Disequilibrium Game Theory

by

William Robert Majure

Submitted to the Department of Economics
in partial fulfillment of the requirements for the degree of

Doctor of Philosophy

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

September 1994

Author ....
Department of Economics
June 16, 1994

Certified by ....
Glenn Ellison
Assistant Professor of Economics
Thesis Supervisor

Accepted by ....
Olivier Blanchard
Chair, Departmental Committee on Graduate Students

Disequilibrium Game Theory

by

William Robert Majure

Submitted to the Department of Economics
on June 16, 1994, in partial fulfillment of the
requirements for the degree of
Doctor of Philosophy

Abstract

This thesis is comprised of three essays that all take different approaches to the subject of disequilibrium in games. The theory of disequilibrium in games is not really new, but has gained a lot of attention recently in the literature on learning and evolution. Chapter 1 exemplifies the use of disequilibrium for estimating the structural form of an Industrial Organization model. Chapter 2 demonstrates the use of experimental data to select among disequilibrium models and to establish benchmarks of behavior. Chapter 3 illustrates the use of numerical simulation to explore the convergence behavior of more complex models and to derive comparative statics.

In chapter 1 we develop a simple model of oligopoly with multi-product, or multi-location, firms. We show that the variable "number of products" (or locations) can be either a strategic substitute or a strategic complement. Implications for comparative statics are drawn. The model is estimated with data from the Portuguese banking industry and with a partial adjustment version of the "cobwebbing" mechanism. The results provide an explanation for the puzzling pattern of branching during the late 80's.

In chapter 2 individual data from an experimental setting is used to estimate a learning model and an evolution model. These models characterize the behavior of boundedly rational individuals in a setting of repeated interaction that mirrors the experimental environment. The models represent two somewhat distinct approaches by game theorists to relaxing the perfect rationality of Nash equilibrium. The estimation reveals that both models are significant determinants of behavior. Moreover, the strengths of each model complement the weaknesses of the other. A mixed model does not reject either element. Some support is also found for focusing attention on convergence behavior and for modelling players in the most heterogeneous fashion possible. The learning model is found to be more sensitive to aggregation issues than evolution, though learning is a better description of individual behavior.

This paper establishes a technique for estimating the probability of each equilibrium for a broad class of static games. The procedure is to simulate learning processes to generate a database of convergence behavior and then estimate the probability of a given outcome as a function of the game played. This prediction will approximate reality to the extent that the theoretical learning process aproximates real players. The predictions are also a useful way to characterize the behavior of the learning process both as direct information about the process and as a basis for comparative statics. In the class of 2x2 symmetric coordination games, those games with extreme risk dominance are almost sure to eventually have the risk dominant equilibrium played by everyone. The rest of the games have non-zero probability of convergence to either equilibrium. Weighting more recent observations more heavily in the calculation of beliefs about others' play did not affect the expected action converged to, but did increase the speed of convergence. Adding a cost of updating caused a dramatic shift in the action converged to for some games, but only for some distributions of the players' propensity to update their behavior after one non-equilibrium outcome relative to another. Making fictitious play more like an evolutionary model by changing the way that players respond to their information from adopting the best response to emulating other players slowed the process down dramatically.

Thesis Supervisor: Glenn Ellison
Title: Assistant Professor of Economics

For Kim

Acknowledgments

I thank the National Science Foundation for their graduate fellowship which paid for part of my graduate education. Chapter 1 is a variant of joint work with Luis Cabral that began while we were both visiting the Banco de Portugal. I am grateful to the research staff at the bank for their help in assembling the data, especially to Margarida Catalão Lopes and to Luis for allowing me to use our work in my thesis. I am also grateful to Douglas DeJong for generously supplying the data for Chapter 2.

Many people at M.I.T. have given me helpful advice and friendly support. The ones that stand out particularly are Amy Ando, Guy Debelle, Peter Klibanoff, Matt Slaughter, Koleman Strumpf, Stacey Tevlin and Beth Anne Wilson. There have also been a lot of professors who helped and influenced me. Glenn Ellison, Frank Fisher, Drew Fudenburg, Matt Rabin and Dick Schmalensee went further out of their way for me than I had any right to expect.

Mostly, though, I thank my wife Kim for always pushing me to aim higher, for helping me through the rough times, and for understanding when I wasn't able to do the same for her. Without Kim this thesis would never have been.

Contents

List of Figures

List of Tables

Chapter 1

A Model of Branching With an Application to Portuguese Banking

1.1 Introduction

How many branch offices should competing banks open in a given area and how many varieties of soda should competing manufacturers produce? While there is an extensive literature on product differentiation, the question of how many locations a given firm will serve do not seem to have been examined for their strategic nature.[1] Salop's (1979) study asked how many locations are served in equilibrium with free entry when each firm serves at most one location. Brander and Eaton (1984), Bhatt (1987) and Klemperer (1992) look at the question of where firms should locate multiple products (i.e. clustered together, interspersed with the rival's products or on top of the rival's products), but consider only cases where firms are locating a fixed number of products.

An important exception to this critique is the symmetric goods case of monopolis-

[1] There is a noteable difference between the strategic question of number of locations or products and questions of the optimal firm size based on cost considerations. See Aron (1988) for an interesting analysis of diversification of products based on the costs of managerial incentives.

tic competition as presented in Dixit and Stiglitz (1977). In that case the symmetry of the goods and the zero profit condition of free entry result in an equilibrium where each firm produces just one product. The intuition is that zero profit determines the number of locations served so each firm is deciding if they want to serve an additional location given that if they don't someone else will. The symmetry ensures that an incumbent firm faces the same zero profit as an entrant, so the incumbent is indifferent. This solution is very sensitive to these two assumptions and when they are relaxed Dixit and Stiglitz note that a firm could produce multiple products, but do not treat this case.

This paper addresses the issue of multi-product, or multi-location, oligopoly competition. We develop a simple model that allows us to focus on the size of a firm's network in either a product space or in ageographic region. One application of this model is to banks' decisions of how many branch offices to operate and in this respect the model will be estimated with data from the Portuguese banking industry. The qualitative analysis of the model, together with the empirical estimates, suggest an explanation for the puzzling pattern of branch expansion in this industry during the late 80's.

From a theoretical point of view, the main question we address is the strategic nature of the variable "number of varieties" or "number of locations." There are essentially two ways one can look at this variable. First, we can think of it as a measure of scale—firms with a larger number of branches are larger than firms with a lower number of branches. In this sense, we expect the variable "number of branches" to be a strategic substitute, that is, we expect firm i's reaction function to be decreasing in firm j's number of branches (just as in a Cournot oligopoly, for example).

However, we can also think of branching as a location decision. When a firm opens a branch in some location, it will receive customers which originate from one of three alternative sources: (i) previously unserved customers, (ii) previously own customers in some other branch, or (iii) previously customers at some rival firm's branch. Now, when the rivals' number of branches increases, it is possible that the importance of the third source increase so much relative to the second one that the marginal revenue of an extra branch goes up; that is, the number of branches may be a strategic complement (just as price in a Bertrand oligopoly or quality in various vertically differentiated oligopolies).

In Section 1.2, we present a simple theoretical model which conforms to this latter intuition. It is shown that the firms' best response functions are in general non-monotonic. Specifically, they are increasing for low levels of the rivals' number of branches and decreasing for high levels of the rivals' number of branches. Some comparative statics implications of the model are also explored in Section 1.2. In particular, it is shown that, in response to entry, small incumbents should reduce the number of branches, while large incumbents, especially large incumbents with a large loyal customer base, should increase the number of branches.

In Section 1.3.2, we estimate a version of the theoretical model using branching data from the Portuguese banking industry. (Before, in Section 1.3.1, we provide a brief overview of the industry.) The theoretical model is extended to allow for the possibility of a partial adjustment to the equilibrium number of branches, which in turn is given by the estimated reaction functions. A counterfactual analysis based on the estimated model indicates that some of the incumbents did, indeed, respond to entry by increasing their branch networks, a possibility which is suggested by the theoretical model. In other words, if there had been no entry, some incumbent banks would have increased their networks by less than they actually did.

Finally, in Section 1.3.3 we use the comparative statics of the theoretical model, as well as the empirical estimates, to advance an interpretation for the pattern of branch expansion in Portuguese banking during the late 80's. We show that efficiency and a large customer base are "substitute" factors in determining the desire to expand. This in turn explains the observed heterogeneity of the banks which expanded the most (large, old, inefficient banks as well as new, efficient ones). It also explains the patterns of expansion across geographical areas.

1.2 The basic model

In this section, we develop a simple static model of oligopoly competition between multi-product, or multi-location, firms.[2] In this model, each firm's decision of the number of varieties/locations plays a central role. We first present the basic structure of the model and then derive some results regarding equilibrium and comparative statics. As in much of the literature on product differentiation, the model can have two interpretations: choice of varieties and spatial location. Since our empirical application refers to the latter, we will in general refer to the firm's decision of *branching*, that is, choosing the number of locations. As we will see, given the assumptions we make, the model is actually better seen as a model of choice of varieties. However, in the end of the section, we compare our model with an alternative, more standard location model, and argue that our model provides a good approximation to reality, both qualitatively and quantitatively.

1.2.1 Primitives

We assume that the market is characterized by a fixed number, N, of identical locational nodes and a mass of consumers, one per node. Consumers are located both at the nodes and along the paths which connect the nodes. Firms, on the other hand, can only locate their branches at the nodes.[3] Conceptually, these nodes are meant to represent something like shopping areas or street corners abstracted from the immense complexities of a realistic street system. Under the choice-of-variety interpretation, nodes represent typical taste preferences. For example, in the market for ready-to-eat breakfast cereals, typical nodes would be raisin bran, corn flakes, natural, etc.

Consumer preferences are lexicographic. Each consumer's first concern is to min-

[2]There is an extensive related literature on the use of product proliferation as a means of entry deterrence, which includes Hay (1976), Prescot and Visscher (1977), Schmalensee (1978), Eaton and Lipsey (1979), Lane (1980), Bonano (1987), Shaked and Sutton (1990) and Judd (1985). This literature typically allows strategic choice of the number of locations to serve, but imposes sequential timing of these decisions and restricts the entrant's choice to a single location.

[3]If the concentration of consumers around the nodes is sufficiently high, then this would be a derived result.

imize transportation costs.[4] Therefore, if some firm is located at the node nearby the consumer's address, then the consumer purchases from some firm located at that node. If no firm is located at that node, then the consumer travels to another node according to some pre-determined preference ordering, until a served node is found. These preference orderings for second-, third-, etc. choices are uniformly distributed across consumers and are *independent of location*. The independence assumption is similar to that in Von Ungern-Sternberg's (1991) model of "monopolistic competition on the pyramid" (see also Deneckere and Rothschild, 1986). In fact, Von Ungern-Sternberg's idea of an $(N-1)$-dimensional pyramid with firms located at each of the N vertices and consumers located along the edges may be a useful one, and will be maintained throughout. The assumption that second-choice preferences are independent of location may not be such a bad approximation for geographic differentiation. A consumer may first look for a branch located close to his home address. If no branch is found, then he will look for a branch close to work; and then close to the child's school; etc. The idea is that there is little correlation between place of residence and place of work, etc. In the case of product varieties, the assumption reflects the fact that there are many dimensions to consumer preferences, so that relative orderings are very subjective.[5]

Each of the m firms simultaneously decide how many branches to have, n_i, and allocate these to one node each. The analysis in Brander and Eaton (1984) and in Bhatt (1987) of which subsets of a small finite set of product lines each firm can pursue in equilibrium is similar to the question of where to locate branches. The multiplicity of cases and equilibria in those papers illustrates the intractability of the *where to locate* problem with many potential sites. Since our focus is on the choice of the *number* of branches, we make the location choice as simple as possible. Specifically, we assume that location is a random process, with each node being given

[4] Anecdotal evidence from banking services, for example, shows that this assumption is not too unrealistic.

[5] From the perspective of firm competition, the assumption implies the absence of "neighborhood effects:" every firm/branch competes with every other firm/branch. In fact, this is the essence of Chamberlin's monopolistic competition. See Páscoa (1993).

equal probability (given the constraint that not more than one branch is allocated to any given node).[6]

The revenue earned by a branch of firm i from the customers located at each node is given by the function $R_i(S)$[7], where S is the set of firms located at that node (S will thus include firm i). It is natural to assume that $R_i(\cdot)$ is monotonic: $S \subset S' \Rightarrow R_i(S) \geq R_i(S')$. An extreme example of functions $R_i(\cdot)$ is given by Bertrand competition: $R_i(S) = 0, \quad \forall S \neq \{i\}$.

The opposite extreme of $R_i(\cdot)$, one which will be considered in the empirical application, is the case when there is no price competition. Under this assumption, when there is more than one firm serving a given node, that is, when a node is contested, consumers are divided between firms according to the following rule. If all firms have a branch at the node, then firm i gets a share q_i of the total. For any set S of firms with branches at a contested node, firm i will get a fraction $q_i / \sum_{j \in S} q_j$.

The shares q_i may reflect quality differences between firms. In the case of product variety, an example is given by ready-to-eat breakfast cereals: a consumer may prefer corn flakes to raisin bran (preferences across nodes); and, given that he chooses to buy corn flakes, there may be a preference between *Kellogs* and some other maker of corn flakes (preference between firms, modeled by q_i).[8] In the case of spatial competition, in particular bank competition, the values q_i could reflect the proportion of the population who are "loyal" to each firm in the sense that they would incur a switching cost to purchase from a different firm. The assumption on transportation costs means that these switching costs are low enough relative to transportation costs

[6]There are several reasons why this assumption may be a sensible one. First, one can show that even for very simple location games there exists no equilibrium in pure strategies. Second, the pattern of randomization we consider does constitute an equilibrium. Third, playing mixed strategies can be interpreted on the basis of Harsanyi's (1973) purification argument. Applied to the case in hand, the idea is that firms have some private information about demand at each node (for example) and choose locations optimally according to this information. If each firm's private information is uncorrelated with the other firms', then equilibrium strategies will look as if firms were randomizing locations.

[7]This is the revenue net of the cost of serving these consumers where the marginal cost of serving an additional consumer is constant.

[8]The difference between firms can be a difference in vertical quality. In that case, interior values of q_i would be justified by a countervailing price difference which is valued differently by different consumers.

to make a customer unwilling to walk away from a branch at the node where he lives to a branch of the bank he is loyal to.

Finally, notice that the potential mass of customers at each node is given by "local residents" plus the consumers who travel from unserved nodes. The fact that second-choice node preferences are independent of location implies that unserved customers are evenly distributed among the served nodes. As a result, the distribution of consumers per served node is identical. We assume that each firm's revenue per branch is given by $QR_i(S)$, where Q is the measure of customers served at the node where the branch is located.[9]

1.2.2 Reaction functions: basic results

We will begin by characterizing the strategic nature of the variable "number of branches" by studying the shape of the reaction functions. In order to avoid trivial solutions, we will need some assumptions regarding the marginal cost of an additional branch, c_i for firm i, and the size of the market.

Assumption 1 *For all* i, $c_i < (N+1)R_i(\{i\})/6$.

Assumption 2 *For all* i, $c_i > R_i(U)$, *where* U *is the set of all firms.*

The meaning of these assumptions can be seen from the following three basic lemmas.

First, suppose that firm i is a monopolist, that is, $n_j = 0$, $\forall j \neq i$. Then, firm i's optimal response is never to open more than one branch, that is, $n_i^r(0) \leq 1$, where $n_i^r(n_{-i})$ is firm i's reaction function to n_{-i}, the vector of the rivals' number of branches. In fact, since demand is inelastic with respect to transportation costs, the only effect of opening more than one branch would be to increase the total cost of branches opened, while total revenue would remain constant. This leads to the following result.

[9] A linear demand function which is scaled up by the number of customers is one example which would naturally lead to this proportionality rule.

Lemma 1 $n_i^r(0) = 1$.

Proof: One can easily check that Assumption 1 implies that the profit from opening one branch is positive, so that $n_i^r(0) > 0$. The argument above implies that $n_i^r(0) \leq 1$. ■

Now suppose that one of firm i's rival firms opens one branch ($n_j = 1$), while all other firms open zero branches ($n_k = 0, \ \forall k \neq i, j$). What is firm i's expected payoff if it opens one branch? With probability $1/N$, two rival branches will be located at the same node, while with probability $(N-1)/N$ the two rival branches will be located at different nodes. In the first case, both branches compete for a total demand of N consumers (or, N times the measure of consumers at each node); while, in the second case, each firm will be a monopolist over one half of the total demand. We thus have

$$\Pi(1, e_j) = \frac{1}{N} N R(\{i,j\}) + \frac{N-1}{N}\frac{N}{2} R(\{i\}) - c_i, \tag{1.1}$$

where $\Pi(n_i, n_{-i})$ is the expected payoff function and e_j is a vector with a 1 on the $j-th$ component, zeros elsewhere.

Now suppose that firm i opens 2 branches, still maintaining the same value of n_{-i}. Expected payoff is now given by

$$\Pi(2, e_j) = \frac{2}{N}\left(\frac{N}{2} R(\{i\}) + \frac{N}{2} R(\{i,j\})\right) + \frac{N-2}{N}\frac{2}{3} N R(\{i\}) - 2c_i. \tag{1.2}$$

Based on this, we can derive the following result.

Lemma 2 $n_i^r(e_j) > 1$, *where the j-th element of e_j is 1 and all other are zero.*

Proof: Computation establishes that Assumption 1 is equivalent to $\Pi_i(2, e_j) > \Pi_i(1, e_j)$. ■

Finally, let us consider the opposite extreme, namely the case when all firms $j \neq i$ choose the maximum possible number of branches, that is, $n_j = N, \ \forall j \neq i$.

Lemma 3 *If n_{-i} is sufficiently large, then $n_i^r(n_{-i}) = 0$.*

Proof: By Assumptionn 2, $n_j = N, \ \ \forall j \neq i$ implies $n_i^r(n_{-i}) = 0$. The result then follows by continuity. ■

Together, Lemmas 1–3 imply the following corollary.

Corollary 1 *n_i^r is (strictly) increasing for low values of n_{-i} and decreasing for large values of n_{-i}.*

In other words, the reaction functions are first increasing and then decreasing. The number of branches are strategic complements for low numbers of the rivals' branches and strategic substitutes for high numbers of the rivals' branches.

The intuition for these results can be seen with reference to the two standard examples of strategic complements and substitutes, namely price and quantity competition.[10] We now consider each of these in turn.

(a) Strategic complementarity. When a firm opens a branch in some location, it will receive customers which originate from one of three alternative sources: (i) previously unserved customers, (ii) previously own customers in some other branch, or (iii) previously customers at some rival firm's branch. Now, when the rivals' number of branches increases, it is possible that the importance of the third source increase so much relative to the second one that the marginal revenue of an extra branch goes up. Since by Lemma 1 $n_i^r(0) = 1$ and by Lemma 2 $n_i^r(e_j) > 1$, this will indeed be the case when n_{-i} is small, that is, the number of branches will be a strategic complement.

This is analogous to the complementarity effects in price competition. If $p_i < p_j$, then the cost to firm i (in terms of lower margins) of a reduction in price are mostly or fully internalized by this firm. However, if p_j is decreased such that $p_j < p_i$, then part of the costs of the reduction of p_i are borne out by the rival firm, in the form of a lower market share. This is roughly the intutition for why prices are strategic complements.

[10] As Bulow, Geanakoplos, and Klemperer (1985b) note (and exemplify), this characterization of price and quantity competition is not general. It is correct for the case of linear demand and in other cases.

(b) Strategic substitutability. When n_{-i} is very large, $P_{-i} \approx 0$, that is, the fraction of unserved nodes is close to zero. Consequently, firm i's expected revenue from an extra branch is proportional to the expected value of $R_i(S)$. But since $R_i(S)$ is a decreasing function of S, and S is "increasing" in n_{-i}, it follows that expected revenue is decreasing in n_{-i}.

In this case, the analogy is with quantity competition. When a Cournot oligopolist increases the quantity it supplies, it reduces not just its marginal revenue but all firms' marginal revenue as well. In fact, marginal revenue is an increasing function of price, which in turn is a decreasing function of total quantity. Likewise, in the case of branching, marginal revenue is a decreasing function of S, which in turn is an increasing function of all rivals' number of branches (although not in the simple way as in quantity competition).

Another immediate consequence of Lemmas 1–3 is the following.

Corollary 2 *In any symmetric equilibrium with $N > 1$, the number of branches per firm is greater under duopoly than under monopoly.*

Section 1.2.5 includes a discussion of this corollary.

1.2.3 Reaction functions: continuous case

Unfortunately, it is not possible to determine the exact reaction function for a general n_{-i}. In order to get around the integer problem we will assume that the number of branches is a continuous varaiable. When the elements of n_{-i} are large numbers this assumption is reasonable.

Denote by $E_i(n_{-i})$ firm i's expected revenue from a branch located at a given node which is originated from the customers located at that node. As we have seen before, firm i's revenue is given by $R_i(S)$, where S is the set of firms located at the node. Therefore, we have $E_i(n_{-i}) = E(R_i(S) \mid n_{-i})$. Notice that $E_i(n_{-i})$ does not depend on n_i, since the firms' randomizations are independant.

The number of customers served at each node is typically greater than the number of customers located at that node. This is so because there are generally nodes which remain unserved. The expected fraction of such nodes is given by

$$P(n) = \prod_{k=1}^{m} \left(1 - \frac{n_k}{N}\right), \tag{1.3}$$

where $n \equiv [n_i]$. The total expected revenue for firm i of a branch is then

$$E_i + E_i \left(\frac{P}{1-P}\right) = E_i \left(\frac{1}{1-P}\right). \tag{1.4}$$

Recalling that firm i's marginal cost of having an additional branch is a constant c_i, firm i's total expected profit is

$$\Pi_i(n) = n_i E_i \left(\frac{1}{1-P}\right) - c_i n_i. \tag{1.5}$$

Equation (1.5) illustrates the tradeoffs to branching in this model. The first and last n_i represent the usual tradeoff between the revenue from an additional branch and its cost. The revenue per branch, which is given by E_i times the term in brackets, is a decreasing function of n_i. Additional branches take customers away from existing branches. Some of these customers will switch between branches of the same firm, and so they will not be a contribution to revenue. The larger a firm's network is, the more of the customers it gets at a new branch will be transfers of this sort, and the lower will be the contribution to revenue.

The general profit function of Equation (1.5) is concave in n_i so we can use the first order condition to derive a reaction function. Let

$$P_{-i} = \prod_{k \neq i}^{m} \left(1 - \frac{n_k}{N}\right). \tag{1.6}$$

It can be shown that firm i's reaction function is given by

$$n_i^r(n_{-i}) = N \left[\frac{P_{-i} - 1 + \sqrt{\frac{E_i}{c_i}(1 - P_{-i})}}{P_{-i}}\right]. \tag{1.7}$$

1.2.4 Equilibrium in duopoly with no price competition

In order to have a better idea of the shape of the reaction curves, we now specialize the model to the case of a duopoly with no price competition. Specifically, we assume there is a potential revenue π originating from the consumers at each node, and that this revenue is divided between firms when the node is contested. The revenue sharing is relatively simple in the duopoly case. If a node is served by only one firm, then that firm gets all of the potential revenue from that node. If a node is served by both firms, firm i gets q_i of the potential revenue. If a node is served by neither firm, the potential revenue from that node is divided among the nodes that are served. Firm i's expected revenue at each node it serves is given by

$$E_i = \pi\left(\frac{N-n_j}{N}\right) + q_i\pi\left(\frac{n_j}{N}\right) = \pi\left(1 - q_j\frac{n_j}{N}\right). \tag{1.8}$$

where $q_j = 1 - q_i$. Profits can thus be written as

$$\Pi_i(n_i, n_j) = n_i\left(1 - q_j\frac{n_j}{N}\right)\left[\frac{N^2}{N^2 - (N-n_i)(N-n_j)}\right]\pi - c_i n_i. \tag{1.9}$$

Taking the derivative with respect to n_j and simplifying yields

$$n_i^r(n_j) = \frac{N}{N-n_j}\left[-n_j + \sqrt{n_j\frac{\pi}{c_i}(N - q_j n_j)}\right]. \tag{1.10}$$

For the duopoly case, we are able to derive a strong result with respect to the slope of the reaction functions at the equilibrium. The following lemma states that, in equilibrium, the number of branches is a strategic substitute for one firm but a strategic complement for the other.[11]

Lemma 4 *In equilibrium,* $(\partial n_i^r/\partial n_j)(\partial n_j^r/\partial n_i) < 0 \quad (i \neq j)$.

Proof: Taking the derivative of firm i's reaction function with respect to n_j, we get

[11] In the words of Bulow, Geanakoplos, and Klemperer (1985a), the number of branches is a strategic substitute *from the perspective of* one firm but a strategic complement *from the perspective of* the other. Bulow, Geanakoplos, and Klemperer (1985b) present an example of duopoly quantity competition which implies a result similar to the following one.

$$
\begin{aligned}
\frac{\partial n_i^r}{\partial n_j} &= \frac{N}{(N-n_j)^2}\left[-n_j + \sqrt{n_j \frac{\pi}{c_i}(N - q_j n_j)}\right] + \\
&+ \frac{N}{N-n_j}\left[-1 + \frac{1}{2}\frac{\frac{\pi}{c_i}(N - 2q_j n_j)}{\sqrt{n_j \frac{\pi}{c_i}(N - q_j n_j)}}\right] \\
&= \frac{n_i}{N-n_j} + \frac{N}{N-n_j}\left[-1 + \frac{1}{2}\frac{\pi}{c_i}(N - 2q_j n_j)\frac{n_i \frac{N-n_j}{N} + n_j}{n_j \frac{\pi}{c_i}(N - q_j n_j)}\right] \\
&= \frac{n_i}{N-n_j} + \frac{N}{N-n_j}\left[-1 + \frac{1}{2}(N - 2q_j n_j)\frac{Nn_i + Nn_j - n_i n_j}{n_j(N - q_j n_j)}\right] \\
&= \Big[n_j(N-n_j)(N-q_j n_j)\Big]^{-1}\Big[(N/2 - q_j n_j)(Nn_i + Nn_j - n_i n_j) - \\
&\quad - (Nn_j - n_i n_j)(N - q_j n_j)\Big] \\
&= \Big[n_j(N-n_j)(N-q_j n_j)\Big]^{-1}\Big[(N/2 - q_j n_j)(Nn_i + Nn_j - n_i n_j) - \\
&\quad - (Nn_j + Nn_i - n_i n_j)(N - q_j n_j) + Nn_i(N - q_j n_j)\Big] \\
&= \Big[n_j(N-n_j)(N-q_j n_j)\Big]^{-1}\Big[-\frac{N}{2}(Nn_i + Nn_j - n_i n_j) + \\
&\quad + Nn_i(N - q_j n_j)\Big] \\
&= \Delta_i\Big[\frac{N}{2}(n_i - n_j) + (\frac{1}{2} - q_j)n_i n_j\Big] \qquad (1.11)
\end{aligned}
$$

where $\Delta_i \equiv N[n_j(N-n_j)(N-q_j n_j)]^{-1}$. Since $\Delta_i > 0$ and $q_i = 1 - q_j$,

$$
\text{sign}\left[\left(\frac{\partial n_i^r}{\partial n_j}\right)\left(\frac{\partial n_j^r}{\partial n_i}\right)\right] = \text{sign}\left[-\left[\frac{N}{2}(n_i - n_j) + (\frac{1}{2} - q_i)\right]^2\right] \qquad (1.12)
$$

which proves the result. ■

Corollary 3 *If*

$$
\begin{aligned}
&\textit{(i)} \quad q_i = q_j \ \ \textit{and} \ \ c_i < c_j, \\
\textit{or} \ \ &\textit{(ii)} \quad c_i = c_j \ \ \textit{and} \ \ q_i > q_j,
\end{aligned}
$$

then

$$\text{(a)} \quad n_i > n_j$$
$$\text{and} \quad \text{(b)} \quad \partial n_i^r / \partial n_j > 0 > \partial n_j^r / \partial n_i.$$

The idea of this result is that there are two main factors which influence a firm's size and whether its number of branches is a strategic substitute or complement. The first one is the cost of opening a new branch. The lower c_i is, the higher firm i's number of branches, other things equal, and the greater the likelihood that n_i is a strategic complement. The second factor is the expected market share in each served node, q_i. The higher q_i is, the higher firm i's number of branches, other things equal, and the greater the likelihood that n_i is a strategic complement.[12] In this sense, one can say that efficiency (as reflected in c_i) and a loyal customer base (as reflected in q_i) are "substitute" factors in determining a firm's desire to expand (and the likelihood that n_i is a strategic complement). (This observation will play a crucial role in the interpretation of the pattern of Portuguese banks' expansion of their branch networks.) Finally, Corollary 3 suggests that it is the larger firms that have a positively sloped reaction function.[13]

The situation of simultaneous strategic complements and substitutes is especially relevant to analyzing the effects of entry. In order to keep the duopoly apparatus, we will consider an expansion in the number of Firm 2's branches resulting from an exogenous decrease in its cost, c_2, from c_2' to c_2''. As illustrated in Figure 1-1, it is conceivable that a dominant firm (a firm with lower c_i and/or higher q_i, Firm 1 in this case) would rationally expand its branch network in response to this form of entry. In the figure, the equilibrium changes from E' to E''. Both Firm 2 *and* Firm 1 increase n_i as a result of a decline in c_2.

As we noted before, the greater q_1 is, the greater the likelihood that n_1 is a strategic complement; and q_1 may be interpreted as the degree of customer loyalty

[12]Notice that the effect of q_i on the derivative of n_i^r with respect to n_j comprises both a direct effect and an effect through the value of n_i; both effects have the same sign.

[13]However, one may find cases with $c_i < c_j$, $q_i < q_j$ such that the order is reversed.

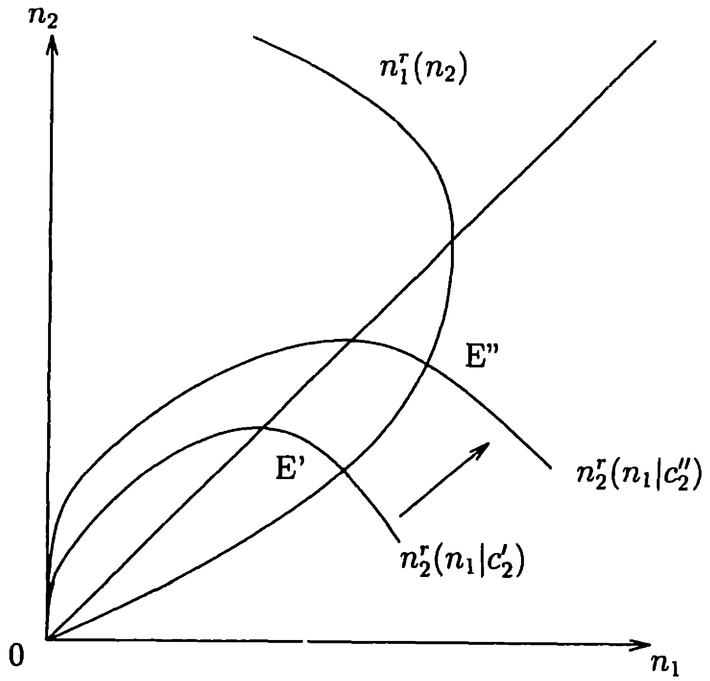

Figure 1-1: Dominant firm's reaction to rival's increased efficiency

that Firm 1 enjoys. Then, the comparative statics on entry indicates that the more loyal customers a firm has, the more likely it is that it will expand its branch network in response to entry. The intuition is that when there are more loyal customers, the incumbent firm has more to lose from reducing its relative density of branches; which is to say, has more to gain from expanding the density of its branch network.

Finally, it is interesting to investigate the implications of the above results in a model of sequential entry of the kind considered by Fudenberg and Tirole (1984). Suppose that firm 1, the incumbent, has the opportunity of making an investment I in cost reduction, such that $c_1 = c_1(I)$ with $c_1'(I) < 0$, before firm 2 enters the market. In addition, suppose that $q_1 = 1_2$ and that c_2 is given. Following Fudenberg and Tirole (1984), the effect I on firm 1's equilibrium profits can be decomposed in a direct effect and a strategic effect. The direct effect is always positive: an increase in I implies a decrease in c_1, which, ceteris paribus, implies an increase in firm 1's profits. The strategic effect is given by

$$\frac{\partial \Pi_1}{\partial n_2}\frac{d\hat{n}_2}{dI}. \tag{1.13}$$

It can be shown that the sign of the strategic effect is given by

$$\text{sign}\left(\frac{\partial \Pi_1}{\partial n_2}\frac{d\hat{n}_2}{dI}\right) = \text{sign}\left(\frac{\partial \Pi_2}{\partial n_1}\frac{d\hat{n}_1}{dI}\right) \times \text{sign}\left(\frac{\partial n_2^r}{\partial n_1}\right). \tag{1.14}$$

The first term on the right hand side is negative. In fact, one can see that firm i's profit is decreasing in firm j's number of branches. In the terminology of Fudenberg and Tirole (1984), investment in cost reduction makes firm 1 "tough." The sign of the strategic effect is therefore opposite of the sign of the slope of firm 2's reaction curve. In particular, if the reaction curve is positively sloped, then the strategic effect is negative. This in turn implies that the firm underinvests in cost reduction

Now, Corollary 3 implies that firm 2's reaction function is positively sloped if and only if $c_1 > c_2$. Therefore, we conclude that 1 will underinvest in cost reduction if $c_2 > c_1$ (in the terminology of Fudenberg and Tirole (1984), follow a "puppy dog" strategy) and overinvest in cost reduction if $c_2 < c_1$ (in the terminology of Fudenberg and Tirole (1984), follow a "top dog" strategy).

1.2.5 Dissipative competition strategies

A closer examination of the previous subsections reveals that the main features of the variable "number of branches" should also be found in the more general context of strategic variables related to dissipative competition. In particular, the marginal gain from investing in a dissipative strategy is positively related to two factors: (i) the rivals' market share (the greater it is, the more there is to transfer), and (ii) the rivals' investment in that strategy (the greater it is, the more difficult it is to transfer market share).

These two factors are consistent with the observations made before, namely: (a) investment in dissipation is a minimum when the number of firms is one (factor (i): there is no market share to steal) or when it is very large (factor (ii): large investments are needed in order to steal market share); (b) at low levels, reaction curves are upward sloping (factor (i): as the rivals' increase their market share, there is more to steal from); and finally, (c) at high levels, reaction curves are downward sloping (factor (ii): as the rivals increase their investments, the required investment to recover becomes higher).

An example to which these ideas seem to apply as well as branching is dissipative advertising. In fact, our model can also be interpreted as a very stylized model of dissipative advertising. Suppose that each node represents one consumer. Each firm decides how many advertising messages to send, n_i, and distributes them randomly across consumers. A consumer who receives only one ad buys from the firm sending that ad. A consumer who receives more than one ad chooses one of the sellers according to the probability vector $[R_i(S)/\sum_{j\in S} R_j(S)]$, where S is the set of firms whose ads that consumer receives. Finally, a consumer who does not receive any ads "travels" to one of her neighbors and imitates his decision. This is the same story of advertising as in Butters (1977), but with a different technology of advertising. Butters assumes that the random allocation of ads to consumers is done with replacement so that a consumer can receive multiple ads from the same firm. In this interpretation of our model the firms are assumed to be able to keep from duplicating their own effort and we add word-of-mouth advertising for those consumers who receive no ads directly.

In this context, it is interesting to address the question of what is the relation between market structure (namely, number of firms) and the equilibrium number of branches (or advertising messages). Numerical simulations reveal that in an m-firm symmetric oligopoly the equilibrium number of branches per firm, $n^*(m)$, is a decreasing function of m for $m > 1$, while the total number of branches, $n^*(m)m$, is an increasing function of m. In addition, the analysis presented above implies that $n^* = 1$ when $m = 1$.[14]

1.2.6 An alternative model of spatial competition

The model we have developed so far may be criticized on the basis that it only allows for localized competition in a very stylized way. Each network branch is in direct

[14] As the model applies to advertising, these results seem consistent with the evidence that advertising/sales ratios are a quasiconcave function of market concentration (cf Martin, 1979).

competition with other branches at the same node; but, other than that, it is equally placed with respect to all other branches in all other nodes. We will now argue that a more standard model of spatial competition may exaggerate in the opposite direction, that is, may involve *too much* localized competition. We will also provide other reasons why our model may provide, both qualitatively and quantitatively, a good approximation to reality.

The most common model of localized competition, in the tradition of Hotelling (1929), is Salop's (1979) circular road model. In order to make the comparison with our model possible, we will consider the extreme case when the density of nodes is so large that the probability of two branches coinciding is close to negligible. Suppose that there are two firms who have to decide how many branches to have and how to distribute them along the circle. As before, we will assume there is no price competition, so that we have a pure location game.

It can be shown that there exist no pure-strategy equilibria in this location game: each firm's best response is always to locate just around the rival's locations and thus maximize the customer base. It can also be shown that it is an equilibrium for firms to locate their branches equally spaced along the circle and randomize the location of this pattern on the circle with a uniform distribution.

Based on this, market shares can be easily computed: they are given by $s_i = n_i/(2n_j)$ and $s_j = (2n_j - n_i)/(2n_j)$, where $n_i \leq n_j$. From market shares, we can derive expected profit and reaction functions. It can be shown that the reaction function $n_i^r(n_j)$ consists of three different sections. For $n_j < n'$, $n_i^r(n_j)$ is increasing, concave and greater than n_j; for $n' < n_j < n''$, $n_i^r(n_j) = n_j$; and for $n_j > n''$, $n_i^r(n_j) = 0$.

The circular road model has both differences and elements in common with the model we presented above, the pyramid model. First, note that in both cases reaction functions are first increasing and then decreasing. However, in the circular model this occurs in a discontinuous way. In particular, the decreasing "region" is the discontinuity at n'' where it is suddenly no longer profitable to participate. Alternatively, the reaction functions for the pyramid model estimated in Section 1.2.4 are smooth. The circular model does not have a region of downward sloped reaction functions, so it is not possible for an equilibrium to occur there. Second, while a unique equilibrium was found for the pyramid model, the circular road model can easily admit multiple equilibria: recall that, when $n' < n_j < n''$, $n_i^r(n_j) = n_j$.

The third point of comparison relates to market shares. As we mentioned before, the pyramid model may be criticized on the basis that it does not admit enough localized competition. However, the circular road model may be criticize for the precise opposite reason. In fact, the model tends to over-estimate the larger firms' market share. The idea is that the larger firm guarantees to itself all of the demand between two of its adjacent branches; the smaller firm is thus confined to a small fraction of the potential market. In fact, the smaller firm's market share, $n_i/(2n_j)$, is less than the fraction of the total number of branches that it owns, $r_i \equiv n_i/(n_i + n_j)$. In the pyramid model, however, it can be shown that when $N \to \infty$, we have $s_i \to n_i/(n_i + n_j)$. In particular, it can be shown that, at $r_i = 0$, $\partial s_i/\partial n_i = 1$ for the pyramid model and $\partial s_i/\partial n_i = 1/2$ for the circular road model. But the data we use in the empirical section seem to reject the second hypothesis against the first.[15]

In summary, the pyramid and the circular road models are qualitatively similar with regards to the shape of reaction functions. However, the former seems better indicated for the purpose of empirical estimation.

1.3 An application to Portuguese banking

Bank branching is one of the most obvious applications of the theoretical model developed in the preceding section. In this section, we consider the case of the Portuguese banking industry during the late 80's. The section begins with a brief overview of the industry.[16] We then present the empirical model used and the estimation results. The section ends with a discussion of the results.

[15] For example, in the district Lisbon and in 1990, a regression of s_i on r_i yielded, for the smaller banks, $s_i = 0.88r_i$, with a standard error of 0.09.

[16] The reader is also referred to Gordy (1991a,b) for interesting analyses of Portuguese banking (namely branching).

1.3.1 The puzzling pattern of Portuguese bank branching

From 1975 to 1984, virtually all banks operating in Portugal were State-owned. Interest rates, credit ceilings, and other variables were directly controlled by the Government and the Bank of Portugal (the central bank). Competition was therefore virtually nonexistent. The 1980's, by contrast, were a decade of overall deregulation and privatization, and the banking sector was not an exception. In 1984, a law was passed allowing entry by new private banks. Later on, a plan of privatization for several of the public banks was drafted; the first privatization was processed by the end of the decade.[17] Beginning in December of 1984, new domestic private banks were created and foreign banks entered the Portuguese market. These new banks gradually increased their number of branches. Together with the expansion by the incumbent banks, this led to a very rapid increase in the total number of branches.[18] Table 1.1 presents the number of branches in continental Portugal by the end of 1988, as well as the increase in the number of branches during the period 1989–91.

The table depicts a pattern of branching that is puzzling. First, the total increase in the number of branches during 1989–91 was fairly large, a growth rate of more than 10% a year. Second, growth was highly concentrated in four banks, *Caixa Geral de Depósitos* (CGD), *Banco Comercial Português* (BCP), *Banco Espírito Santo e Comercial de Lisboa* (BESCL), and *Banco Pinto e Sotto Mayor* (BPSM), which accounted for more than two thirds of the increase in the number of branches. Of these four, CGD and BCP alone accounted for roughly half. Third, the increase in the number of branches was more or less evenly distributed between public and private banks (although, since there were more public banks initially, the growth rate was much higher for private banks than for public banks). Finally, while private banks have mainly invested in "urban areas" (Lisbon and Oporto districts), public banks seem to have focussed on "rural areas" (all districts except Lisbon and Oporto).

In addition to the data, there is some anecdotal evidence that the new, private

[17] To be more precise, banks were re-privatized, since most of them had been nationalized in 1975.

[18] The rate of expansion might have been higher if it were not for the bureaucratic costs of getting new branches approved by the government—by 1992, entry has still not been completely liberalized.

Table 1.1: Number of bank branches in continental Portugal. (Source: Bank of Portugal.)

Concept	By end 88	New in 89+90
Four largest	628	225
BCP	18	113
CGD	339	45
Private	191	174
Urban	102	118
Rural	89	56
Public	1325	159
Urban	477	65
Rural	848	94
Total	1516	333

banks are more efficient than the old, public banks; in particular, there is a common perception that the private banks' marginal profitability of an extra branch is larger than the public banks', especially because operating costs are lower. On the other hand, public banks, especially CGD, have the advantage of controlling a fairly large customer base consisting mainly of very "loyal" depositors, that is, depositors with a high cost of switching to other banks.

If we were to base our analysis solely on comparisons of bank profitability, we would expect private banks to be the only ones expanding. But, perhaps a large and loyal customer base is a "substitute" for high margins as a determinant of high propensity to open new branches. This would help explain the difference between "urban" and "rural" areas, but the traditional analysis would still have difficulty explaining an incumbent that responds to entry by expanding.

All these stylized facts and conjectures motivate the estimation of the theoretical model of Setion 1.2 using data from Portuguese bank branching.[19]

[19]The theoretical model does not contemplate the possibility of direct consumer benefits from network size. We believe that these benefits are small in the Portuguese market. For a typical depositor, branches other than the "home" branch are used for the sole purpose of obtaining cash. Now, the network of ATM machines is common to all Portuguese banks, so that depositors attach relatively little importance to the network size of the bank they patronize. Models which explicitly consider network effects include Gale (1992), Matutes and Padilla (1993), Nakamura and Parigi (1992). Other models which explicitly address issues of banking, although not the same as ours, include Neven (1990), Chiappori, Perez-Castrillo and Verdier (1992). Fuentelsaz and Salas (1992), in particular, estimate the empirical determinants of the *total* number of bank branches in Spain.

1.3.2 Empirical model

Unfortunately, the extremely high and sustained growth rate of branches in the Portuguese banking industry during 1989–1991 indicates that the industry was not in an equilibrium during this period. It is possible that year to year changes could produce an equilibrium one year that was much higher than the equilibrium in another year. It is also possible that another wave of changes could cause growth in multiple periods. But, a pattern of very high growth rates over a period of several years requires a somewhat particular sequence of significant and unforseen shocks. Despite the high rate of legislative change in this industry, a more compelling explanation seems to be that the industry is undergoing a non-instantaneous transition to a higher equilibrium. This target equilibrium may be subject to some fluctuation from year to year, but these changes seem less significant than the difference between the current state and the equilibrium.

From an empirical perspective, disequilibrium is not very attractive. To completely model the situation, we should expand the model of competition in branches to include a constraint on the acquisition of new branches and to include dynamic interaction through an infinite-period game with discounting. However, this would introduce additional complications since there would exist multiple equilibria. For the purpose of empirical estimation, we have chosen to assume that the players follow a particular strategy. In each period, they increase the number of branches by some fixed proportion of the difference between the current level and the static reaction function evaluated at that date—the game-theoretic correspondent to the partial adjustment model, a model which is frequently used in applied econometrics. This strategy is not completely arbitrary in that it is similar to a set of strategies that Fudenberg and Tirole (1983) found to be a Markov perfect equilibrium of a capacity accumulation game with infinitely patient players. The strategy has the advantage that the accumulation is proportional to the myopic "need," so the reaction function of the static game, and therefore the accumulation path, can be estimated with

In this paper, we are interested in the number of branches per bank.

dynamic data.

The empirical model is thus

$$n_{ij}^{t+1} - n_{ij}^{t} = \alpha_i^t \left(n_{ij}^r(n_{-ij}^t) - n_{ij}^t \right) + u_{ij}^t \tag{1.15}$$

where n_{ij}^t is the number of branches of bank i in geographic region j at time t. The geographic regions of Portugal are considered to be separate markets, since it would be unreasonable to believe that consumers in the Lisbon metropolitan area travel to Oporto to do their banking. Therefore, the reaction function of player i, $n_{ij}^r(\cdot)$, is also indexed by the region, j. This allows for variation across regions in the number of nodes, N, the expected share of revenue at a node, Q_i, and the potential profit of a node relative to the branching cost, $\frac{\pi_i}{c_i}$. The reaction function also depends, of course, on the number of branches that opponents have, but only in that region, the vector n_{-ij}^t. The adjustment rate α_i^t is allowed to vary across banks and time. This variance accomodates differences in the implicit constraints that keep the banks from fully adjusting. The errors, u_{ij}^t, are assumed to be independent and identically distributed.

Since players are determining their reactions to the actions of all players in the previous period, there is no problem of simultaneity in the estimation. Requiring that all players assume this stationarity in the play of their opponents also has implications for Q_{ij}. In a forward looking framework, Q_{ij} would be a random variable depending on the current actions of the opponents, but in this model Q_{ij} is fully determined. In fact, Q_{ij} can be obtained from the market shares in deposits, MS_{ij}, since market share is equal to the ratio of firm i's revenue to the total revenue earned by firms in this market.

$$\begin{aligned} MS_{ij} &= \frac{n_{ij}Q_{ij}\pi\frac{1}{1-P}}{\sum_k(n_{kj}Q_{kj}\pi\frac{1}{1-P})} \\ &= \frac{n_{ij}Q_{ij}}{\sum_k(n_{kj}Q_{kj})} \end{aligned} \tag{1.16}$$

$$Q_{ij} = \frac{MS_{ij}}{1 - MS_{ij}} \frac{\sum_{k \neq i}(n_{kj} Q_{kj})}{n_{ij}}. \tag{1.17}$$

This deterministic Q_{ij} is a function of each other firm's Q_{kj}, but they can all be approximated together by an itterative method. In the first round $Q^1_{ij} = MS_{ij}$ where the superscript is an index of itteration. Subsequent values are given by

$$Q^{s+1}_{ij} = \frac{MS_{ij}}{1 - MS_{ij}} \frac{\sum_{k \neq i} n_{kj} Q^s_{kj}}{n_{ij}}. \tag{1.18}$$

This procedure is continued until convergence.

Rearranging Equation (1.15) and substituting the general reaction function for $n^r(\cdot)$ yields

$$n^{t+1}_{ij} = \alpha^t_i N^t_j \left[\frac{P_{-i} - 1 + \sqrt{Q^t_{ij} \frac{\pi^t_{ij}}{c_i}(1 - P_{-i})}}{P_{-i}} \right] + (1 - \alpha^t_i) n^t_{ij}. \tag{1.19}$$

Since N^t_j , π^t_{ij} , c_i and α^t_i are unknown, we must also specify

$$N^t_j = \beta_1 + \beta_2 Pop^{\,t}_{\,j} + \beta_3 (Pop^{\,t}_{\,j})^2 + \beta_4 (Dense^{\,t}_{\,j}) + \beta_5 (ATMC^{\,t}_{\,j}), \tag{1.20}$$

$$\pi^t_{ij} = \delta_1 (Prv_{\,i}) + \delta_2 (Pub_{\,i}) + \delta_3 (\frac{VAM_{\,j}}{N^t_j}) + \delta_4 (ATM^{\,t}_{\,j}), \tag{1.21}$$

$$c_i = \delta_5 (Prv_{\,i}) + \delta_6 (Pub_{\,i}) \tag{1.22}$$

and

$$\alpha^t_i = \alpha_0 (Big^{\,t}_{\,i}) + \alpha_1 (89) + \alpha_2 (90) + \alpha_3 (91). \tag{1.23}$$

Since π^t_{ij} and c_i do not enter the equation seperately, their coefficients cannot be estimated seperately. Combining Equations 1.21 and 1.22 and simplifying produces a single equation for the ratio of expected profits to startup costs:

$$K_{ij}^t = Q_{ij}^t \frac{\pi_j^t}{c_i} = Q_{ij}^t \quad [(Prv_{\,i})(\gamma_1 + \gamma_2(\frac{VAM_{\,j}}{N_j^t}) + \gamma_3(ATM_{\,j}^{\,t}))$$
$$(Pub_{\,i})(\gamma_4 + \gamma_5(\frac{VAM_{\,j}}{N_j^t}) + \gamma_6(ATM_{\,j}^{\,t}))]. \qquad (1.24)$$

In these approximations *Pop* is the population, *Dense* is population density, *ATMC* is the number of ATM machines, *VAM* is the Value Added in Manufacturing (which is computed for 1990), *ATM* is the average daily transactions per ATM and *MS* is the regional market share (calculated from deposit data); *Pub* is 1 if the bank is one of the old public banks, and *Prv* is 1 if *Pub* is 0. Those banks that were privatized during the period are still classified as public since the intention is to explore the effects of incumbency.[20] The variable *Big* is the total number of branches opened by the bank in a year. This variable serves as a normalization so that the number of branches opened locally is compared to the global number. Finally, the errors from these estimations are assumed to have been already accounted for in the u_{ij}^t.

1.3.3 Empirical results and discussion

Substituting the linear approximations (1.20), (1.23) and (1.24) into the nonlinear equation (1.19), the model can be estimated by Nonlinear Least Squares. The results, with aproximate *t*-statistics in parentheses, are

$$\alpha = \underset{(9.031)}{0.002289}\ (Big) \ \underset{(-0.9960)}{-0.001236}\ (89) \ \underset{(-0.9895)}{-0.001263}\ (90) \ \underset{(-1.686)}{-0.002443}\ (91)$$

$$N = \underset{(5.589)}{19.18} \ \underset{(-0.1347)}{-0.001395}\ (Pop) \ \underset{(6.559)}{+4.102e-05}\ (Pop^2)$$
$$\underset{(-1.092)}{-1.591e+05}\ (Dense) \ \underset{(0.9245)}{+0.03504}\ (ATMC)$$

[20]This is probably a good place to mention that we assume profit maximization by both the private *and* the public banks. We believe this is a sensible assumption since there is a significant flow of managerial talent between the two types of banks, so that career concerns motivate profit maximization by all managers.

$$
\begin{aligned}
K \;=\; & (Q_i)[(Prv)(\quad 15.04 \underset{(0.9854)}{} \quad \underset{(7.991)}{+0.02034} \; (VAM) \quad \underset{(0.5257)}{+0.04485} \; (ATM)) \\
& +(Pub)(\quad \underset{(1.949)}{18.99} \quad \underset{(0.5828)}{+0.03959} \; (VAM) \quad \underset{(-0.5646)}{-0.05328} \; (ATM))]
\end{aligned}
\tag{1.25}
$$

The traditional measure of goodness of fit, R^2, is the ratio of estimated sum of squares to total sum of squares and in this estimation is equal to 0.9831. However, in the linear model R^2 has an interpretation as part of an F test of the hypothesis that all of the coefficients are 0 except the constant term. Setting all of the coefficients to be 0 in this nonlinear model generates a restricted model where n_{ij}^t is predicted by n_{ij}^{t-1} as opposed to the sample mean as in a linear model. Define a generalized R^2 based on this generalized F test as

$$
R_*^2 = 1 - \frac{\sum(n_{ij}^t - \hat{n}_{ij}^t)^2}{\sum(n_{ij}^t - n_{ij}^{t-1})^2}, \tag{1.26}
$$

then R_*^2 is 0.8111. The generalized F test of the hypothesis that all of the coefficients are 0 rejects overwhelmingly.

Although R^2 and R_*^2 indicate that this specific model fits the data very well, there is still a question of whether or not another model might explain it better. In order to test the specification of the model various coefficients were introduced on the computed variables. The first test places a coefficient on P_{-i} wherever it appears in $n^r(\cdot)$. The test statistic of the t test that this coefficient is equal to 1 is -0.0006687 which gives it a p-value of 0.9995. Thus, this specification restriction is not rejected by the data for any typical significance level.

The next test expands the first. Since P_{-i} appears 3 times, let the coefficient be different for each one and test that all 3 are equal to one. The coefficients are given by

$$
n_{ij}^{t+1} = \alpha_i^t N_j^t \left[\frac{\beta_1 P_{-i} - 1 + \sqrt{Q_{ij}^t \frac{\pi_{ij}^t}{c_i}(1 - \beta_2 P_{-i})}}{\beta_3 P_{-i}} \right] + (1 - \alpha_i^t) n_{ij}^t. \tag{1.27}
$$

Table 1.2: P_{-i} Coefficients

	t statistic	p-value
β_1	-2.495	0.01280
β_2	-1.712	0.08725
β_3	1.857	0.06361

Table 1.3: N Coefficients

	t statistic	p-value
β_4	-0.1912	0.8484
β_5	-24.36	3.470e-98

The test statistics are reported in Table 1.2 for the t tests of the seperate hypotheses that each of the coefficients equal 1. These tests are not as overwhelming as the first test. In fact, for some reasonable significance levels these tests will reject. Still, the results are somewhat ambiguous as the test statistics are borderline. In addition to the seperate tests it is also possible to estimate a joint hypothesis that all three of the coefficients are equal to 1.

The last specification test uses the fact that N enters $n^r(\cdot)$ both directly and through the calculation of P_{-i}. Putting a coefficient on the direct N, β_4, and another on the indirect, β_5, allows tests of the specification restriction that each of these equals 1. The test statistics are reported in Table 1.3. This time one of the coefficients is overwhelmingly likely to be equal to 1, but the other is overwhelmingly unlikely. Taken with the test results from above the indication seems to be that the data would like more freedom in predicting the individual firms' beliefs about the likelihood of confrontation, but that if this belief is to come from a consistent model, then the one specified is a good one. i.e. A better estimate could be made if each firm were allowed to use a different value of N to predict the likelihood of confrontation from the value used in determining their own actions. Also if the N's are taken to be the same then a better estimate might be made if the firms could be inconsistent about there use of the prediction of the probability of confrontation. Still, if consistency is imposed the specified model is valid.

The coefficient on $ATMC$ is positive, so more ATM machines in a region implies more potential sites for bank branches. This would follow if ATM's and branches are providing complementary services. Then, more ATM's would indicate more demand for banking and, thus, more potential sites for branches. On the other hand, if they are substitutes, then an ATM would act as a branch at that location for every bank since all banks use the one ATM network. Thus, a region with a given number of nodes initially would seem to have fewer nodes for actual branches the more of them were taken out of competition by an ATM machine. We can, therefore, conclude that bank branches and ATM's are providing complementary services.

The coefficient on *Pop* is negative and that on Pop^2 is positive, which would indicate a U-shaped curve. This shape and the fact that the regions of Lisbon and Oporto are high outliers in the distribution of population indicates that these regions constitute much larger markets than the other regions. The negative coefficient on *Dense* indicates that in denser sections there are fewer nodes with larger populations. These results accord well with a perception of the urban markets as large and profitable.

The last set of coefficients that offers an interesting interpretation are those on *ATM* for the private and public banks. The fact that this is positive for private banks and negative for public banks can be understood by viewing this variable as an indication of depositor sophistication. Customers who use the ATM network extensively probably have lower switching costs, so in a region where there is a large proportion of this type of customer the established banks have less of an advantage.

We are interested in comparing the explanation of our model with the alternative story of capacity expansion. A pure capacity expansion situation can be represented by letting P_{-i} tend to 1. In such a story, a new branch serves mainly customers who were not being served before, and in the extreme is purely market expanding; in other words, the probability of such a new branch being at an uncontested node is close to 1. The parameters estimated in (1.25) produce values of P_{-i} that range from 1.042e-16 to .1240 with a mean of only .05591. A test of the hypothesis $P_{-i} = 1$ based on the generalized Wald statistic rejects overwhelmingly. Moreover, the fact that the distribution of the estimated values of P_{-i} is so tightly packed toward 0 means that

Table 1.4: Average Values of Explanatory Variables

	1988	1989	1990
Population	522.0	521.4	520.5
Urban	1843.	1844.	1843.
Rural	356.8	356.2	355.1
ATM (average daily use)	133.9	159.2	161.4
Urban	219.9	219.3	205.2
Rural	123.2	151.6	156.0
ATMC (number of machines)	17.94	27.61	42.78
Urban	100.0	154.0	227.0
Rural	7.688	11.81	19.75
Area	4.933e+07		
Urban	2.550e+07		
Rural	5.231e+07		
Value Added Manufacturing			6.688e+04
Urban			2.605e+05
Rural			4.268e+04
Market Share	0.02857	0.02857	0.02857
Public Urban	0.06496	0.06354	0.06050
Private Urban	0.004309	0.005256	0.007283
Public Rural	0.07064	0.07067	0.07030
Private Rural	0.0005258	0.0005066	0.0007552

new branches are seen to be very likely to be at contested nodes; in other words, new branches play a predominantly business stealing role in the market.

We are particularly interested in comparing private and public banks and their behavior in the two types of regions. Table 1.4 provides some average values of the explanatory variables and Table 1.5 provides averages of the computed variables for the relevant groups. Recall that K is the ratio of expected profits to startup costs. Dividing this term by Q, the share of total profits at a node that a firm can expect, yields the firm's perception of the ratio of profits to startup costs that the firm would receive if it was a monopolist, π/c. Examining the average values of this variable seems to contradict the popular conception that urban markets are more profitable than rural ones and that private banks are more efficient than public banks other things equal. But, other things are not equal. In particular, there is a significant difference in the distributions of the variable Q. Public banks have more customer loyalty, so a higher Q on average. The significance of this loyalty is demonstrated

Table 1.5: Average Values of Computed Variables by Group

	1988	1989	1990
N			
Urban	152.3	154.8	157.0
Rural	24.62	24.90	25.45
K			
Public	8.025	8.566	7.279
Urban	7.298	7.403	8.556
Rural	8.119	8.717	7.106
Private	3.972	3.801	2.713
Urban	4.240	4.241	3.662
Rural	3.506	3.262	1.765
Q			
Public	0.1175	0.1320	0.1190
Urban	0.09412	0.09663	0.1152
Rural	0.1205	0.1366	0.1195
Private	0.06146	0.05930	0.04397
Urban	0.07005	0.06978	0.06117
Rural	0.04657	0.04643	0.02677
π/c			
Public	70.43	68.52	67.49
Private	64.52	64.98	66.66
Urban	69.27	68.38	66.92
Rural	69.55	67.74	67.40
α			
Public	0.01357	0.01791	0.03404
Private	0.01495	0.04409	0.06569

Table 1.6: Descriptive Statistics for Computed Variables

Variable	Mean	Min	Max	Std Dev
α	0.02634	-0.007021	0.2356	0.04347
N	50.93	19.47	190.6	54.64
K	7.084	0.002006	63.14	7.135
$n^r(\cdot)$	45.75	1.000	190.6	52.80
$N - n^r(\cdot)$	5.172	0.000	189.6	24.74
$n^r(\cdot) - n^t$	39.56	0.000	189.6	49.52
P_{-i}	0.05591	1.042e-16	0.1240	0.03777
Q_i	0.1096	3.178e-05	0.9527	0.07494

by the observation that the expected ratio of profits to costs, K, is higher for public banks than for private banks. Customer loyalty, then, is a significant "substitute" for efficiency. Further examination of K reveals that an average public bank should not open a branch in an urban region, since it has a higher return to opening one in a rural region. Also, an average private bank should not open a branch in a rural region. Thus, the branching pattern is explained. The private banks, having less customer loyalty, focused on the lucrative urban markets. The public banks have also expanded slightly in these lucrative markets, but have acted more in the rural markets where their loyal customer base is larger. The reversal of this story for public banks in the last period is probably a result of privatization, but could represent a successful reaction to entry, or diminishing returns in the rural regions.

Besides the traditional goodness of fit measure, R^2, we are also concerned with the reasonableness of the computed variables. Descriptive statistics for these have been reported in Table 1.6. As a whole these do seem to be reasonable. There are some players who seek a saturation of the market, $n^r(\cdot) = N$, and there are others who are barely participating. The average of 90% saturation seems reasonable when the average region has 51 locations.

However, there is obviously a problem at the low end of the distribution of the adjustment rate, α. A negative adjustment rate is clearly wrong. One explanation for this shortcoming is that Nonlinear Least Squares does not do well when the error distribution has thick tales. In this data there is a large disparity between the growers and the non-growers. If this difference is not accounted for sufficiently in the

Table 1.7: Predicted and Actual Growth

	Urban areas		Rural Areas	
	Pred.	Actual	Pred.	Actual
Public	197	172	212	202
BESCL	34	26	41.5	36
BPSM	23.4	22	27.9	21
BTA	26	20	31.3	27
CGD	44.5	44	25.9	51
Private	183	240	107	133
BCI	45.7	46	19.5	23
BCP	121	125	86.1	92
Total	380	412	319	335

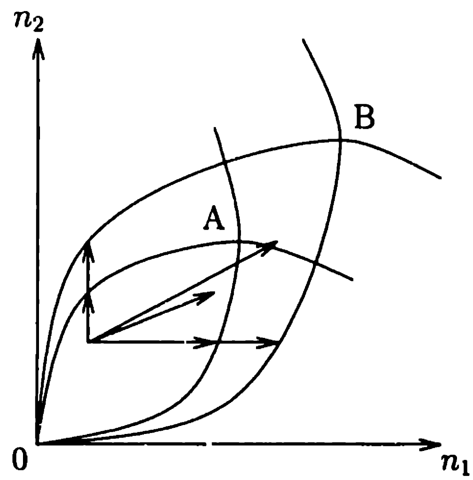

Figure 1-2: Effect of Lower Reaction Curves

explanatory variables, then the errors will have thick tales. Banks either expanded a lot or very little. This disparity probably arises from different business practices. One bank may consider depositors an essential source of funds and compete heavily in branches while another prefers the interbank market and uses branches for loans only. Since this analysis has focused exclusively on the deposits side of banking, and therefore assumed that all banks had the same business practice, the data have a problem explaining banks that do not expand. Such a problem would lead to low values of α and is indicated by the fact that the fitted values of growth in Table 1.7 are low for the private banks since that group would contain most of the banks with business practices that are not focussed on deposit taking in Portugal.

Another explanation of the low adjustment rate, however, is also interesting. If banks perceive currently high profitability as a temporary state, then they will react not to the current levels but to their expected levels discounted over time. This behavior would mean that banks are not adjusting to today's reaction curves but to lower curves. Figure 1-2 illustrates how a large adjustment toward reaction curves that are lower than the current ones can look like a small adjustment rate. The short vectors are the laws of motion based on the reaction curves with equilibrium at A. The longer ones are those for the curves that intersect at B. The adjustment rate is the ratio of the actual distance travelled along the diagonal vector to the length of that vector, so if the longer vector is used the adjustment rate looks smaller. It is possible that banks have expectations of future profitability that make them adjust in a slightly different direction than today's reaction curves would indicate. If that is true, then there will be some mis-estimation of the reaction functions. The difference in the direction of adjustment can also be caused by an asymmetry in the starting point. In this case, the estimated reaction functions will be slightly larger than the true ones for the firms that have lower initial numbers of branches.

While the low end of the distribution of estimated adjustment rates looks bad, the high end looks good. It would be difficult to say that the adjustment rate should be much higher than 25% since the growth patterns do not taper off that much. Thus, the problem seems to be that more of the adjustment rates should be at the higher end of the observed range, rather than that the observed range is too low. This property lends credence to the first explanation, different business practices. This model is, perhaps, not appropriate for some of the participants in Portuguese banking.

Nevertheless, the estimation of the empirical model allows the numerical computation of an interesting counterfactual. What would the public banks have done if the private banks had not entered the market? By using the estimated parameter values on a data set that contains only the public banks we can address a question that is very close to this one. The difference arises from the fact that the parametric form of K cannot be decomposed to fit the counterfactual. We cannot say what the market shares would have been without entry, but this problem is less worrisome than it appears. First, a new bank has very little customer loyalty. In this model a

Table 1.8: New Branches In All Regions 1988-1991

Bank	Actual Number	Predicted Number	Predicted Number If No New Banks	% Difference In Predictions
BBI	3.00	1.88	1.85	0.798
BCA	0.00	-0.460	-0.460	
BESCL	62.0	75.5	76.1	-0.977
BFB	15.0	18.3	18.3	-0.00823
BNU	7.00	6.83	6.95	-1.62
BPSM	43.0	51.4	51.9	-1.19
BPA	19.0	22.3	22.4	-0.279
BTA	47.0	57.3	57.5	-0.531
UBP	13.0	14.7	14.9	-1.09
CGD	95.0	70.4	70.9	-0.503
CPP	38.0	52.4	52.5	-0.280
BFE	8.00	9.98	9.96	0.292
BMELL	2.00	0.00	0.00	0.00
CEL/MG	22.0	29.2	29.2	-0.0719

new bank is, therefore, not a significant factor in the expected share of revenue, Q_i, of other banks. Note that a new bank is still important to the level of their profits, however, through the increased density of the system (i.e. through the P terms). Second, a case can be made that since the people most likely to have switched to the new banks are the people with the lowest switching costs. This self selection means that the fact that we do not know who these customers would have been loyal to in the counterfactual is not very important since they would not have been very loyal. We have chosen to allocate the loyal customers of the new banks to the old banks proportionately to the old banks' current loyalty. This is somewhat ad hoc but the above arguments indicate it is not overly significant.

Table 1.8 presents the results of the counterfactual analysis. While entry seems to have had a slight stifling affect on the branching of incumbents, there is evidence of increased branching by some incumbents as a result of entry. The increase is small in terms of the total number of branches opened, but it is of the same magnitude as the downward shifts. The fact that the increases are as significant as the decreases indicates that the case of simultaneous strategic complements and substitutes is more than just an interesting theoretical diversion. It is, quite possibly, the reality in this industry.

1.4 Conclusion

In this paper, we have developed a general model of oligopoly competition between multi-product, or multi-location, firms. In particular, the model indicates two important factors in the analysis of banks that compete by establishing branch networks. First, that switching costs and consumer loyalty are substitutes of efficiency in determining the profitability of branching. And second, that it may be optimal for a bank to increase its own size in response to an increase in its rival's size. The analysis of the data reveals that both of these characteristics are important in explaining the pattern of branch network expansion in Portugal.

While we have mainly focused on issues of positive analysis, the results of the paper have interesting policy implications as well. It is possible that the large customer loyalty in rural regions inhibits social efficiency, in that this loyalty prompts a less efficient bank to be the one opening a new branch. Moreover, pushing the new banks to open branches in the rural regions could lead to even more branching by the incumbents, an effect which would likely be socially wasteful. Efforts to reduce customer switching costs would then be beneficial both directly *and indirectly.* The incentives to open new branches would place greater weight on efficiency, and entry would be promoted in a less risky manner.

Chapter 2

Fitting Learning And Evolution Models to Experimental Data

2.1 Introduction

The traditional solution concept of game theory, Nash equilibrium, relies on some stringent assumptions. One is that all players have perfect information about their opponents' strategies. Another is that each player is perfectly rational and responds to the opponents' strategies with a best response. Many authors have explored the results of relaxing these assumptions by replacing the single game with an infinite repetition of the same game where the only connection between rounds is a dynamic process on the information or the response of the players.[1] There are two primary strains to this literature. Learning models assume that players are still rational, but initially lack information which they learn over time. Evolution models assume that the players can have all of the relevant information but not choose the fully rational action. "Natural selection" works over time to eliminate players who choose worse actions. These categories are not mutually exclusive nor are they complete.[2]

[1]See Young (1993), Matsui and Rob (1991) and Ellison (1993a) for example.

[2]A hybrid type of model, for example, is the myopic best response. These models assume that players use only information from the most recent repetition and try to make a best response to this information but the players also randomly "mutate" and play some other action. An example of these models is Kandori, Mailath and Rob (1992).

Researchers in this dynamic-process approach have typically focused on convergence results as the number of repetitions approaches infinity. When information is eventually learned or evolutionary pressure has removed all of the players who do not play the best response, the systems converge and the distribution of play approaches a Nash equilibrium. In this sense, certain Nash equilibria might be said to be justified on the grounds that with even looser assumptions players would learn (or evolve or whatever) to play according to the equilibrium. Unfortunately, this equilibrium-selection aspect of the literature suffers from the same problems as the explicit equilibrium-selection literature. The equilibrium selected can differ depending on the mechanism used, and the modeller is left with no basis for choosing one selection mechanism over another other than his or her own faith in the assumptions. Ideally, such faith should follow from objective empirical evidence that the assumptions are a good approximation of reality.

While allowing for more realistic decision makers than under the Nash assumptions, the models introduce another distortion from reality in terms of the structure. There are very few games of interest that have been played infinitely often, let alone being played by restricted pools of players with intertemporal effects limited to the specified dynamic. In response to this fact, it is often conjectured that players will recognize games which are technically different as being similar enough to apply learning from one to the other.[3] Still, there are likely to be many situations of interest that are new and different or which are not similar to any game that has been played enough to make convergence results relevant. Thus, there is no *ex ante* justification for evaluating the models on the basis of convergence results.

While game theorists have been developing models of learning and evolution, experimental economists have typically viewed learning as a hurdle to be overcome in collecting their data.[4] The typical experimental procedure is to have the subjects

[3] Li Calzi (1992) illustrates some of the difficulties involved in modelling learning in similar games but shows that it is possible. Brandts and Holt (1992) find that experience in very similar experimental games is significant.

[4] For example see Camerer and Weigelt (1991). For an exception see Van Huyck, et al (1991). Also, Crawford (1993) and Selten and Stoecker (1986) use the entire experiment to estimate Markov processes that represent learning models. Roth and Erev (1993) use early period data to calibrate

play the same game a large number of times and to focus on the behavior toward the end of this process. Just as in the theoretical literature, players are led to treat each repetition as a separate instance of a one period game through the devices of random matching and anonymous opponents so there is a negligible effect on the rest of the experiment from choices in any period. The focus on final periods in the experimental procedure comes naturally from observing that in early periods behavior tends to be widely scattered and tends not to support any prediction over any other, while the play in later periods tends to converge to something and some predictions can be rejected.[5] The situation created in the laboratory is a very close match to the situation postulated in the theory and a cursory examination of some of the results indicates some support for the notion of an underlying dynamic system of behavioral adjustment that tends to converge.

While the data on experimental play in early rounds is not very useful in distinguishing between static predictions, the path of play in these experiments may be ideal for comparing and evaluating the models proposed by theorists. Since two theoretical models will use different dynamic systems, they will have different predictions about the path of the observations from an experiment even when they agree on the convergence behavior. Specifically, each dynamic system corresponds to a particular Markov process in the behavioral state space and to a unique likelihood function for evaluating a path of play. The theoretical models produce parameterized functional forms for the processes. The data from experimental subjects playing the same game repeatedly against randomly selected, anonymous opponents are observations on these processes. Thus, the definition of a good relaxation of perfect rationality will be that the model fits the data from a finitely repeated experiment.

Section 2.2 provides a summary of the experiment that will be used as a source of data in this analysis. Section 2.3 describes the particular theoretical models to be compared and section 2.4 describes the statistical methods used as well as reporting the results obtained. The particular feedback from past play that each model predicts

---

the starting point of a numerical approximation of learning.

[5]This focus on final periods also takes the form of discussing the effect of experience on play as in Van Huyck, et al (1990).

will be important in determining future play turn out to be significant and to act in the predicted fashion. However, each model has weaknesses. Learning insists that players will not play a dominated action, but enough of them do that this failure is significant. Evolution requires that the "gene pool" be an important variable, but how the "gene pool" affects decisions requires some interpretation. Fortunately, the models are complementary in their strengths and weaknesses. Evolution can be a good description of players choosing a dominated action and learning is a good interpretation of the "gene pool's" role. A mixed model of evolution and learning is estimated and the complementarity is found to be significant.

2.2 The Experimental Data

While the optimal results would be obtained from conducting an experiment specifically designed to generate data for comparing the theories, such an experiment is not truly necessary. The theories are applicable to the games already studied in the experimental literature and the data from these studies has not been used in this way before. Douglas DeJong was gracious enough to provide the extensive data from the experiment used in "Selection Criteria in Coordination Games: Some Experimental Results" by Cooper, DeJong, Forsythe and Ross (1990). The experimental method is described fully there, and comparisons are made of various equilibrium-selection predictions based on their fit to the end of experiment behavior. The extensive form of the data allows tracking each individual's choices in each period. Thus, the econometric techniques of discrete choice can be applied rather than relying on an analysis of aggregate data.

The experiment involved eleven subjects who played a symmetric 3x3 game twenty times. The subjects were randomly matched each period to anonymous opponents and were told afterward the action of their particular opponent only. The experiment itself was then repeated six times with a new pool of subjects each time and a slight variation of the game. The random matching of the experiment was not the ideal of the theorist. Each player knew that there were only ten possible opponents and was told that he would face each one exactly twice. Nevertheless, this experiment is fairly close to the theoretical assumptions.

Table 2.1: Experimental Games

Game 1

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

Game 2

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

Game 3

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

Game 4

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

Game 5

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

Game 6

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

Table 2.1 are the payoff matrices for the six games.[6] The experimental payoffs were points (out of 1000) for a lottery of $1 after the game so they should represent utility. The six games are all a variant on a theme. There is a particular symmetric 2x2 coordination game common to all of the games. The games differ only in the payoffs to the third action which is always dominated. While traditional non-cooperative game theory says that such a variation does not affect the game, other approaches disagree. In Games 1-4 the combined payoff maximizing outcome (which may be a cooperative outcome) occurs at (3,3). Thus, players could have an approach to the game that is fundamentally different from the maximization of the stated payoffs and which causes them to view the games as different.

Figures 1-6 characterize the true distribution of play in each period in each game and figure 7 does so for the games taken as a whole. In general each game has an

[6]Note that the numbering of the games here differs from that in Cooper, et al. since a dominant strategy equilibrium game and an asymmetric game have been omitted from the present treatment.

action that is played by more people over time. In games 1 and 2 this action is action 1 and in the rest of the games it is action 2. Since the pure strategy Nash equilibria of these games are (1,1) and (2,2), the population distributions are converging to a Nash equilibrium. However, the rate of convergence differs across the games whether that rate is measured as the rate at which the distribution accumulates on one action or as the first period in which all of the weight is on one action. Games 4 and 5 appear to have converged to everyone playing action 2, while game 3 would appear to have done the same if the last round were omitted. This is not to say that the last round of game 3 should be ignored. Indeed, it is particularly interesting that someone would choose in the last round of a game to play a dominated action when no one else had played that action for at least seven rounds including the person who did it at the end. What kind of signal could that person have gotten to induce him to switch his play in such a way? This person was not alone. In games 1 and 6 there is a person who plays action 3 once after an even longer spell and then doesn't play it again. Game 2 has three people switching to action 3 at the same time after two rounds with no one playing it and then at least one of them stops right away. It is such peculiarities arising out of considering the path of play as a connected stream that will serve to distinguish the fits of the dynamic models.

Figure 7 does not show a significant growth of the probability of one action, certainly not to the extent of one action becoming overwhelmingly more likely than any other. This "failure to converge" indicates that a person with no knowledge of population specific history could make a good prediction of 30% action 1, 65% action 2 and that if this person were told the number of times the people had played the game the only significant change would be in adjusting the probability of seeing action 3. While there is a mixed strategy equilibrium, it has every player putting probability 2/3 on action 1 and 1/3 on action 2. In fact, this figure does not even represent convergence to something other than a Nash equilibrium. Rather, the figure represents a problem of aggregation. The individual game figures indicate that the mixture of actions comes from aggregating games that move toward action 1 with games that move toward action 2, so a considerable improvement could be made by

Table 2.2: Transition Probabilities By Game

	Game 1		
	1	2	3
1	.8844	.3023	.4737
2	.0816	.6279	.1579
3	.0340	.0698	.3684
#	147	43	19

	Game 2		
	1	2	3
1	.8219	.4516	.5313
2	.0685	.3548	.1875
3	.1096	.1935	.2813
#	146	31	32

	Game 3		
	1	2	3
1	.2000	.0156	0
2	.8000	.9323	.8333
3	0	.0521	.1667
#	5	192	12

	Game 4		
	1	2	3
1	.6000	.0185	.0938
2	.2667	.9753	.1563
3	.1333	.0062	.7500
#	15	162	32

	Game 5		
	1	2	3
1	.5000	.0449	.6667
2	.4643	.9551	.3333
3	.0357	0	0
#	28	178	3

	Game 6		
	1	2	3
1	.6769	.1007	.6000
2	.2615	.8921	.4000
3	.0615	.0072	0
#	65	139	5

knowing more details about the relevant players.

The simple Markov transition matrices from last period's action to the current period's are estimated in Table 2.2 by game (so averaging across players and periods) and in Table 2.3 by period. The possible Nash equilibria are not absorbing states but they are fairly close. For example, in game 3 few players stopped playing action 2 once they started and many players switched from the other actions to action 2. It is interesting to note that what distinguishes action 3 from the other two is not necessarily an inability to retain players but an inability to attract new players. Table 2.3 indicates that while the three actions started out relatively equal, actions 1 and 2 increased their retention rate and the number of people playing them while action 3 lost players and only had repeat players until round 9. This pattern implies a dynamic system that eventually eliminates unprofitable play and which leads to coordination on one of the possible Nash equilibria.

The mechanics of this dynamic system are illustrated in Table 2.4 which reports the sample probability of each action conditional on the previous outcome. The outcomes are listed with the player's own action first. Two observations are apparrent for actions 1 and 2. The first is that once a player plays 1 or 2 he is much more likely to do so again than to play any other action no matter what his opponent played. The second observation is that if a player's opponent plays 1 or 2, then the player

Table 2.3: Transition Probabilities By Period

	Period 2		
	1	2	3
1	.6500	.1250	.2143
2	.1500	.7500	.1429
3	.2000	.1250	.6429
#	20	32	14

	Period 3		
	1	2	3
1	.8000	.0345	.3529
2	.1500	.8621	.2353
3	.0500	.1034	.4118
#	20	29	17

	Period 4		
	1	2	3
1	.7391	.0625	.0909
2	.1304	.8438	.2727
3	.1304	.0938	.6364
#	23	32	11

	Period 5		
	1	2	3
1	.7500	.1515	.3077
2	.2500	.7879	.3077
3	0	.0606	.3846
#	20	33	13

	Period 6		
	1	2	3
1	.5833	.1714	.1429
2	.2500	.8286	0
3	.1667	0	.8571
#	24	35	7

	Period 7		
	1	2	3
1	.6667	.2286	.6000
2	.3333	.7429	.1000
3	0	.0286	.3000
#	21	35	10

	Period 8		
	1	2	3
1	.8214	.1471	0
2	.1429	.8529	.2500
3	.0357	0	.7500
#	28	34	4

	Period 9		
	1	2	3
1	.8214	.0588	.2500
2	.1429	.9412	.2500
3	.0357	0	.5000
#	28	34	4

	Period 10		
	1	2	3
1	.6923	.0811	.6667
2	.2692	.8649	.3333
3	.0385	.0541	0
#	26	37	3

	Period 11		
	1	2	3
1	.6957	.0250	.3333
2	.1304	.9750	.6667
3	.1739	0	0
#	23	40	3

	Period 12		
	1	2	3
1	1	0	.2500
2	0	.9773	.7500
3	0	.0227	0
#	18	44	4

	Period 13		
	1	2	3
1	.7895	.0435	0
2	.1053	.9348	1
3	.1053	.0217	0
#	19	46	1

	Period 14		
	1	2	3
1	.7647	.0870	.3333
2	.2353	.8913	.6667
3	0	.0217	0
#	17	46	3

	Period 15		
	1	2	3
1	.8333	.1064	0
2	.1667	.8723	1
3	0	.0213	0
#	18	47	1

	Period 16		
	1	2	3
1	.9500	.0889	1
2	.0500	.9111	0
3	0	0	0
#	20	45	1

	Period 17		
	1	2	3
1	.7500	0	NA
2	.0833	1	NA
3	.1667	0	NA
#	24	42	0

	Period 18		
	1	2	3
1	.8333	.0227	.7500
2	.1667	.9545	.2500
3	0	.0227	0
#	18	44	4

	Period 19		
	1	2	3
1	.8947	.0435	1
2	0	.9565	0
3	.1053	0	0
#	19	46	1

	Period 20		
	1	2	3
1	.9500	0	1
2	0	.9773	0
3	.0500	.0227	0
#	20	44	2

Table 2.4: Action Distribution Conditional on Previous Outcome

Previous Outcome	Current Action 1	2	3	Number
(1,1)	0.8348	0.0870	0.0783	230
(1,2)	0.6423	0.2920	0.0657	137
(1,3)	0.9744	0.0000	0.0256	39
(2,1)	0.2628	0.6861	0.0511	137
(2,2)	0.0248	0.9592	0.0160	564
(2,3)	0.1136	0.7727	0.1136	44
(3,1)	0.5641	0.1795	0.2564	39
(3,2)	0.2045	0.3636	0.4318	44
(3,3)	0.1500	0.2000	0.6500	20

is more likely to adopt that action than if his opponent played another action. For action 3, however, neither of these is true. After a (1,3) or a (3,1), both players are more likely to play action 1. Playing 3 after (3,3) has only the same probability as 2 after (2,1) or 1 after (1,2). Both of these facts indicate the importance of previous period's profits in determining a player's propensity to play action 3 repeatedly.

In short, the data indicate that there is some underlying dynamic process pushing play toward one of the two Nash equilibria in each game. While this fact supports the broad research program of modelling such a process we are also interested in how closely particular models fit to the true process.

2.3 The Theoretical Models

Fictitious play will serve as the paragon of the learning models, but since the actual estimation will introduce errors it will look a bit strange. The replicator dynamic would be the natural candidate to represent evolution models, except that it is a continuous-time model with dubious applicability to the experimental setting. Instead, the economic muddler model which Binmore and Samuelson (1992) present as a discrete time extension of the replicator dynamic for economic modelling will be used. Table 2.5 represents the games in a generalized fashion which will be used in this section.

Table 2.5: Generalized Game

	1	2	3
1	(350,350)	(350,250)	(a,0)
2	(250,350)	(550,550)	(b,0)
3	(0,a)	(0,b)	(c,c)

a>c in all games, so action 1 always dominates action 3.

2.3.1 The Learning Model

Fictitious play assumes that each player believes his opponents to be playing a particular mixed strategy each period. The player does not know these strategies, so he is assumed to have Dirichelet priors over the possible mixed strategies. Each period all players play a best response to their current beliefs about the actions of their opponents and update their beliefs on the basis of what their opponent that period played. Assuming Dirichelet priors over the possible mixed strategies and Bayesian updating implies that the beliefs over current opponent's actions have a conveniently simple form. The priors over current opponent's actions of player i can be represented by n_{i1}, n_{i2}, n_{i3}. If at round t the number of times player i has seen action j is x^t_{ij}, then

$$P^t_{ij} = \frac{n_{ij} + x^t_{ij}}{\sum_k (n_{ik} + x^t_{ik})} \tag{2.1}$$

where P^t_{ij} is the probability that player i assigns to the event that his period t opponent is playing action j. Note that the assumption of a common mixed strategy and updating beliefs about what that mixed strategy is over time ignores the possibility that opponents' strategies change. Assuming a constant strategy by all opponents is particularly limited thinking for a player who is doing the updating in order to change his own behavior. Thus, strategies are fully rational given beliefs, but belief formation is not assumed to be consistent with common knowledge of the game, or of the learning process itself.[7]

[7]For a more formal and extensive justification of fictitious play as a learning model see Fudenberg and Kreps (1991). For a learning model that is fully rational see Kahlai and Lehrer (?).

Since 3 is dominated, a fictitious-play player would never choose it. Therefore, the decision is only between 1 and 2. A player (i) will choose to play action 1 in period t if

$$(100)P_{i1}^t - (200)P_{i2}^t + (a-b)P_{i3}^t > 0, \tag{2.2}$$

where a and b are the parameters in Table 2.5. Substituting for the P_{ij}^t condition (2.2) becomes

$$-100n_{i1} + 200n_{i2} - (a-b)n_{i3} < 100x_{i1}^t - 200x_{i2}^t + (a-b)x_{i3}^t. \tag{2.3}$$

Even ignoring the fact that some players play action 3, it is not possible to construct beliefs of this form so that condition (2.3) is satisfied if and only if a player plays action 1. This failure is not surprising since any number of factors could interfere between the realization of the condition and the realization of the action. For example the players could have non-material concerns over the outcomes which contribute in an unobserved way to their utility functions.[8] The players could also be experimenting with actions that they would not otherwise play to be sure of the consequences.[9] We will say that fictitious-play learning is an important determinant of behavior, then, if this interference is such that when condition (2.3) is close to equality the interference may be strong, but when the term on the left (N_i) is much smaller than the term on the right (h_i^t) it is much more likely that the player will play action 1. Letting $G(\cdot)$ be the distribution of an independant error term incorporating the interference, the distribution of play is

$$\begin{aligned} Prob(a_i^t = 1|h_i^t) &= 1 - G(N_i - h_i^t) \\ Prob(a_i^t = 2|h_i^t) &= G(N_i - h_i^t). \end{aligned} \tag{2.4}$$

[8]See Rabin (1992) for an example of incorporating such concerns into game theory.

[9]Fudenberg and Kreps (1991) demonstrate how such experimentation can be important to learning in extensive-form games.

For lack of any more compelling model, let $G(\cdot)$ be logit so

$$G(N_i - h_i^t) = \frac{exp(L_0 + L_1 N_i - L_1 h_i^t)}{1 + exp(L_0 + L_1 N_i - L_1 h_i^t)}. \tag{2.5}$$

The term $L_0 + L_1 N_i$ represents the effect on the path of play of each player having unobserved and different priors. In order to examine the significance of this effect we will estimate the model under three specifications. $N_i = N$ for all i assumes no heterogeneity of priors. $N_i = N_j$ if i and j are in the same game assumes heterogeneity only across games. The final specification places no restrictions on N_i which allows for complete heterogeneity of priors.

As noted above, we also observe some instances of players playing action 3. Such observations cause an immediate rejection of the pure fictitious-play model. Unlike the interference between condition (2.3) and playing action 1, this failure cannot be reconciled to the assumptions of fictitious play. In order to isolate this failure from the question of whether or not fictitious play is an important determinant of behavior we will say that players play according to (2.4) except that with probability ϵ they play action 3, so the distribution of play is

$$\begin{aligned} Prob(a_i^t = 1 | h_i^t) &= (1-\epsilon)(1 - G(N_i - h_i^t)) \\ Prob(a_i^t = 2 | h_i^t) &= (1-\epsilon)(G(N_i - h_i^t)) \\ Prob(a_i^t = 3 | h_i^t) &= \epsilon. \end{aligned} \tag{2.6}$$

Since we observe that most of the instances of action 3 are in early periods, we will let ϵ depend on the number of rounds a player has participated in,

$$\epsilon^t = \frac{exp(\epsilon_0 + \epsilon_1 t)}{1 + exp(\epsilon_0 + \epsilon_1 t)}. \tag{2.7}$$

Within the context of fictitious play this can be loosely interpreted as allowing for the possibility that players learn to be rational as well as learning about their opponents. Since, however, we also noted that many of the action 3's were by players who had played action 3 in the previous round we will also be concerned that ϵ will be a source of heteroskedasticity.

Another factor which should influence ϵ is the tendency to play action 3 which arises from concerns other than material self-interest. In the case of ϵ some of these concerns can be modeled in such a way that differences between the experimental games can be used as a variable to control for them. Rabin (1992) provides a formal way to incorporate concerns over fairness into game theory. When Rabin's fairness equilibria of the games studied here are computed (3,3) is an equilibrium if the weighting of fairness in the player's utility is above a threshold that varies across the games. Without any reason to suspect that one pool of players has a different distribution of fairness weights than another, including the thresholds as variables to determine ϵ should improve the estimation. Letting T_i be the threshold of the game player i plays, the specification

$$\epsilon_i^t = \frac{exp(\epsilon_0 + \epsilon_1 t + \rho_0 T_i + \rho_1 t T_i)}{1 + exp(\epsilon_0 + \epsilon_1 t + \rho_0 T_i + \rho_1 t T_i)}, \tag{2.8}$$

will be used for comparison to determine if the players are choosing to play action 3 for rational reasons which are just not a part of fictitious play. This alternative specification allows the non-material concerns to affect not only the probability of players choosing action 3, but also the persistence of that rate over time. The minimum fairness weights for (3,3) to be an equilibrium of games 1-6 respectively are 800, 200, 800, 200, 400 and 1000.

2.3.2 The Evolution Model

The assumptions of the replicator dynamic come from evolution rather than rational optimization. Each player is "genetically coded" to play a particular action. Those players whose action is profitable against the current population reproduce the most. That generation is then replaced by their offspring who tend to have the same action as their predecessor but with some slight mutation to playing a different action. Fortunately no one actually dies during these experiments, but that means that there is no natural selection so the pure replicator dynamic does not apply.

Assume, instead, that the players are boundedly rational in that they play a particular strategy until it provides them with a payoff below some threshold level at which point they realize that it is not very successful and choose another strategy. If the thresholds are drawn from independent, identical distributions each period, such an agent could be described by the economic muddler model. Binmore and Samuelson (1992) specify that an agent whose payoffs are below the threshold tries to adopt the strategy of a randomly chosen member of the population. With probability $1-\lambda$ this attempt is succesful and with probability λ the agent is a "mutant" and adopts another action. A full characterization requires specifying the distribution used by the mutant in determining this other action, so let the probability of choosing any action after a mutation be 1/3. Binmore and Samuelson show that as λ goes to zero, and other variables are taken to their limits, the muddler model converges to the replicator dynamic.

This story, however, requires that the players have information that they do not actually have in these experiments. Specifically, the players only observe the actions of their own opponents each period so cannot choose just any player in the prior period to emulate. A natural adaptation is to assume that when a player updates he chooses one of his former opponents to emulate.

Let F(·) be a logit distribution for the thresholds, so

$$F(x) = \frac{exp(E_0 + E_1 x)}{1 + exp(E_0 + E_1 x)}. \tag{2.9}$$

And let π_i^t be i's payoff in period t, and $c_i^t(a)$ be the proportion of player i's opponents in periods 1 to $t-1$ who played action a. Then

$$\begin{aligned} Prob(a_i^t|\pi_i^{t-1}, a_i^t = a_i^{t-1}) = \quad & (F(\pi_i^{t-1})) + \quad (1 - F(\pi_i^{t-1}))P(a_i^t) \\ Prob(a_i^t|p_i^{t-1}, a_i^t \neq a_i^{t-1}) = \quad & \qquad\qquad\qquad (1 - F(\pi_i^{t-1}))P(a_i^t) \end{aligned} \tag{2.10}$$

where

$$P(a_i^t) = (\lambda \frac{1}{3} + (1-\lambda) c_i^t(a_i^t)) \tag{2.11}$$

is the probabilty that a player who switches chooses the action a_i^t. Note that after deciding to switch actions a player can be "reassured" of his or her original action by observing other players choosing it or can end up playing the original action again by mutating back to it.

As in the learning model the distribution on errors could depend on the player. In order to allow for heterogeneity and for comparability with the learning model we will estimate the three specifications: E_0 is the same for all players; E_0 is the same for players in a given game; and E_0 is not necessarily the same for any players.

Binmore and Samuelson first fix λ and compute the limit as $t \to \infty$ and then compute the limit of this fixed λ limit as $\lambda \to 0$. However, the value of λ is still a constant within any instance of the process. As noted above it is interesting to test the assumption that λ is unaffected by players' experience. The specification

$$\lambda^t = \frac{exp(\lambda_0 + \lambda_1 t)}{1 + exp(\lambda_0 + \lambda_1 t)} \qquad (2.12)$$

allows such a test while restricting λ to take on values in (0,1).

2.4 The Empirical Results

Tables 2.6-2.8 provide the maximum-likelihood estimated coefficients for the relevant specifications of each model with t-statistics in parentheses where available. The exceptions to this format are the columns for the mutation rates which have the rate at t=1 on the first line and for t=20 on the second and the columns for the constant terms of the logit distribution. The latter has entries of "ind" and "game" to represent that the coefficient was estimated separately for each individual and for each experimental pool respectively. For the former case the variance-covariance matrix could not be computed so t-statistics are not reported.

The learning model specifications discussed in section 2.3.1 are estimated in Tables 2.6 and 2.7. For all the specifications L_1 is positive so the more likely condition (2.2) is to fail the more likely that players will play action 2 rather than action

Table 2.6: Learning Specifications

Logit Constant (L_0)	Logit $E[\pi(1) - \pi(2)]$ (L_1)	ϵ Constant (ϵ_0)	ϵ Time (ϵ_1)	ϵ(t=1) ϵ(t=20)	lnL/n
ind	.001801	-.987233	-.175157	.23823 .01109	-.425053
game	.000859 (7.806)	-.987233 (-4.723)	-.175157 (-7.037)	.23823 .01109	-.564191
.632670 (7.386)	.001364 (17.116)	-.987233 (-4.723)	-.175157 (-7.037)	.23823 .01109	-.584685

Table 2.7: Learning Specifications With Fairness

Logit Constant (L_0)	Logit $E[\pi(1) - \pi(2)]$ (L_1)	ϵ Constant (ϵ_0)	ϵ Time (ϵ_1)	ϵ Fairness (ρ_0)	ϵ Fair*Time (ρ_1)	lnL/n
ind	.001801	-.426565	-.130084	-.000952	-.000119	-.415640
game	.000859 (7.431)	-.426566 (-1.101)	-.130084 (-2.918)	-.000952 (-1.315)	-.000119 (-1.256)	-.554778
.632670 (7.389)	.001364 (17.117)	-.426565 (-1.101)	-.130084 (-2.918)	-.000952 (-1.315)	-.000119 (-1.256)	-.575272

1. Since condition (2.2) was the condition for a rational player to choose action 1 over action 2, this result is as expected. In every specification the inconsistency parameter ϵ (the probability of playing the dominated action 3) depends significantly on experience as represented by the number of repetitions the player has participated in (coefficient ϵ_1). The starting value for the inconsistency (ϵ_0) is significantly less than 0 which in the specification used corresponds to the probability of playing action 3 initially being less than 1/2. The rate quickly drops off, so the probability at the end of the 20 periods is less than 1/20. The thresholds for an individual's fairness weight in utility to make (3,3) a fairness equilibrium also significantly affect both the initial probability of playing action 3 and the persistence of that probability. Since ρ_0 is negative, players in games with lower thresholds (so the weights are more likely to be on the (3,3) side) are more likely to play action 3 initially. Since ρ_1 (the coefficient on the threshold interactive with experience) is negative, fairness also increases the persistence of playing action 3. Note that while the coefficients ρ_0 and ρ_1

Table 2.8: Evolution Specifications

Logit Constant (E_0)	Logit π^{t-1} (E_1)	λ Constant (λ_0)	λ Time (λ_1)	λ(t=1) λ(t=20)	lnL/n
ind	.012888	.876985	-.195849	.66399 .045651	-.390520
game	.007891 (6.578)	0.628426 (-1.809)	-.224361 (-2.400)	.59966 .020658	-.453011
-2.711573 (-6.326)	0.007417 (8.065)	0.649036 (2.118)	-.220352 (-5.774)	.60556 .022799	-.460698

do not have significant t-statistics the likelihood ratio test of their joint significance overwhelmingly rejects the hypothesis that both are 0.

The evolution model specifications are estimated in Table 2.8. The coefficient E_1 is always positive so high profits in the previous period make playing the same action as in the previous period more likely. This accords with the assumption that players are switching strategies only when they fall below a threshold. The mutation rate (λ) depends on experience. The initial level of mutation (λ_0) is positive and significantly different from 0, so the probability of mutation starts well above 1/2. The dependence on experience (λ_1) is significantly negative so the probability declines over time.

Why would a mutation rate be affected by the experience of the player? Even dropping the evolution metaphor and considering λ as the mixture by a boundedly rational individual between a uniform distribution and a more reasoned distribution the dependence on experience is puzzling. Perhaps the players are becoming more rational in their approach to the game. This evolution of the evolutionary process could be accomplished by players remembering what happened to them in previous periods of playing different actions so that the computational costs of deciding which action to employ decrease with experience and we observe more rational play. This model is only a proxy, though, for the way people process information. It is likely that the dependence of λ on experience is due to the proxy becoming better than the uniform distribution as players gain experience and as play converges. Thus, the decreasing λ could indicate that the model is incomplete as there are other factors (such as

Table 2.9: Evolution - Population Distribution *vs* Observed Distribution

		1	2	3
Satisficing	Constant	-2.711572	-1.361150	-2.788616
		(-6.329)	(-5.018)	(-5.893)
	Profit Last Period	0.007417	0.004582	0.006477
		(8.069)	(7.022)	(6.382)
Mix Of Observed (vs Population) Model	Constant			0.905910
				(1.332)
	Experience			-0.067685
				(-1.239)
Probability Of Mutating (Observed)	Constant	0.649035		0.864645
		(2.121)		(1.245)
	Experience	-0.220352		-0.447879
		(-5.789)		(-1.993)
Probability Of Mutating (Population)	Constant		-0.607765	-2.981663
			(-1.303)	(-0.496)
	Experience		-0.094059	-0.061806
			(-2.540)	(-0.173)
Mean Log Likelihood		-0.460698	-0.464673	-0.434959

learning) which compel the population to converge and which so are correlated with experience.

A natural way to explore this issue is to return to the original model of Binmore and Samuelson. Recall that this model put great weight on the population distribution over actions in the previous period to determine current play. We can think of this last period population distribution as a proxy for the forces of convergence outside of the model. We re-estimated the no heterogeneity of constant term specification of evolution with the last period's population distribution replacing the sample distribution from the player's observation of past opponents' behavior. The two estimates are reported in Table 2.9 for comparison along with the estimates from a model that mixes the two distributions. The mutation rate in the new estimation (specification 2) is insignificantly below 1/2 initially, but declines slower to a much higher final value. This change in the pattern of the mutation rate's dependence on time indicates that the population distribution is a reasonable proxy for the outside convergence forces. The mixed model (specification 3) has a less significant depen-

Table 2.10: Merged Model

Probability	Cond'l Probability	Cond'l Cond'l Probability	Action
$F(\pi_i^{t-1})$			Last Period's Action
$1 - F(\pi_i^{t-1})$	μ	$G(h_i^t)$	Action 1
	μ	$1 - G(h_i^t)$	Action 2
	$1 - \mu$	$\lambda\frac{1}{3} + (1-\lambda)c^t(a)$	Any Action a

dence of mutation on experience. Likelihood ratio tests overwhelmingly reject the hypothesis that the model is complete without the proxy.

The learning model is also a simplified description of the way that players process information and could be improved by adding the possibility that players play their last period's action again unless the profits were particularly low. Such an extension would explain much of the observations of action 3. The implication is that the two models have a great deal of complementarity. In order to explore this possibility a merged model was estimated. The story of this model is that players first decide whether or not to update their action depending on whether profits were high or low last period. If the action is updated, then the new action is chosen on the basis of condition (2.3) with probability μ and is determined on the basis of the evolution updating scheme with probability $1 - \mu$. This model is outlined in Table 2.10 and the estimates for it are reported in Table 2.11. Because there may be some concern that the most recent observations are being weighted differently in the two formulations of beliefs (i.e. the learning update rule treats all past observations the same and the evolution update rule only uses last period's population distribution) we will allow the most recent observations to enter the calculation of condition (2.3) with a different coefficient.[10] We will also estimate this model substituting sample distributions from observed opponents for the previous period's population distribution. In all of these specifications both models are significant and the parameters have the correct signs. In the specifications where the evolution component uses last period's population distribution, however, it is no longer the case that the mutation rate depends

[10]This allows for both all past observations are the same and for a somewhat hyperbolic discounting of observations in the updating of priors.

Table 2.11: Merged Model Estimates

		1	2	3	4
Satisficing	Constant	-1.601022	-1.534988	-1.532919	-1.464930
		(-5.041)	(-4.882)	(-5.037)	(-4.898)
	Profit Last Period	0.004465	0.004324	0.004440	0.004298
		(5.614)	(5.484)	(5.792)	(5.658)
Learning	Constant	-1.030387	-0.868225	-0.970586	-0.853118
		(-1.857)	(-1.771)	(-2.064)	(-1.798)
	$E_t[\pi(1) - \pi(2)]$	0.002000	0.000795	0.001464	0.000335
		(3.531)	(0.851)	(4.117)	(0.507)
	$E_t[\pi(1) - \pi(2)] -$ $E_{t-1}[\pi(1) - \pi(2)]$		0.002203		0.002441
			(1.052)		(1.472)
Mix Of Learning (vs Evolution) Updating	Constant	-1.550398	-1.486832	-1.406879	-1.413977
		(-2.294)	(-2.130)	(-2.188)	(-2.175)
	Experience	0.117992	0.116409	0.101499	0.100762
		(1.922)	(1.833)	(2.036)	(2.044)
Mix Of Uniform (vs Population) Distribution	Constant	-0.042674	-0.049692		
		(-0.097)	(-0.059)		
	Experience	-0.042674	-0.044765		
		(-0.623)	(-0.566)		
Mix Of Uniform (vs Observed) Distribution	Constant			2.600734	2.705044
				(1.724)	(1.818)
	Experience			-0.537321	-0.563272
				(-1.845)	(-2.003)
Mean Log Likelihood		-0.456408	-0.455988	-0.455813	-0.454569

significantly on experience. When the evolution component uses the sample distribution based on the player's observations the dependence on experience is significant, but the initial value is even higher and the decline even steeper than in the evolution model alone with this distribution. Such a change in the path indicates that the uniform distribution is representing priors that are stronger initially but quickly updated away. These results reaffirm the implication that the proxy interpretation of evolution is correct. With this merged model it is also possible to make restrictions on the parameters to get back each of the separate models. The likelihood ratio tests of these restrictions reject for both of the models. This result indicates that the inferred complementarity of these models is significant.

The evolution model best approximates the underlying dynamic when mutation rates decrease with experience. Since in this model the mutation rate slows convergence, this result indicates that the speed at which the true system accumulates probability on an action increases over time. Some of the theoretical literature incorporates these rates as declining over time, but much of it imposes the assumption of constant rates.[11] Since the mutation rates generally act as a means of escaping a basin of attraction of the dynamic system without mutation, a decrease in these rates over time will lead to a larger role for the initial conditions in determining the state of the system at any given time than if the mutation rates had stayed at their initial levels. As this effect increases the rate of convergence, it indicates that analyzing the limits of the theoretical models may not be such an abstraction from reality as earlier rounds will be more like the limit than with a constant rate.

Both the learning model and the evolution model demonstrate an importance of heterogeneity in the constant term of the logit distribution (L_0 and E_0). The likelihood ratio tests in each case on each level of restriction produce p-values below 10^{-8}. Thus, the restrictions are rejected and the specification that allows complete heterogeneity is strongly implied to be acceptable despite the dramatic increase in the number of parameters. Recall that for the learning model this result has the

[11] For example, the experimentation rate in Fudenberg and Kreps (1991) declines over time. The mutation rate in Binmore and Samuelson (1992) does not.

interpretation of rejecting the assumption of common priors. The question arises of what heterogeneity of priors means. The likelihood ratio test does not distinguish between a small number of extreme outliers and a smooth distribution. However, in asking whether an assumption of common priors is valid it may be that 65 players with exactly the same priors and one player with very different priors is not very different from common priors. To address this issue we will need to compute N_i for each player. If $L_0 = 0$ then a statistic which converges to N_i is the ratio of the estimated constant term over the estimated L_1. Even if $L_0 \neq 0$ this statistic will still be informative of the distribution of the N_i as long as the true L_0 does not vary across individuals. Note that N_i is player i's expectation of the net profitability of action 2 versus action 1 where the expectation is with regard to player i's priors. Figure 8 is a histogram of the N_i computed from the coefficients in the full heterogeneity specification. The histogram displays a three-peaked distribution. Slightly more than half the players are in the central peak and are nicely distributed around 0. The remainder are at one of the two extremes 12,000 and -12,000. These magnitudes are large enough to prevent a fictitious-play player from switching strategies for the entire 20 periods. While there are no observations between the peaks this does seem to be a case of heterogeneous priors.

Since the nested comparison of the models rejected both in favor of the merged model, we need to introduce a mechanism for comparing non-nested models in order to look at the relative fits. One such mechanism is Akaike's Information Criterion (AIC). Under this criterion each model is assigned a score equal to the maximized log likelihood minus a parameter penalty of two times the number of parameters. The model with the highest score is chosen. AIC is an approximation to Bayesian choice with a particular loss function, but the parameter penalty has drawn criticism for arbitrariness.[12] As long as the models to be compared have the same number of parameters, however, any parameter penalty will drop out in the comparison of the AIC scores. Therefore, we will only compare models with the same number of parameters (i.e. the same degree of heterogeneity). For each level of heterogeneity

[12] See Amemiya (1985).

the evolution model has higher log likelihood and so higher AIC score than learning.

Figures 9 and 10 aggregate the individual predictions for the learning and evolution models. These graphs are representative of the other games for these models in game 1 and can usefully be compared to Figure 1. The learning model has excessive smoothness especially in the prediction of play on the dominated action 3. Since the model itself had nothing to say about this action the indication is that some exploration of a less arbitrary characterization of the reasons for playing such an action would be warranted. Since fairness turned out to be a significant determinant of this probability, a rational motivation for playing action 3 can be found even though such a motivation does not appear in this learning model. In comparison the evolution model is less smooth but is still more smooth than the data. Moreover, the evolution model reacts strongly to the previous period's behavior and can be exactly off a cycle in the data. For example, in periods 7 to 10 there is a cycle in the proportion playing action 1 (from .8 to .7 to .9 and back to .7). The evolution prediction also has a cycle in this period but it lags behind by one period so it moves in the wrong direction. The prediction in period 9 (the height of the observed cycle) is lower than in periods 8 and 10. Both models demonstrate excessive smoothness, the learning model does not handle dominated actions well, but, the evolution model places too much importance on last period's proportions for this period's predictions.

The excessive smoothness problem may be a result of another concern. The previous predictions allowed some heterogeneity, but they assumed a great deal of homogeneity. Particularly, they assumed that all of the individuals were following the same underlying dynamic process (i.e. that everyone was learning). There is no compelling reason why this should be the case. In fact, there is no reason to assume that two individuals who are following a particular model should have the same coefficients (i.e. that everyone learns at the same rate). Since the number of parameters is small it is possible to break the data into individual time series and estimate them separately. One drawback of this approach is that some individuals may not demonstrate enough change in behavior to distinguish the competing explanations of how their behavior should change. Specifically, about one third of the people never change their

Table 2.12: Results of Individual Estimations

Parameter	All Players	Non Constant	"Learners"	"Evolvers"
Learning Constant	-1.859 (58.13)	-7.115 (67.05)	-6.303 (72.2)	-11.17 (33.5)
Learning $E[\pi(1) - \pi(2)]$	.02158 (.06313)	.02938 (.07267)	.03231 (.07815)	.01473 (.03371)
Learning ϵ Constant	-6.859 (93.07)	-2.086 (109.1)	6.586 (105.2)	-45.44 (125.2)
Learning ϵ Time	-3.876 (13.02)	-5.288 (15.06)	-5.586 (14.22)	-3.802 (19.85)
Evolution Constant	-22.8 (85.66)	-33.99 (98.32)	-35.2 (107.0)	-27.99 (32.22)
Evolution π^{t-1}	.07325 (.2081)	.09107 (.2420)	.09541 (.2637)	.06934 (.07311)
Evolution λ Constant	-20.85 (202.1)	-21.04 (237.2)	-3.538 (236.1)	-108.6 (238.1)
Evolution λ Time	-.4173 (13.17)	-.9084 (15.42)	-2.642 (14.92)	7.758 (15.93)
mean lnL(Learn)	-.2538	-.3490	-.3137	-.5252
mean lnL(Evol)	-.3207	-.4410	-.4441	-.4256
# of players	66	48	40	8

Mean values of parameters across players
(Standard Dev. of parameters across players)

action. These people can all be perfectly explained by either model in an individual estimation by setting the priors or the distribution of thresholds at extreme values.

Table 2.12 lists the population mean and standard deviation for each parameter for several interesting groups of players. The parameters are estimated for each individual separately and then the estimates are grouped in this way for tractability. The first group is all of the players. The rest of the groups exclude those players who never change their action. For the last two groups a player is identified as a "learner" if the likelihood of the learning model is larger than the likelihood of the evolution model when estimated for that player alone. At the bottom of the table are the mean log likelihoods for the populations. Since the two models were estimated with the same number of parameters, a comparison of AIC scores holding the level of heterogeneity constant again reduces to comparison of the likelihoods for any parameter penalty. Thus, allowing for extreme heterogeneity across individuals reverses the earlier ordering of the two models. All of the parameters have the same signs on average as before, but the magnitudes are much larger. Since we don't have any intuition for the correct magnitude for any of these parameters interpreting this change is difficult. However, note that the most extreme values seem to be for the sub-population that is best described by the other model. The exception to that observation is the mutation rate in evolution. Here, the extreme values are realized by the evolvers and the effect of experience is reversed. Combined with the number of learners relative to the number of evolvers, the indication is that evolvers are really the outliers of the distribution which learning cannot be stretched as easily to explain. Thus, learning is a better story of individual behavior, but since individuals differ greatly in terms of how concerns outside of the scope of this particular learning model affect behavior (e.g. non-material preferences over outcomes), learning is much more sensitive to aggregation.

2.5 Conclusion

Theorists have proposed a number of competing models of how agents who do not meet the requirements of Nash equilibrium would adjust their behavior in response to certain feedback resulting from repeated play of a particular game. This phenomenon is observable in experimental games and such games have the same structure as is assumed in the theories. Since the theoretical models are essentially functional form assumptions, they can be fit to the data from experiments. Doing so allows us to explore the consequences of relaxing Nash equilibrium in two somewhat orthogonal ways and reveals that the best way incorporates the features of both. Real players in the sittuation assumed by theorists will play a given action until it gives them profits which they consider to be too low. When they do change actions, the new action is determined by the best response to what the players think their next opponent will be doing based on what the players have seen other opponents do.

The estimation also has implications for additional modelling. Both of the separate models have a dependence on experience that should be incorporated in new models.

Also, the models reject assumptions of common priors and of common aspiration levels across players. A useful extension that is not possible with this sample would be to incorporate demographic data on the players and look for some regularities of the parameters by identifiable individual characteristics. Mason and Phillips (1990) have found that experience affects men and women differently, so such an extension should prove rewarding.

As far as direct comparison of the two models goes learning is probably a better description of individual behavior than evolution. However, learning suffers from a greater sensitivity to heterogeneous factors and so is more negatively affected by aggregation. Learning does better when the parameters are distinct for individuals while evolution does better when most of the parameters are held to be the same for the population. Another usefull extension would be to examine the aggregate prediction of a learning model that allows rational players to play a dominated action by formally incorporating some non-material concerns.

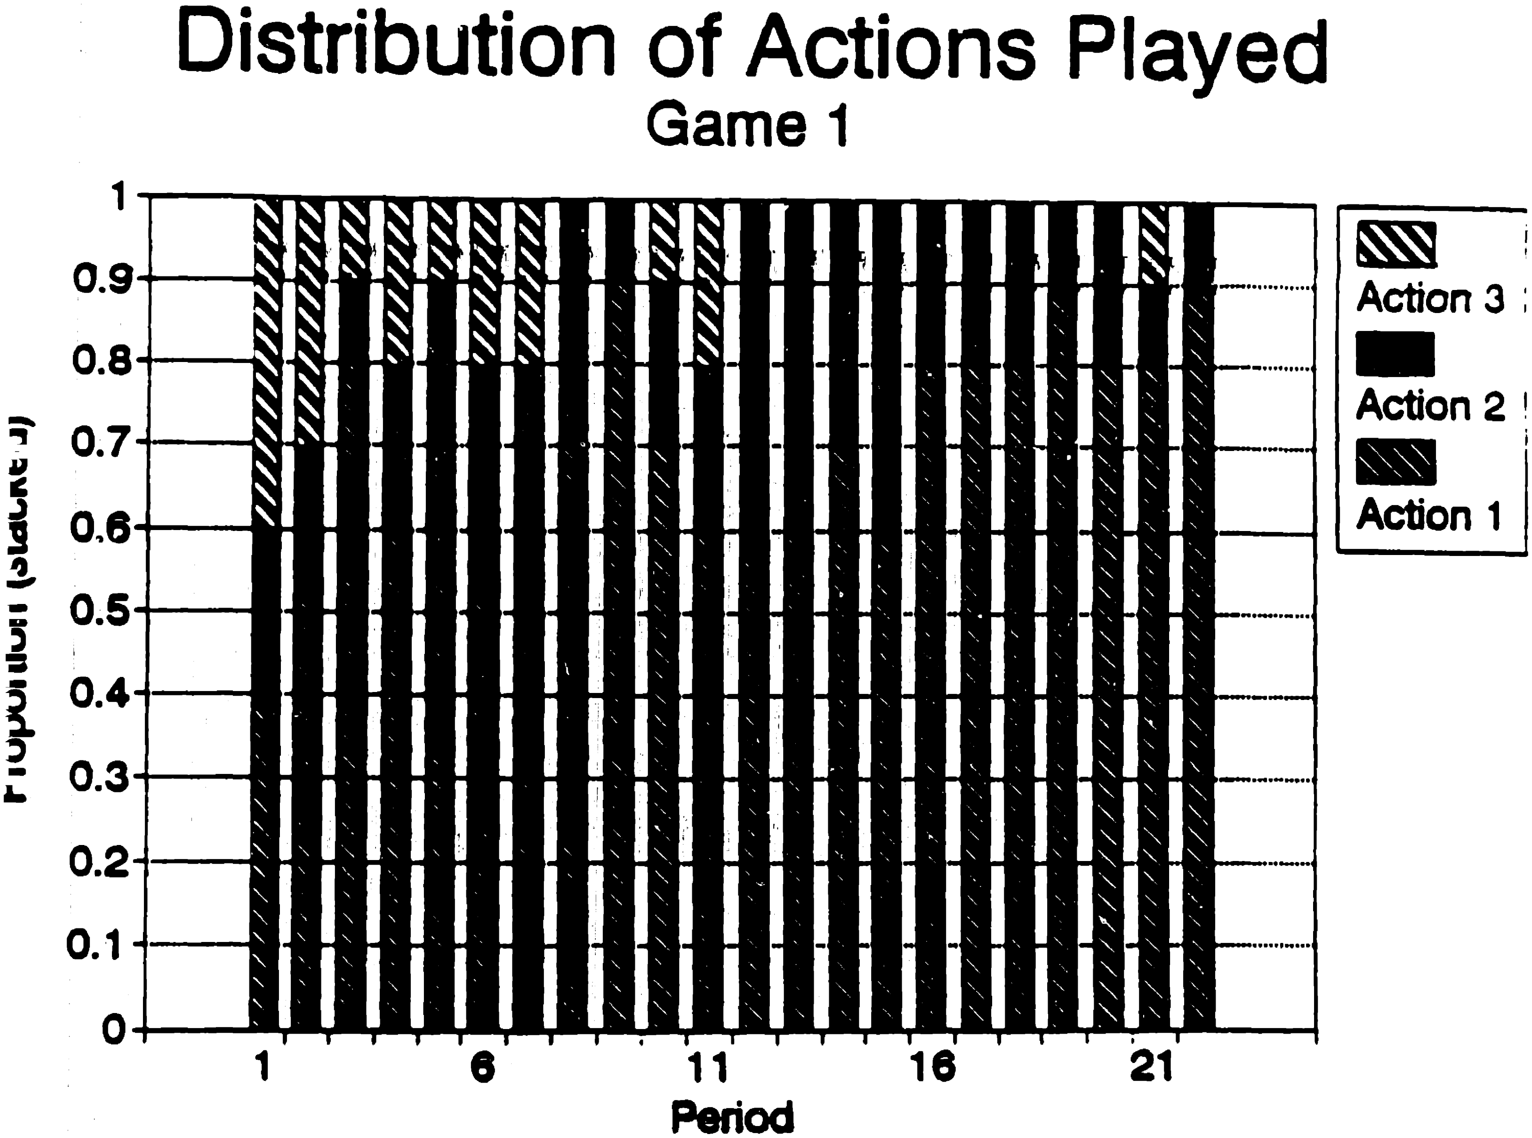

Distribution of Actions Played
Game 1
Proportion (stacked)
1
0.9
0.8
0.7
0.6
0.5
0.4
0.3
0.2
0.1
0
1
6
11
16
21
Period
Action 3
Action 2
Action 1

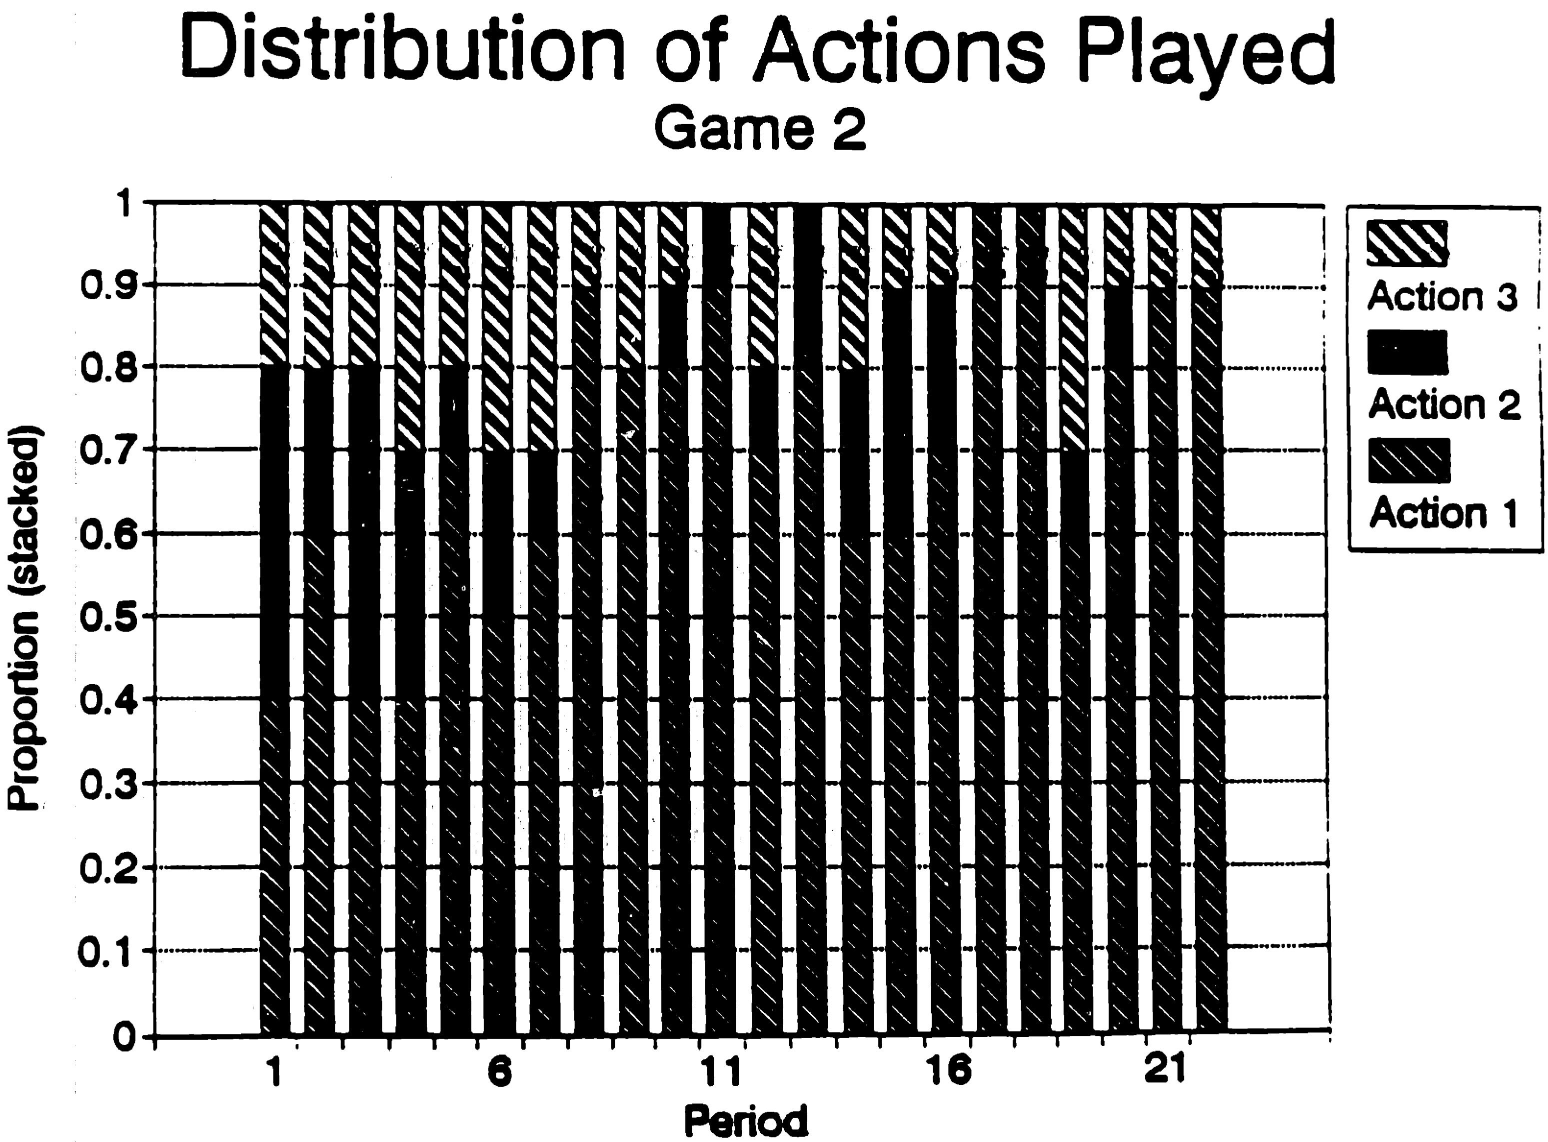

Distribution of Actions Played
Game 2
Proportion (stacked)
1
0.9
0.8
0.7
0.6
0.5
0.4
0.3
0.2
0.1
0
1
6
11
16
21
Period
Action 3
Action 2
Action 1

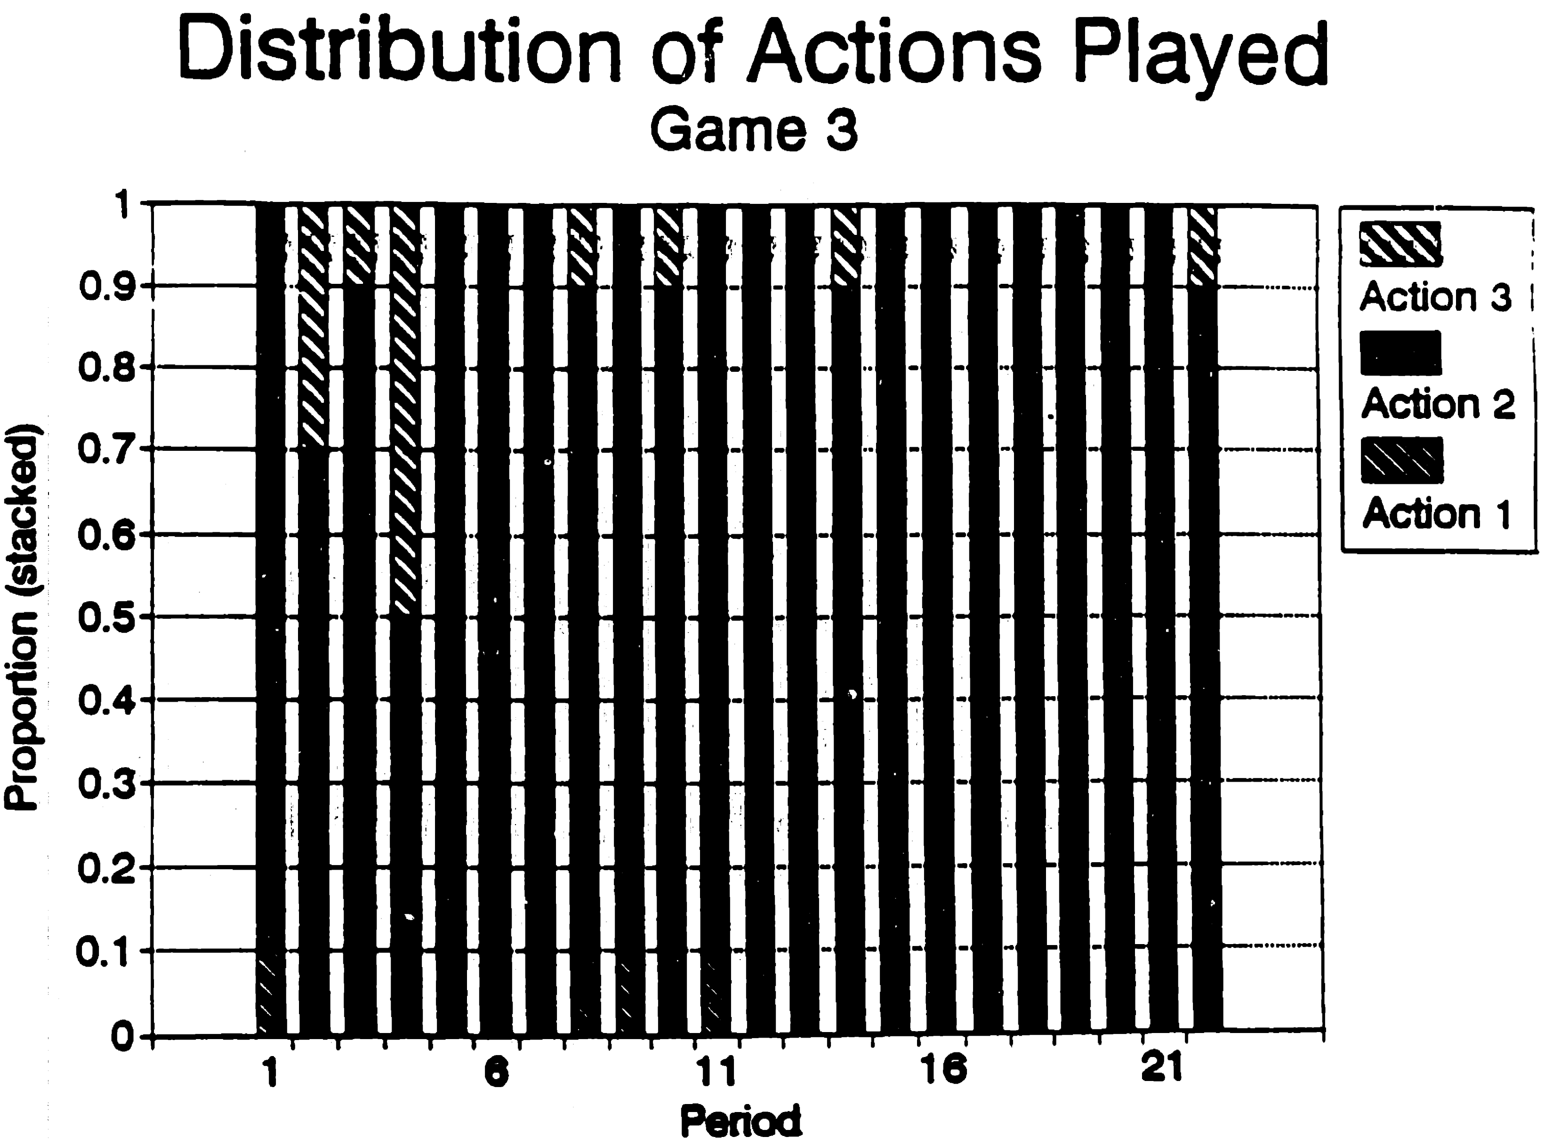
Distribution of Actions Played
Game 3
Proportion (stacked)
1
0.9
0.8
0.7
0.6
0.5
0.4
0.3
0.2
0.1
0
1
6
11
16
21
Period
Action 3
Action 2
Action 1

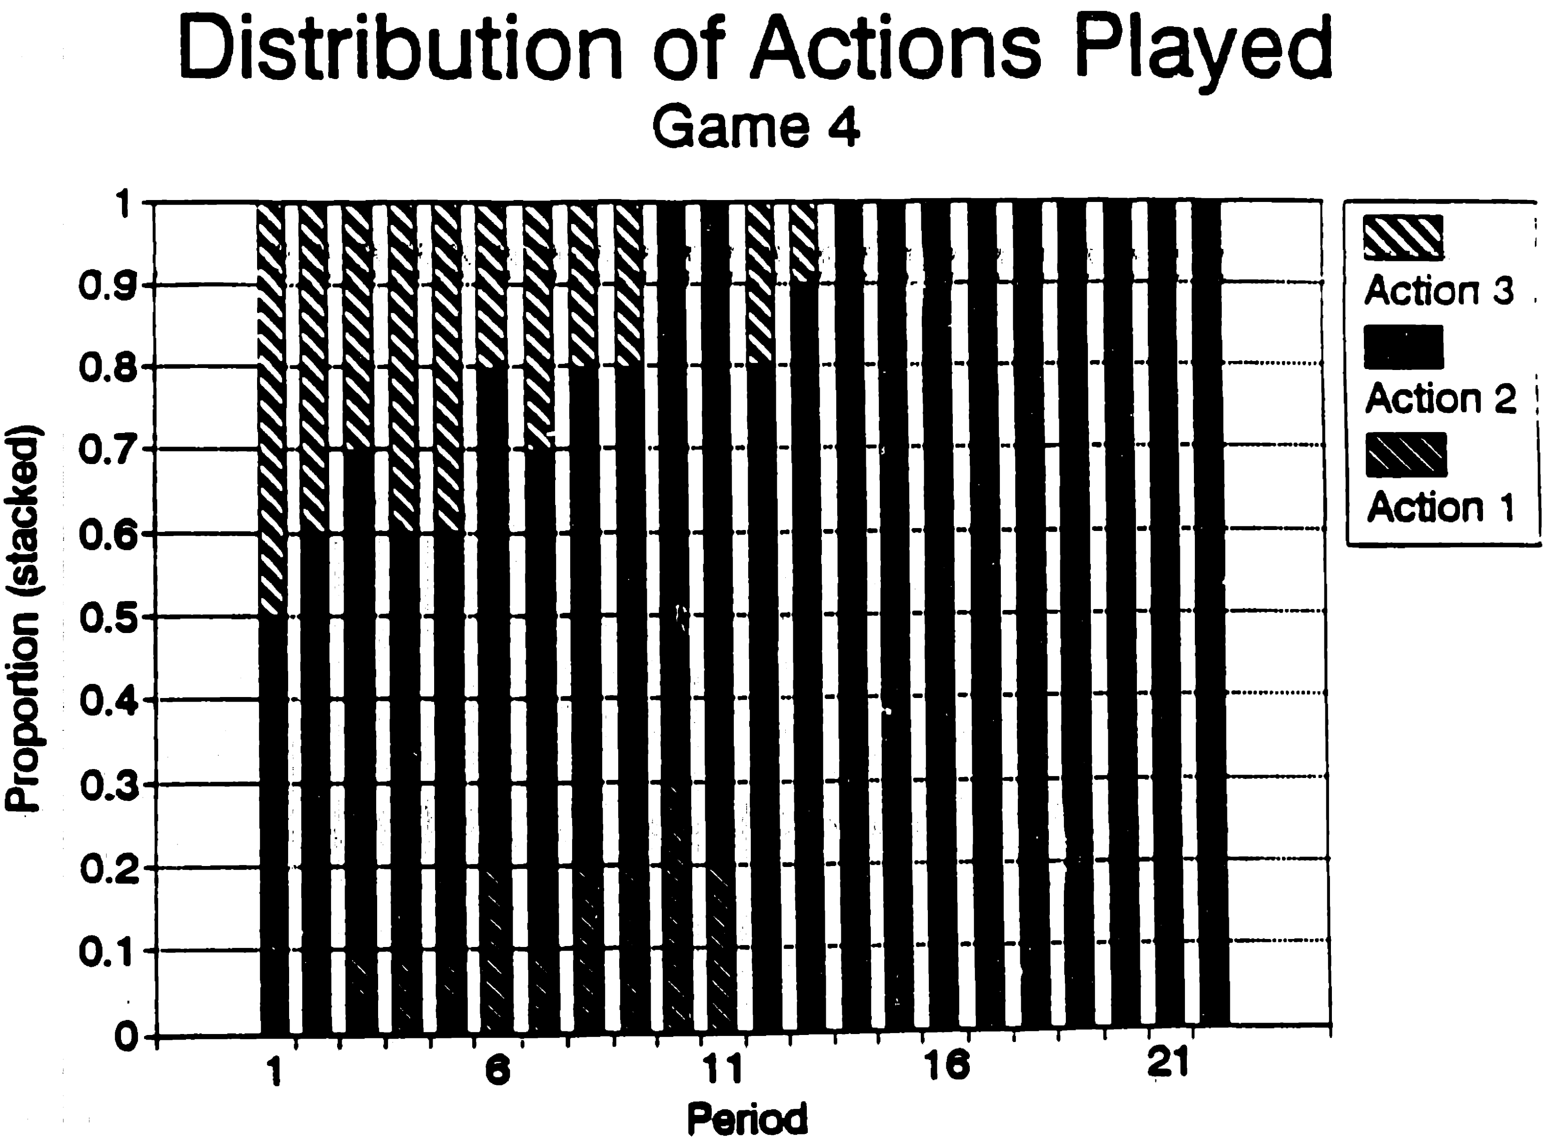

Distribution of Actions Played
Game 4
Proportion (stacked)
1
0.9
0.8
0.7
0.6
0.5
0.4
0.3
0.2
0.1
0
1
6
11
16
21
Period
Action 3
Action 2
Action 1

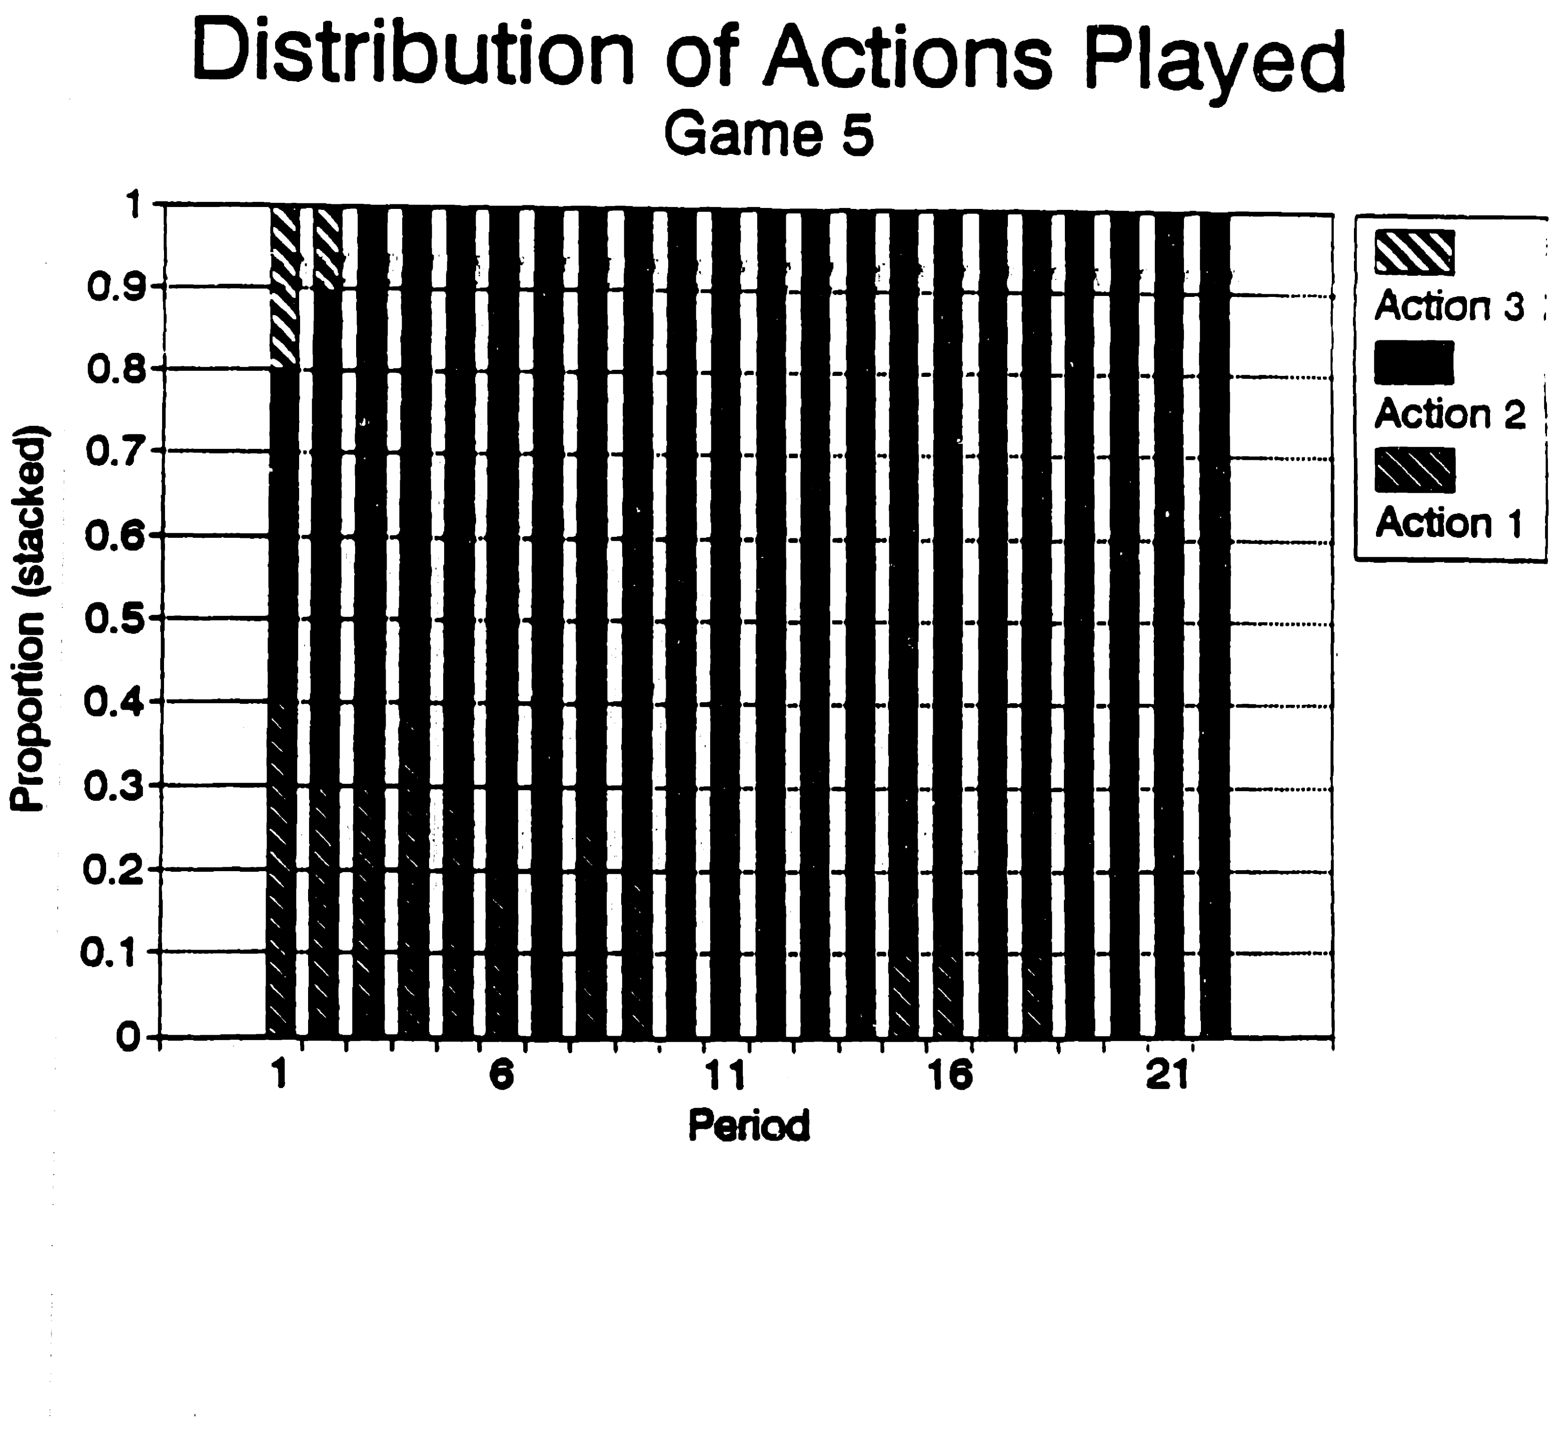

Distribution of Actions Played
Game 5
1
0.9
0.8
0.7
0.6
0.5
0.4
0.3
0.2
0.1
0
Proportion (stacked)
1
6
11
16
21
Period
Action 3
Action 2
Action 1

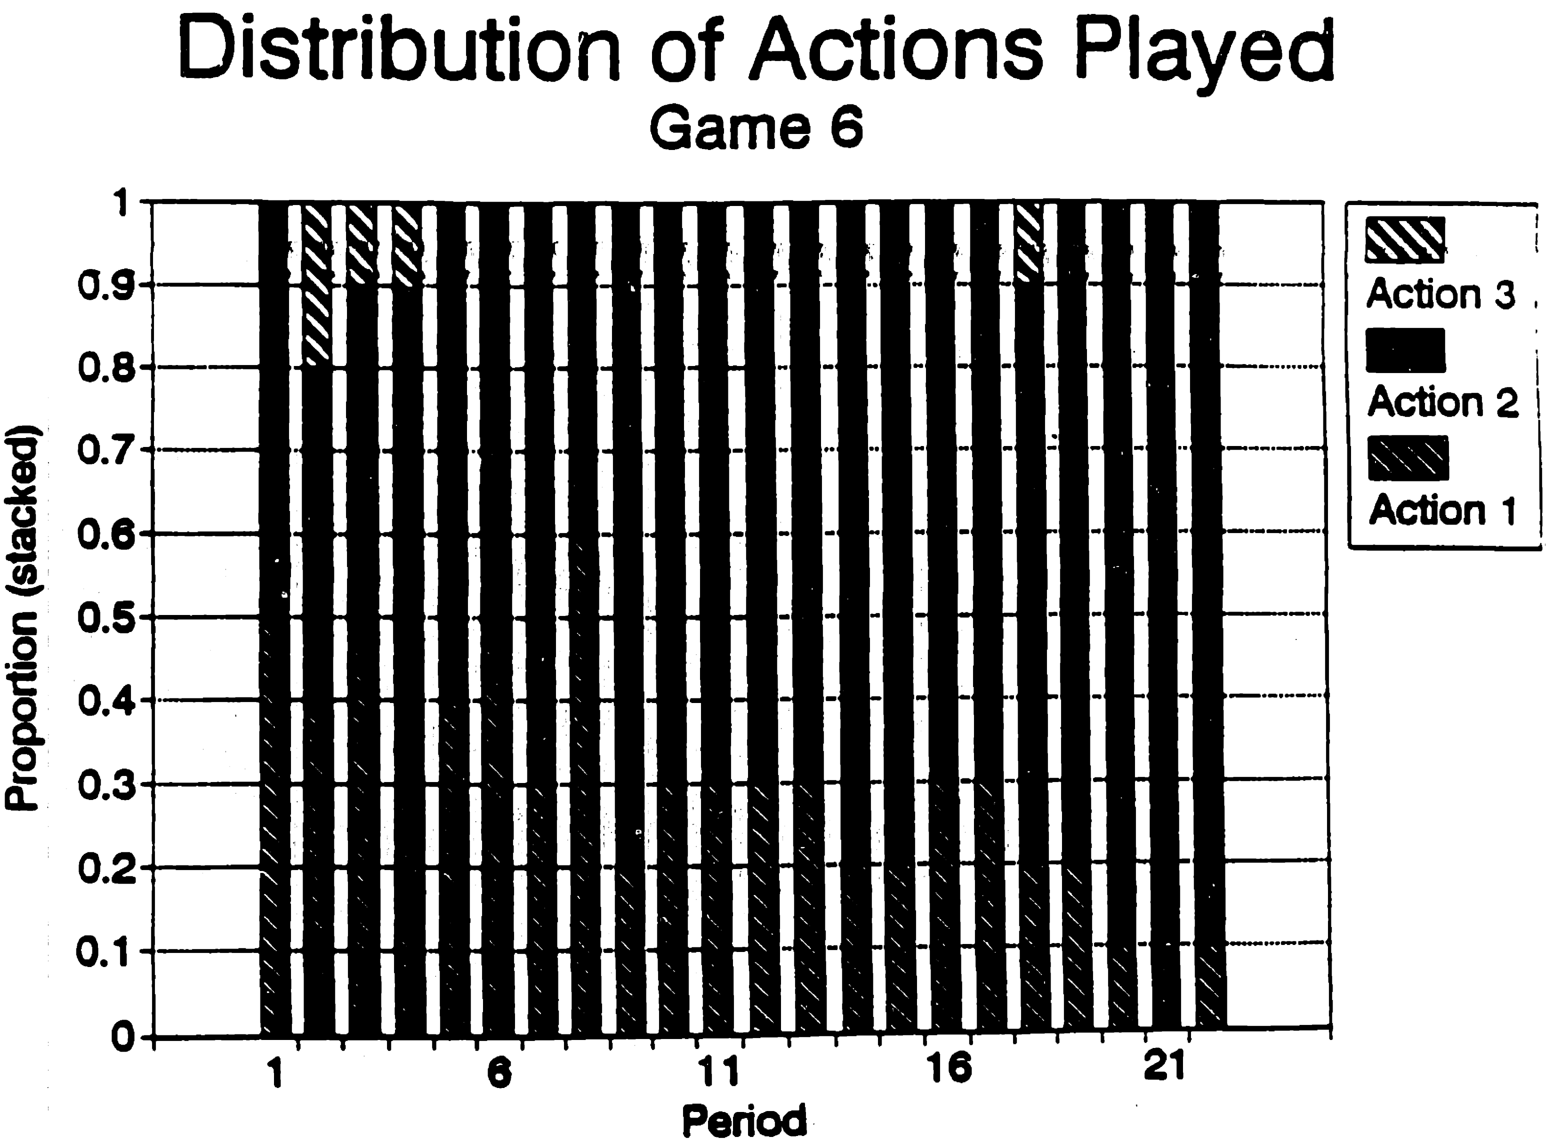

Distribution of Actions Played
Game 6
Proportion (stacked)
1
0.9
0.8
0.7
0.6
0.5
0.4
0.3
0.2
0.1
0
1
6
11
16
21
Period
Action 3
Action 2
Action 1

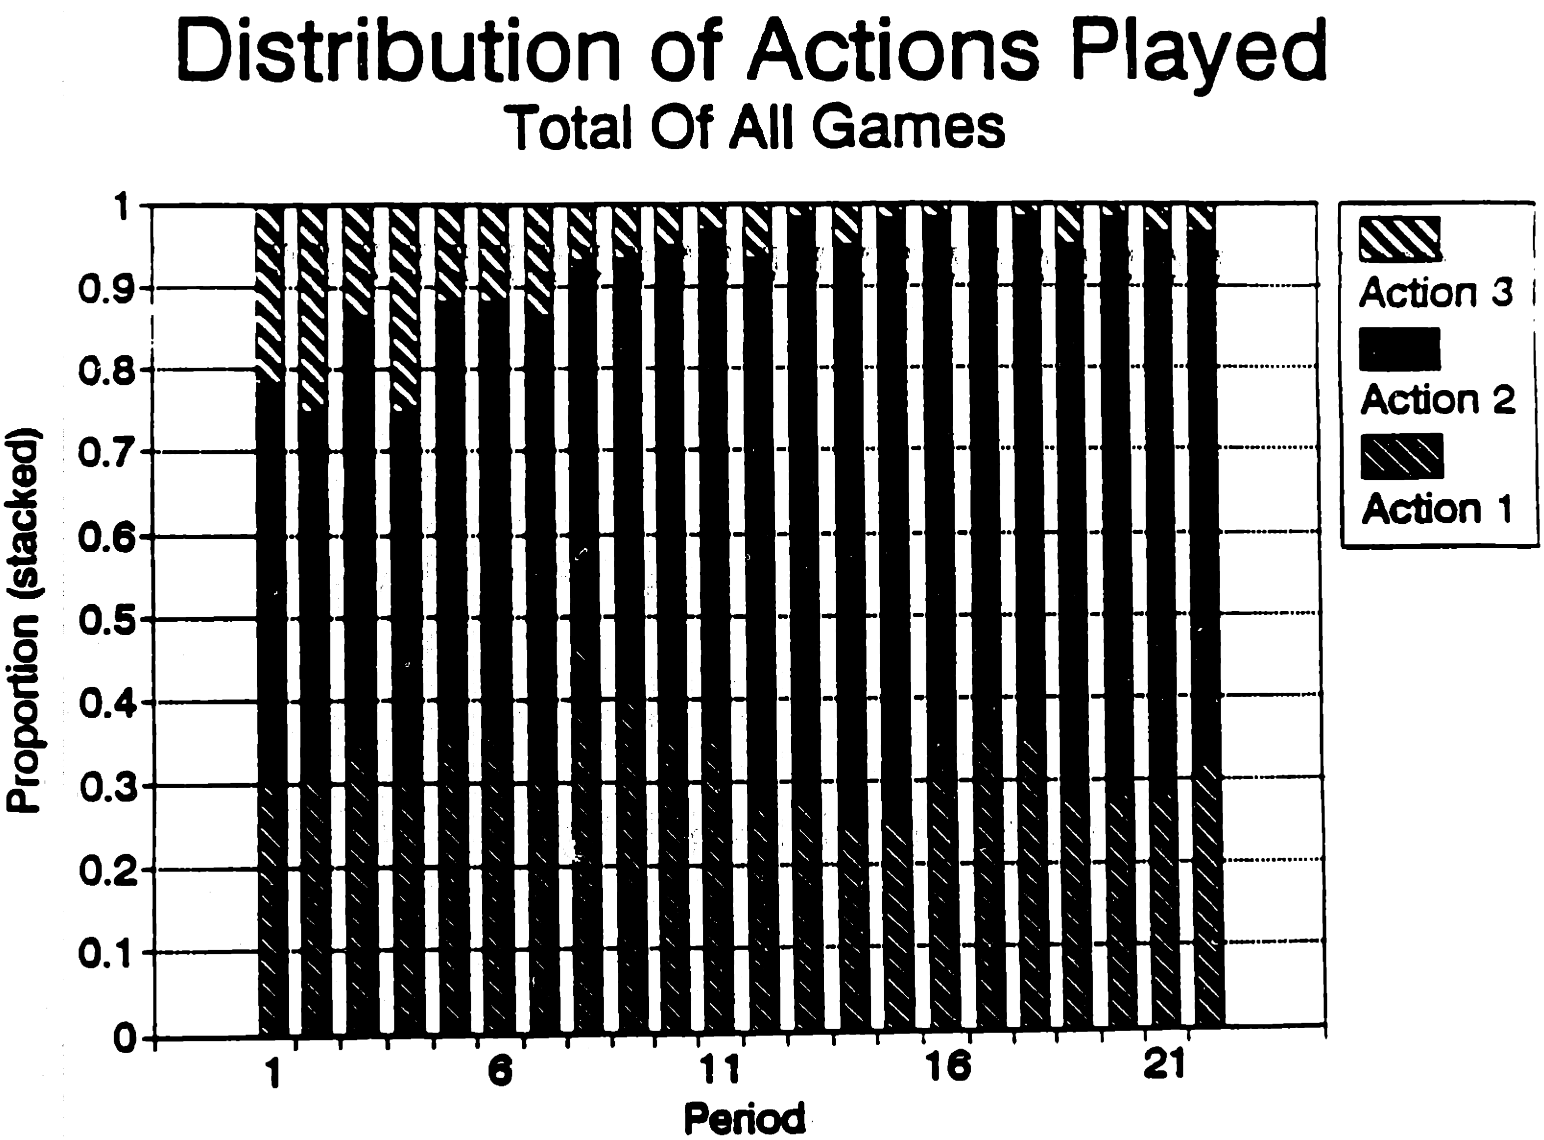
Distribution of Actions Played
Total Of All Games
1
0.9
0.8
0.7
0.6
0.5
0.4
0.3
0.2
0.1
0
Proportion (stacked)
1
6
11
16
21
Period
Action 3
Action 2
Action 1

Figure 8: Histogram of N's From Full Heterogeneity Estimation

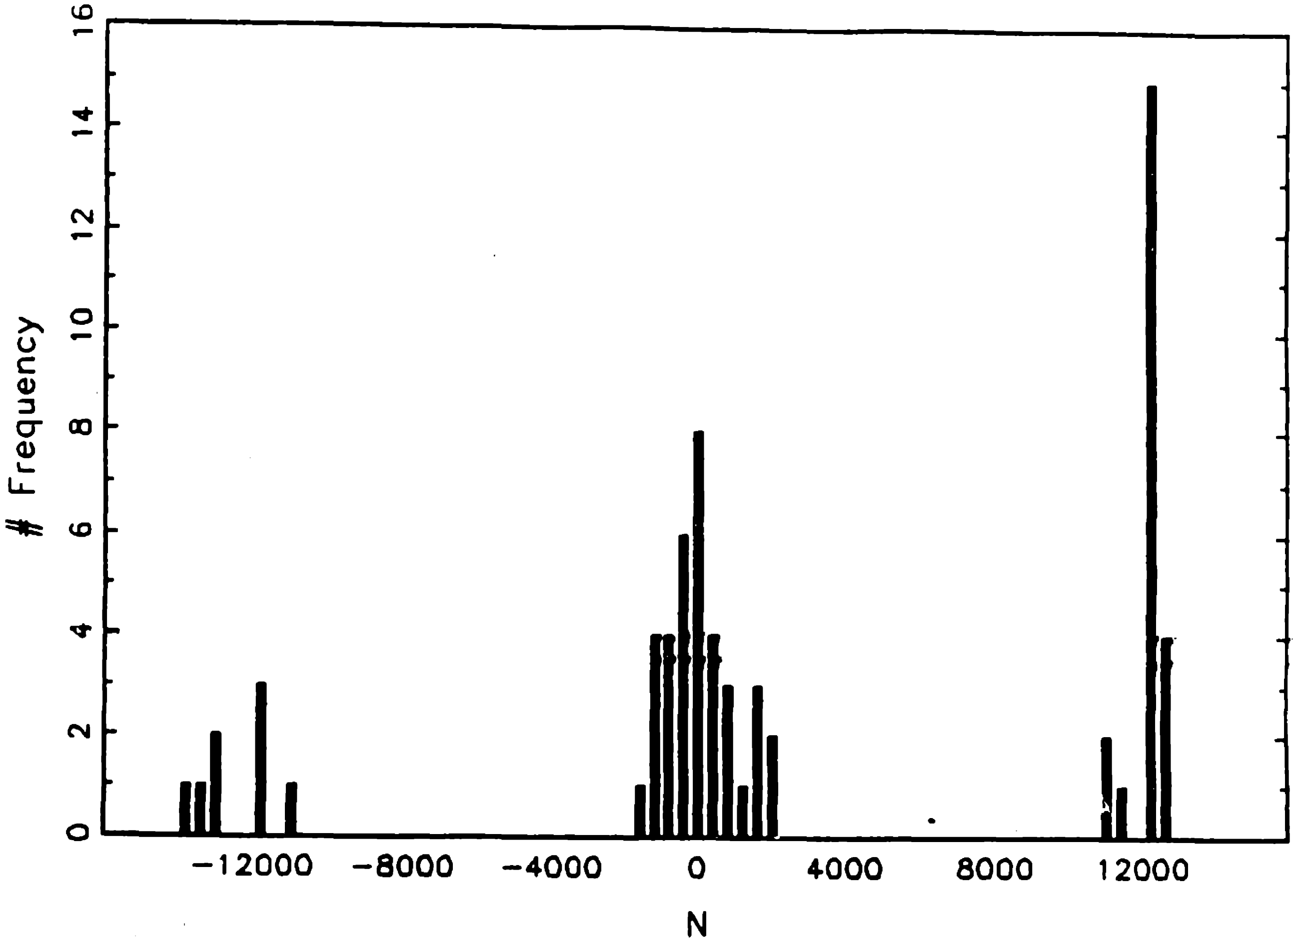

Figure 9: Estimated Distribution

Learning Model -- Game 1

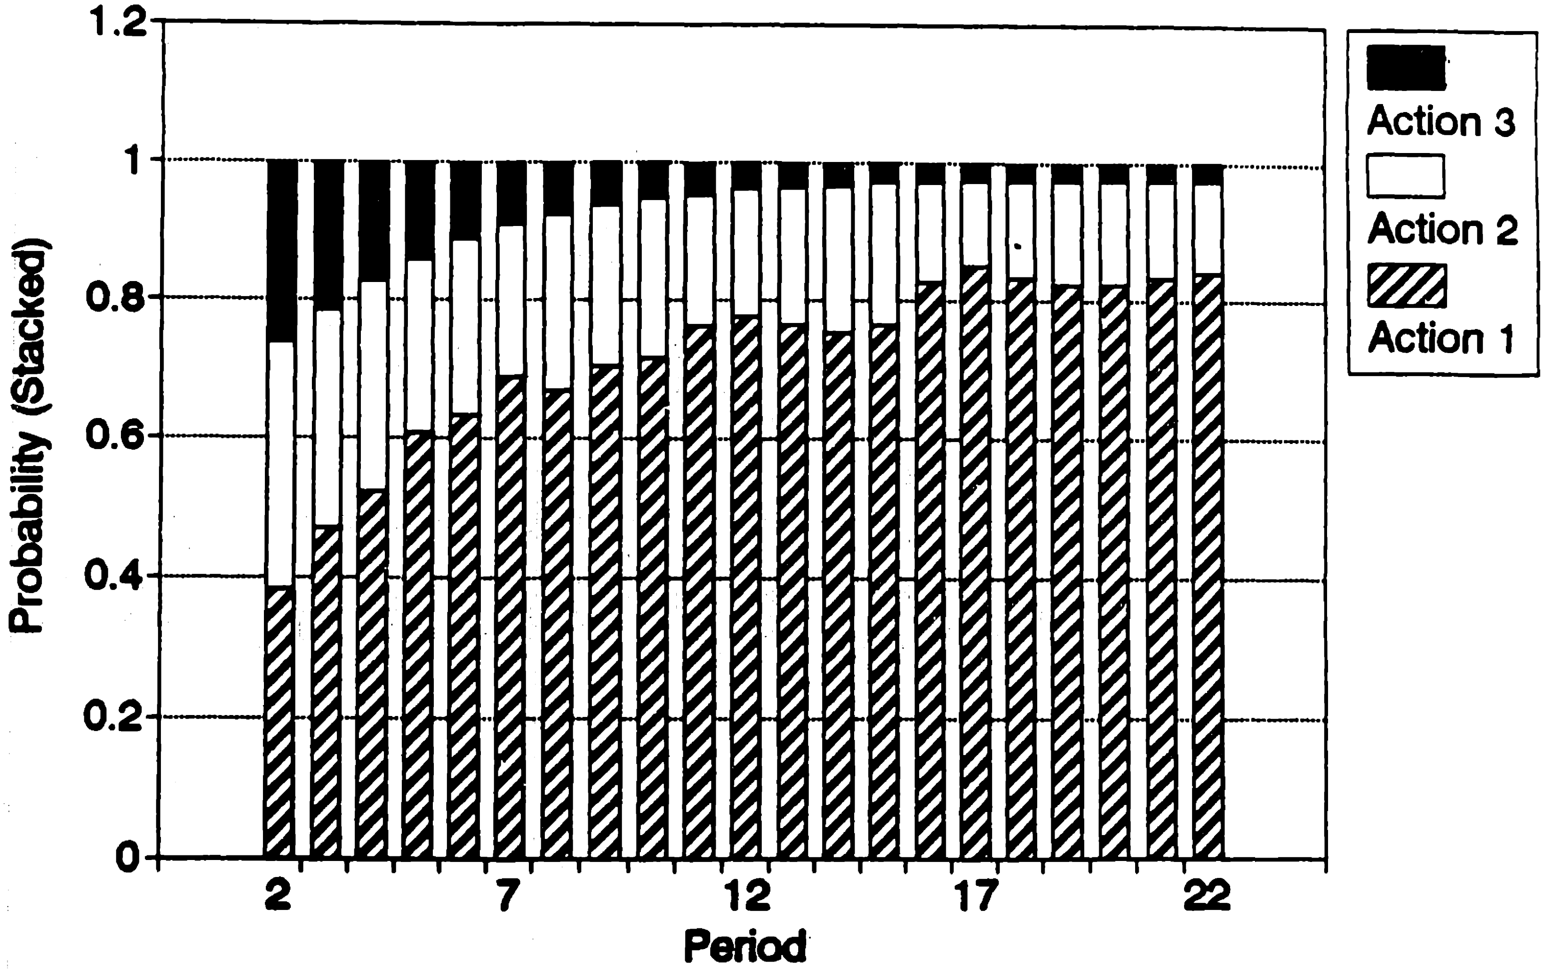

Figure 10: Estimated Distribution
Evolution Model -- Game 1

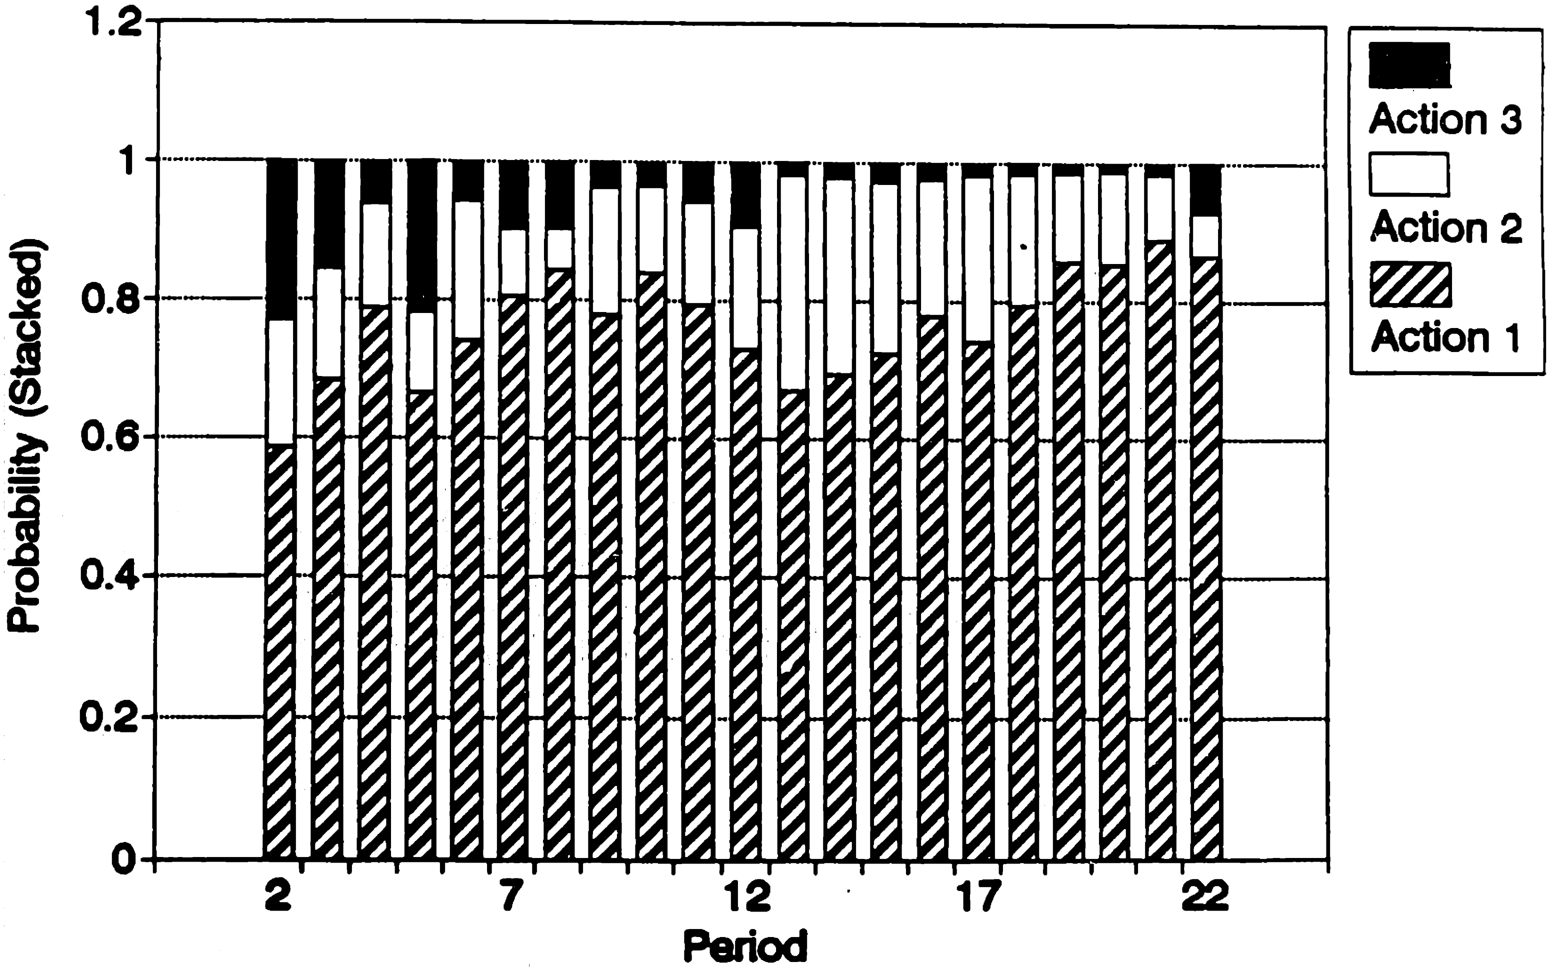

Chapter 3

Simulating Players Through Learning

3.1 Introduction

There is a large and growing literature that models the process by which players learn to play games. This literature relaxes the assumptions of Nash equilibrium to more realistically model the players. Specifically, the players have some uncertainty about what everyone else is doing in the game. The players are also boundedly rational in some of their decisions. These realistic players are set in an environment of repeatedly playing the same game and given a simple rule for updating their behavior from one period to the next based on their experience at that point. The behavior rules have to be relatively simple because results are obtained by computing the limit of behavior as the number of repetitions becomes infinite. Even for these simple behavior rules proving propositions about the limit behavior can be computationally difficult.

In this paper we will simulate a number of variations on the basic fictitious play model of learning. Simulation offers several advantages over analytic computation of the limits. For one thing it will be computationally feasible to compare many different learning models for a broad spectrum of games. Some of the variations that will be simulated have not been solved analytically. Another advantage of simulation is that the complexity of the behavior rule is not constrained by computability. With more complicated behavior rules we can more closely approximate the learning process of real players. Majure (1994) showed that fictitious play was not sufficiently complex to fully describe the learning behavior of players in an experimental setting. A third advantage to simulation is that our attention is not restricted to the limit as the number of repetitions goes to infinity. Very few games are well characterized as having been played so many times that this limit would be approached. It may be the case that the limit behavior is relevant even to games that are played infrequently,[1] but without some approach, such as simulation, that allows us to look at finite repetition behavior of the process, there is no way to know.

While the simulation procedure allows us to look at any aspect of behavior in finite time for any learning model, in order to have a basis of comparison for the learning models considered here, we will focus on convergence of the population to all playing the same action. In the pure fictitious play model this type of convergence implies that the action will be the limit behavior as well. The short-term convergence need not imply limit behavior in the other models, though. Obviously, we can't simulate an infinite repetition. One interpetation of the short-term convergence behavior is as descriptive statistics of the finite repetition of the processes. In this interpretation we would use the fact that in a given game the short-term convergence is much more likely to be to action 1 than to action 2, say, to infer that this process predicts that action 1 is much more likely in this game. Another interpretation is that we are measuring the relative sizes of the basins of attraction. If, for randomly selected initial conditions, action 1 is much more likely to be converged to in the short term than action 2, then we can infer that action 1 has a much larger basin of attraction.

A caveat to the first interpretation is that to take these simulations as a Monte Carlo estimation of the probability that real players will play a given action we have to believe that we have used a distribution on the initial conditions that accurately reflects the population. Majure (1994) estimates parameter values for several learning processes in games similar to the class studied here. That work demonstrates how we

[1]For example we will see that in some games the expected time until fictitious play achieves the limit behavior is only four repetitions.

can address this question and justifies for this interpretation some of the parametric assumptions that we will make in this paper. All of the parametric assumptions in this analysis, though, are subject to further experimental testing. Doubtless, further experimental and theoretical work will refine our knowledge of learning and the prediction interpretation will become more compelling, especially if a particular learning process can be shown to accurately reflect the learning process of real players. If we have correctly specified the learning process and we have accurately represented the population in our parameters, then under this interpretation of simulation we have a Monte Carlo approach to making actual predictions in games. The observed proportion of simulated players playing an action will approximate the probability that real players will play that action. These probabilities will allow us to make predictions that say more than just what *can* happen in a game.

Perhaps a more direct interpretation of the short-term convergence behavior is as a measure of the size of the basins of attraction and their relative pull. Counting the number of times that an even distribution of initial conditions produce a given action is akin to measuring the size of some bowls by evenly dropping marbles over them and counting the number in each bowl. It is, admittedly, a rough measure, but in this case it is much easier than direct measurement. Determining the relative sizes of the basins of attraction is the approach to computing the limit behavior of some learning models.[2] There is a possibility of such a connection in the models studied here, but I do not want to compare the processes on the basis of limit behavior. Rather, I want to compare them on the basis of their propensity to produce a certain result for realistically finite repetitions.

The comparisons yield a few characterizations that hold true across all of the specifications. In games with very strong risk dominance, the risk dominant action is converged to almost every time. The risk dominant outcome has such a large basin of attraction that if initial conditions are independent and uniform, then the probability of observing another equilibrium is effectively zero. Between these regions of prediction with certainty is a region for each specification where the basins of

[2] See for example Kandori, Mailath and Rob (1992)

attraction are more even. In these games the prediction is a non-zero probability for each action. Both of these results deserve further exploration as predictions. If the former is borne out, then we have a set of games in which this approach mirrors the refinement approach by completely eliminating an outcome. Conversely, the second demonstrates the existence of a set of games that should not be refined to a unique prediction. In this second set of games predictions of the relative probabilities of the different equilibria are essential for the application of game theory to real situations.

The comparisons also yield some comparative statics that should be useful in the search for a learning process that more closely describes reality in a broad spectrum of games. We will see that weighting more recent observations more heavily in fictitious play's calculation of beliefs does not significantly affect the action converged to, but does increase the speed of convergence. Adding a cost of updating behavior to fictitious play causes a dramatic shift in the action converged to for some games and for some parametric distributions but not for others. We will also see that changing the behavior rule of fictitious play from choosing a best response to trying to emulate other players will slow the convergence down. There will also be an indication that emulation changes the action converged to, but the new process is so slow to converge that this effect may be confused with a failure to converge.

While these predictions and comparative statics are useful, the most important contribution of this paper is the methodology. By taking the learning process seriously as a description of how players come to have the beliefs that they do, this paper demonstrates the usefulness of simulating the learning process to acquire enough data to be able to predict how people will play in a broad class of static games.

3.2 The Games and the Learning Processes

The class of games that will be used in this paper is the simplest class that has multiple equilibria – symmetric, 2-player, 2-action coordination games. The procedure can easily be generalized to more complicated games. The games in this class can all be represented, through some rescaling of payoffs and relabeling of the actions, by

Table 3.1: The General Form of Games in the Class

	1	2
1	1,1	0,c
2	c,0	d,d

the game in table 3.1 where the parameters c and d are on the interval [0,1]. These games have a Pareto dominant equilibrium at (1,1) and another equilibrium at (2,2). Each also has a mixed strategy equilibrium which we will essentially ignore since the learning processes that we are considering do not converge to mixed strategy equilibria in a manner that is useful for prediction. As the games are all symmetric, we will act as if all players are player 1, so the outcome (1,2) is where the player is playing action 1 and her opponent is playing action 2.

The general learning process assumes that a game is not just played once. Rather, the current game is just one in an infinite repetition of the same game. At each repetition players are randomly assigned opponents from a large population. Each player uses their history to determine their current play and then updates their history on the basis of the outcome. The manner in which the players use their information determines the particular model. There are roughly two strains of models in the literature – "learning" and "evolution." When these models converge, they converge to a Nash equilibrium.[3] But not all Nash equilibria are equally likely to be converged to. In fact, some Nash equilibria are so unlikely to be converged to that equilibrium refinements have been developed on this basis. The evolution model in particular has been used in this fashion to refine the set of predicted equilibria.[4] We will start with

[3]There is really nothing special about equilibria in this simulation approach. The process is completely generalizeable to estimating the probabilities of all outcomes. The learning literature does ascribe importance to convergence behavior and I have maintained that focus here. Thus, we will focus on equilibria only because these learning processes converge to equilibria.

[4]ESS is a primary example of such work. Ellison (1993b) indicates how fictitious play can generate a similar refinement. If fictitious play stability is with regard to the addition of a single fully rational player, then Ellison's work would have the result that Pareto dominated equilibria that are also risk dominated are not stable. We will show a similar result in that strongly risk dominated equilibria are not converged to by fictitious play.

the traditional learning model — fictitious play — and introduce variations that are suggested by other models in the literature. One variation will make fictitious play similar to an evolution model.

The basic fictitious play story is that players do not know the behavior of their opponents, but assume that all of their potential opponents are playing the same strategy and that this strategy is not changing over time so that they can learn what the strategy is by updating prior beliefs on the basis of current observation. Each player has initial beliefs that can be represented by two parameters (n_1, n_2). If the player has observed a history with c_1 observations of action 1 and c_2 of action 2, then the player believes that the probability of action 1 in the next period is

$$P_1 = \frac{n_1 + c_1}{n_1 + n_2 + c_1 + c_2}. \tag{3.1}$$

The players will choose their action in each period as a best response to their current belief about their opponent's action (P_1).

The first variation on fictitious play that we will consider will allow players to put greater weight on more recent observations. If $I(t)$ is 1 if the player observed action 1 in period t and 0 otherwise, then this type of updating can be represented by

$$P_1(T) = \frac{n_1 + \sum_{t=1}^{T} t^{\alpha} I(t)}{n_1 + n_2 + \sum_{t=1}^{T} t^{\alpha}}. \tag{3.2}$$

Note that the updating rule of equation 3.1 is a special case of this discounting rule where $\alpha = 0$.[5]

The second variation that we will consider is to give players a cost of updating their actions. We will make this change by saying that players follow a heuristic rule for when to incur this cost. The rule that we will use is that each player has an aspiration level in the game and the player chooses a new action only when his payoff to playing his old action falls below this level. Note that the traditional fictitious play model is a special case of this variation as well. For the games that we are

[5]Weighting more recent observations more heavily is suggested by the model of Young (1993) where only a fixed number of the most recent observations are kept.

considering if every player has an aspiration level of 1, then they always update as in the traditional fictitious play model.[6]

The last variation that we will consider is to change the way in which a new action is chosen. In fictitious play the players adopt the best response. Therefore, as a player's beliefs change, her action will not change until the beliefs cross a threshold and then suddenly the player changes actions. In this variant beliefs will have a more gradual effect on actions. This variant is based on the economic muddling model of Binmore and Samuelson which is a discrete time version of the familiar replicator dynamic.

The model presented by Binmore and Samuelson has players who only update their behavior when their payoff in a previous period falls below some personal aspiration level. Updating takes the form of drawing another player randomly to emulate. This updating is subject to a random shock, mutation, where instead of emulating another player the updating player just switches actions. As the time between periods and the probability of mutation go to zero this model converges to the familiar replicator dynamic. In order to be consistent with the informational assumption that each player only observes the play of their chosen opponent each period we will replace the emulation of play in the most recent period with emulation of a player drawn from all past opponents. We want to allow the possibility that this choice is biased toward more recent opponents. For consistency with the rest of the analysis we will use the distribution such that

$$P_1(T) = \frac{\sum_{t=1}^{T} t^{\alpha} I(t)}{\sum_{t=1}^{T} t^{\alpha}} \tag{3.3}$$

is the probability that emulating another player leads to playing action 1.

3.3 The Simulation and Estimation Techniques

We will simulate the models of section 3.2 for a large number of randomly drawn initial conditions. In order to describe the convergence behavior of these processes

[6]This aspiration level heuristic comes from the evolution model of Binmore and Samuelson (1992) discussed below.

we will estimate the relationship between the two game parameters (c and d) and the descriptive statistics of convergence — the action converged to and the time of convergence. The random parameters will all be drawn from uniform distributions so that for a large number of simulations this procedure will approximate a grid search simulation.

An advantage to using random parameters is that a large number of parameters can be varied without necessarily having to increase the number of simulations by orders of magnitude. It might be possible to make some simplifying assumptions to reduce the number of parameters, but the more obvious ones are not consistent with earlier results on learning processes. In particular, it is possible to reduce the pool of players by assuming that there are just two players. Ellison has demonstrated that the possibility of "contagion" (meeting someone who's actions have been influenced by your earlier interaction with another player) is significant in large populations Reducing the population, then, would lose some of the interesting dynamics of the process. The number of parameters needed to define the processes could also, potentially, be reduced by assuming homogeneity across players. This approach would maintain the dynamics of a large population since each player would be reacting to a different history of observed repetitions. However, in fitting these processes to the actual behavior of players in an experimental coordination game, Majure statistically rejected the hypothesis that these same parameters could be made homogeneous. In fact, to the extent that the parameters were recoverable in that estimation, the indication is that they come from a wide distribution. Another alternative that was considered was using these estimated parameters. The results obtained from such simulations were found to be extremely sensitive to the particular game in the class chosen to simulate. Since the parameters were estimated for only one game in the class the fact that the processes behaved significantly differently for the same parameters in a slightly different game is an overwhelming rejection of the applicability of the estimates for other games.

Another advantage of random parameters is that if we specify the distribution of the parameters, then we can change it and evaluate the consequences. Because we

Table 3.2: Distributions of Parameter Values

Parameter	Uniform Over Low	Uniform Over High	Comments
Payoff to (2,1) (c)	0	1	
Payoff to (2,2) (d)	0	1	
Aspiration Level	0	1	
	.5	1	
	1	1	always updating
History Weighting (α)	0	0	all history the same
	1	1	linear weights
	2	2	convex weights
	0	2	heterogeneous weights
Unupdateability ($n_1 + n_2$)	0	20	
Initial Belief ($P_1(0)$)	0	1	

want to interpret these simulations as predictions of play, we have to be concerned with the accuracy of our distributional assumptions. The conventional grid-search method of simulation implicitly imposes a uniform distribution on the parameters. It turns out that we have reason to believe that the parameters in this case are uniformly distributed.[7] Nevertheless, we want to be able to examine the robustness of this assumption.

At the start of each simulation all of the parameters are drawn independantly from the distributions in table 3.2. Where appropriate a separate parameter is drawn for each player. If the same parameter is used in two variants of a model or in two different models, then for each player the same value of the parameter is used for each process. Each period a pool of 20 players are randomly matched to play each other where the probability of meeting any two opponents is the same. Play is determined by the process and the players' personal histories of the play in prior periods. The histories do not include the identity of previous opponents as we are imposing the condition that the players treat the population as large. The players also cannot make their play contingent on whether they are a row or column player. Convergence is defined as a run of ten periods which starts with every player playing the same

[7]See Chapter 2.

action and which has at most one player deviating from that action in each of the remaining nine periods. The time of convergence is the first period of the run. If there is no convergence after 100 periods, then the process is stopped and marked as non-convergent. This simulation procedure is repeated 10,000 times.

We are focusing on the action converged to and the time until this convergence. For predicting the outcome of a static game we are most interested in the action converged to, but our confidence in the distribution of these actions as the prediction of play in a given game is reduced if the learning process needs to have been operating for a particularly long time to get to that point. In terms of using these estimations to understand the learning processes themselves, these descriptive statistics are both important characteristics. The action converged to is the basin that attracted this instance of the process. The time to convergence indicates the amount of pull from that basin.

Since we do not want to impose any functional form assumptions on the relationship between the game and these convergence properties, the appropriate estimation technique will be non-parametric estimation. We use a standard kernel estimator.[8] An advantage of this approach is that while imposing some smoothness it estimates the expected value of the dependant variable for a grid on the explanatory variables. In this case that allows us to look for a very non-linear relationship between the game being played and the behavior of the learning process. The estimation technique also allows us to identify games for which the learning process changes when a parameter is changed and types of games for which it does not.

The kernel estimator essentially computes the sample average of the convergence behavior across simulations of the same game. The smoothness comes from the fact that this sample average is actually a weighted average of the behavior in all the simulations with the weight decreasing in the distance of the game played from the game being estimated. The actual estimates are given by a surface over the plane that defines the games. Since all of the games in this class are defined by c and d and since these payoffs are constrained to be in [0,1], the relevant payoff space is the

[8]See Bierens (1987) and also Härdle (1990)

unit square. As we have defined them, every game in this class is a point on the unit square. The kernel estimator actually estimates the expected value of the behavior conditional on the game being played. When the action converged to is the behavior to be estimated there are only two possibilities, so the expectation can be interpreted as the probability of action 2. Thus, if a particular game converges to action 2 in 3/4 of the simulations, the estimation should be 1.75 at that point in the game space. The only way to concisely report these estimates is graphically. Since we have a two dimensional payoff space the graphs are three-dimensional. In order to help illustrate how the surface (the estimation) changes with the game I have also provided contour plots. The contour plots are in the game space, so a contour is the set of games that have the same expected convergence behavior.

3.4 The Estimations and Comparisons

All of the estimations share a common pattern. In each model there is a region in the game space where (2,2) is converged to every time and a region where (1,1) is. On the surface plots these are the regions for which the surface is flat at 1 and at 2. On the contour plots we see that these are the games above the 1.9 contour and also the games below the 1.1 contour. When action 2 guarantees the player a payoff of 1, the prediction is always (2,2). Likewise, when action 2 guarantees the player a payoff of 0, the prediction is (1,1). The certainty of the convergence target should not be surprising in these games. What is surprising is that the two regions of games around these, where the prediction is also certain, are quite large. They are separated only by a relatively narrow band where convergence can occur to either action.

It is possible for any of these learning processes to specify initial conditions for which the process converges to the other action in any of these games where the convergence is certain. For example, it is always possible for a population of fictitious play players to be convinced that their opponents are playing according to a given equilibrium. Since everyone has the same correct beliefs, fictitious play immediately converges to Nash equilibrium. What this result implies is that the probability of drawing any initial conditions that would drive the process to the "wrong" action in these games is so small that it is effectively zero. The basin of attraction for the "wrong" equilibrium is so small that none of our uniformly drawn initial conditions are inside it.

3.4.1 Standard Fictitious Play

The standard fictitious play model (no satisficing and no discounting) is estimated in figure 3-1. The uncertainty region for the action converged to is centered on games with no risk dominance. In the graphs the games with no risk dominance are on the line in the payoff space between (1,0) and (0,1). Games farther from this line have more risk dominance with action 1 risk dominating action 2 above the line and action 2 risk dominating action 1 below it. Note that as risk dominance increases the expected action converged to quickly approaches the risk dominant action. Thus, as risk dominance increases the relative size of the basin of attraction for the risk dominant action increases so quickly that we soon stop seeing the risk dominated action.

The expected time to convergence is decreasing in risk dominance with the extreme cases taking an average of only four periods to converge. Holding the degree of risk dominance constant (i.e. moving parallel to the line between (0,1) and (1,0)) the time to convergence is increasing in d (the payoff to (2,2)). Thus, net of the risk dominance effect, there is a similarity effect that the more similar action 1 is to action 2, the longer convergence takes. This effect implies that the speed of convergence is increasing in the degree of Pareto dominance. Recall that the (1,1) equilibrium always Pareto dominates the (2,2) equilibrium and that the payoff to (1,1) is fixed at 1. Therefore, increasing d decreases the degree of Pareto dominance.

3.4.2 Weighted Histories

We add discounting to the fictitious play model in figures 3-2 and 3-3 In figure 3-2 the players all have $\alpha = 1$. More recent observations are weighted more heavily

Table 3.3: The Game With Action 2 Identical To Action 1

	1	2
1	1,1	0,0
2	0,0	1,1

in computing the player's beliefs and the weighting is linear. Linear weighting is contrasted with the convex weighting of figure 3-3. In figure 3-3 the players all have $\alpha = 2$, so not only are more recent observations weighted more heavily, the size of the weights increases with repetition. Recall that figure 3-1 has $\alpha = 0$ as there is no weighting.

Comparing figure 3-1 with figures 3-2 and 3-3 we see very little change in the expected convergence behavior. In fact, there is no noticeable difference in the expected action converged to for any part of the game space. This similarity of expected action converged to is even more surprising when we note that discounting does not leave the process unaffected in the sense that for the same initial conditions and the same matching of players the discounted and undiscounted models are likely to converge to different actions. Figures 3-4 and 3-5 estimate the expectation of the absolute value of the difference in the convergence behavior between the undiscounted fictitious play and linear and convex discounting respectively. This measure of the disagreement between the models is, of course, zero in the games where the models always converge to one action. But, in the region of games where both actions have non-negligible basins of attraction the probability of disagreement can be as high as 50%. This disagreement probability is increasing as risk dominance decreases (i.e. as the estimated sizes of the basins of attraction become closer to equal) and as the Pareto dominance of action 1 over action 2 is decreasing. The disagreement is at a maximum in the game of table 3.4.2 which has the trait that there is no difference between the actions as they could be re-labeled without affecting the game. The fact that the expected action is not affected by these differences implies an averaging effect across initial conditions that is not biased by discounting.

While the expected action converged to does not change with discounting, the expected time to convergence does. The shape of the surfaces does not change much from figure 3-1 to figures 3-2 and 3-3, but the level does. Both discounted models converge more quickly than the undiscounted model. The fact that the shapes of the surfaces do not change indicates that the same effects across games are present with discounting as without. The increased speed, then, is attributable entirely to the fact that the initial beliefs are also discounted away like early observations. In particular, the player's conviction of the correctness of his initial priors is eroded more quickly. This erosion is necessary for the players to overcome an initial conviction that the game will be played one way when the system is converging the other. It is interesting to note that there is little difference between the linear and the convex discounting. Perhaps a little discounting is sufficient for this effect and more discounting does nothing more.

Since heterogeneity was found to be important for other parameters, we are interested in looking at the robustness of the homogeneity assumption we have implicitly made on α. Figure 3-6 estimates behavior for the model with α drawn independently for each player from U(0,2). It is still the case that the expected action is the same for all of the games. The expected time to convergence is slightly longer than in figures 3-2 and 3-3. This relative slowness is probably due to some players getting α values that are too small to effectively discount wrong initial convictions quickly.

3.4.3 Updating Costs

Updating costs are added to the basic fictitious play model in figures 3-7 and 3-8. In figure 3-7 the players draw their aspiration levels from the distribution U(0,1), while in figure 3-8 they draw from U(.5,1). With the wider distribution on aspiration levels the uncertainty region shifts downward on the contour plot. This shift means that the set of games where action 2 is always converged to grows while the action 1 region shrinks. The implication is that the basin of attraction for action 2 has grown (or action 1's has shrunken). In some games this makes action 1's basin negligible and in others it converts action 2's basin to being non-negligible, but in the rest of the games the change is not enough to affect the expectation. The shift is most noticeable around the point where action 2 guarantees the players the mean (.5) of their aspiration distribution. The players who choose 1 when their opponent chooses 2 always update, but in this case there are some players who choose 2 when their opponent chooses 1 who do not update. This bias makes convergence to action 2 more likely.

When aspiration levels are U(.5,1), there is only a slight change in the uncertainty region from the standard model. This change is only apparent for payoffs of less than .5 to (2,2). The reason for the different magnitude of change is that people in the first case are not updating from action 2 when the system would be converging to action 1 while in the second case the failure to update from action 2 is only likely in games that are converging to action 2 anyway.

This analysis of the effect of updating costs is borne out by looking at the time till convergence. Comparing the estimates in figure 3-7 with figure 3-1 we see that a "hill" has risen above the region of games where the two models disagree on the likelihood of convergence to each action. On the contour plot this region is roughly identified as the games inside the contour of 28 periods. In this region of games updating costs produce much slower convergence. The continued unwillingness to update of a small but significant element can alter the course of the process in these games but it takes a long time to do so.

3.4.4 The "Fictitious Play" Version of Binmore-Samuelson

The last variant that we wanted to consider was the hybrid of Binmore and Samuelson's evolution model with fictitious play. This model with aspiration levels distributed U(0,1) and is estimated in figure 3-9 and with the distribution U(.5,1) in figure 3-10. The uncertainty region of the predicted action is determined by the distribution of individuals' aspirations and is something of an L-shape around the mean of that distribution. Any game where action 2 guarantees at least this mean is certain to converge to action 2. When the payoff to (2,2) is below the mean aspiration, the probability of action 2 quickly drops to zero. When the payoff to playing 2 in response to 1 is below the mean aspiration the probability of action 2 also drops to zero, but the effect can be mitigated by a high payoff to (2,2). This is the same effect of the bias in updating between action 1 and action 2 that we noticed in fictitious play with updating costs, but now it is determining the course of play to a greater extent. Comparing figure 3-9 and 3-10 illustrates the significance of the fact that the convergence action depends so heavily on the mean of the aspiration distribution. If we want to use this model for prediction, then at least for all of the games between the two different uncertainty bands our prediction will depend critically on the distribution assumption. Figure 3-11 demonstrates that this difference creates a region where the two processes are almost certain to diverge for the same initial conditions.

The time till convergence can roughly be described as decreasing in the payoff to (2,2) and increasing in the payoff to playing 2 in response to 1, with noteable exception for games where action 2 guarantees less than the minimum aspiration. Games in the uncertainty region take slightly longer to converge. These times to convergence are much longer than in fictitious play. This causes concern that a sizeable portion of the simulations are not converging in the finite time allowed. When discounting was added to this model there was very low probability of disagreement between the undiscounted and discounted models even for games where the separate estimations gave opposite results. This peculiarity is due to the fact that the estimations here are all conditional on the simulation converging and the congruence estimates are conditional on both models converging for given initial conditions. In standard fictitious play the probability of non-convergence was negligible, but in this model it is a problem.

Figure 3-12 compares this model to fictitious play with both models having a cost of updating with aspiration levels distributed U(0,1). The congruence is startling. Recall that these estimations are conditional on both of the processes having converged so some areas of disagreement may not be represented. Still, the set of games where the two process always converge to the same action for the same initial conditions is about 3/4 of all of the games in the class. In the remaining games the probability of disagreement is only greater than 1/4 in a narrow band and only ever gets as high as

1/2. The biggest difference in the models is that fictitious play converges as much as 55 periods earlier out of a maximum convergence time of 90 periods.

3.4.5 Robustness

As mentioned above an advantage of using random parameter distributions is that we can check the robustness of our results to changes in the distribution. Figure 3-13 estimates the standard fictitious play model with priors that are correlated. The correlation is generated by drawing a base prior for the population (p^*) from U(0,1) and defining each individual's prior to be

$$P_1(0) = p^* + \rho(1 - p^*) \quad \text{or} \quad P_1(0) = p^* - \rho p^*. \tag{3.4}$$

The probability of adjusting upward or downward is 1/2. Each player draws a ρ from the distrbution U(0,1). These correlated priors should make it more likely that the type of initial conditions necessary for the process to converge to the "wrong" outcome are drawn. Therefore we would expect that there is less certainty of convergence in this model. This is not necessarily the case, though. Since the only beliefs that we know go the "wrong" way have absolute correlation, it may be the case that even one player with the opposite beliefs is enough to contaminate the population. Looking at figure 3-13 we see that the band where both actions are converged to some of the time is a little wider than in figure 3-1. The characterization, however, is the same. We conclude, therefore, that correlated priors will weaken our results but won't change them.

3.5 Conclusion and Extensions

We have shown that simulating the learning process is a useful tool for comparison. This approach made it computationally feasible to compare a number of variations of the basic fictitious play model. We were also able to alter the model by incorporating some of the features of an evolutionary model without having to worry that the new, hybrid model would not be analytically tractable. Hopefully, this procedure will lead to more complicated models of learning that better approximate real learning.

Our basis for comparing these models was their short-term convergence behavior. We argued that this behavior could be interpreted as the prediction of how likely particular equilibria were. A natural extension is to try to identify learning models that more accurately depict the learning behavior of real people and to use a distribution on parameters that more closely fits a particular population. In that case we can hope that the methodology of simulation will allow us to say conclusively that a certain outcome in a game will happen with a particular probability. In the interim the predictions found here will serve as a first pass. The general prediction is that games with a great deal of risk dominance are almost sure to eventually have the risk dominant action played by everyone. Conversely, the games between these regions of prediction with certainty have a non-zero probability of each action. These are the games where a conclusive prediction of the probability of each outcome offers a lot to the applied user of game theory.

The comparisons also gave some results that should be useful in developing a model of learning that more closely approximates reality. Weighting more recent observations more heavily in the calculation of beliefs about others' play did not affect the expected action converged to, but did increase the speed of convergence. Adding a cost of updating caused a dramatic shift in the action converged to for some games, but only for some distributions of the players' propensity to update their behavior after one non-equilibrium outcome relative to another. Making fictitious play more like an evolutionary model by changing the way that players respond to their information from adopting the best response to emulating other players slowed the process down dramatically.

Figure 3-1: Standard Fictitious Play

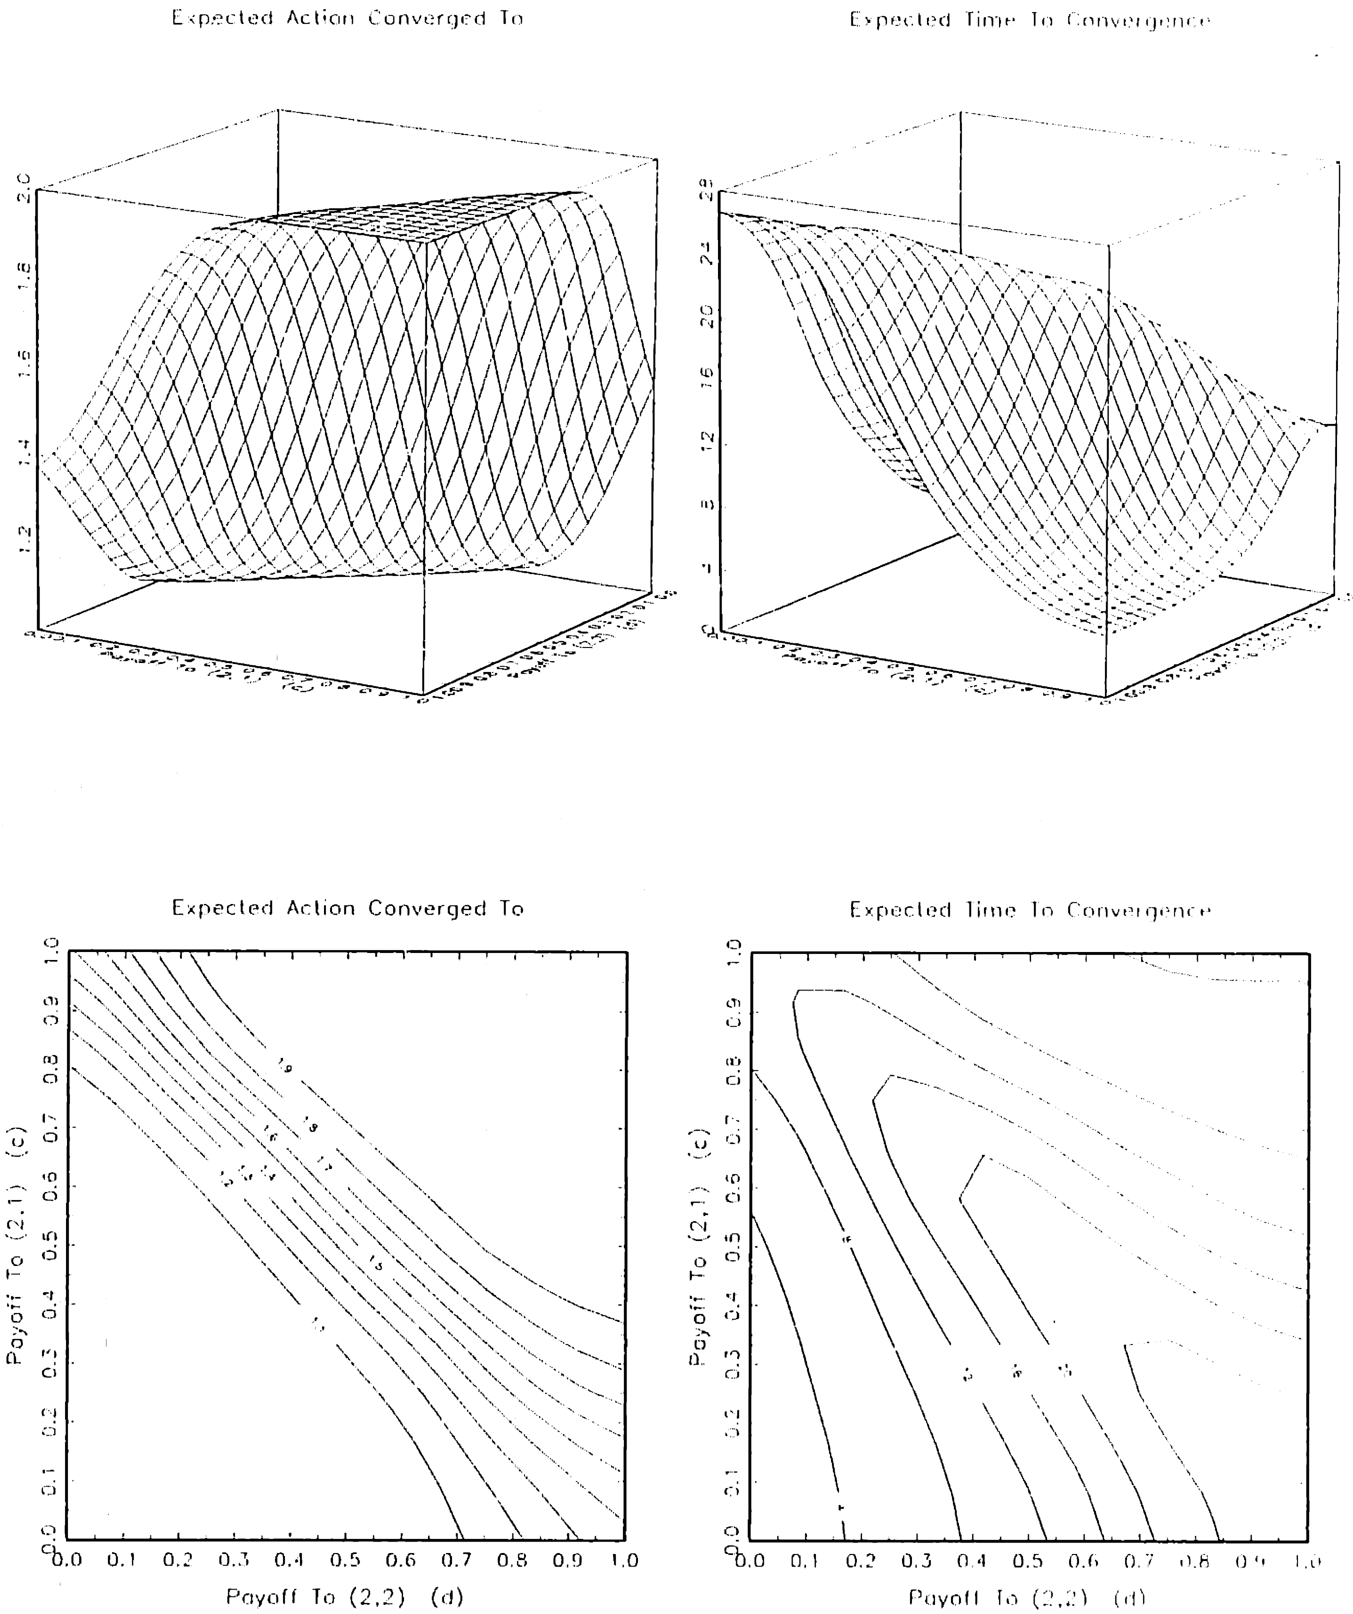

Figure 3-2: Fictitious Play With Linear Discounting

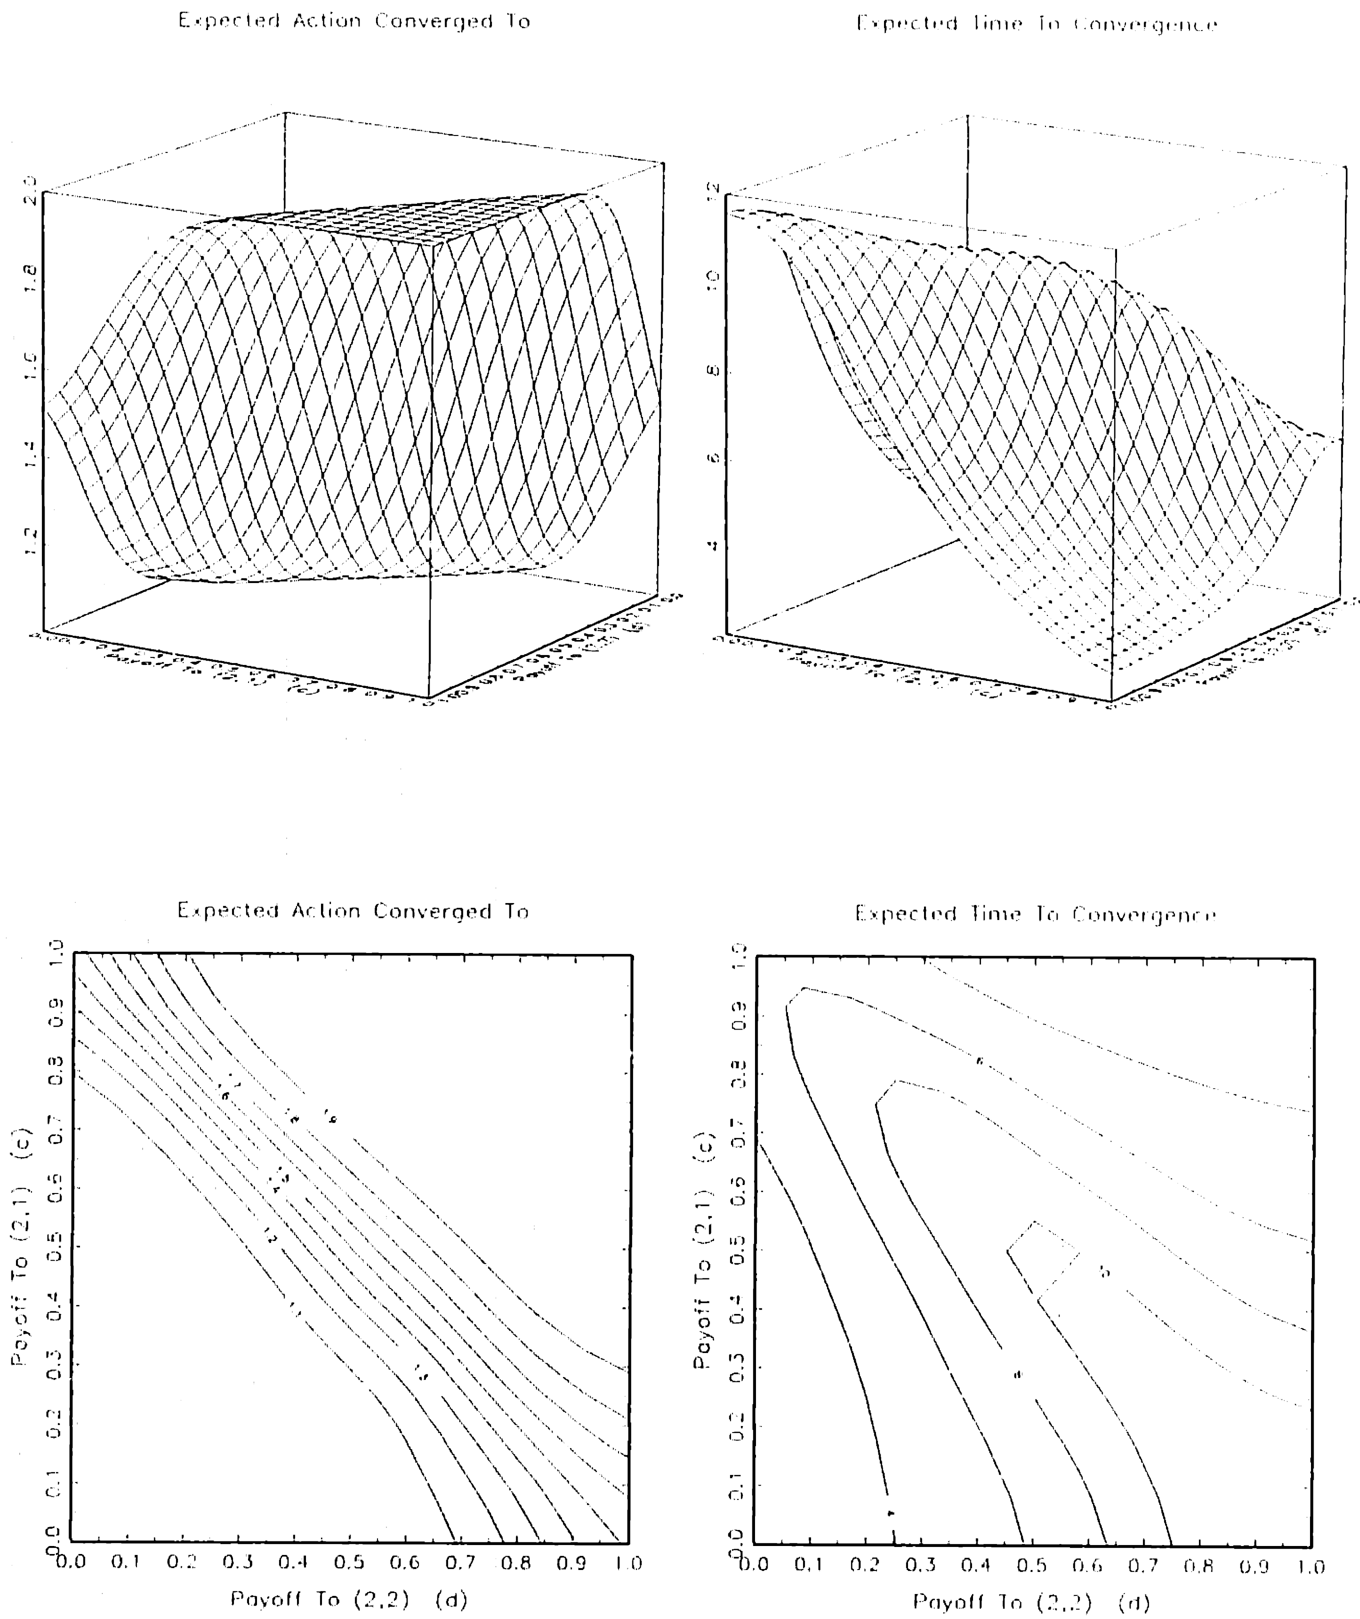

Figure 3-3: Fictitious Play With Convex Discounting

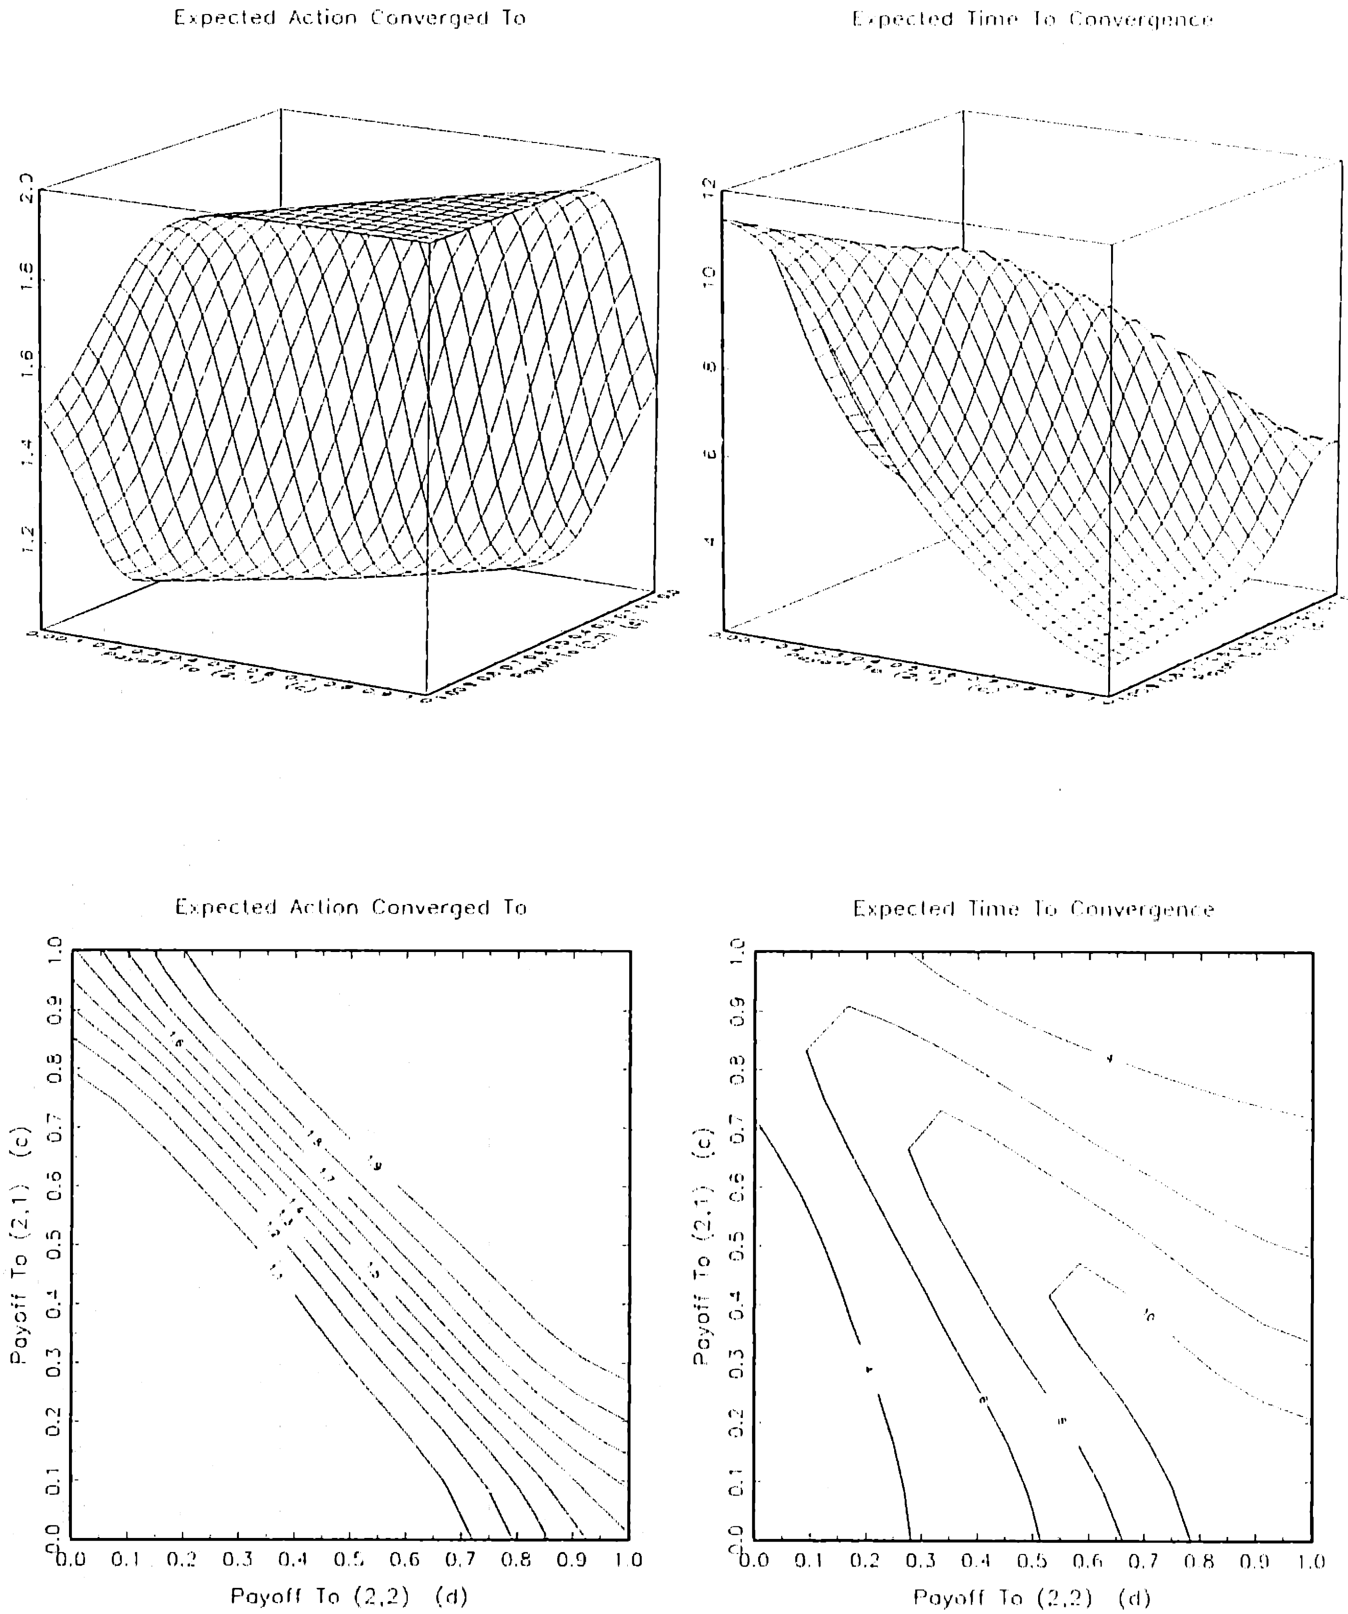

Figure 3-4: Congruence of No Weighting and Linear Weighting Models

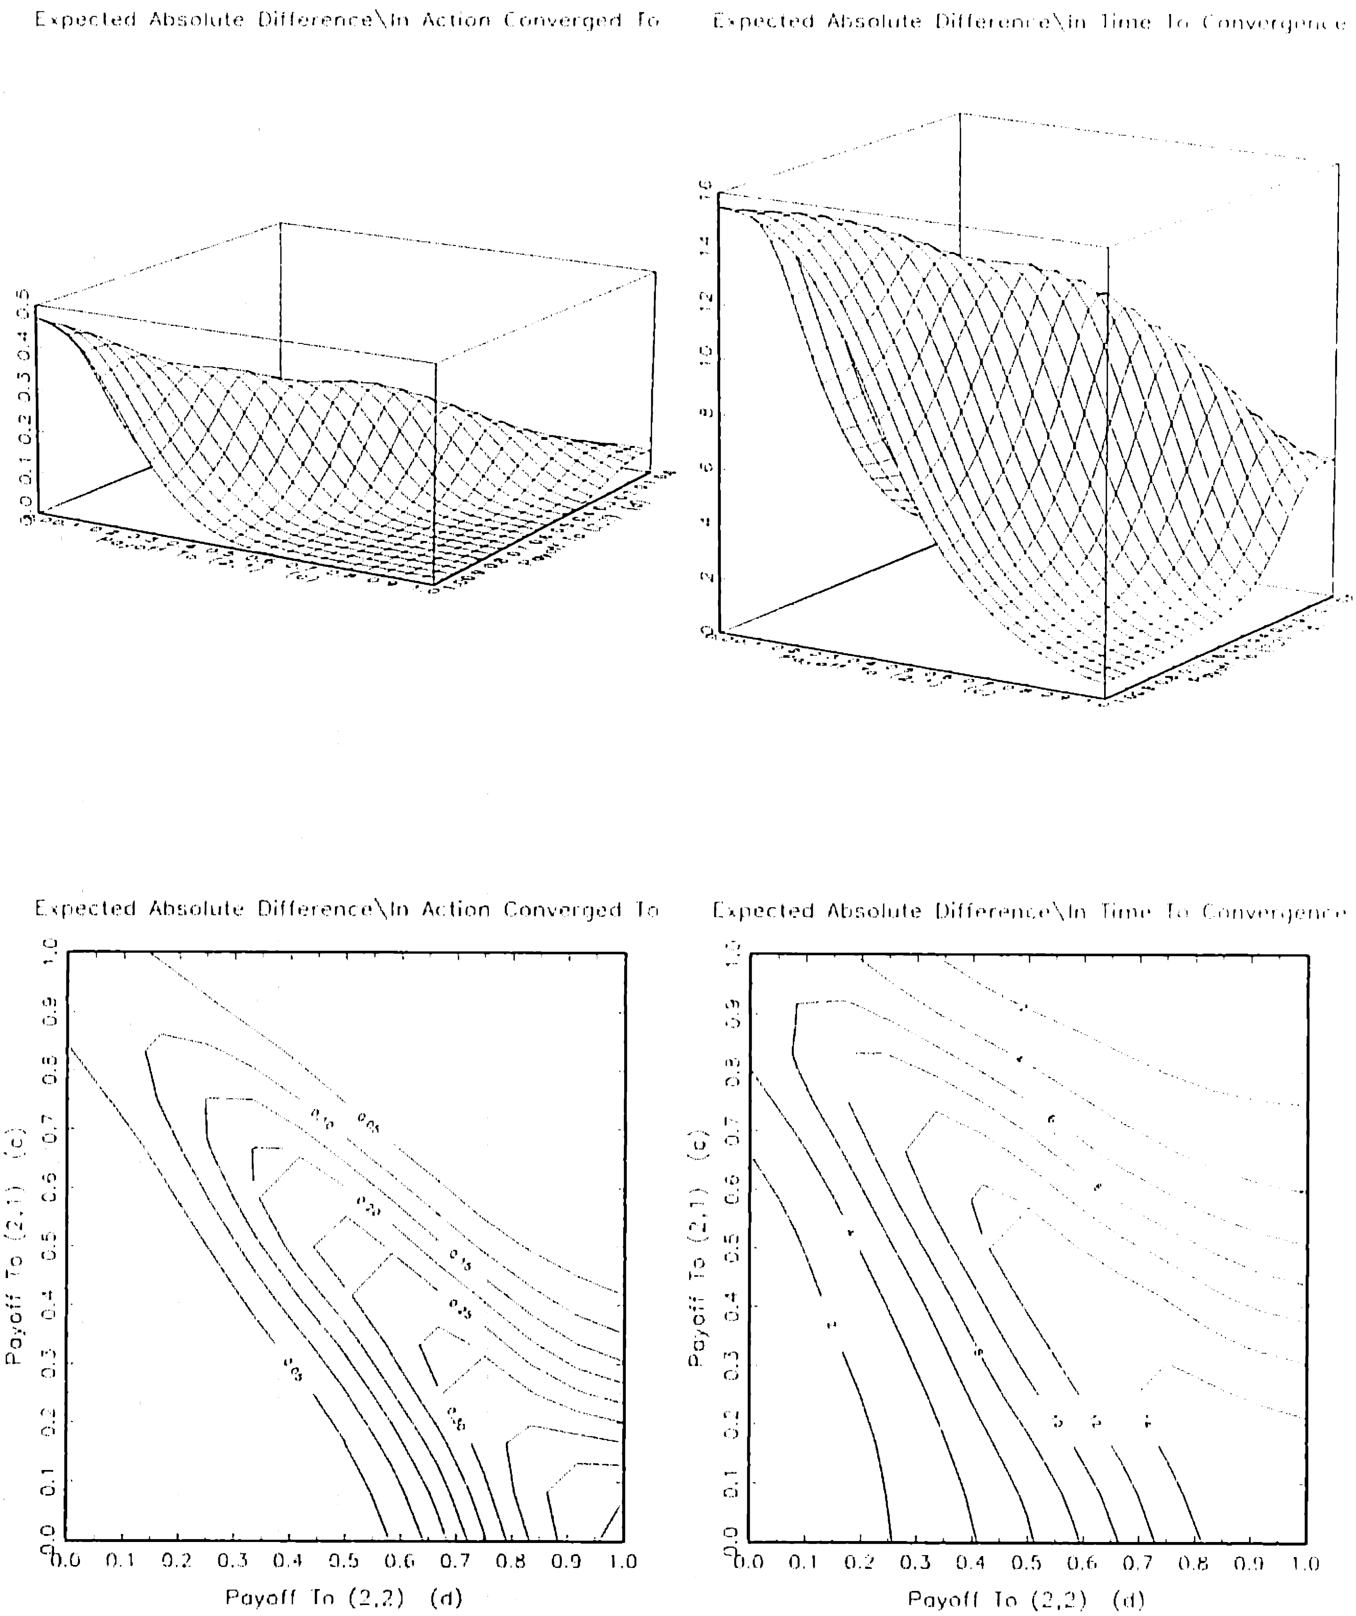

Figure 3-5: Congruence of No Weighting and Convex Weighting Models

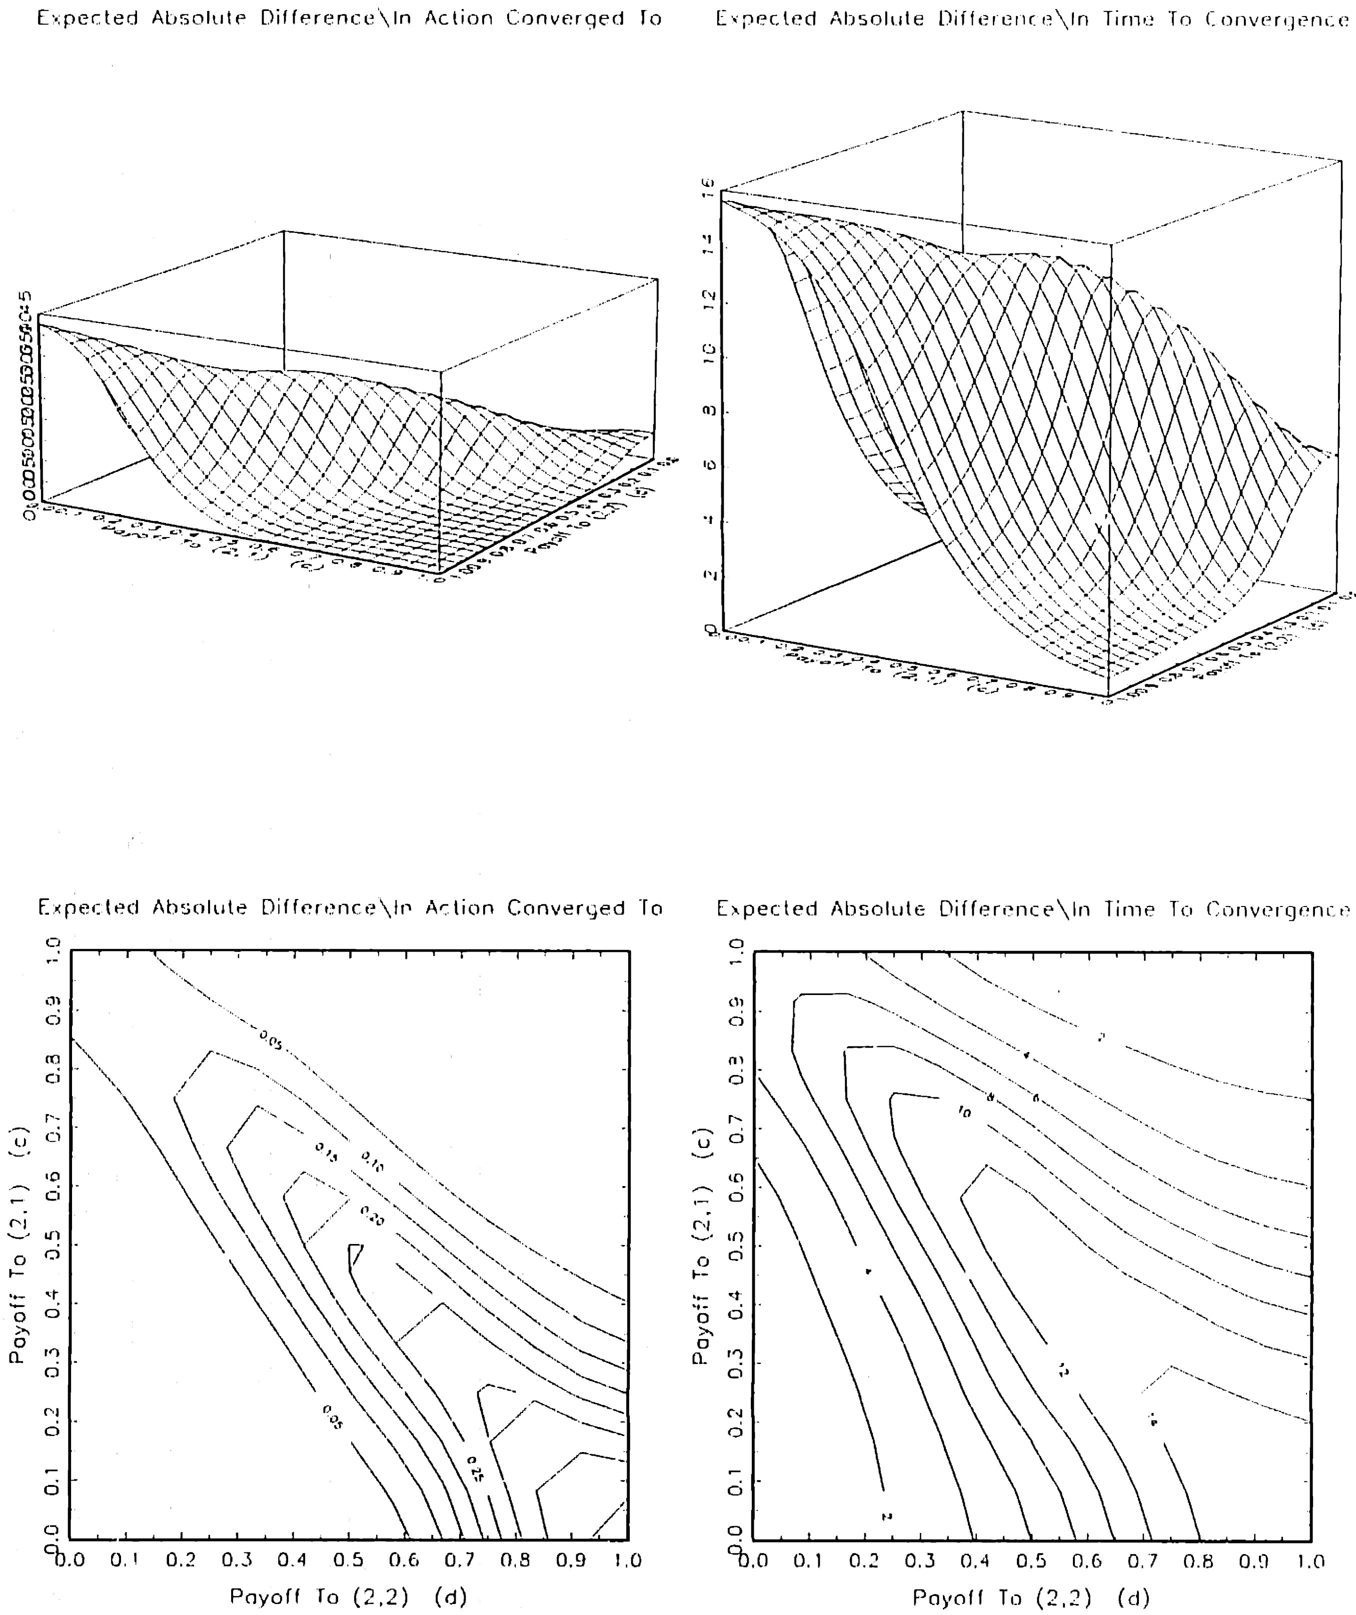

Figure 3-6: Fictitious Play With Heterogeneous Discounting

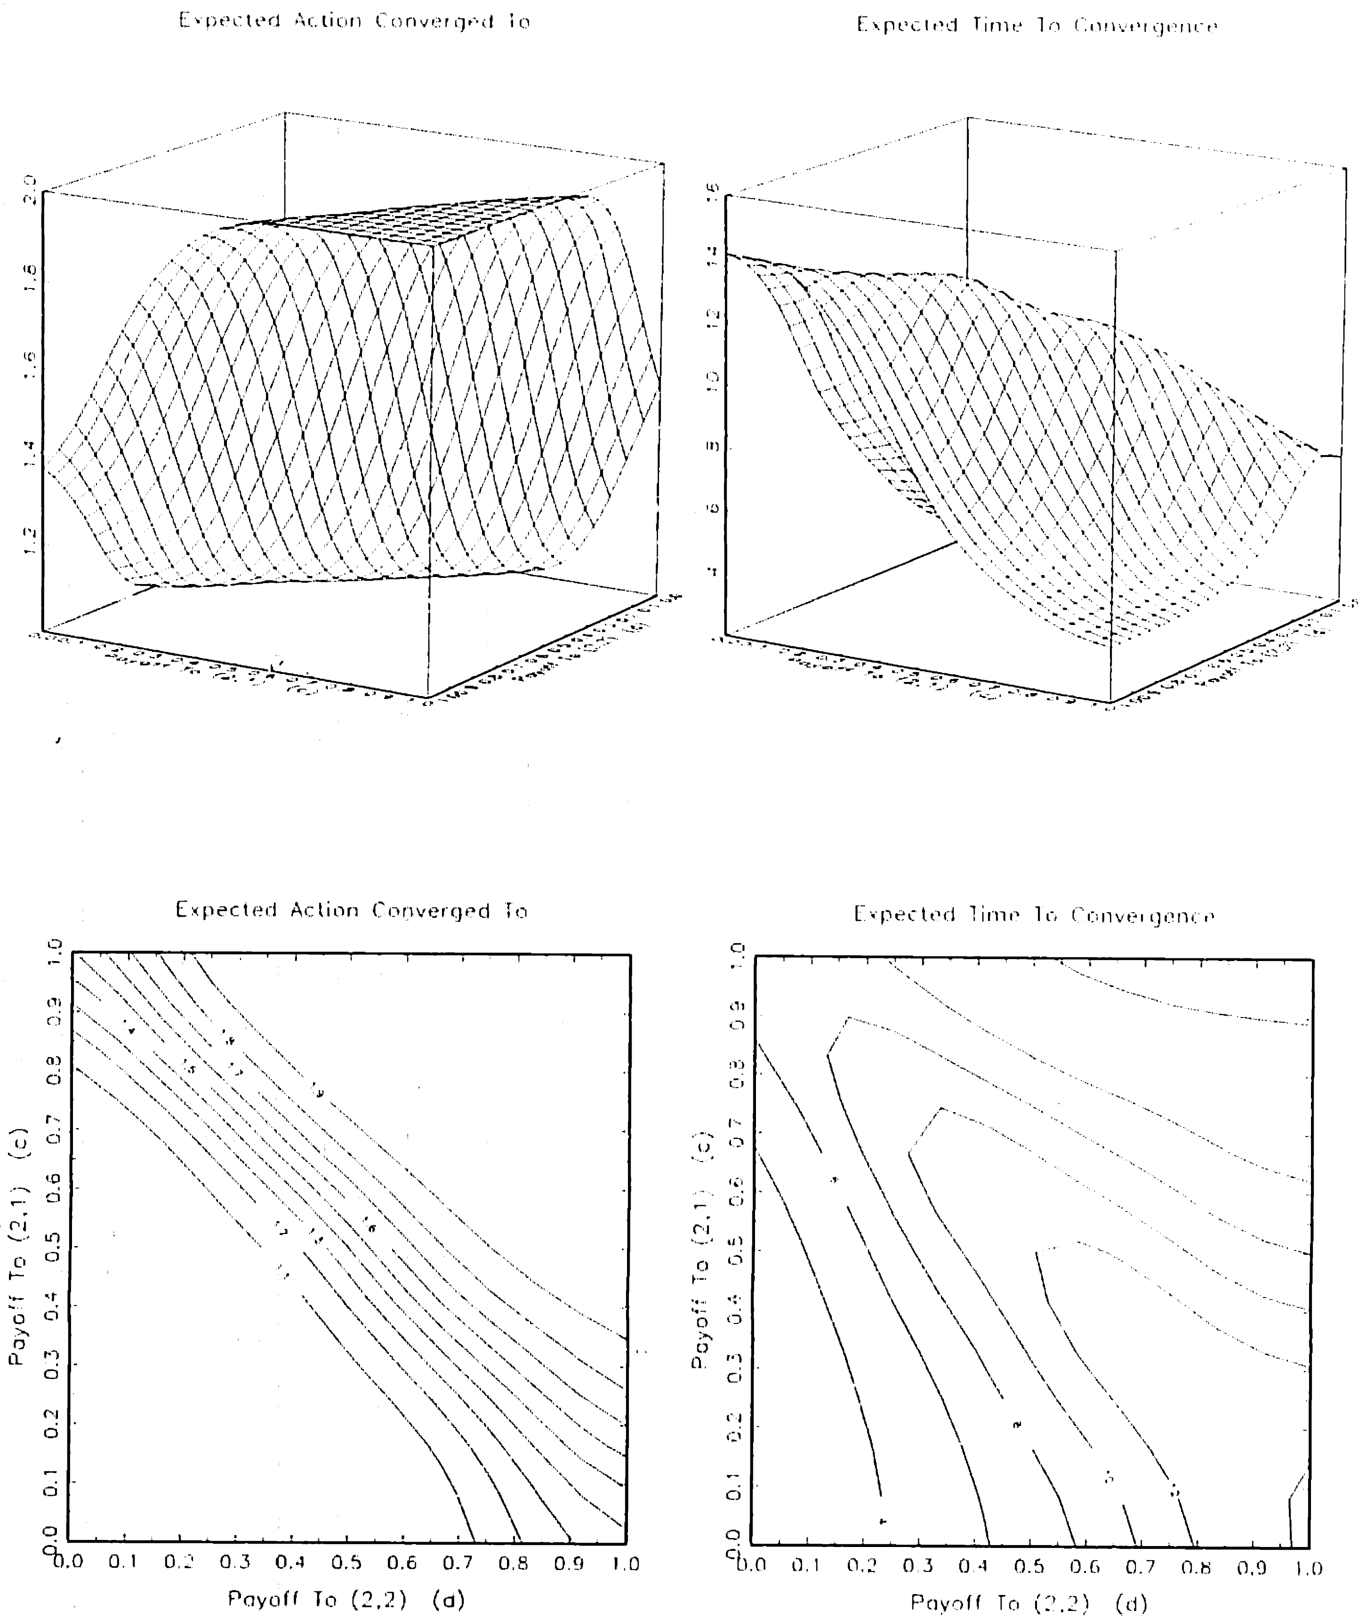

Figure 3-7: Fictitious Play With Widely Distributed Aspiration Levels

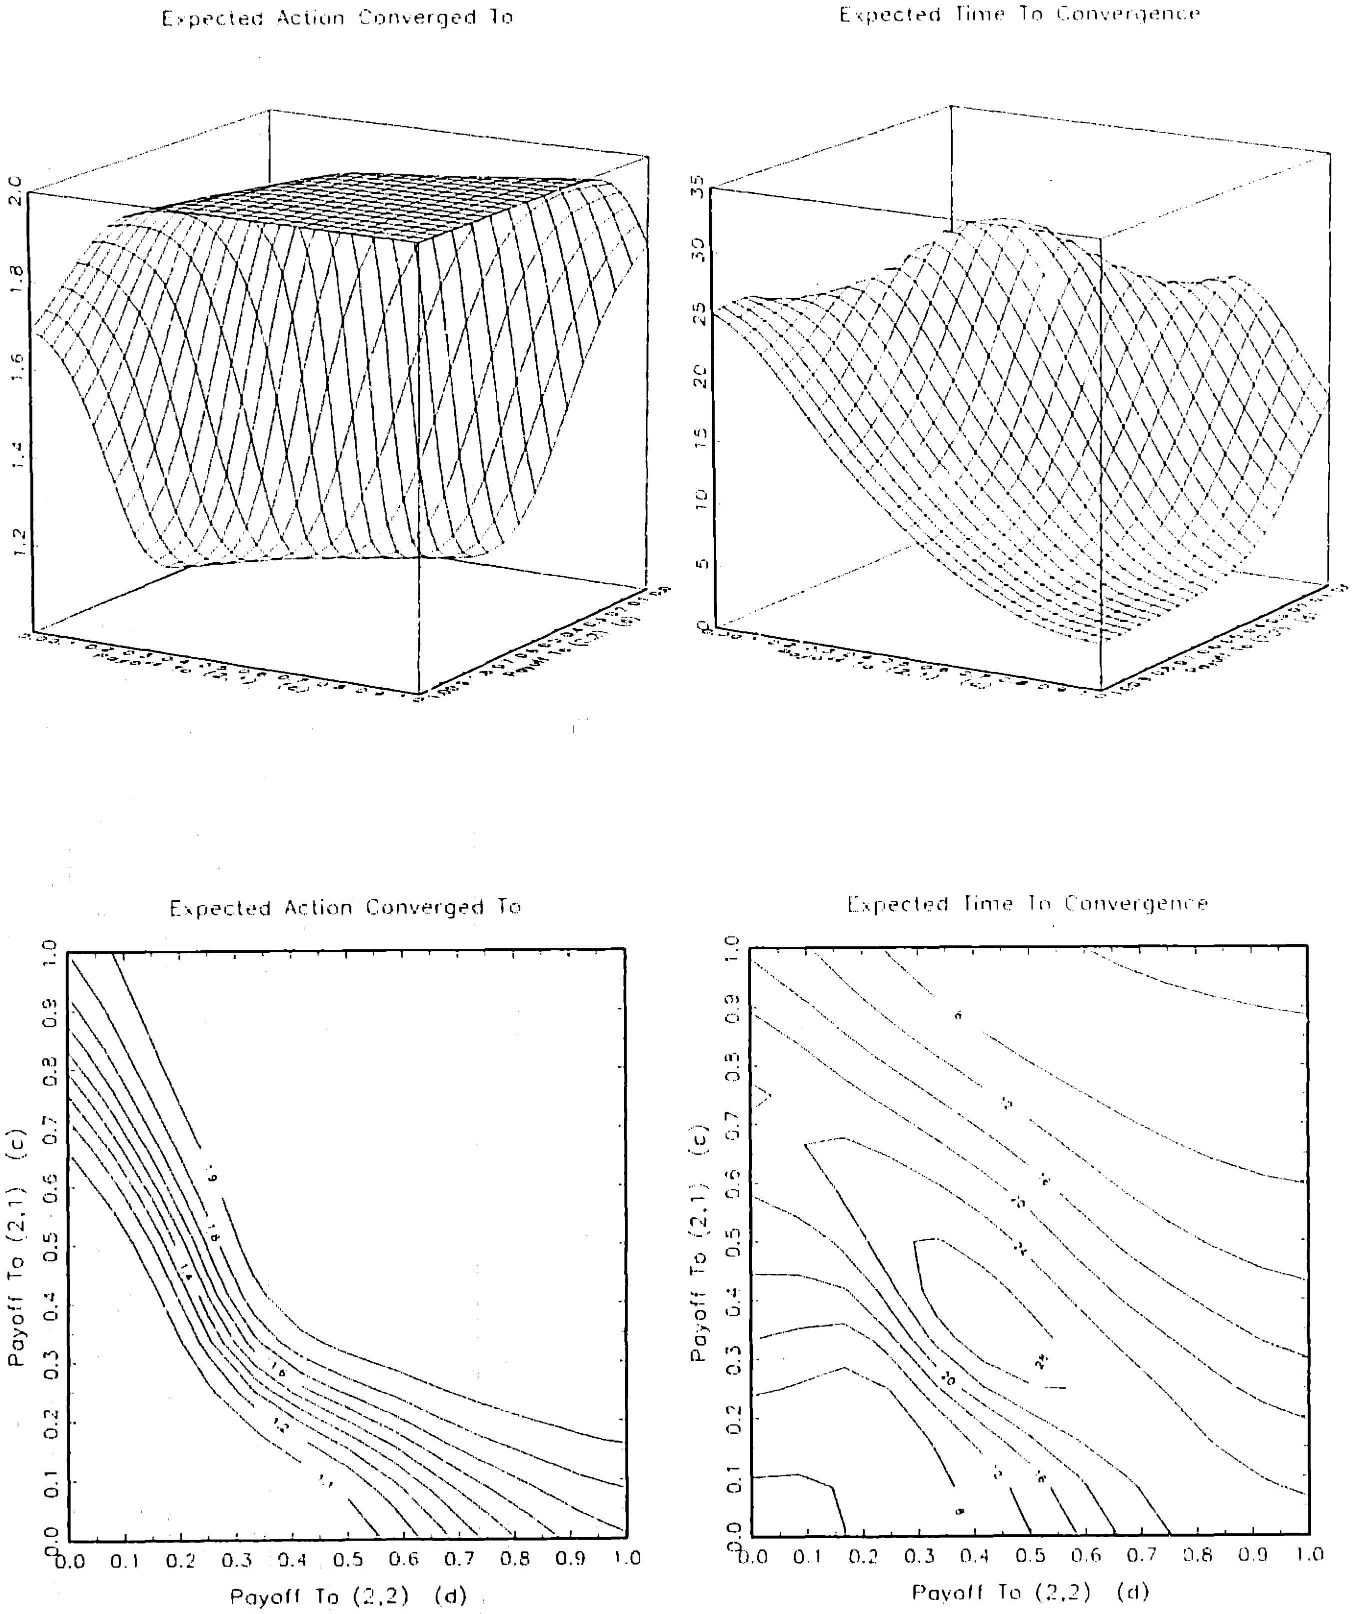

Figure 3-8: Fictitious Play With Higher Aspiration Levels

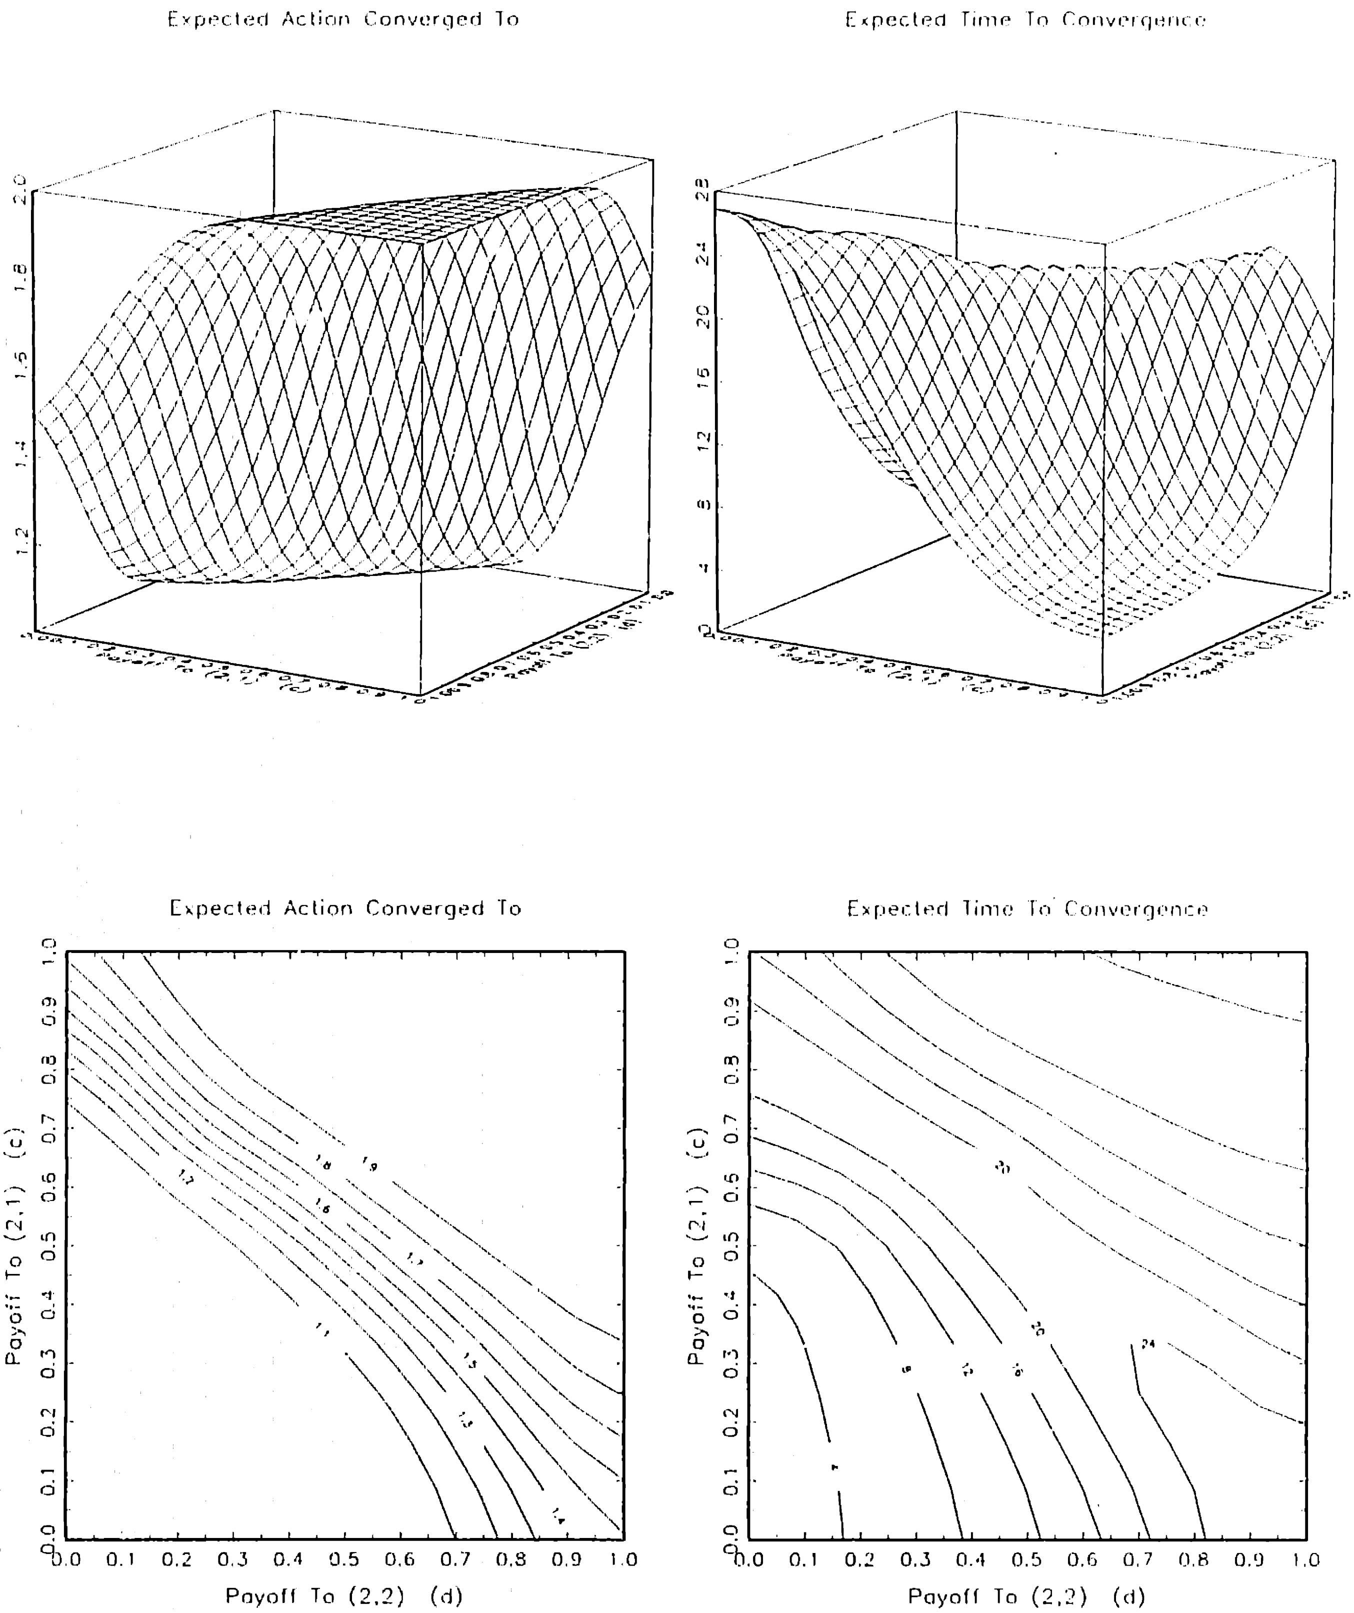

Figure 3-9: "Fictitious Play" Version of Binmore-Samuelson

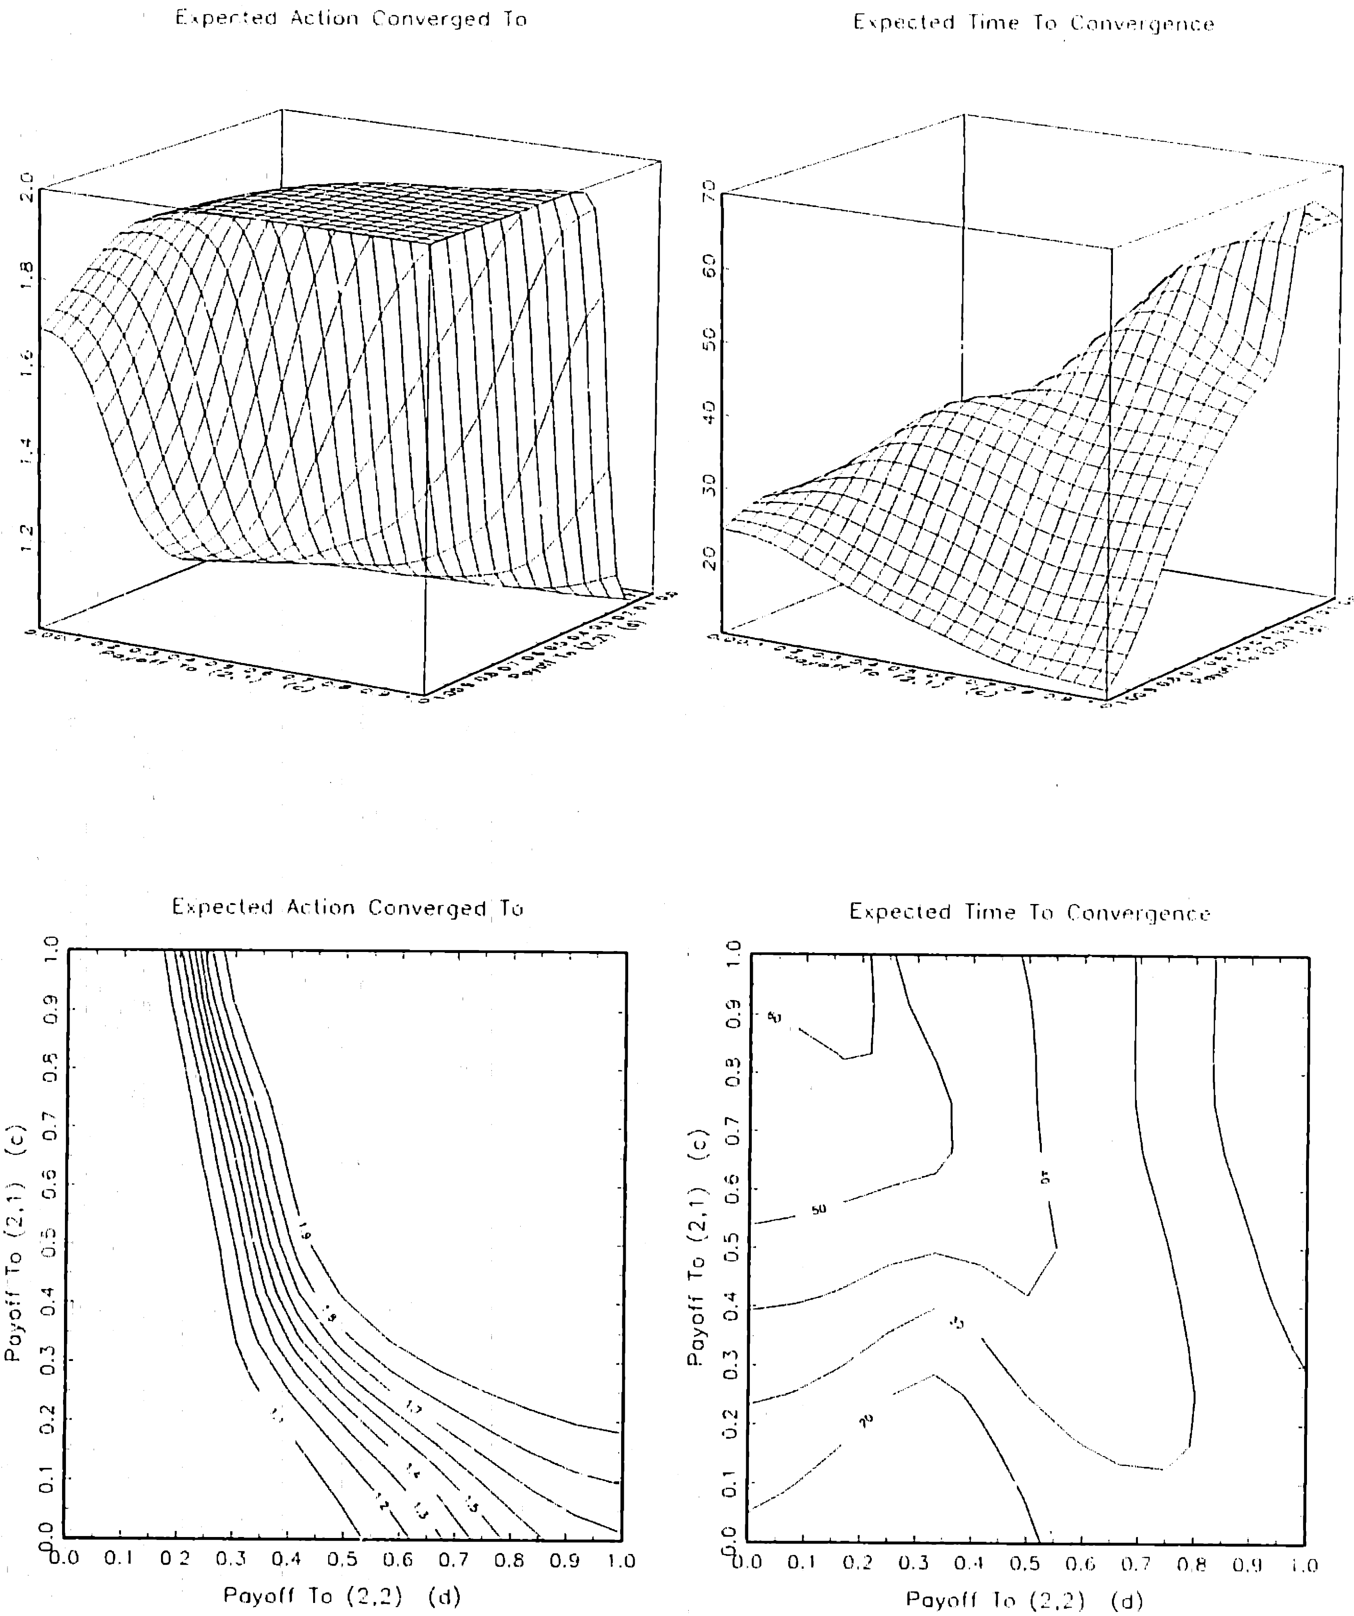

Figure 3-10: Adapted Binmore-Samuelson With Higher Aspiration Levels

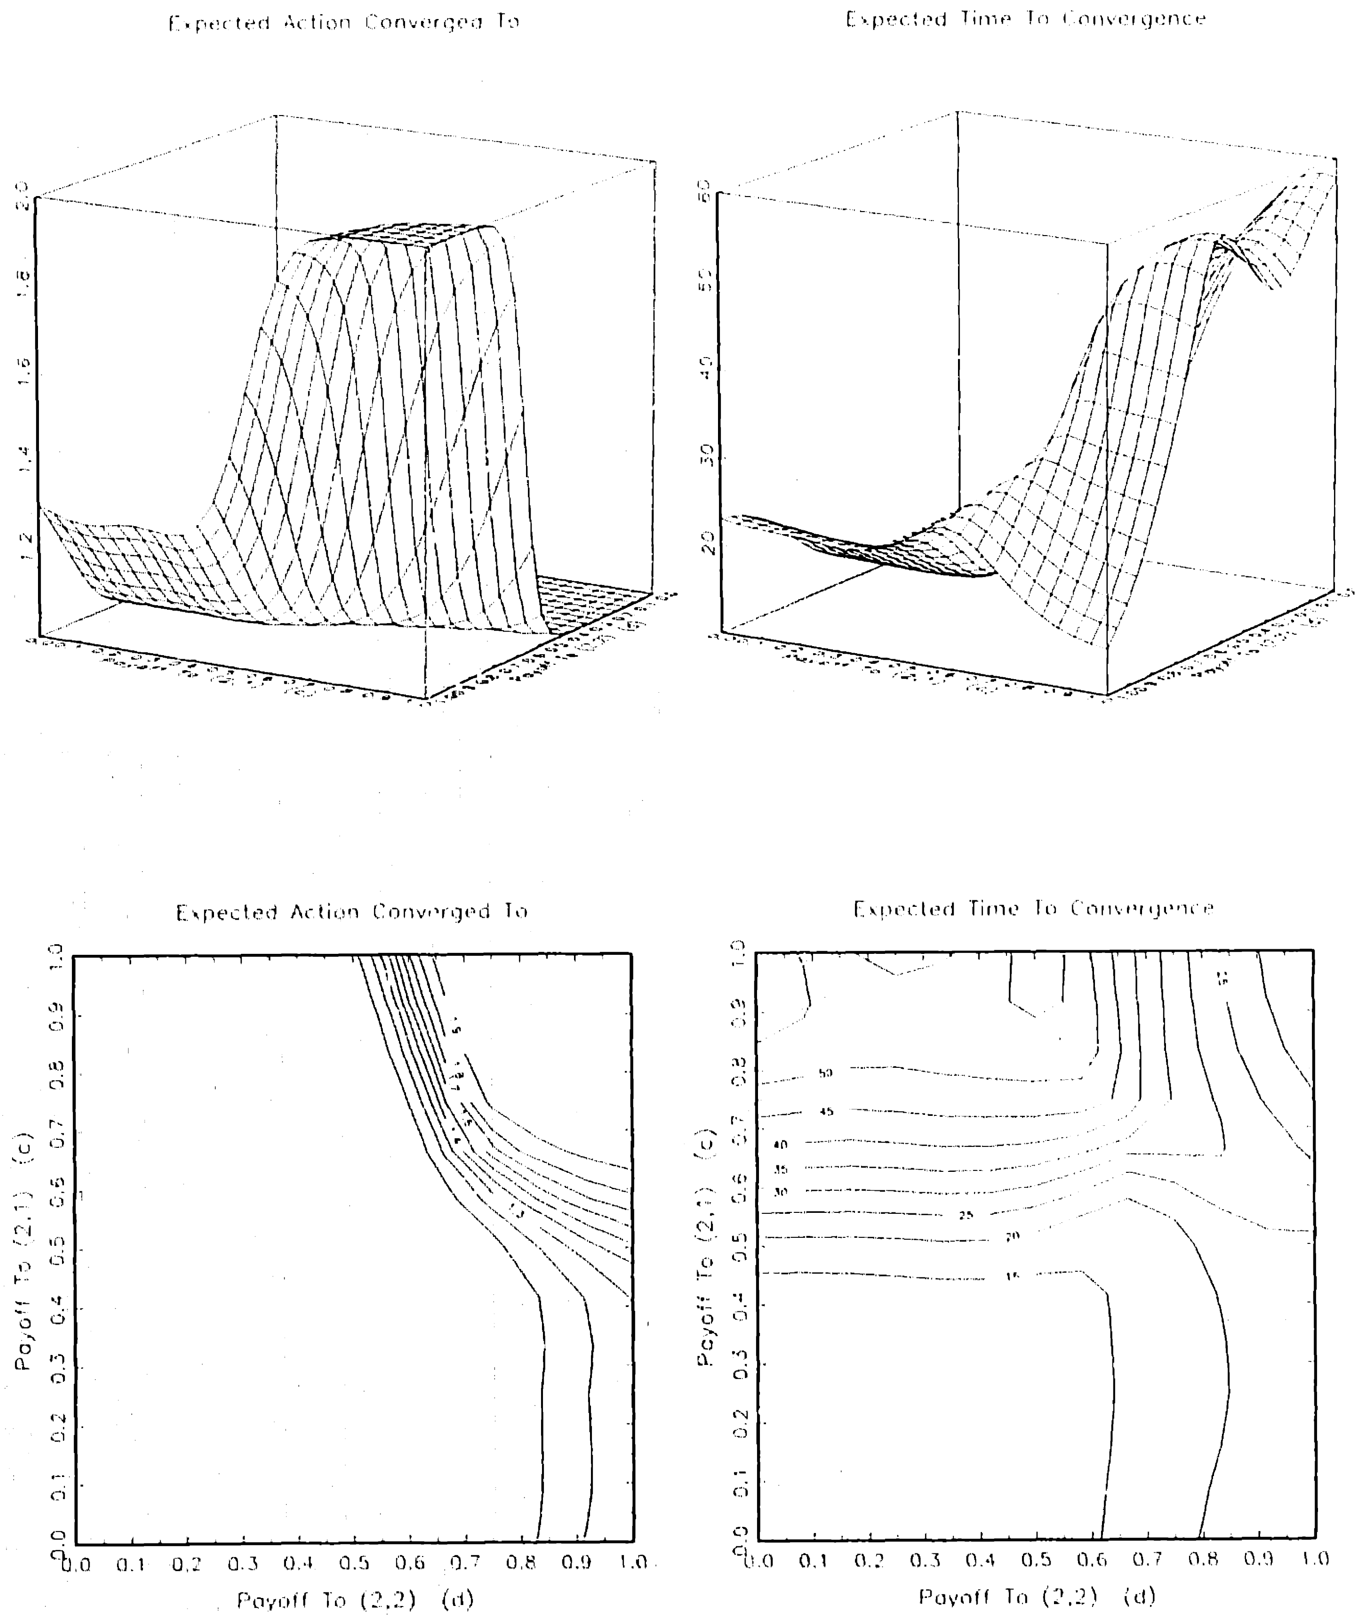

Figure 3-11: Congruence of Binmore-Samuelson with Different Aspiration Dist's

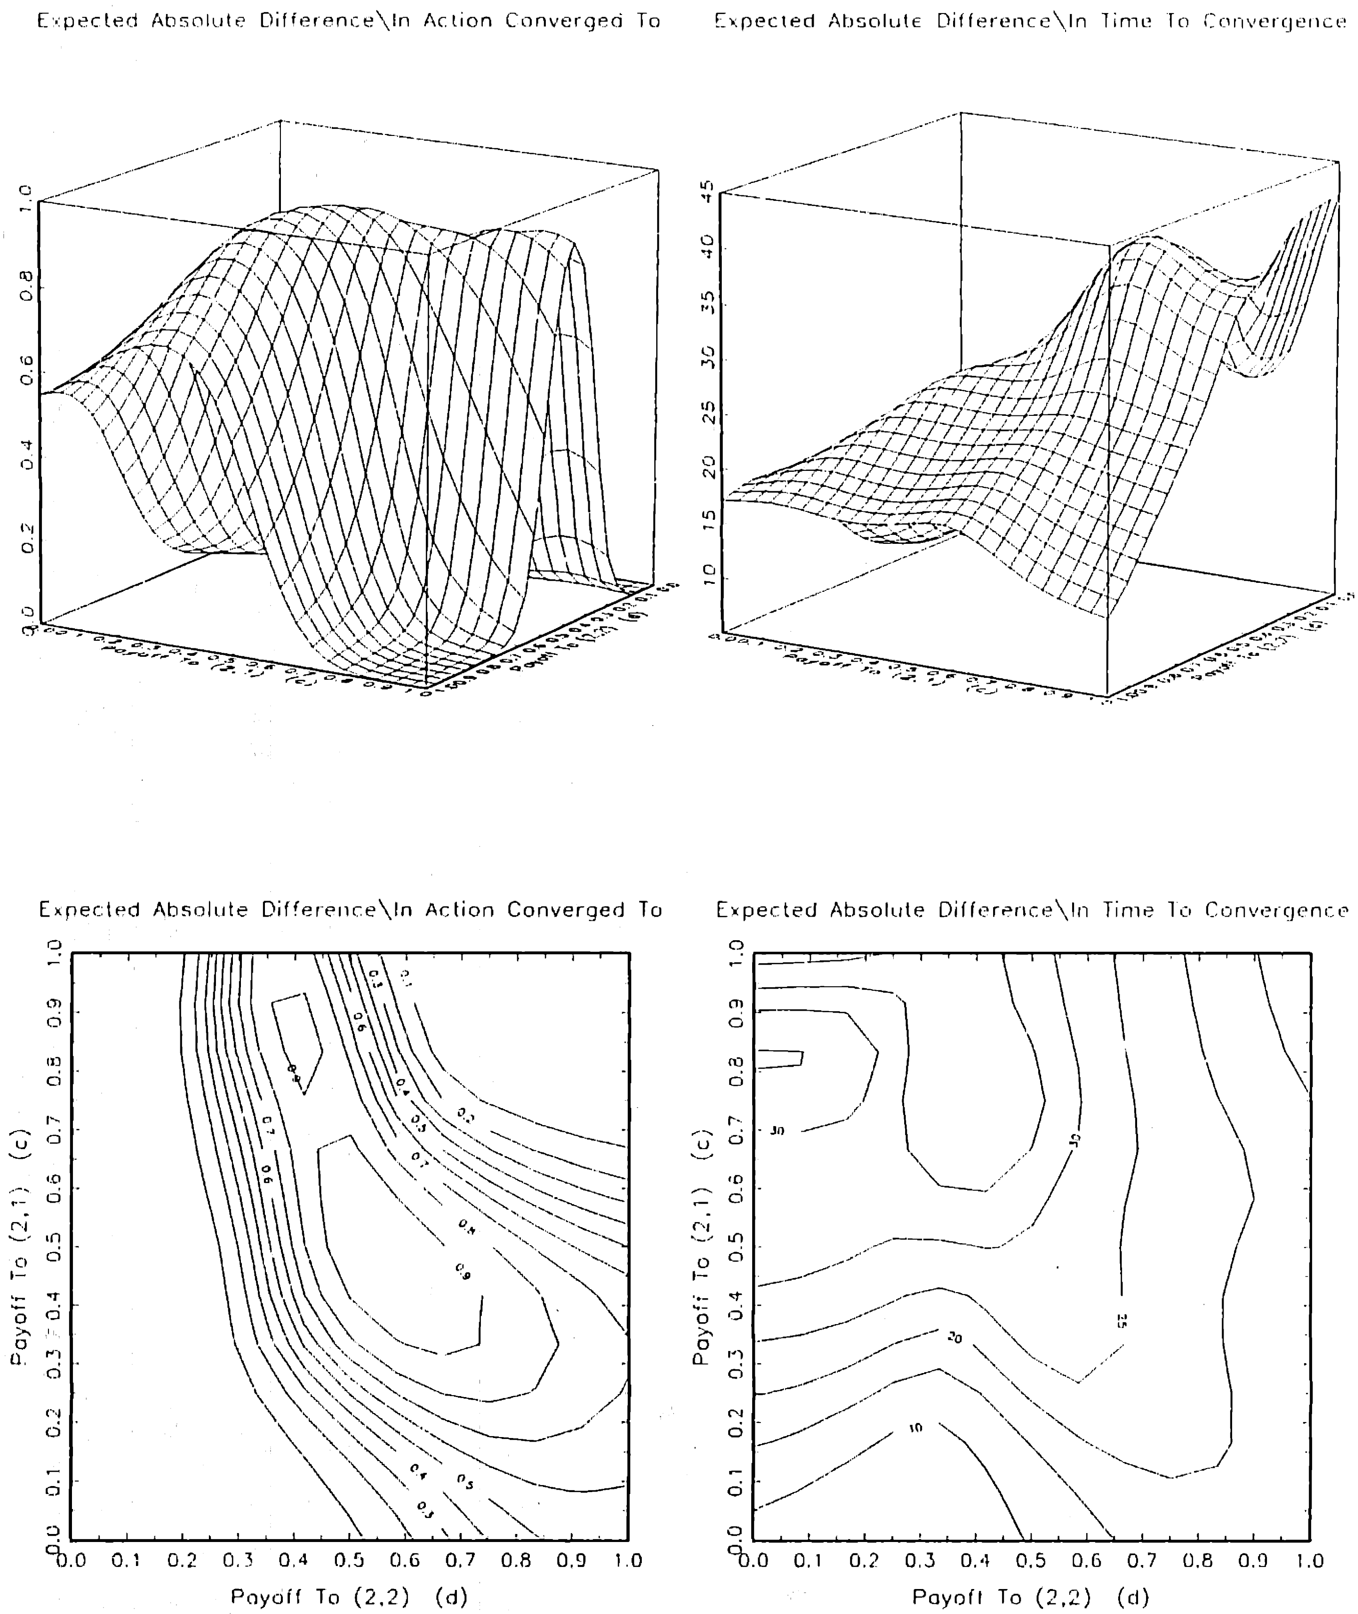

Figure 3-12: Congruence of Fictitious Play With Adapted Binmore-Samuelson

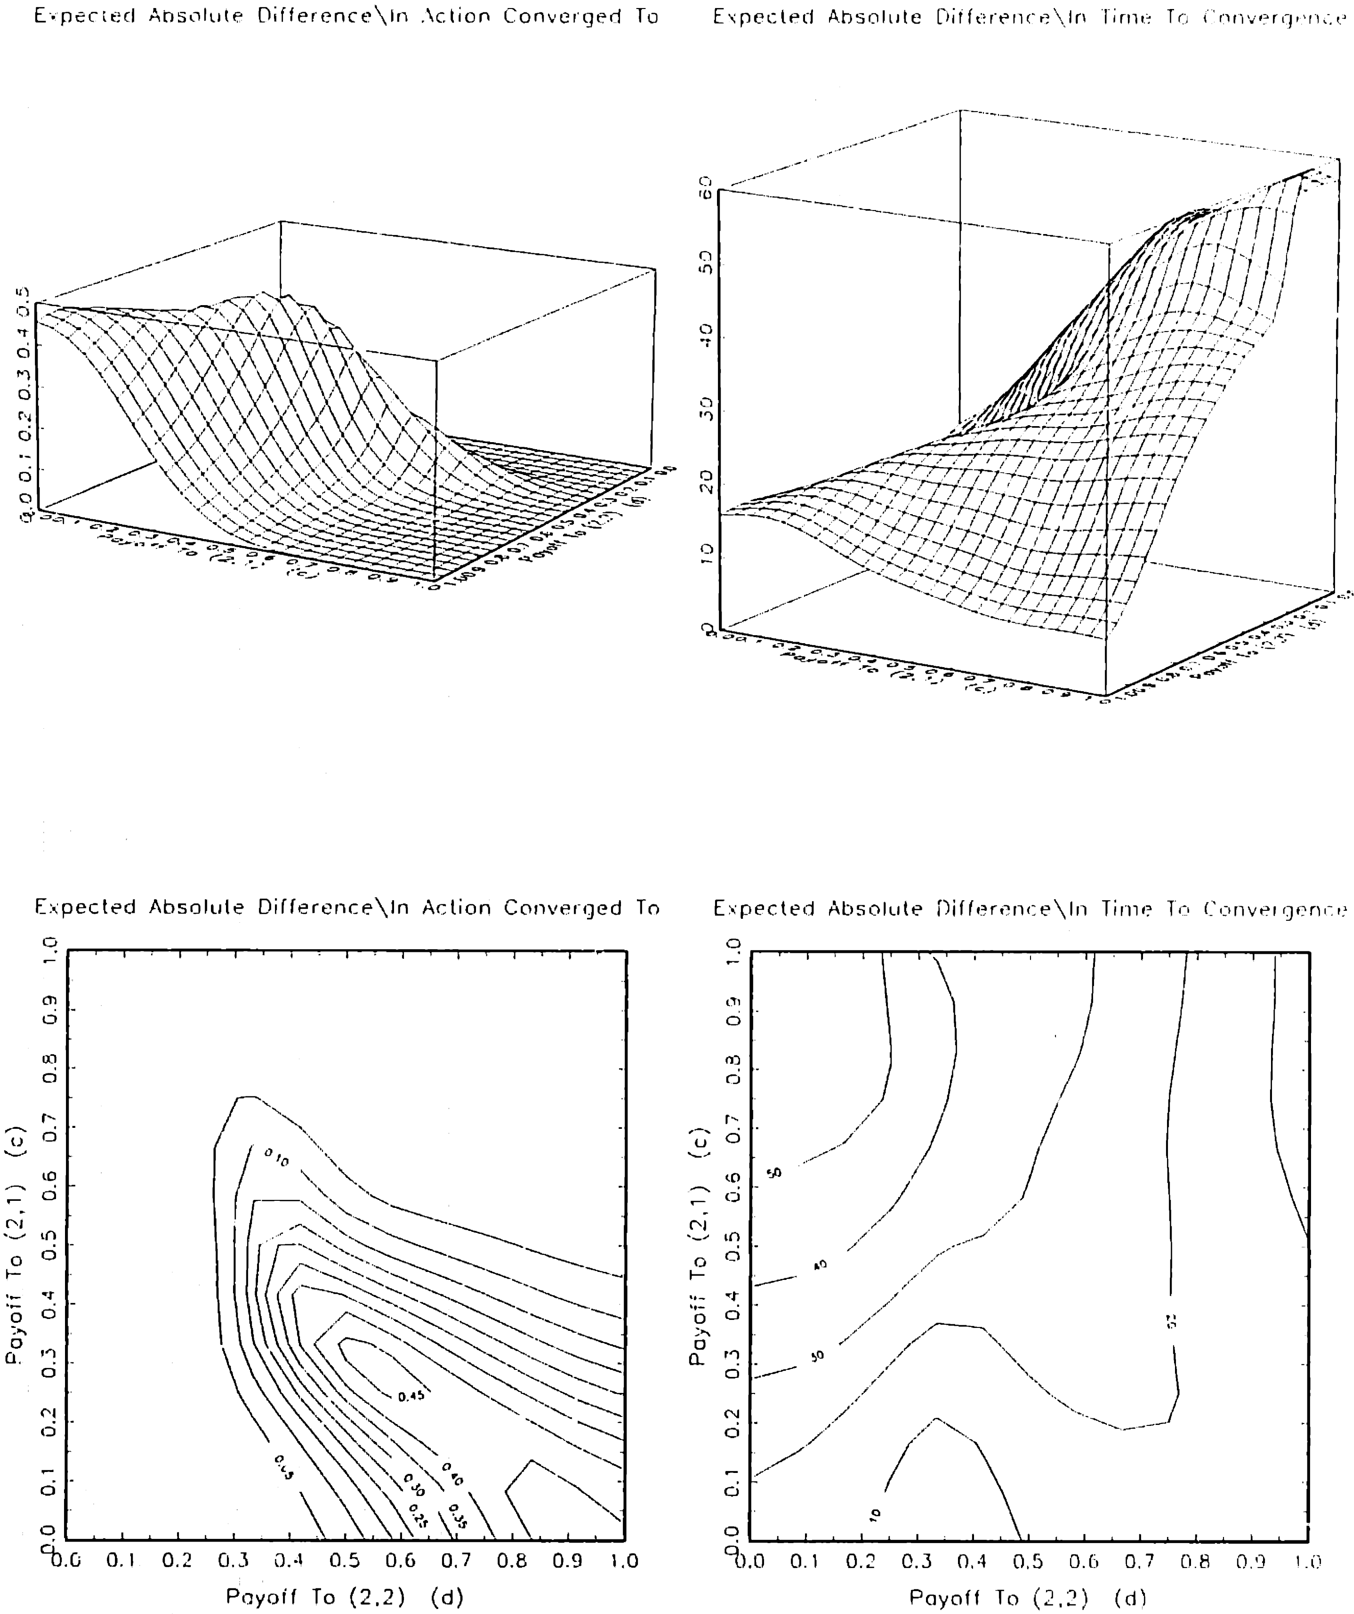

Figure 3-13: Fictitious Play With Correlated Priors

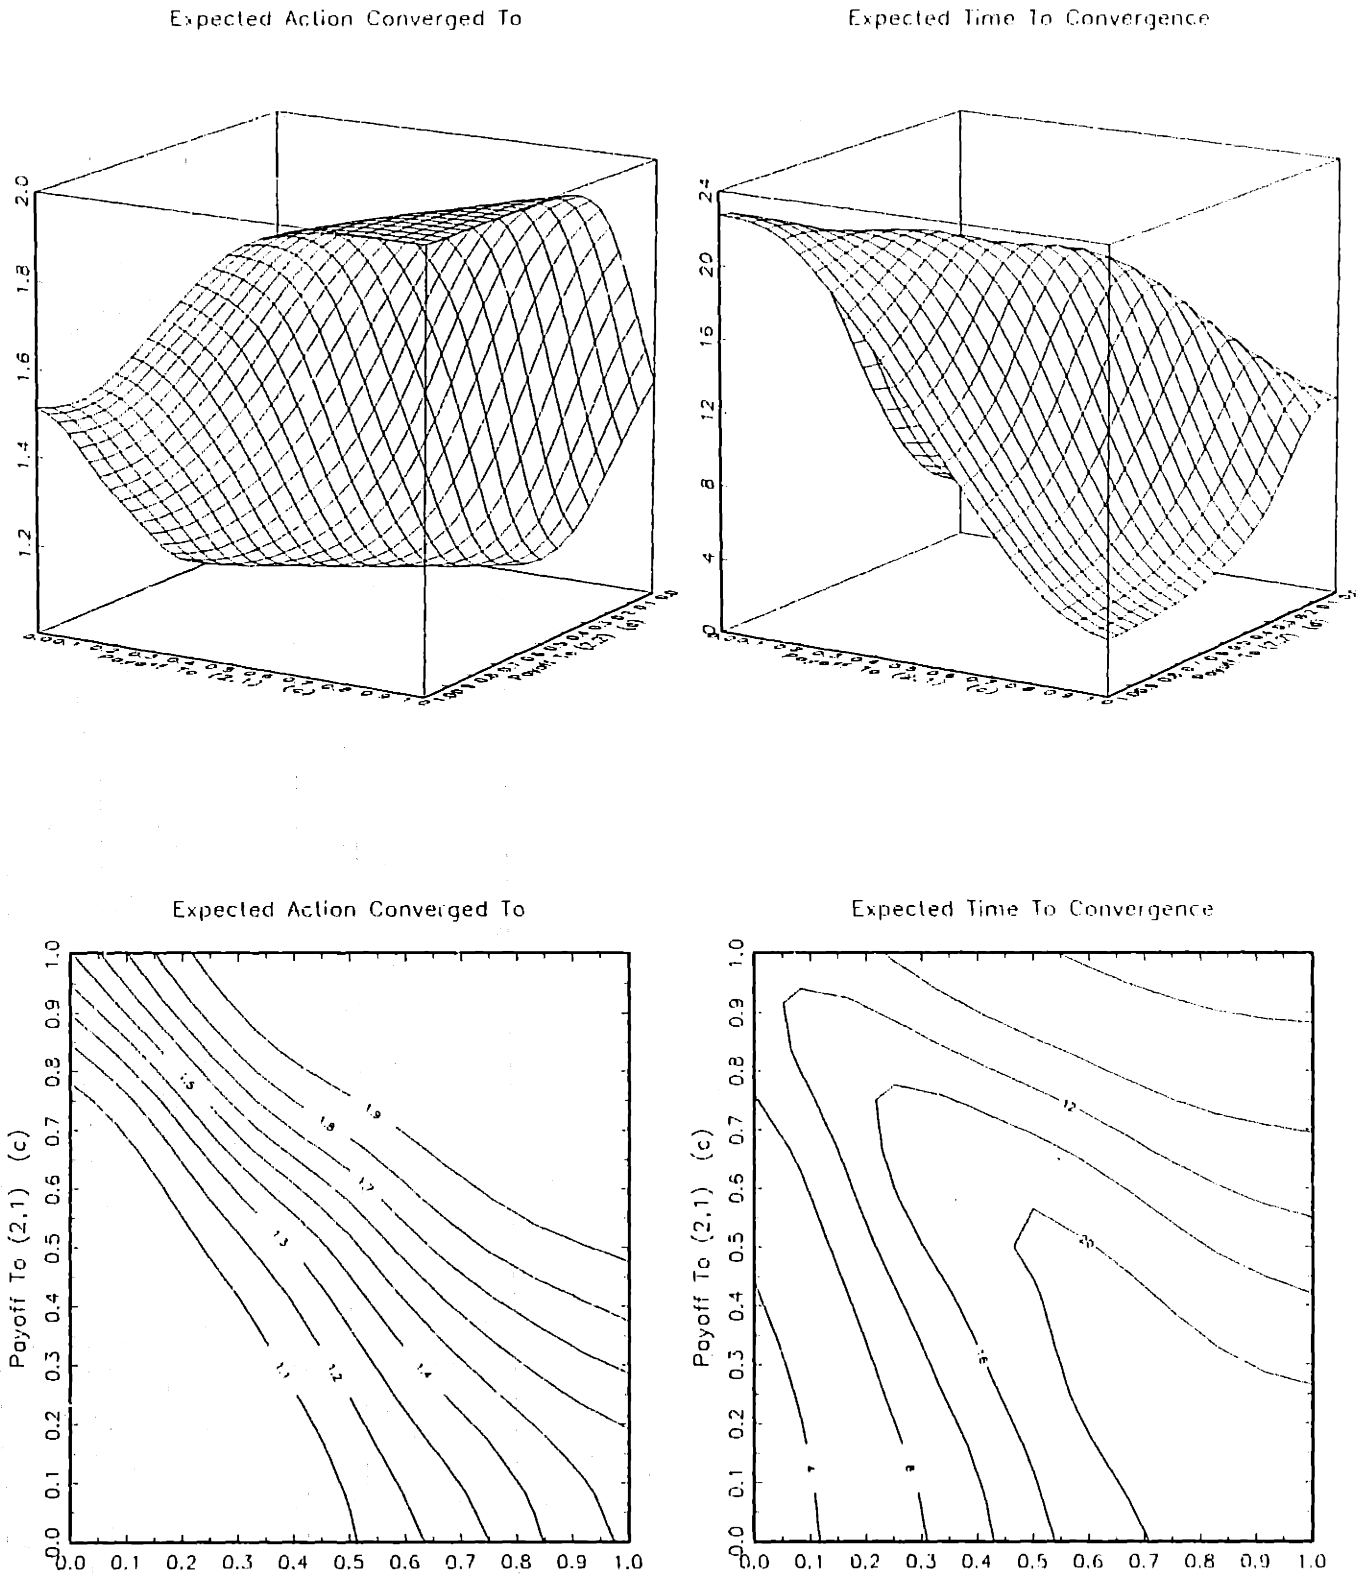

References

Bhatt, S. (1987), "Strategic Product Choice in Differentiated Markets," *Journal of Industrial Eonomics* **36**, 207–216.

Bierens, Herman J., "Kernel estimators of regression functions," in Truman F. Bewley, ed., *Advances in Econometrics–Fifth World Congress*, Volume 1, Cambridge University Press, 1987.

Binmore, Kenneth, and Larry Samuelson (1992), "Muddling Through: Noisy Equilibrium Selection," mimeo.

Bonanno, G. (1987), "Location Choice, Product Proliferation and Entry deterrence," *Review of Economic Studies* **43**, 37–45.

Brander, James, and Jonathan Eaton (1984), "Product Line Rivalry," *American Economic Review* **74**, 323–334.

Brandts, Jordi, and Charles Holt (1992), "An Experimental Test of Equilibrium Dominance in Signalling Games," *American Economic Review* **82**, 1350-65.

Bulow, Jeremy, John Geanakoplos, and Paul Klemperer (1985a), "Multimarket Oligopoly: Strategic Substitutes and Complements," *Journal of Political Economy* **93**, 488–511.

Bulow, Jeremy, John Geanakoplos, and Paul Klemperer (1985b), "Holding Idle Capacity to Deter Entry," *Economic Journal* **95**, 178–182.

Butters, Gerard R. (1977), "Equilibrium Distribution of Prices and Advertising," *Review of Economic Studies* **44**, 465-492.

Camerer, Colin, and Keith Weigelt (1991), "Expeimental Tests of a Sequential Equilibrium Reputation Model," *Econometrica* **56**, 1-36.

Cooper, Russell, Douglas DeJong, Robert Forsythe and Thomas Ross (1989), "Communication in the Battle of the Sexes Game: Some Experimental Results," *Rand Journal of Economics* **20**, 568-587.

Cooper, Russell, Douglas DeJong, Robert Forsythe and Thomas Ross (1990), "Selection Criteria in Coordination Games: Some Experimental Results," *American Economic Review* **80**, 218-233.

Crawford, Vincent (1991), "An 'Evolutionary' Interpretation of Van Huyck, Battalio, and Beil's Experimental Results on Coordination," *Games And Economic Behavior* **3**, 25-59.

Crawford, Vincent (1993), "Adaptive Dynamics in Coordination Games," mimeo.

Deneckere, Raymond, and Michael Rothschild (1986), "Monopolistic Competition and Preference Diversity," Discussion Paper No. 684, Northwestern University.

Eaton, B. C., and R. Lipsey 1979), "The Theory of Market Preemption: Barriers to Entry in a Growing Spatial Market," *Economics* **46**, 149--158.

Ellison, Glenn (1993a), "Learning and Local Interaction," *Econometrica* **61**, 1047-1071.

Ellison, Glenn (1993b), "A Little Rationality and Learning From Personal Experience," Harvard Institute of Economic Research Discussion Paper No. 13.

Fudenberg, Drew, and David Kreps (1991) "Learning and Equilibrium in Games," mimeo.

Fudenberg, Drew, and Jean Tirole (1983), "Capital as a Commitment: Strategic Investment to Deter Mobility," *Journal of Economic Theory* **31**, 227–256.

Fudenberg, Drew, and Jean Tirole (1984), "The Fat Cat Effect, the Puppy Dog Ploy and the Lean and Hungry Look," *American Economic Review, Papers and Proceedings* bf 74, 361–368.

Gale, Douglas, "Branch Banking, Unitary Banking and Competition," mimeo, Boston University.

Gordy, Michael (1991a), "An Empirical Note on Price Competition in the Portuguese Banking Industry," mimeo, Banco de Portugal.

Gordy, Michael (1991b), "On the Recent Expansion of Bank Branch Networks in Portugal," mimeo, Banco de Portugal.

Härdle, Wolfgang, *Applied Nonparametric Regression*, Cambridge University Press, 1990.

Harsanyi, John (1973), "Games With Randomly Disturbed Payoff: A New Rationale for Mixed-Strategy Equiibrium Payoff," *International Journal of Game Theory* **2**, 1–23.

Hay, D. (1976), "Sequential Entry and Entry-Deterring Strategies in Spatial Competition," *Oxford Economics Papers* **28**, 240–257.

Hotelling, Harold (1929), "Stability in Competition," *Economic Journal* **39**, 41–57.

Judd, Kenneth (1985), "Credible Spatial Preemption," *Rand Journal of Economics 16*, 153–166.

Kandori, Michihiro, George Mailath and Rafael Rob (1992), "Learning, Mutation, and Long Run Equilibrium in Games," *Econometrica* **61**, 29-56.

Klemperer, Paul (1992), "Equilibrium Product Lines: Competing Head-to-Head May Be Less Competitive," *American Economic Review* **82**, 740–755.

Lane, W. (1980), "Product Differentiation in a Market with Endogenous Sequential Entry," *Bell Journal of Economics* **11**, 237–260.

Li Calzi, Marco (1992), "Fictitious Play by Similarities," mimeo, Universita di Parma.

Martin, Stephen (1979), "Advertising, Concentration, and Profitability: The Simultaneity Problem," *Bell Journal of Economics* **10**, 639–647.

Mason, Charles, and Owen Phillips (1990), "The Role of Gender in a Non-cooperative Game," *Journal of Economic Behavior and Organizations* **15**, 215-235.

Matsui, Akihiko, and Rafael Rob (1991), "Evolution, Rationality and Equilibrium Selection in Societal Games," CARESS Working Paper No. 91-31.

Matutues, Carmen, and Jorge Padilla (1993), "Shared ATM Networks and Banking Competition," *Journal of Industrial Economics*, forthcoming.

Miller, John and James Andreoni (1991), "Can Evolutionary Dynamics Explain Free Riding in Experiments," *Economic Letters* **36**, 9-15.

Nakamura, Leonard I., and Bruno M. Parigi (1992), "Bank Branching," mimeo, F.R.B. Philadelphia and Virginia P.I. & S.U., August.

Páscoa, Mário (1993), "A General Model of Chamberlin Competition," *Journal of Economic Theory*, forthcoming.

Prescott, E., and M. Visscher (1977), "Sequential Location Among Firms With Foresight," *Bell Journal of Economics* **8**, 378–393.

Rabin, Matthew (1992), "Incorporating Fairness Into Game Theory and Economics," U.C. Berkeley Working Paper No. 92-199.

Roth, Alvin, and Ido Erev (1993), "Learning in Extensive Form Games: Experimental Data and Simple Dynamic Models in the Intermediate Term," mimeo, U. of Pittsburgh.

Salop, Steven, "Monopolistic Competition with Outside Goods," *Bell Journal of Economics* **10**, 141-156.

Schmalensee, Richard (1978), "Entry Deterrence in the Ready-to-eat Breakfast Cereal Industry," *Bell Journal of Economics* **9**, 305–327.

Selten, Reinhard, and Rolf Stoecker (1986), "End Behavior in Sequences of Finite Prisoner's Dilemma Supergames," *Journal of Economic Behavior and Organization* **7**, 47-70.

Shaked, A., and J. Sutton (1990), "Multiproduct Firms and Market Structure," *Rand Journal of Economics* **21**, 45–62.

Van Huyck, John, Raymond Battalio, and Richard Beil (1991), "Strategic Uncertainty, Equilibrium Selection, and Coordination Failure in Average Opinion Games," *Quarterly Journal of Economics* **106**, 885-909.

Von Ungern-Sternberg, Thomas (1991), "Monopolistic Competition on the Pyramid," *Journal of Industrial Economics* **34**, 355–368.

Young, Peyton (1993), "The Evolution of Conventions," *Econometrica* **61**, 57-84.