GAMES AND ECONOMIC BEHAVIOR 5, 368–386 (1993)

Three Problems in Learning Mixed-Strategy Nash Equilibria

J. S. JORDAN

Department of Economics, University of Minnesota, Minneapolis, Minnesota 55455

Received August 11, 1992

This paper discusses three problems that can prevent the convergence of learning mechanisms to mixed-strategy Nash equilibria. First, while players' expectations may converge to a mixed equilibrium, the strategies played typically fail to converge. Second, even in 2×2 games, fictitious play can produce a sequence of frequency distributions in which the marginal frequencies converge to equilibrium mixed strategies but the joint frequencies violate independence. Third, in a three-player matching-pennies game with a unique equilibrium, it is shown that if players learn as Bayesian statisticians then the equilibrium is locally unstable. *Journal of Economic Literature* Classification Numbers: C72, C73, D83. © 1993 Academic Press, Inc.

1. INTRODUCTION

The current literature on learning in games is, in many respects, a natural successor to the earlier literature on learning rational expectations (e.g., Blume *et al.*, 1982). Both literatures address the question of whether decision-makers can, through repeated experience, learn to make optimal or equilibrium decisions. At the level of economic interpretation, the issues studied by the two literatures are as similar as all of the parallels between game theory and general equilibrium theory would suggest. At the level of formal analysis, however, important distinctions arise. Models of learning in games are typically much simpler because of the common assumption that players have only a fixed finite number of pure strategies, which virtually eliminates the measure theoretic difficulties commonly found in rational expectations models. Unfortunately, the cost of this simplifying assumption is the need to include mixed strategies in order to ensure the general existence of equilibrium. Moreover, mixed-strategy

0899-8256/93 $5.00
Copyright © 1993 by Academic Press, Inc.
All rights of reproduction in any form reserved.

equilibria present learning difficulties which are not presented by pure-strategy equilibria and are not found in rational expectations learning models. This paper is devoted to three such problems.

The first difficulty with mixed-strategy equilibria arises from the fact that mixed strategies occur at points where a player's optimal response correspondence is not lower semicontinuous. As a result, a player's expectations of the other players' strategies may converge to the other players' Nash equilibrium strategies but the player's optimal response to the converging expectations may be a unique pure strategy at every point of the expectations sequence. Thus the convergence of expectations to equilibrium need not imply the convergence of actual strategies to equilibrium. Given a typical payoff matrix, a mixed strategy is expected-payoff maximizing only when the expected mixed strategies of the other players lie in a subspace of positive codimension. Since the set of pure strategies is finite, the set of all possible *finite* histories of play is countable. Hence, for any given learning mechanism, it is unlikely that a player's expectations will ever lie in the subspace for which a mixed strategy is an optimal response. Section 2 of this paper constructs a continuum of 2×2 (two players, each with two pure strategies) games, with the properties that (a) each game has a mixed-strategy equilibrium as its unique Nash equilibrium and (b) any mechanism for forming expectations based on the history of play will lead exclusively to pure-strategy responses in all but a countable subset of games. Thus, while it is possible to construct learning mechanisms which support the general convergence of expectations to Nash equilibrium (Jordan, 1991a), the general convergence of expected-payoff maximizing strategies is not possible.

Section 3 of this paper exposes an additional reason why the convergence of expectations can be a less satisfying result if the limit is a mixed-strategy equilibrium than if the limit is a pure-strategy equilibrium. The definition of a mixed-strategy Nash equilibrium includes the requirement that players randomize independently of one another. However, the process of learning from repeated play can induce correlation among the actual strategy choices which may persist in the limit even if expectations converge to the Nash equilibrium marginal distributions. Section 3 borrows from Young (1993) an example of a 2×2 "battle of the sexes" game in which the players form expectations via fictitious play, that is, each player uses the frequency distribution of past plays to forecast the other player's future play. In this example, the two players exactly miscoordinate, so that each player receives the zero payoff in every period, even though expectations converge to the mixed-strategy equilibrium.

Nash equilibrium is analogous to a Walrasian equilibrium in the sense that each player is assumed to ignore the possibility that other players might change their strategies in response to a change in his or her own strategy. In a mixed-strategy equilibrium players regard the random plays of others as though they were exogenously generated by nature. This view of Nash equilibrium motivates the concept of "naive Bayesian learning" (e.g., Eichberger *et al.* 1990), of which fictitious play is the best-known example. Suppose that each player assumes that the plays of others are independent random draws from a fixed but unknown distribution. Each player has a prior probability distribution over the possible strategy distributions of the other players, and forecasts future plays according to the posterior expectations determined by the prior distribution and observed past plays. If the support of the prior distribution contains all mixed strategies, then the posterior expectations are asymptotically close to the empirical frequency distribution.

Shapley (1964) constructed a family of two-player 3×3 games for which he demonstrated that fictitious play, if started at certain strategies, would drive expectations to a limit cycle rather than the unique mixed-strategy Nash equilibrium. This proved that fictitious play is not in general a globally convergent learning process. One might still hope that naive Bayesian learning is at least locally stable, in the sense that if initial expectations are sufficiently near the unique equilibrium, then convergence to equilibrium is assured. This property is obviously satisfied at strict pure-strategy equilibria. However, Corollary 4.14 in Section 4 shows that even local stability can fail when the unique equilibrium involves mixed strategies. This result is based on an example of a three-person matching pennies game. In this example, the space of expectations is a three-dimensional cube, as opposed to the product of two two-dimensional simplices in Shapley's example. The lower dimension makes the learning dynamics easier to analyze and to visualize. For this example we obtain a stronger version of Shapley's nonconvergence result. In particular, even if initial beliefs are concentrated arbitrarily close to the unique Nash equilibrium, naive Bayesian learning can lead to a limit cycle. Thus a mixed-strategy equilibrium can lead to the same local instability of naive disequilibrium dynamics as is found in the tatonnement price-adjustment process in general equilibrium theory (e.g., Scarf, 1960).

Sections 2–4 are each formally self-contained and may be read independently of one another. Section 5 contains some concluding remarks on the implications of the preceding results for the theory of learning in games.

The learning processes studied below all assume best-response dynamics. That is, each player, at each iteration, chooses a strategy which maximizes the player's own expected payoff given the player's current expectations about the strategies of the other players. Best-response dynamics differ from gradient-like dynamics, in which strategies are partially adjusted at each iteration in a payoff increasing direction. A very general instability result for mixed-strategy equilibria under gradient-like dynamics has already been obtained by Crawford (1985). However, the source of instability discovered by Crawford does not appear to extend to best-response dynamics, and is thus quite different from the impediments to convergence described below.

2. Nonconvergence of Strategies

2.1. Definitions. Consider the class of 2×2 games defined by the payoff bimatrix

	L	R
T	$1-\alpha, \beta-1$	$-\alpha, 0$
B	$0, \beta$	$0, 0$

for $\alpha, \beta \in (0, 1)$. Let p_T denote the probability that the row player, player 1, plays T, and let p_L denote the probability that the column player, player 2, plays L. Let $S^1 = \{T, B\}$, $S^2 = \{L, R\}$, and $S = S^1 \times S^2$. For each integer $t \geq 1$, define the set of *t-period histories* by $H_t = \prod_{r=1}^{t} S$, and let H_0 denote the one-point set $\{*\}$. A generic element of H_t is denoted $h_t = (*, s_1, \ldots, s_t) = (*; s_{11}, s_{21}; \ldots; s_{1t}, s_{2t})$. A *learning process* is a pair (e^1, e^2), where each e^i is a sequence of functions e_t^i: $H_t \rightarrow [0, 1]$, $t = 0, 1, \ldots$. For each $h_t \in H_t$, $e^1(h_t)$ represents player 1's believed probability that player 2 will play L in period $t+1$, and $e^2(h_t)$ represents player 2's believed probability that player 1 will play T in period $t+1$.

Given $\alpha, \beta \in (0, 1)$ and a learning process, we say that *strategies are forever pure* if

(1) for every $t \geq 0$, $e_t^1(h_t) \neq \alpha$ and $e_t^2(h_t) \neq \beta$; where for each $t \geq 0$,

$$s_{t+1}^1 = \begin{cases} T & \text{if } e_t^1(h_t) > \alpha; \\ B & \text{if } e_t^1(h_t) < \alpha; \end{cases} \tag{2.1}$$

and

$$s_{t+1}^2 = \begin{cases} L & \text{if } e_t^2(h_t) < \beta; \\ R & \text{if } e_t^2(h_t) > \beta. \end{cases} \tag{2.2}$$

In other words, strategies are forever pure if, for every period t, each player i's optimal response to the expectation $e_t^i(h_t)$ is a unique pure strategy.

2.2. THEOREM. *For each $\alpha, \beta \in (0, 1)$, the unique Nash equilibrium is the mixed-strategy Nash equilibrium defined by $p_L = \alpha$ and $p_T = \beta$. However, given any learning process (e^1, e^2) there are countable sets C_1 and C_2 such that if $\alpha \in (0, 1)\backslash C_1$ and $\beta \in (0, 1)\backslash C_2$, then strategies are forever pure.*

Proof. The first assertion is direct. The second assertion is proved by defining $C_1 = \bigcup_{t=0}^{\infty} e^1(H_t)$ and $C_2 = \bigcup_{t=0}^{\infty} e^2(H_t)$. ∎

2.3. *Remarks.* Theorem 2.2 indicates that learning to *play* equilibrium mixed strategies is more problematic than learning mixed-strategy equilibrium *expectations*. Mixed strategies are controversial in game theory (e.g., Rubenstein 1991), and theorists who dislike mixed strategies as a model of player behavior would probably not be disappointed by this result. However, such prominent learning theorists as Kalai and Lehrer (1991) and Fudenberg and Kreps (1993) have asserted the importance of learning to play equilibrium strategies. Most of Kalai and Lehrer's work concerns repeated games with discounted payoffs in which players learn by forming Bayesian expectations based on prior beliefs over possible histories of play. Beliefs are assumed to satisfy a certain mutual consistency condition across players. As a corollary to their main result, they show that a Bayesian learning process based on prior beliefs over player types will ensure the convergence of actual strategies to the Nash equilibrium set with probability one, assuming that the support of the type distribution is countable (Kalai and Lehrer, 1991, Theorem 2.1). Theorem 2.2 shows that the countable support assumption is essential. Fudenberg and Kreps (1993, Section 8) show that for 2×2 games with a unique mixed-strategy equilibrium, the convergence of strategies played can be obtained by augmenting the game with random perturbations of the payoff characteristics of each player. This device purifies the mixed strategies over the realizations of payoff characteristics.

It may be worth mentioning that Crawford's instability result (Crawford, 1985) applies to games for which both Kalai and Lehrer (1991) and Fudenberg and Kreps (1993) obtain the convergence of both expectations and strategies. This contrast illustrates how differently learning and gradient-like dynamics can behave. In particular, the reader should be cautious in attempting to draw inferences about evolutionary models from the results of the present paper.

3. MARGINAL VERSUS JOINT FREQUENCIES

In 2×2 games, fictitious play is known to converge to Nash equilibrium (e.g., Rosenmüller 1971). That is, if each player uses the empirical frequency distribution of the opponent's plays to estimate the opponent's mixed strategy, then the pair of frequency distributions will converge to the set of Nash equilibrium strategy pairs. However, we borrow an example of Young (1993) to show that the empirical joint-frequency distribution over pairs of pure strategies need not converge to the Nash equilibrium set. Thus the joint-frequency distribution can be quite different, in the limit, from the product of the two marginal-frequency distributions.

Consider the "battle of the sexes" game represented by the payoff bimatrix:

	L	R
T	0, 0	1, $\sqrt{2}$
B	$\sqrt{2}$, 1	0, 0

For each i, let $\hat{p}_{it}$ denote the frequency with which player i's first pure strategy (T for player 1 and L for player 2) has occurred during the first t plays. Suppose that player i's expectation of the mixed strategy to be used by player j in period $t + 1$ is $\hat{p}_{jt}$. Then play in each period $t > 1$ is determined by

$$s_t^1 = \begin{cases} T & \text{as } \hat{p}_{2(t-1)} \lesseqgtr (1 + \sqrt{2})^{-1} \\ B & \end{cases}; \qquad (*)$$

and

$$s_t^2 = \begin{cases} L & \text{as } \hat{p}_{1(t-1)} \lesseqgtr (1 + \sqrt{2})^{-1} \\ R & \end{cases}.$$

Since empirical frequencies are rational numbers, only strict inequalities need be included on the right-hand side of (*). Play in period 1 is arbitrary, so let $s_1 = (s_1^1, s_1^2) = (T, L)$. Then we have $(\hat{p}_{11}, \hat{p}_{21}) = (1, 1)$, $s_2 = (B, R)$; $(\hat{p}_{12}, \hat{p}_{22}) = (\frac{1}{2}, \frac{1}{2})$, $s_3 = (B, R)$; etc. More generally, it is easily verified that every period t satisfies either

Case (1). $\hat{p}_{1t} = \hat{p}_{2t} < (1 + \sqrt{2})^{-1}$ and $s_t = (T, L)$; or

Case (2). $\hat{p}_{1t} = \hat{p}_{2t} > (1 + \sqrt{2})^{-1}$ and $s_t = (B, R)$.

In Case 1, $\hat{p}_{i(t+1)} = (1 + t\hat{p}_{it})/(t + 1) > \hat{p}_{it}$, and in Case 2, $\hat{p}_{i(t+1)} = t\hat{p}_{it}/$

$(t + 1) < \hat{p}_{it}$. Therefore $\hat{p}_{it} \to (1 + \sqrt{2})^{-1}$ for each i, so the frequency pair $(\hat{p}_{1t}, \hat{p}_{2t})$ converges to the mixed-strategy Nash equlibrium. The joint-frequency distribution, however, is concentrated on the two strategy pairs (T, L) and (B, R), and thus does not converge to the Nash equilibrium joint distribution. In fact, both players receive the zero payoff in every period, so it is clear that the limiting joint-frequency distribution also fails to be a correlated Nash equilibrium.

It should be mentioned that this example is not robust because it depends on the symmetry of the payoffs. A slight perturbation away from symmetry will induce convergence to one of the two pure-strategy equilibria.

As far as I am aware, Young (1993) was the first to note that fictitious play could lead to persistent miscoordination. A similar example is described by Fudenberg and Kreps (1993, Section 5). The pathological behavior of the joint-frequency distribution is possible in such examples because the players ignore the correlation between their strategies, which is caused by the learning process. In the Bayesian learning processes described by Jordan (1991a,b, 1992a), Kalai and Lehrer (1991), and Nyarko (1992a,b), expectations are defined as joint distributions, so the convergence of expectations to Nash equilibrium entails the convergence of expected probabilities over joint strategies. In fact, Nyarko (1992a) and Jordan (1992a) have shown that the joint frequencies converge to the Nash equilibrium set under Bayesian learning. In this sense, Bayesian learning is a device for purifying a mixed-strategy equilibrium over time.

4. Nonconvergence of Naive Bayesian Learning

Shapley (1964, pp. 24–27) exhibited a family of two-player 3×3 games with the property that expectations formed via fictitious play can cycle rather than converge to equilibrium. The purpose of this section is to exhibit a simple three-player game with the same property. More importantly, it will be shown that expectations paths which begin outside a one-dimensional stable manifold will converge to the limit cycle, and that this property extends to other methods of expectations formation, naive Bayesian learning in particular, which behave like fictitious play asymptotically.

The game itself is a three-person version of matching pennies. Player 1 seeks to match player 2, player 2 seeks to match player 3, and player 3 seeks *not* to match player 1. In particular, each player i is concerned only with predicting the actions of player $i + 1$ (mod 3). It is easily seen that the unique Nash equilibrium requires each player to mix equally between *heads* and *tails*.

4.1. Definitions. There are three players and each player has two pure strategies, which we will denote *heads* and *tails*. Thus for each i, $S^i = \{heads, tails\}$, $1 \le i \le 3$. Player i's payoff function, $\pi^i\colon S \to R$, is defined for each i as

$$\pi^1(s^1, s^2, s^3) = \begin{cases} 1 & \text{if } s^1 = s^2; \\ -1 & \text{if } s^1 \neq s^2; \end{cases}$$

$$\pi^2(s^1, s^2, s^3) = \begin{cases} 1 & \text{if } s^2 = s^3; \\ -1 & \text{if } s^2 \neq s^3; \end{cases}$$

and

$$\pi^3(s^1, s^2, s^3) = \begin{cases} 1 & \text{if } s^3 \neq s^1; \\ -1 & \text{if } s^3 = s^1. \end{cases}$$

For each player i, a mixed strategy is represented by a number $p_i \in [0, 1]$ which represents the probability that $s^i = heads$.

4.2. PROPOSITION. *The above game has a unique Nash equilibrium, which is given by* $p_i^* = \frac{1}{2}$ *for each* $1 \le i \le 3$.

Proof. Direct. ■

4.3. *Stable Manifold and Limit Cycle.* Let Q denote the closed unit cube $[0, 1]^3$ in R^3, with generic element $p = (p_1, p_2, p_3)$. Let M denote the closed line segment between (0, 1, 0) and (1, 0, 1). Let g denote the "golden ratio," $g = (1 + \sqrt{5})/2$ (g is characterized as the positive solution to the equation $x^{-1} = x - 1$). Let

$$x^\alpha = (2g)^{-1} = (\sqrt{5} - 1)/4$$
$$x^\beta = 1 - (g/2) = (3 - \sqrt{5})/4$$
$$x^\gamma = 1 - (2g)^{-1} = (5 - \sqrt{5})/4$$
$$x^\delta = g/2 = (1 + \sqrt{5})/4.$$

Let

$$p^a = (x^\alpha, x^\beta, \tfrac{1}{2})$$
$$p^b = (x^\beta, \tfrac{1}{2}, x^\gamma)$$

$$p^c = (\tfrac{1}{2}, x^\gamma, x^\delta)$$
$$p^d = (x^\gamma, x^\delta, \tfrac{1}{2})$$
$$p^e = (x^\delta, \tfrac{1}{2}, x^\alpha)$$
$$p^f = (\tfrac{1}{2}, x^\alpha, x^\beta).$$

Define the set $C \subset Q$ by

$$C = p^a \to p^b \to p^c \to p^d \to p^e \to p^f \to p^a,$$

where the arrows represent the closed line segments between the respective points.

4.4. *Remarks.* Suppose that each player i uses the observed frequency

$$\hat{p}_{(i+1)t} = {}^{\#}\{\tau \leq t: s_\tau^{i+1} = heads\}/t$$

to estimate the probability that player $i + 1$ (mod 3) will play *heads* at iteration $t + 1$. The resulting path of expectations $(\hat{p}_{1t}, \hat{p}_{2t}, \hat{p}_{3t})$ in the unit cube Q is fairly easy to visualize. First, note that the three planes $\hat{p}_i = \frac{1}{2}$ partition Q into eight octants, in the interior of which each player has a unique optimal pure strategy. The optimal responses to the expectations $(\hat{p}_i)_i$ in the eight octants are given by

$\hat{p}_3 < \frac{1}{2}$

	$\hat{p}_2 < \frac{1}{2}$	$\hat{p}_2 > \frac{1}{2}$
$\hat{p}_1 < \frac{1}{2}$	T, T, H	H, T, H
$\hat{p}_1 > \frac{1}{2}$	T, T, T	H, T, T

$\hat{p}_3 > \frac{1}{2}$

	$\hat{p}_2 < \frac{1}{2}$	$\hat{p}_2 > \frac{1}{2}$
$\hat{p}_1 < \frac{1}{2}$	T, H, H	H, H, H
$\hat{p}_1 > \frac{1}{2}$	T, H, T	H, H, T

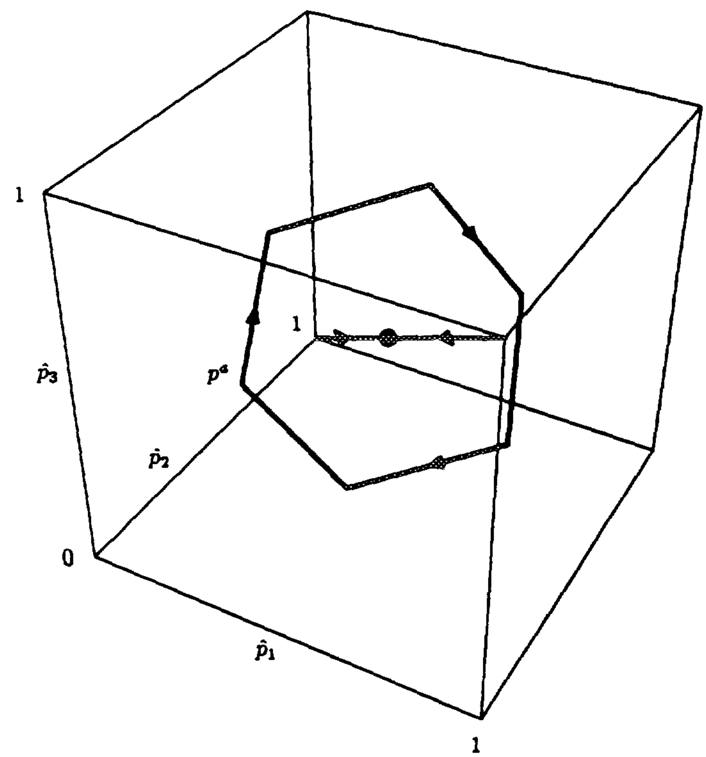

FIGURE 1

where T denotes *tails* and H denotes *heads.*

If $\hat{p}_t$ lies in one such octant, then $\hat{p}$ will proceed from $\hat{p}_t$ in the direction of the corner of Q which represents the optimal pure-strategy triple. For example, if $\hat{p}_t$ satisfies $\hat{p}_{it} < \frac{1}{2}$ for each i, then $\hat{p}$ proceeds along the line segment from $\hat{p}_t$ to the corner (0, 0, 1) until the first iteration $t' > t$ with $\hat{p}_{3t'} > \frac{1}{2}$ (neglecting, for the moment, the case $\hat{p}_{3t'} = \frac{1}{2}$), and thence from $\hat{p}_{t'}$ toward (0, 1, 1); and so on. The resulting path, which is shown in Fig. 1,[1] is a sequence of line segments with kinks near the planes $\hat{p}_i = \frac{1}{2}$, unless $\hat{p}_t \in M$. In this case the relevant corners are the two endpoints of M, so $\hat{p}$ stays within M, converging in an oscillating fashion to the equilibrium $(\frac{1}{2}, \frac{1}{2}, \frac{1}{2})$. Of course, if $\hat{p}_{t'} = (\frac{1}{2}, \frac{1}{2}, \frac{1}{2})$ for some t', the strategies played at iteration $t' + 1$ are not uniquely defined, so $\hat{p}_{t'+1}$ may lie outside M (this

[1] I am indebted to Julio Escalano for this drawing.

can be avoided by imposing the choices (H, T, H) or (T, H, T) at $(\frac{1}{2}, \frac{1}{2}, \frac{1}{2})$).

From any $\hat{p}_t \notin M$, the expectations path spirals toward the limit cycle C. This fact is a consequence of Theorem 4.11 below, and depends on t being large enough so that the effect of fortuitous strategy choices at the $\hat{p}_i = \frac{1}{2}$ planes can be ignored. The set C, which lies in the six octants of Q which do not intersect M, is derived as follows: $\hat{p}^b$ is the point at which the line segment from $\hat{p}^a$ to (0, 1, 1) intersects the $\hat{p}_2 = \frac{1}{2}$ plane, $\hat{p}^c$ is the point at which the line segment from $\hat{p}^b$ to (1, 1, 1) intersects the $\hat{p}_1 = \frac{1}{2}$ plane, etc.

Figure 1 depicts the sets M and C. The black ball in the center of the cube indicates the unique Nash equilibrium, the diagonal line segment from (0, 1, 0) to (1, 0, 1) depicts the stable manifold M, and the six hexagonal line segments represent the limit cycle C. The corner point p^a is indicated in Fig. 1, and the points p^b, p^c, etc., lie at the clockwise successive corners of C.

4.5. Definition. An *expectations process* is a triple of functions e^i: $[0, 1] \times \mathbf{Z}_{++} \to [0, 1]$, $i = 1, 2, 3$, where $\mathbf{Z}_{++}$ denotes the strictly positive integers.

4.6. *Assumption.* For each i and each $\varepsilon > 0$ there exists $t^0 > 0$ such that $|e^i(\hat{p}_{i+1}, t) - \hat{p}_{i+1}| < \varepsilon$ for all $\hat{p}_{i+1} \in [0, 1]$ and all $t > t^0$.

4.7. *Remarks.* The number $e^i(\hat{p}_{i+1}, t)$ is interpreted as player i's expectation of the probability that player $i + 1$ will play *heads* in period $t + 1$ as a function of the empirical frequency $\hat{p}_{i+1}$ of *heads* during the previous t periods. Naive Bayesian expectations conform to this definition because the posterior expectation of the probability of *heads* depends on an observed history $(s_1^{i+1}, \ldots, s_t^{i+1})$ only through the frequency of *heads* and the number, t, of observations. A formal definition of naive Bayesian expectations is given below. Any expectations process satisfying Assumption 4.6 is *asymptotically empirical*, as this term is defined by Fudenberg and Kreps (1903, Section 4) but the converse is not generally true because the latter property does not require expectations to depend solely on the frequency of past plays, or to be uniformly close to the frequency for large t. Proposition 4.9 below shows that naive Bayesian learning satisfies Assumption 4.6.

4.8. *Naive Bayesian Learning.* Let $\Delta([0, 1])$ denote the set of Borel probability measures on [0, 1]. Define the function β: $\Delta([0, 1]) \times [0, 1] \times \mathbf{Z}_{++} \to [0, 1]$ by

$$\beta(\mu, x, t) = \left(\int_0^1 q[q^{tx}(1-q)^{t(1-x)}]\mu(dq)\right) \bigg/ \left(\int_0^1 q^{tx}(1-q)^{t(1-x)}\mu(dq)\right).$$

For each i, let $\mu^i \in \Delta([0, 1])$. Then the expectations process $(\beta(\mu^i, \cdot))_{i=1}^3$ is a *naive Bayesian* expectations process.

The number $\beta(\mu, x, t)$ defined above is simply the expected value of q conditional on observing xt heads in t observations, provided that xt is an integer. The combinatorial coefficient $t!/xt!(t - xt)!$ cancels from the numerator and denominator in the above expression. Nonintegral values of xt would never by observed, but including them simplifies the definition.

4.9. PROPOSITION. Let $\mu \in \Delta([0, 1])$ and suppose that supp $\mu = [0, 1]$. Then for each $\varepsilon > 0$ there exists $t^0 > 0$ *such that* $|\beta(\mu, x, t) - x| < \varepsilon$ *for all* $t > t^0$ *and all* $x \in [0, 1]$. In particular if supp $\mu^i = [0, 1]$ for each $i = 1, 2, 3$, then the naive Bayesian expectations process $(\beta(\mu^i, \cdot))_{i=1}^3$ satisfies Assumption 4.6.

Proof. Let $\varepsilon > 0$. Since supp $\mu = [0, 1]$ and $[0, 1]$ is compact,

$$\inf\{\mu([x - \varepsilon/2, x + \varepsilon/2]) : x \in [0, 1]\} > 0. \tag{$*$}$$

Also, it is straightforward to show that for any $k > 0$ there is some $t^0 > 0$ such that

$$x^{tx}(1 - x)^{t(1-x)}/q^{tx}(1 - q)^{t(1-x)} > k, \tag{$**$}$$

for all, $x, q \in [0, 1]$ with $|q - x| \geq \varepsilon/2$ and all $t > t^0$. The desired conclusion follows from ($*$) and ($**$) and the definition of β. ∎

4.10. *Dynamics.* For each i, define the function f^i: $[0, 1] \times S^i \times \mathbf{Z}_+ \to [0, 1]$ by

$$f^i(\hat{p}_i, s_i, t) = \begin{cases} (t\hat{p}_i + 1)/(t + 1) & \text{if } s_i = \textit{heads}; \\ t\hat{p}_i/(t + 1) & \text{otherwise}. \end{cases}$$

Let $Q = [0, 1]^3$ and define f: $Q \times S \times \mathbf{Z}_+ \to Q$ by $f(\hat{p}, s, t) = (f^i(\hat{p}^i, s_i, t))_{i=1}^3$. In what follows, for any $\hat{p} \in Q$ and any t, i, we conserve notation by writing $e^i(q, t)$ to mean $e^i(q_{i+1}, t)$. For each $(\hat{p}, t^0) \in Q \times \mathbf{Z}_+$, a *frequency path from* $(\hat{p}, t^0)$ is defined to be a sequence $\{\hat{p}_t\}_{t=t^0}^\infty$ satisfying $\hat{p}_{t^0} = \hat{p}$ and, for each $t > t^0$,

$$\hat{p}_t = f(\hat{p}_{t-1}, s_t, t - 1),$$

where, for each i, s_t^i is an optimal response to the expectation $e^i(\hat{p}_{t-1}, t - 1)$, that is, for $i = 1, 2$,

$$s_t^i = heads \text{ if } e^i(\hat{p}_{t-1}, t-1) > \tfrac{1}{2}$$

and

$$s_t^i = tails \text{ if } e^i(\hat{p}_{t-1}, t-1) < \tfrac{1}{2},$$

with the inequalities reversed for $i = 3$. For each $t > t^0$ and each i, let $p_{(i+1)t} = e^i(\hat{p}_{t-1}, t-1)$. The sequence $\{p_t\}_{t=t^0+1}^{\infty} \subset Q$ is defined to be an *expectations path from* $(\hat{p}, t^0)$.

4.11. THEOREM. *Suppose that the expectations process satisfies Assumption 4.6, and let $\hat{p} \in Q \backslash M$. Then there exists $t > 0$ such that for every $t^0 > t$, the set of cluster points of every expectations path from $(\hat{p}, t^0)$ is equal to C.*

4.12. *Remarks.* The proof of Theorem 4.11, which is quite tedious, is given in the appendix. The proof consists largely of showing that C is a limit cycle for fictitious play, and that the convergence of fictitious play to C is sufficiently robust that C is also a limit cycle for any process that behaves like fictitious play asymptotically in the sense of Assumption 4.6. The proof also suggests that the theorem is qualitatively robust to small perturbations of the payoff matrices, although we will not state or prove this claim formally. Of course, the equilibrium and the sets M and C will be perturbed as well. If player i's payoff becomes slightly sensitive to s^{i-1}, then player i's plane of strategy indifference is a slight perturbation of the $p_{i+1} = \frac{1}{2}$ plane within the cube Q. Assumption 4.6 and Proposition 4.9 extend directly to bivariate expectations. The appendix also contains a more detailed discussion of the behavior of fictitious play itself.

The requirement that t be sufficiently large is needed for two reasons. First, since the adjustment process is discrete and the strategy s^i is not uniquely defined if player i's expectation $p_{i+1} = \frac{1}{2}$, a fortuitous choice of s^i might place next period's expectations in M if t is small enough for this to be accomplished in one step. This could not occur in a continuous time version of the expectations adjustment process. Second, Assumption 4.6 only restricts the expectations functions f^i asymptotically in t. Thus expectations can be forced artificially to the stable manifold, or to the equilibrium itself, for any given value of t without violating Assumption 4.6.

For a naive Bayesian expectations process, it is more natural to initialize expectations using the prior distributions $(\mu^i)_{i=1}^3$, so Definition 4.13 below defines a naive Bayesian expectations path accordingly. Proposition 4.9 and Theorem 4.11 imply that the Nash equilibrium is not locally stable under naive Bayesian expectations. More precisely, even if the initial beliefs $(\mu^i)_{i=1}^3$ are concentrated arbitrarily near the Nash equilibrium ($\frac{1}{2}$, $\frac{1}{2}$,

$\frac{1}{2}$), the expectations path may converge to the limit cycle C. This follows from Proposition 4.9 and Theorem 4.11 because beliefs concentrated near the equilibrium can be obtained from arbitrary prior distributions $(\mu_0^i)_{i=1}^3$, satisfying supp $\mu_0^i = [0, 1]$ for each i, by updating in response to a t^0-period history in which the associated frequency vector $\hat{p}$ is near the equlibrium but outside M. This local instability result is stated formally as Corollary 4.14.

4.13. DEFINITION. For each i, let $\mu^i \in \Delta([0, 1])$, a *naive Bayesian expectations path* from $(\mu^i)_{i=1}^3$ is defined as a sequence $\{p_t\}_{t=1}^\infty \subset Q$ constructed as follows. For each i, let $p_{(i+1)1} = \int_0^1 q\mu^i(dq)$, and let

$$\hat{p}_{i1} = \begin{cases} 1 & \text{if } s_1^i = \textit{heads}; \\ 0 & \text{if } s_1^i = \textit{tails}, \end{cases}$$

where s_1^i is an optimal response to $p_{(i+1)1}$. Now let $\{p_t\}_{t=2}^\infty$ be an expectations path from $(\hat{p}_1, 1)$ for the naive Bayesian expectations process $(\beta(\mu^i, \cdot))_{i=1}^3$.

4.14. COROLLARY. For any $\varepsilon > 0$ there exists $\mu^i \in \Delta([0, 1])$, $i = 1, 2, 3$, satisfying

$$\mu^i(\{q: |(\tfrac{1}{2}) - q| < \varepsilon) > 1 - \varepsilon, \qquad (*)$$

such that the set of cluster points of every naive Bayesian expectations path from $(\mu^i)_{i=1}^3$ is equal to C.

Proof. For each i, let $\mu_0^i \in \Delta([0, 1])$ with supp $\mu_0^i = [0, 1]$. Then by Proposition 4.9, the naive Bayesian expectations process $(\beta(\mu_0^i, \cdot))_{i=1}^3$ satisfies Assumption 4.6. Let $\varepsilon > 0$ and let $\hat{p} \in Q\backslash M$ such that for each i, $\hat{p}_i$ is a rational number and $|(\frac{1}{2}) - \hat{p}_i| < \varepsilon$. Let t^0 be given by Theorem 4.11 and, choosing t^0 larger if necessary, let $\{s_t\}_{t=1}^{t^0}$ be a t^0-period history such that for each i, $^\#\{t \le t^0: s_t^i = \textit{heads}\}/t^0 = \hat{p}_i$. For each i, let $\mu^i \in \Delta([0, 1])$ be the "posterior" distribution on [0, 1] determined by "prior" distribution μ_0^i and the "observations" $\{s_t^{i+1}\}_{t=1}^{t^0}$. By the same argument used to prove Proposition 4.9, t^0 can be chosen large enough so that μ^i satisfies (*) for each i. The corollary now follows from the fact that the naive Bayesian expectations path from $(\mu^i)_{i=1}^3$ is identical to the expectations path from $(\hat{p}, t^0)$ for the expectations process $(\beta^i(\mu_0^i, \cdot))_{i=1}^3$. ■

4.15. *Remarks.* Proposition 6.4 of Fudenberg and Kreps (1993, Section 6) shows that *asymptotically empirical* expectations, together with the players' behavior strategies, can converge to the equilibrium of the above matching-pennies game, provided that either: (1) the players' expectations can be fixed at the equilibrium as long as the equilibrium expectations are, in a certain sense, approximately empirical or (2) the players' strategies can be fixed at the equilibrium strategies as long as the equilibrium strategies are, in a certain sense, approximately best responses to expectations. In both cases the adjustment process is tailored to the particular equilibrium. Theorem 4.11 and Corollary 4.14 show that some such tailoring is essential to obtain even the convergence of expectations alone. Corollary 4.14 indicates that naive Bayesian learning can drive expectations to the limit cycle even when initial beliefs are arbitrarily close to the equilibrium expectations.

4.16. THE SHAPLEY EXAMPLE. The 3×3 game constructed by Shapley (1964) can be represented by the following payoff bimatrix:

	L	C	R
T	1, 0	0, 0	0, 1
M	0, 1	1, 0	0, 0
B	0, 0	0, 1	1, 0

Shapley demonstrated that if the first play is (T, L) then expectations formed via fictitious play will approach a hexagonal limit cycle C_S as the sequence of plays follows the cycle

$$(T, L) \to (T, R) \to (B, R) \to (B, C) \to (M, C) \to (M, L) \to (T, L).$$

Moreover, the number of periods spent at each successive play pair increases exponentially. Shapley also demonstrated that these properties are not dependent on the particular payoff numbers, but are satisfied as long as the payoffs satisfy certain inequalities.

I am not aware of any formal analysis of the fictitious play dynamics from arbitrary initial conditions in this example, but some numerical analysis suggests that the Shapley example behaves much the same as the matching-pennies game studied above. The space of expectations is $\Delta^2 \times \Delta^2$, where Δ^2 is the unit simplex in R^3. Within this space there is a two-dimensional subset $M_S = \{(p_1, p_2) \in \Delta^2 \times \Delta^2 \colon p_{2L} = p_{1B}, p_{2C} = p_{1T}$, and $p_{2R} = p_{1M}\}$ such that if initial expectations lie in M_S, fictitious play converges to the unique Nash equilibrium $p_1^* = (\frac{1}{3}, \frac{1}{3}, \frac{1}{3})$, $p_2^* = (\frac{1}{3}, \frac{1}{3}, \frac{1}{3})$. Convergence takes the form of an inward spiral in M_S as the sequence of plays follows the cycle

$$(T, C) \rightarrow (M, R) \rightarrow (B, L) \rightarrow (T, C).$$

If initial expectations lie outside of M_S, fictitious play appears to diverge to the limit cycle C_S. This behavior seems sufficiently robust that a result very similar to Theorem 4.11 may hold for the Shapley example as well, but I have not attempted to prove this.

5. Conclusion

The problems described above represent obstacles to learning which must be addressed by any general theory of learning in games. To date, the only learning mechanisms for which expectations have been shown to be asymptotically Nash for all normal form games with a finite number of players and a finite number of pure strategies are the Bayesian learning processes described by Jordan (1991a, Corollary 3.10). The Bayesian learning processes have two notable shortcomings. First, convergence occurs at the level of expectations but not necessarily at the level of actual strategies. Second, the sophistication and implicit coordination required of players in order to form the specified expectations seems excessive. Theorem 2.2 shows that the first limitation is inevitable unless one relaxes either the requirement that expectations are functions of the history of play or the requirement that plays are chosen as optimal responses to expectations. With respect to the second shortcoming, the example in Section 3 and Theorem 4.11 indicate that any generally convergent learning process is likely to require, at least implicitly, that the players possess more awareness of their strategic interactions than one might wish to require in a theory of Nash disequilibrium dynamics.

Appendix

Theorem 4.11. *Suppose that the expectations process satisfies Assumption 4.6, and let $\hat{p} \in Q \backslash M$. Then there exists $t > 0$ such that for every $t^0 > t$, the set of cluster points of every expectations path from $(\hat{p}, t^0)$ is equal to C.*

Proof. Let $Q_0 = \{p \in Q: p_1 \leq \frac{1}{2}, p_2 \geq \frac{1}{2}, p_3 \leq \frac{1}{2}\}$ and $Q_{00} = \{p \in Q: p_1 \geq \frac{1}{2}, p_2 \leq \frac{1}{2}, p_3 \geq \frac{1}{2}\}$ ($Q_0 \cup Q_{00}$ contains the stable manifold M). Let $\hat{p} = (\hat{p}_1, \hat{p}_2, \hat{p}_3)$, and suppose that $\hat{p} \notin Q_0 \cup Q_{00}$. Without loss of generality, assume $\hat{p}_1 < \frac{1}{2}$, $\hat{p}_2 \leq \frac{1}{2}$, and $\hat{p}_3 \leq \frac{1}{2}$. Let $\varepsilon < \frac{1}{2} - \hat{p}_1$, and choose ε^0 and ε' with $0 < \varepsilon^0 < \varepsilon' < \varepsilon/3$ so that for any $0 < \lambda < 1$,

(1) if $\lambda + (1 - \lambda)(\frac{1}{2} - \varepsilon) \geq (\frac{1}{2}) - (\frac{2}{3})\varepsilon$ then $\lambda + (1 - \lambda)(\frac{1}{2} - \varepsilon^0) > \frac{1}{2} + \varepsilon'$.

Let t^* satisfy

(2) for all i and all $q \in Q$, $|e^i(q, t) - q_{i+1}| < \varepsilon^0$ for all $t \geq t^*$, and

(3) $t^* > 1/\varepsilon^0$.

Suppose that $t^0 > t^*$, and let $\{\hat{p}_t\}_{t=t^0+1}^{\infty}$ be a frequency path from $(\hat{p}, t^0)$. By (2), s_t^3 = *heads* for every t such that $\hat{p}_{1t} < \frac{1}{2} - \varepsilon^0$. Hence, for all such t, $\hat{p}_{3(t+1)} > \hat{p}_{3t}$. Similarly, if $\hat{p}_{2t} < \frac{1}{2} - \varepsilon^0$, s_{t+1}^1 = *tails*, and if $\hat{p}_{3t} < \frac{1}{2} - \varepsilon^0$, s_{t+1}^2 = *tails*. Thus, if $\hat{p}_2 < \frac{1}{2} - \varepsilon^0$, the sequence $\{\hat{p}_t\}_{t=t^0+1}^{\infty}$ moves toward the corner (0, 0, 1) until $\hat{p}_{3t} \geq \frac{1}{2} - \varepsilon^0$. Alternatively, suppose $\hat{p}_2 \geq \frac{1}{2} - \varepsilon^0$. If $\hat{p}_3 < \frac{1}{2} - \varepsilon^0$, then $\hat{p}_{2t}$, as before, will decrease at least until the first time t' that $\hat{p}_{3t'} \geq \frac{1}{2} - \varepsilon^0$, at which time $\hat{p}_{1t'} < \hat{p}_1 = (\frac{1}{2}) - \varepsilon$. Hence, there are two possible cases, differing only in $\hat{p}_{2t}$:

Case (1). $\hat{p}_{1t} \leq \frac{1}{2} - \varepsilon$, $\hat{p}_{2t} < \frac{1}{2} - \varepsilon^0$, $(\frac{1}{2}) - \varepsilon^0 \leq \hat{p}_{3t} \leq \frac{1}{2}$; and

Case (2). $\hat{p}_{1t} \leq \frac{1}{2} - \varepsilon$, $\frac{1}{2} - \varepsilon^0 \leq \hat{p}_{2t} \leq \frac{1}{2}$, $(\frac{1}{2}) - \varepsilon^0 \leq \hat{p}_{3t} \leq \frac{1}{2}$.

In Case 2, s_{t+1}^3 = *heads*, but s_{t+1}^1 and s_{t+1}^2 are ambiguous. However, by (1) and (2), there will be some $t' > t$ with $\hat{p}_{1t'} < \frac{1}{2} - (\frac{3}{4})\varepsilon$ and $\hat{p}_{3t'} > \frac{1}{2} + \varepsilon'$. For $t > t'$, we have s_{t+1}^2 = *heads* and s_{t+1}^3 = *heads* until some t with $\hat{p}_{1t} \geq \frac{1}{2} - \varepsilon^0$, at which point, again by (1) and (3), we have

Case (3). $\hat{p}_{1t} \geq \frac{1}{2} - \varepsilon^0$, $\hat{p}_{2t} > \frac{1}{2} + \varepsilon'$, $\hat{p}_{3t} > \frac{1}{2} + \varepsilon'$.

By a similar argument, Case 1 also leads eventually to Case 3.

Now suppose $\hat{p} \in Q_0 \cup Q_{00}$, say $\hat{p} \in Q_0$. Let $\overline{p}$ denote the point at which the line segment from $\hat{p}$ to the corner (1, 0, 1) intersects the boundary of Q_0. Since $\hat{p} \notin M$, there is some i with $\overline{p}_i \neq \frac{1}{2}$. Without loss of generality suppose that $\overline{p}_1 < \frac{1}{2}$. Let $\varepsilon = (\frac{1}{2} - \overline{p}_1)$ and let δ and t^* satisfy (1–3) above. Then the frequency path $\{\hat{p}_t\}_{t=t^0+1}^{\infty}$ proceeds along the line segment from $\hat{p}$ to (1, 0, 1) until some period t' in which $|\hat{p}_{it'} - \frac{1}{2}| < \delta$ for $i = 2$, or $i = 3$. From this point, virtually the same reasoning as in the previous paragraph leads to Case 3.

Proceeding from Case 3, we have s_{t+1}^1 = *heads*, s_{t+1}^2 = *heads*, and s_{t+1}^3 is ambiguous. Let t^c denote the first period $t' > t$ with $\hat{p}_{1t'} > \frac{1}{2} + \varepsilon^0$. From $\hat{p}_{t^c}$, the frequency path will proceed along the line segment from $\hat{p}_{t^c}$ to (1, 1, 0) until $\hat{p}_3 \leq \frac{1}{2} + \varepsilon^0$. Let p'^c denote the point at which this line segment, if extended back from $\hat{p}_{t^c}$, would intersect the plane $\hat{p}_1 = \frac{1}{2}$, and let p'^d denote the point at which the line segment intersects the plane $\hat{p}_3 = \frac{1}{2}$. Recall the points $p^c = (\frac{1}{2}, x^\gamma, x^\delta)$ at $p^d = (x^\gamma, x^\delta, \frac{1}{2})$ defined in 4.3 above, and note that the line segment from p^c to (1, 1, 0) intersects the plane $\hat{p}_3 = \frac{1}{2}$ at p^d. Then a geometric argument shows that

(4) $|p_2^d - p_2'^d| < |p_2^c - p_2'^c|/2x^\delta + [(\frac{1}{2} - 2\delta)/(\frac{1}{2} + 2\delta)]|p_3^c - p_3'^c|/2x^\delta$; and

(5) $|p_1^d - p_1'^d| < |p_3^c - p_3'^c|/2x^\delta$.

Let $\eta = (1 - \varepsilon')/(1 + \varepsilon')$. Then, in matrix notation, the right-hand side of (4) and (5) can be written as

$$\begin{bmatrix} \eta(2x^\delta)^{-1} & (2x^\delta)^{-1} \\ (2x^\delta)^{-1} & 0 \end{bmatrix} \begin{bmatrix} |p_2^c - p_2'^c| \\ |p_3^c - p_3'^c| \end{bmatrix}.$$

The 2×2 matrix has the characteristic polynomial $r(r - \eta/2x^\delta) - (2x^\delta)^{-2} = 0$. Since $\eta < 1$ and $(2x^\delta)^{-2} = 1 - (2x^\delta)^{-1}$ ($2x^\delta$ is the golden ratio), both characteristic roots have absolute value less than unity. Let r_1 be the positive root, which also has the larger absolute value. Then (4) and (5) imply that $\|p^d - p'^d\| < r_1\|p^c - p'^c\|$; then since $|p_3^d - p_3'^d| = |p_1^c - p_1'^c| = 0$, (4) and (5) imply

(6) $\|p^d - p'^d\| < r_1\|p^c - p'^c\|$,

where $\|\cdot\|$ denotes the Euclidean norm. Let t^d be the first period $t' > t^c$ with $\hat{p}^3_{t'} < \frac{1}{2} - \varepsilon^0$; then $\|p^d - \hat{p}_{t^d}\| < 6\varepsilon^0$. Similarly, $\|p^c - \hat{p}_{t^c}\| < 6\varepsilon^0$, so

(7) $\|p^d - \hat{p}_{t^d}\| < r_1(\|p^c - \hat{p}_{t^c}\| + 6\varepsilon^0) + 6\varepsilon^0.$

From $\hat{p}_{t^d}$, the frequency path proceeds along the line segment from $\hat{p}_{t^d}$ to $(1, 0, 0)$ until some period t' with $\hat{p}_{2t'} \leq \frac{1}{2} + \varepsilon^0$. Then, using $p^e = (x^\delta, \frac{1}{2}, x^\alpha)$, the same argument can be used to show that $\|p^e - \hat{p}_{t^e}\| < r_1(\|p^d - \hat{p}_{t^d}\| + 6\varepsilon^0) + 6\varepsilon^0$, where $\hat{p}_{t^e}$ is defined analogously to $\hat{p}_{t^d}$ and $\hat{p}_{t^c}$. Continuing in this fashion leads, via (7) and its analogues for the other five turning points of C, to the implication that

(8) $\limsup_{t\to\infty} (\min\{|\hat{p}_t - \hat{p}| : \hat{p} \in C\}) \leq 6\varepsilon^0((1 + r_1)/(1 - r_1)).$

As $t \to \infty$, ε^0 can be made arbitrarily small, independently of r_1, so the right-hand side of (8) can be replaced by zero. Along the frequency path the step size is on the order of $1/t$, and thus approaches zero as $t \to \infty$. Hence, for each $\hat{p} \in C$, lin $\inf_{t\to\infty} |\hat{p}_t - \hat{p}| = 0$. Let $\{p_t\}_{t=t^0+1}^{\infty}$ be the expectations path corresponding to the frequency path $\{\hat{p}_t\}_{t=t^0}^{\infty}$. By Assumption 4.6, $\lim_{t\to\infty} |p_{(i+1)t} - e^i(\hat{p}_t, t)| = 0$, so the proof is complete. ■

The asymptotic behavior of fictitious play in the three-person matching-pennies game can be described more precisely than is done in Theorem 4.11. This appendix mentions some further results which were obtained numerically via a program available from the author on request. The program selects initial expectations $p^0 \in Q$ and an initial time t^0 pseudo-randomly, and then computes the successive "turning points" of the path from (p^0, t^0) according to fictitious play. A turning point occurs when some player i's pure strategy switches from *heads* to *tails*, or vice versa, that is, when p_{i+1} crosses the $p_{i+1} = \frac{1}{2}$ plane. The expectations path from (p^0, t^0) is a sequence of points within the line segments connecting successive turning points, including the turning points themselves. The line segments connecting the successive turning points form a hexagonal spiral converging to the limit cycle C. The program terminates when it reaches a turning point which lies within a prespecified distance ε of the corresponding turning point in C. The program also computes a "continuous-time" version of fictitious play, which puts the turning points exactly on the respective $p_i = \frac{1}{2}$ planes.

The golden ratio, $g = (1 + \sqrt{5})/2 \approx 1.618$, plays a large role in the asymptotic behavior of fictitious play. First, the number of periods spent between successive turning points increases at the rate g. That is, given successive turning points p_{t_n}, $p_{t_{n+1}}$, and $p_{t_{n+2}}$, we have $(t_{n+2} - t_{n+1})/(t_{n+1} - t_n) \to g$ as $n \to \infty$. Second, g is also involved in the rate at which the path approaches C. Let p_{t_n} be a turning point on the expectations path, and let p^C_n be the corresponding turning point in C. Suppose that $p^C_{i_n} = \frac{1}{2} \approx p_{it_n}$, and let $x_n = |p^C_{(i+1)n} - p_{(i+1)t_n}|$ and $y_n = |p^C_{(i+2)n} - p_{(i+2)t_n}|$. Then, along a "typical" path, as $n \to \infty$, $(x_{n+1} + y_{n+1})/(x_n + y_n) \to g^{-1}$, $y_n/x_n \to g$, $y_{n+1}/x_n \to 1$, and $x_{n+1}/x_n \to g^{-1}$. An exception to this behavior occurs along paths starting from initial conditions such as $t^0 = 3$ and p^0 $(\frac{1}{3}, \frac{1}{3}, \frac{1}{3})$, which have the property that $x_n = y_n$ for all n and $x_{n+1}/x_n \to g^{-2}$. This case depends on the equality $x_n = y_n$ for all n, and is therefore disrupted by rounding errors in floating point arithmetic.

References

Blume, L., Bray, M., and Easley, D. (1982). "Introduction to the Stability of Rational Expectations Equilibrium," *J. Econ. Theory* **26**, 313–317.

Crawford, V. (1985). "Learning Behavior and Mixed Strategy Nash Equilibria," *J. Econ. Behav. Organ.* **6**, 69–78.

Eichberger, J., Haller, H., and Milne, F. (1990). "Naive Bayesian Learning in 2 × 2 Matrix Games," *J. Econ. Behav. Organ.*, in press.

Fudenberg, D., and Kreps, K. (1993). "Learning Mixed Equilibria," *Games. Econ. Behav.* **5**, 320–367.

Jordan, J. S. (1991a). "Bayesian Learning in Normal Form Games," *Games Econ. Behav.* **3**, 60–81.

Jordan, J. S. (1991b). "The Exponential Convergence of Bayesian Learning in Normal Form Games," *Games Econ. Behav.* **4**, 202–217.

Jordan, J. S. (1992a). "Bayesian Learning in Games: A NonBayesian Perspective," mimeo.

Jordan, J. A. (1992b). "Bayesian Learning in Repeated Games," mimeo.

Kalai, E., and Lehrer, E. (1991). "Rational Learning Leads to Nash Equilibrium," *Econometrica*, in press.

Nyarko, Y. (1992a). "Bayesian Learning in Repeated Games Leads to Correlated Equilibria," mimeo.

Nyarko, Y. (1992b). "Bayesian Learning without Common Priors and Convergence to Nash Equilibrium," mimeo.

Rosenmüller, J. (1971). "Über Perioizitätseigenschaften spieltheoretischer Lernprozesse," *Z. Wahrscheinlichkeitstheorie verwandte Gebeite* **17**, 259–308.

Rubenstein, A. (1991). "Comments on the Interpretation of Game Theory," *Econometrica* **4**, 909–924.

Scarf, H. (1960). "Some Examples of Global Instability of the Competitive Equilibrium," *Int. Econ. Rev.* **1**, 157–172.

Shapley, L. (1964). "Some Topics in Two-Person Games," in *Advances in Game Theory, Annals of Mathematical Studies* (M. Dresher, *et al.*, Eds.), Vol. 52, pp. 1–28.

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