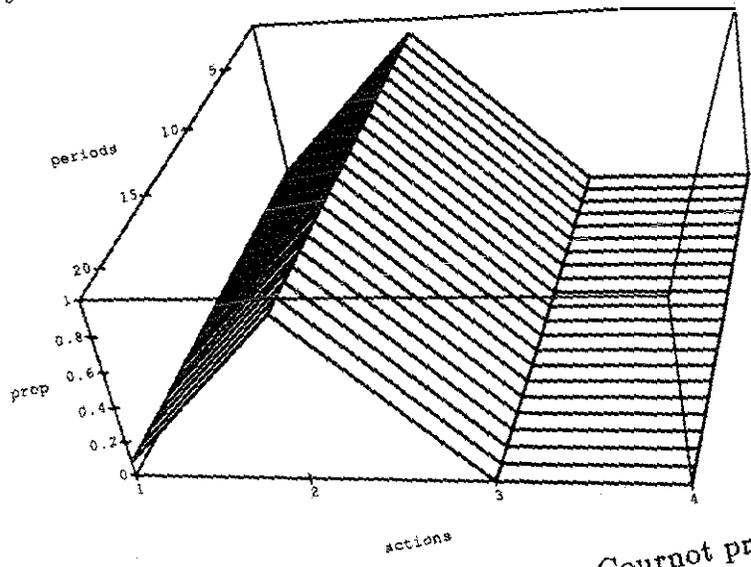


Figure 31: Simulation of Cooper's experiment 5 with the Cournot process.

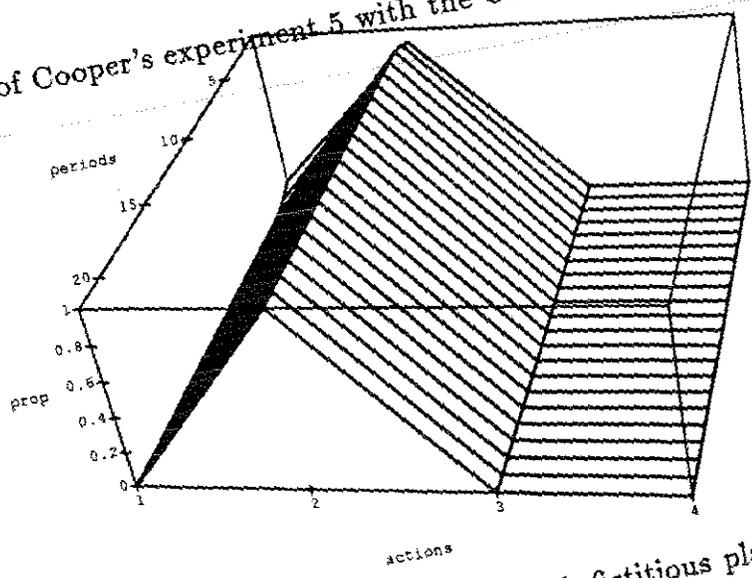


Figure 32: Simulation of Cooper's experiment 5 with fictitious play.

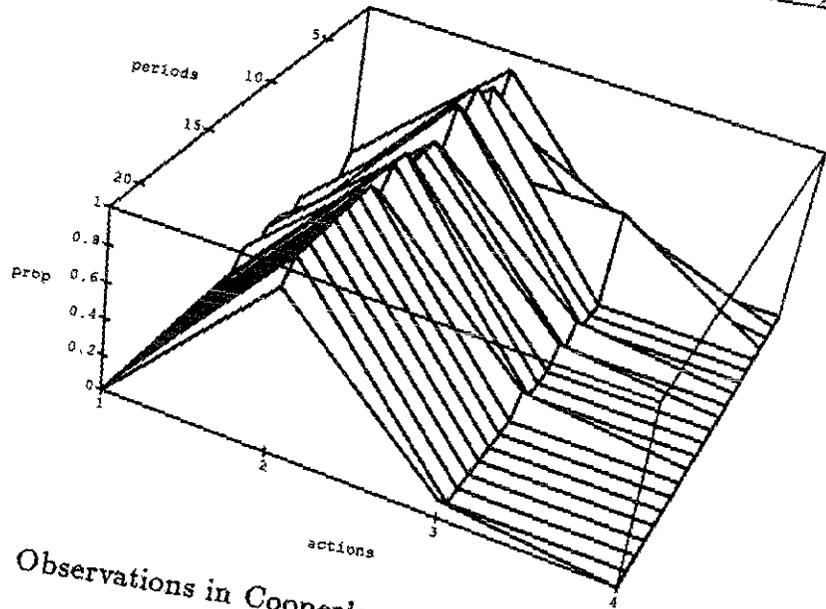


Figure 33: Observations in Cooper's experiment 5.

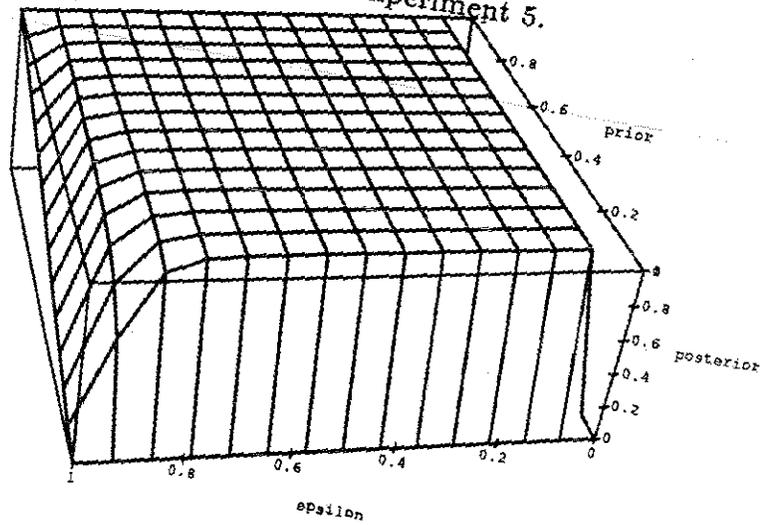


Figure 34: Posterior probability of agents playing according to fictitious play in Cooper's experiment 5.

Appendix 6

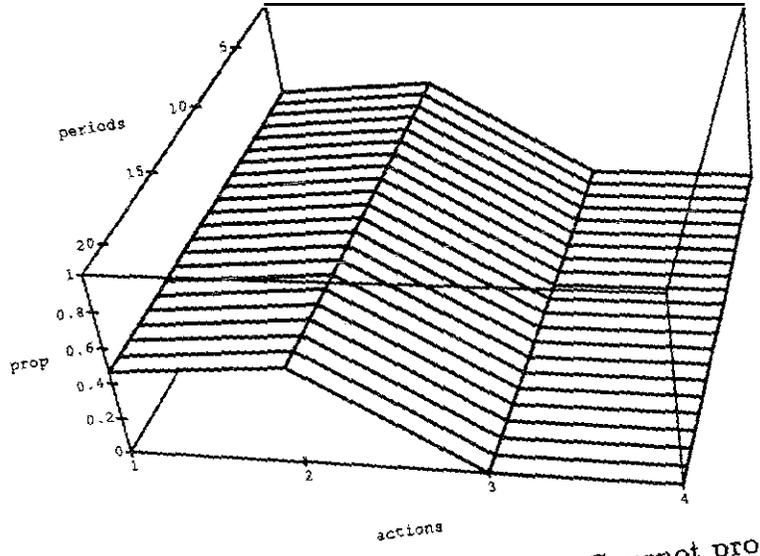


Figure 35: Simulation of Cooper's experiment 6 with the Cournot process.

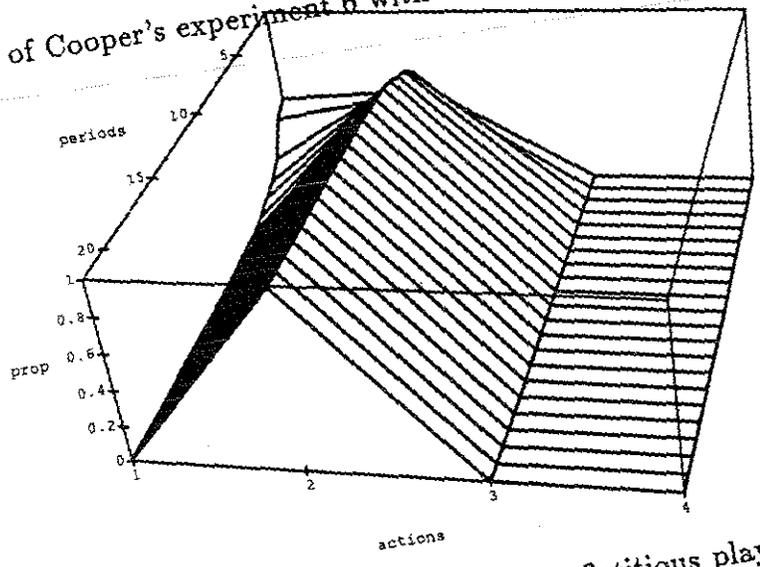


Figure 36: Simulation of Cooper's experiment 6 with fictitious play.

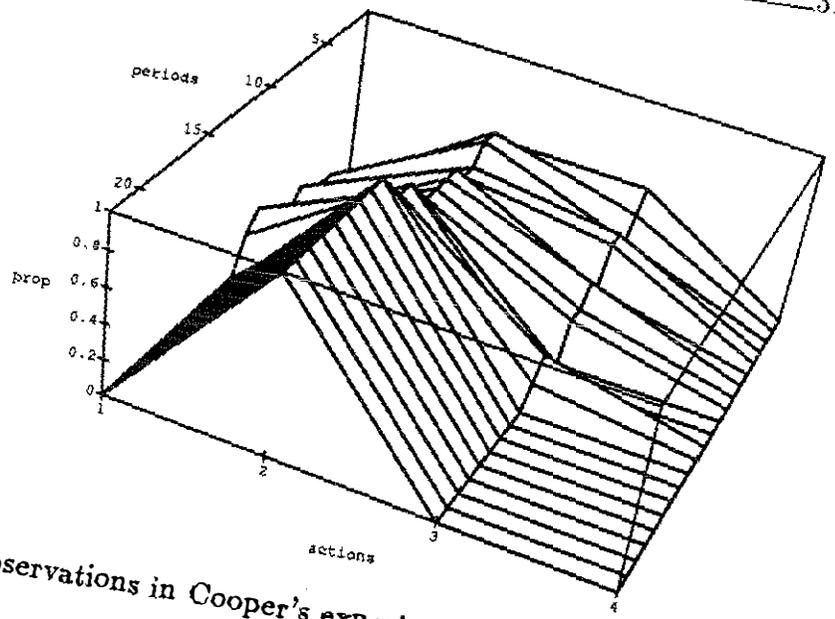


Figure 37: Observations in Cooper's experiment 6, Process.

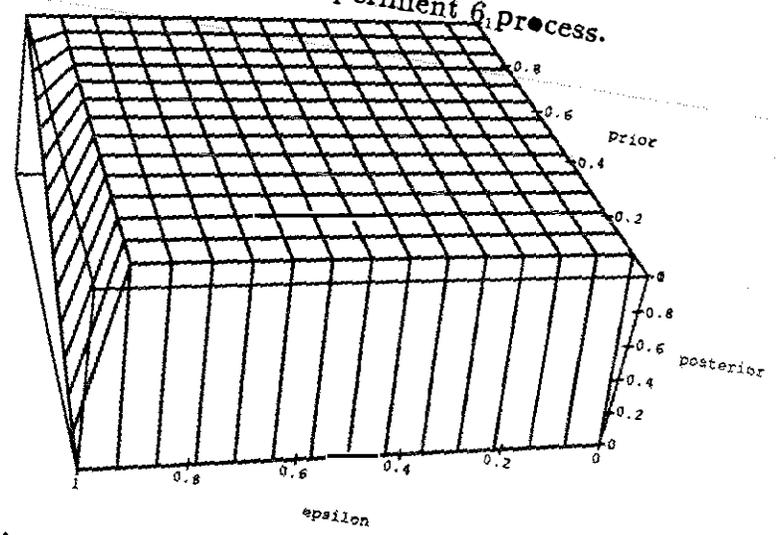


Figure 38: Posterior probability of agents playing according to fictitious play in Cooper's experiment 6.

Boylan & El-Gamal, May 31, 1990

Appendix 7

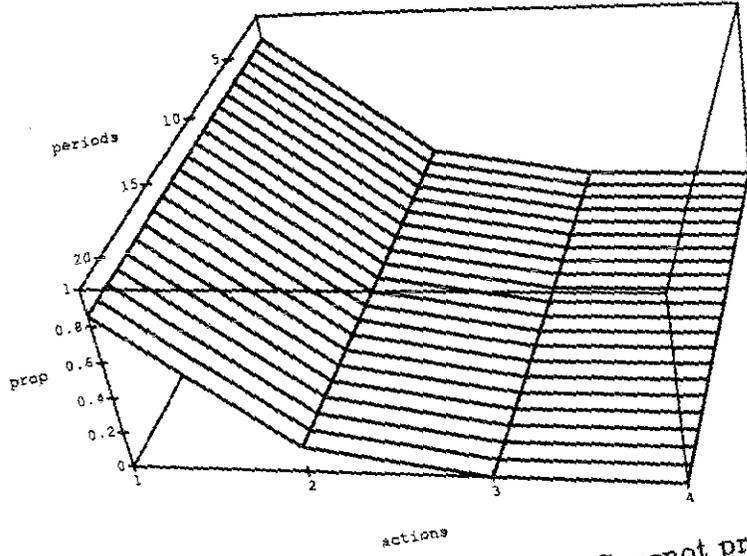


Figure 39: Simulation of Cooper's experiment 7 with the Cournot process.

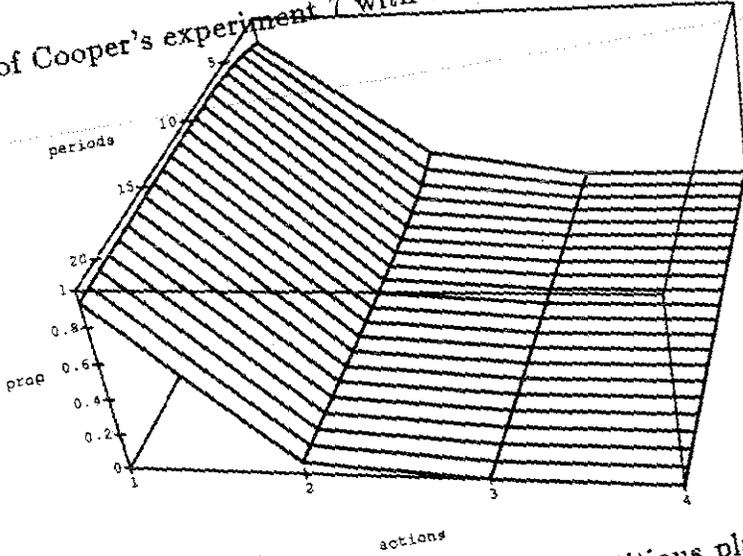


Figure 40: Simulation of Cooper's experiment 7 with fictitious play.

Appendix 7

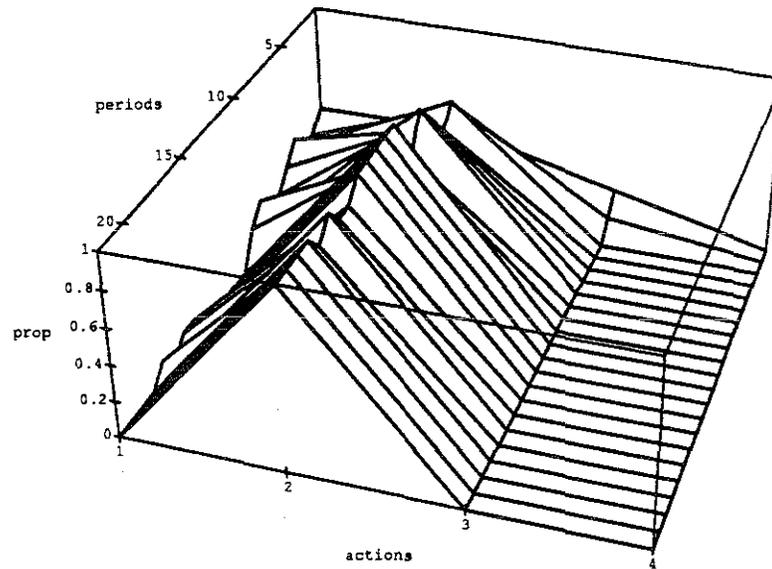


Figure 41: Observations in Cooper's experiment 7₁ process.

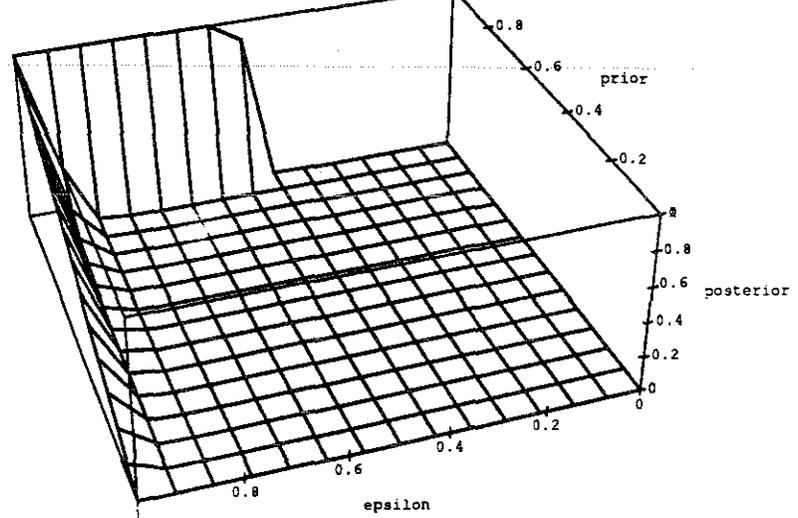


Figure 42: Posterior probability of agents playing according to fictitious play in Cooper's experiment 7.

Appendix 8

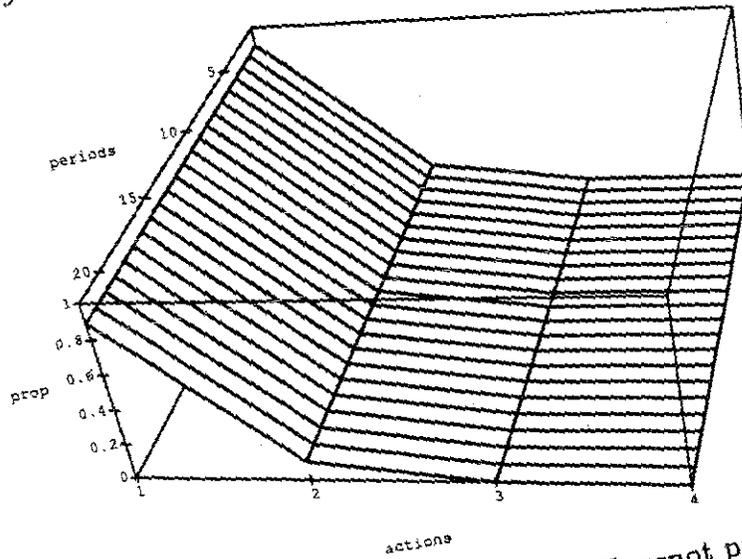


Figure 43: Simulation of Cooper's experiment 8 with the Cournot process.

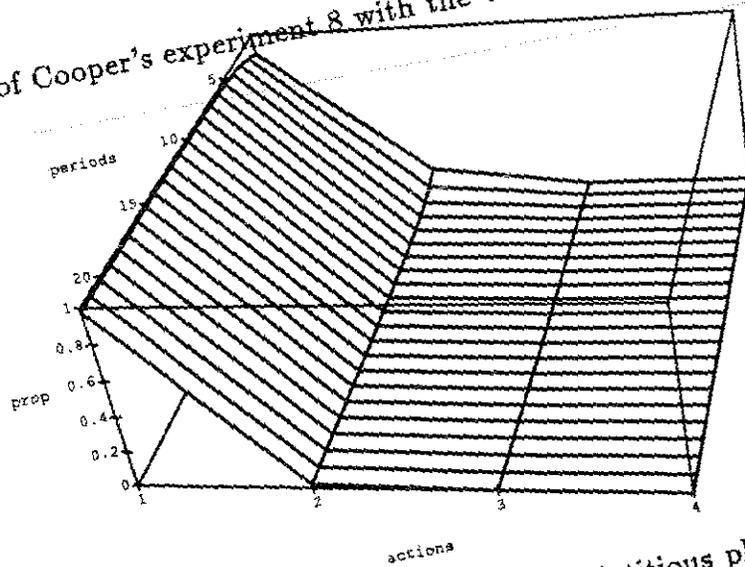


Figure 44: Simulation of Cooper's experiment 8 with fictitious play.

Appendix 8

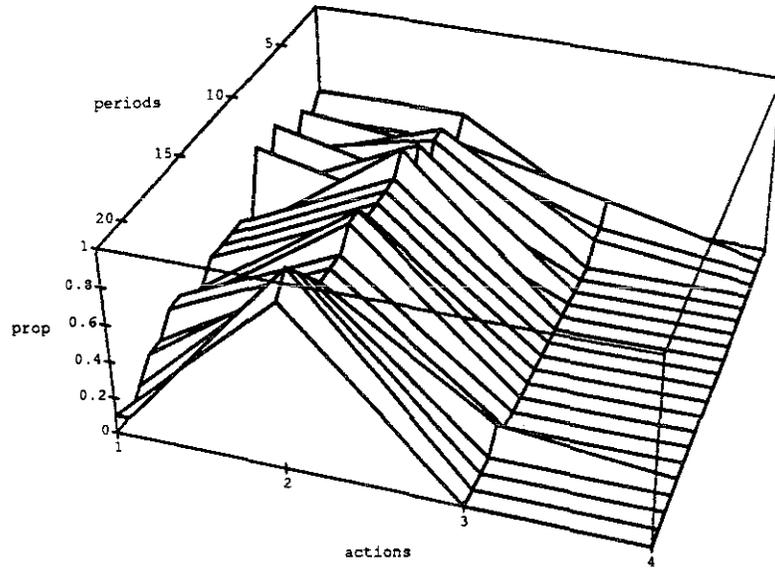


Figure 45: Observations in Cooper's experiment 8₁ process.

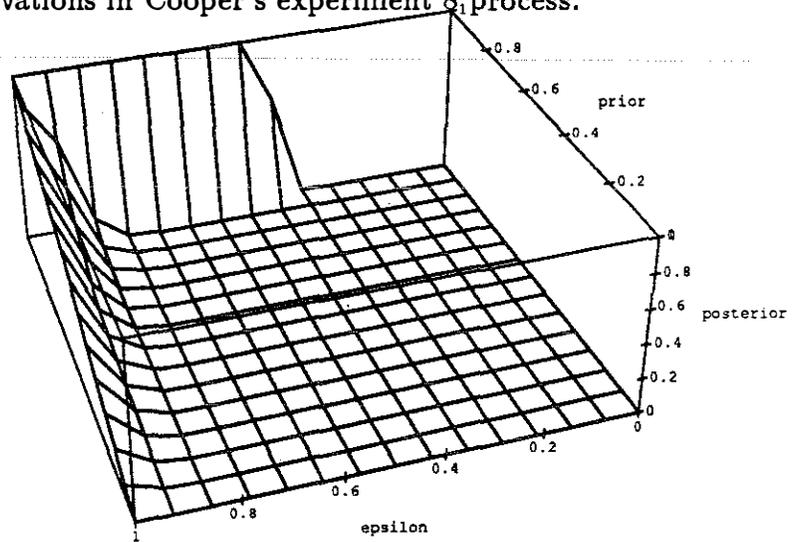


Figure 46: Posterior probability of agents playing according to fictitious play in Cooper's experiment 8.