Econometrica, Vol. 64, No. 6 (November, 1996), 1375–1393

A PROBABILISTIC MODEL OF LEARNING IN GAMES

BY CHRIS WILLIAM SANCHIRICO[1]

This paper presents a new, probabilistic model of learning in games which investigates the often stated intuition that common knowledge of strategic intent may arise from repeated interaction. The model is set in the usual repeated game framework, but the two key assumptions are framed in terms of the likelihood of beliefs and actions conditional on the history of play. The first assumption formalizes the basic intuition of the learning approach; the second, the indeterminacy that inspired resort to learning models in the first place. Together the assumptions imply that, almost surely, play will remain almost always within one of the stage game's "minimal inclusive sets." In important classes of games, including those with strategic complementarities, potential functions, and bandwagon effects, all such sets are singleton Nash.

KEYWORDS: Learning in games, rationalizability.

FOR ALMOST HALF A CENTURY Nash equilibrium has been game theory's predominant solution concept. Yet in recent years foundational research on games has focused on the need to shore up the justification for equilibrium's fundamental assumption: in plain terms, that players correctly guess their opponents' strategies. Figuring large in this new literature is a resurgence in research on learning in games. The learning approach rests on a simple intuition: namely that players who play together repeatedly will eventually reach a common understanding of their strategic intentions. Translating this intuition into concrete convergence results, however, has proven to be no simple matter.

Consider, for example, "fictitious play." In this leading model agents' beliefs about their opponents' current actions are assumed to equal the empirical frequency of past opponent play. Agents, moreover, act "myopically" in choosing to play a best response to such beliefs without regard to the effect on their opponents' future beliefs. Shapley's (1964) well known example showed that the resulting empirical frequencies of play do not generally converge. More recent research, however, has focused on the model's proclivity to generate nonconvergent sequences of *actual play* even when frequencies do converge. Such is the case in Game 1.[2] A simple geometric argument confirms that for a range of initial conditions the generated sequence of actions fails to converge to the game's unique pure equilibrium (Heads, Out), but instead behaves as if the game consisted solely of its Matching Pennies component (shaded). Beliefs cycle

[1] I wish to thank Truman Bewley, Margaret Bray, Joseph Chang, John Geanakoplos, Roger Myerson, Phil Reny, Rafael Rob, Susan Rose-Ackerman, the members of Yale's game theory reading group, and seminar participants at the Econometric Society's 1994 North American Summer Meetings, SITE's 1994 Summer Workshop, U. C. Berkeley Economics, Yale Law School, Yale Economics, and the Universidad Carlos III de Madrid's 1995 Summer Conference in Economic Theory. I especially wish to thank David Pearce for his advice and encouragement.

[2] Game 1 is essentially the same as an example in Fudenberg and Kreps (1993).

	Heads	Tails	Out
Heads	1, -1	-1, 1	1, 2
Tails	-1, 1	1, -1	-1, -4

FIGURE 1.—Game 1.

toward the Matching Pennies mixed equilibrium while actual play cycles among the four action profiles.

Convergence of actual play, here and in general, requires the coordination of two complementary phenomena. First, (Heads, Out) must be in some sense "absorbing." Second, there must be sufficient "entropy" in the system to insure that play reaches (Heads, Out) often enough for this absorption to occur. Under fictitious play, strict equilibria like (Heads, Out) are in fact *immediately* absorbing in the sense that if they are ever played, they are played forever after. Yet fictitious play fails to converge because the same rigid structure that produces absorption makes the process prone to eternally ignore its absorptive states.

The central problem of learning in games is to simultaneously generate both of these seemingly contradictory forces—absorption and entropy—in general games and from reasonable assumptions about how the history of play affects current beliefs. Such is the design here. The paper introduces a class of models defined by two mutually consistent assumptions, one each for absorption and entropy. Impotant on their own, the assumptions combine to imply convergence of actual play to one of the stage game's "minimal inclusive sets" (Basu and Weibull (1991)). Roughly, a minimal inclusive set is one that includes all its own best responses, and no other sets with the same property. In Game 1, for instance, {(Heads, Out)} is the only minimal inclusive set. More generally, such sets are generically singleton, Nash in games with strategic complementarities, potential functions, identical interests, or bandwagon effects. At the other extreme, in Shapley's example and Matching Pennies, the whole set of profiles is minimal inclusive; the paper leaves open the question of whether and how a common understanding of intent could develop in such "irreducible" games.

An example of a process satisfying the two assumptions for Game 1 will help to introduce the general approach, describe the assumptions, and explain how they imply convergence. Given $0 < \lambda \leq 1$, define for each player i and each subset of stage game action profiles,[3] $E \subseteq A$ the set $\Delta_i^{\lambda}(E) = \{\psi_{-i} \in \Delta(A_{-i}) | \psi_{-i}(E_{-i}) \geq \lambda\}$, where $\Delta(A_{-i})$ represents the set of probability measures on the set of opponent actions, A_{-i}. The set $\Delta_i^{\lambda}(E)$ represents the event that i thinks her opponent likely to play within E_{-i}, with "likely" defined by λ. For

[3]All symbols used here are defined formally in the first paragraph of Section 2.

each subset $B \subseteq \Delta(A_{-i})$ of stage game beliefs for i, let $u_i(B)$ represent the measure placing unit weight uniformly on B and zero weight elsewhere. Combining notation, consider for a given history $\{a^1, \ldots, a^{t-1}\}$, the measure[4] $u_i(\Delta_i^\lambda(a^{t-r}, \ldots, a^{t-1}))$. The measure puts probability one on the event that i believes her opponent likely to repeat an action he has taken within the last r periods.[5] All beliefs consistent with this event are regarded as equally likely. The graph marked "A" in Figure 2, for instance, shows the frequency distribution of $u_{col}(\Delta_{col}^\lambda(\cdots))$ when $r = 2$ and Row has just played Tails twice in a row.

In this example the measure describing i's current beliefs after the partial history $\{a^1, \ldots, a^{t-1}\}$ is a convex combination of the measure $u_i(\Delta_i^\lambda(a^{t-r}, \ldots, a^{t-1}))$ and a one period lag:[6]

$$(1) \qquad p_i(a^1, \ldots, a^{t-1}) = \alpha p_i(a^1, \ldots, a^{t-2}) + (1 - \alpha) u_i\big(\Delta_i^\lambda(a^{t-r}, \ldots, a^{t-1})\big).$$

Thus, the measure after the history $\{a^1, \ldots, a^\tau, \ldots, a^{t-1}\}$ is just a geometric average of past $u_i(\Delta_i^\lambda(a^{\tau-r}, \ldots a^{\tau-1}))$'s. Figure 2, for example, shows the progression of $u_{col}(\Delta_{col}^\lambda(a^{t-r}, \ldots, a^{t-1}))$ and $p_{col}(a^1, \ldots, a^{t-1})$ for a given history of play for Row, starting from given initial conditions, with $\alpha = 1/2$ and $r = 2$. Finally, assume that players' current beliefs are drawn independently at each partial history and that both players play a myopic best response to their current beliefs.

These specifications define a probability measure P on the set of all sequences $\{\psi^t, a^t\}$ of beliefs and actions.[7] The theorems of this paper imply that if we set $\lambda \geq 3/4$ (which we do hereinafter), then P converges to (Heads, Out) in the sense that almost surely (Heads, Out) will be played almost always. In other words, P assigns probability one to the set of all sequences $\{\psi^t, a^t\}$ in which a^t remains at (Heads, Out) forever after some point.

Before explaining how this result obtains, it is worth emphasizing an important difference between this and more traditional learning models. Notice that the measure $p_i(a^1, \ldots, a^{t-1})$ is not player i's belief about her opponent's choice at t, but rather, a probability measure on such beliefs. While learning rules such as fictitious play stipulate for each history what players' beliefs *certainly* are, this model specifies a measure describing what such beliefs *tend* to be. This probabilistic approach rests on three general assertions:

(i) *Placing "extra-rational" restrictions on beliefs is unavoidable given the indeterminacy of rational strategic interaction.* A corollary to the literature on extensive form rationalizability (Pearce (1984)) is that rationality alone will not yield

[4] I use sequential notation for both the sequence and its range.

[5] The parameter r should not be confused with the parameter ρ in Assumption 2.

[6] All statements made about this system of measures applies to any system which approximates it in a particular sense: the approximating measures always assign the same weights as (1) to each cell in a canonical simplicial subdivision of the simplex (of i's beliefs), whose cells have sides with length no greater than $1 - \lambda$.

[7] First, specify arbitrary tie-breaking measures for the case where two or more actions are best response to agent's stage game beliefs. Then apply Ash (1972, Theorem 2.7.2) to the product $(\Delta(A_{Col}) \times \Delta(A_{Row}) \times A_{Col} \times A_{Row}) \times (\Delta(A_{Col}) \times \Delta(A_{Row}) \times A_{Col} \times A_{Row}) \times \ldots.$

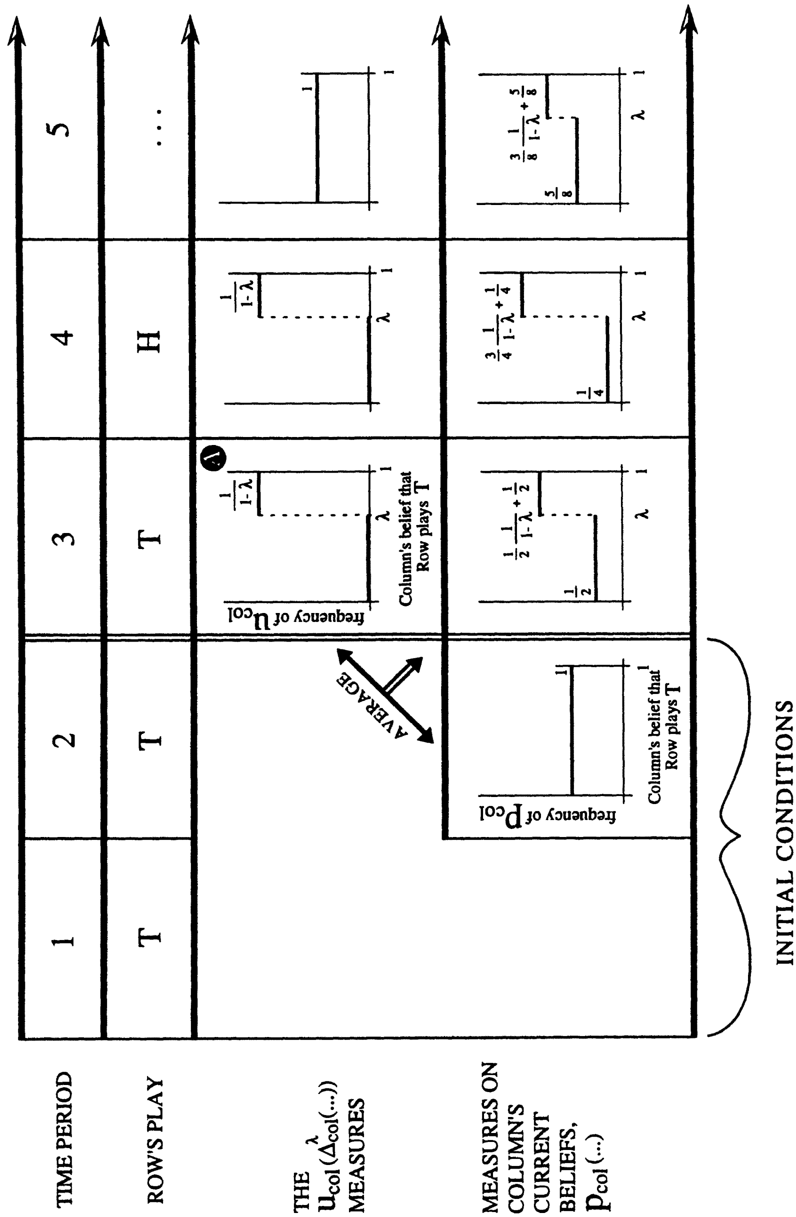
TIME PERIOD
ROW'S PLAY
THE $u_{col}(\Delta_{col}^{\lambda}(\ldots))$ MEASURES
MEASURES ON COLUMN'S CURRENT BELIEFS, $p_{col}(\ldots)$
1
2
3
4
5
T
T
T
H
. . .
AVERAGE
frequency of u_{col}
frequency of p_{col}
Column's belief that Row plays T
INITIAL CONDITIONS

FIGURE 2

convergence. In Game 1, for instance, wherein all actions are stage game rationalizable, *any* sequence of action profiles $\{a^t\}$ is consistent with common knowledge of rationality in the extensive form of the repeated game.[8]

(ii) *Such "extra-rational" restrictions should reflect the indeterminacy that necessitates their use*. The usual way to impose extra-rational restrictions is to stipulate that players form their beliefs according to a particular learning rule, such as fictitious play. In contrast, the probabilistic approach introduced herein eschews rigid, formulaic algorithms in favor of probabilistic statements about the tendency of players to think the past repeats itself. The necessary restrictions are thus cast in a manner that acknowledges and incorporates our agnosticism about how players form their beliefs (c.f., Gul (1991) and Milgrom and Roberts (1991)).

(iii) *Formalizing indeterminacy in a probabilistic framework is a natural way to generate entropy*. If we believe that many things are possible after each partial history of play, then we must also believe it likely that many things *will* happen over time. Thus, to the extent that convergence requires both absorption *and* entropy, the use of probability theory makes the model not only more palatable, but also more effective in generating convergence.

Returning to Game 1, convergence here follows from two intermediate results, which follow in turn from two properties of the process[9] (1). The first result is that (Heads, Out) is *eventually absorbing*: conditional on the event that (Heads, Out) is played infinitely often, it will be played almost always (i.e., always, after some point) with probability 1. Eventual absorption is in turn a consequence of the manner in which (Heads, Out) feeds back on itself under this process.

Feedback follows jointly from (1) and the best response properties of (Heads, Out). The more (Heads, Out) is played, the more often and predominant it is in recent history, $(a^{\tau-r},\dots,a^{\tau-1})$, and so the more probable it is that both players think their opponent likely to play (Heads, Out) again (since each player i's current beliefs $p_i(a^1,\dots,a^{t-1})$ are a geometric average of past $u_i(\Delta_i^{\lambda}(a^{\tau-r},\dots a^{\tau-1}))$ measures). This much is true of all action profiles. Since (Heads, Out) is an equilibrium (and λ is large enough), each player does in fact repeat her part of (Heads, Out) when she thinks her opponent likely to do the

[8] To "rationalize" $\{a^t_{Row}\}$, for example, construct the history-independent strategy s_{Row} prescribing a^t_{Row} at *all* time-t information sets/stage games. Since all Row's actions are stage game rationalizable, a^t_{Row} is stage game best response to some stage game beliefs ψ^t_{-Row} on *Column's* rationalizable actions. By the classic measure theorem (see, e.g., Ash (1972, Corollary 2.7.3)) we can then find a repeated game prior μ_{-Row} that is supported on similarly history-independent, stage game rationalizable strategies for *Column*, and induces ψ^t_{-Row} at each time-t information set. Since Row believes his current actions do not affect Column's future play, s_{Row} is a perfect best response to μ_{-Row}, whatever Row's discount factor $\delta \in [0,1)$. Thus, the sequence $\{a^t_{Row}\}$ is generated by a perfect best response to a prior that is supported on similarly structured strategies for Column, which may be similarly rationalized.

[9] Convergence is no easier to prove for this example than for the general case and so, to save space, the following analysis is confined to explanation of these properties rather than proof.

same. Hence, the feedback: the more (Heads, Out) is played, the more likely it is to be played yet again.

Feedback per se, however, is not enough to generate eventual absorption. For example, the probability that (Heads, Out) is repeated might increase to an asymptote of $\frac{1}{2}$, implying that almost surely (Heads, Out) is infinitely often *not* played. Indeed, the bare fact that the probability of (Heads, Out) approaches 1 is also not sufficient. Theorem 1, however, establishes that eventual absorption *does* follow if the feedback is *uniformly summable*: i.e., the probability that players think (Heads, Out) likely to recur always increases (in the number of times in a row n that it has played) faster than $1-x_n$ for some uniformly chosen (across t), nonnegative, summable sequence $\{x_n\}$. Such is the case here, where, as the reader can check, after $n+r-1$ plays of (Heads, Out), the chance that the players jointly think it likely to recur is always at least

$$[(1-\alpha)(1+\alpha+\cdots+\alpha^{n-1})]^2=[1-\alpha^n]^2=1-(2\alpha^n-\alpha^{2n}).$$

The proof that eventual absorption follows from uniformly summable feedback has two steps. By an argument related to the second Borel-Cantelli Lemma, the summability condition just discussed implies that every time we arrive at (Heads, Out) there is some chance we stay there forever. Since the summability is uniform, this chance is uniform as well. An argument from the *first* Borel-Cantelli Lemma then establishes that this small chance of absorption on each arrival translates into a long run certainty, so long as we arrive at (Heads, Out) sufficiently often.

Arriving at (Heads, Out) sufficiently often is the role of this process' second key property, *best response entropy*: namely, if the profile a is a best response (for both players) to $\tilde{a}$, and $\tilde{a}$ is played at t, then the chance that a will be played at $t+1$ is, in this example, never less than $((1/16)(1-\alpha))^2$. Importantly, the chance is uniformly (over t) bounded away from zero. The indeterminacy of rational strategic interaction might suggest that anything is possible after all partial histories. Applied literally, this precludes any form of convergence. Best response entropy insists only that best responses to actions played recently—actions that are in a sense still "in play"—be regarded as possible. The degree of entropy over current play is thus a function of whether recent history has many or few best responses, which depends in turn on whether recent history is itself diffuse or concentrated.

That best response entropy implies infinite plays of (Heads, Out) is the content of Theorem 2. For intuition, note that if (Heads, Out) is *not* played infinitely many times, some other profile in this finite game, say (Tails, Tails), must be. But, because (Tails, Heads) is a joint best response to (Tails, Tails), each time the former is played, there is a chance the latter is played in the following period, implying that (Tails, Heads) is *also* played infinitely often. Continuing the argument to (Heads, Heads) and then (Heads, Out) we obtain a contradiction. There being such a "best response chain" from all profiles to (Heads, Out), the result follows.

Taken alone, uniformly summable feedback is consistent with the performance of fictitious play in Game 1. Best response entropy on its own is consistent with drawing action profiles in i.i.d. fashion. Together, however, the properties imply almost sure, almost always convergence to (Heads, Out). Best response entropy gives Pr(Heads, Out) i.o.) = 1, where i.o. means infinitely often. Uniformly summable feedback yields Pr((Heads, Out) a.a.|(Heads, Out) i.o.) = 1, where a.a. means almost always. The product of these is the convergence result.

It is worth noting that this convergence is not a special case of Kalai and Lehrer's (1993) model of "rational learning." Indeed, so long as $\lambda < 1$, one can show that players will almost surely *not* put positive weight on the true path of play—what is required by rational learning's absolute continuity assumption in this context. Yet the main point of comparison with rational learning is perhaps more methodological than technical. Arguably, rational learning's assumption that i puts positive weight on the true path of play[10] is really just another way of saying that as play unfolds, i becomes more and more certain and correct in her beliefs about the future course of the game[11]—thus begging the question of why this might occur. In contrast, this paper is an explicit attempt to *explain* convergence. The result is not in any sense a mathematical restatement of the fact of convergence, but rather a mathematical formalization of an explanation for convergence that is fundamentally behavioral: namely, that the right combination of indeterminacy, and a self-intensifying tendency for players to think that history repeats itself, will lead players over time to a common understanding of their strategic intentions.

The theorems and lemmas of this paper generalize this introductory example along several dimensions. First, the model captures the two highlighted properties of this example in two assumptions on general measures over paths of play and beliefs. *Any* measure satisfying these assumptions is shown to converge. Second, convergence is shown for general games, to one of the stage game's minimal inclusive sets. Third, in the general model it is the probability that a subset is "salient" that increases as the subset is played repeatedly, not necessarily the probability that players think it likely to recur. As explained within, salience generalizes the latter to full hierarchies of beliefs. A final generalization —not included here—is that players need not be myopic. It is enough that

[10] To be sure Kalai and Lehrer (1993) allow for behavioral strategies, in which case there is no one true path of play. This makes their theorems more interesting and difficult than the results reported here, but it does not defeat the basic criticism.

[11] Let $\{a^t\}$ be the true path of play induced by the players' repeated game strategy profile. Let p_t be the probability that i places on the true action profile a^t in the stage game following the true history $\{a^1, \dots, a^{t-1}\}$. (We may derive this from i's strategy and prior.) The probability that i places on the *true* path is just $\prod_{t=1}^{\infty} p_t$. A basic result on infinite products says that $\prod_{t=\tau}^{\infty} p_t > 0$ for some $\tau \geq 1$ is equivalent to $\lim_{\tau \to \infty} \prod_{t=\tau}^{\infty} p_t = 1$. But $\prod_{t=\tau}^{\infty} p_t$ is just the probability that i's prior assigns to the *continuation* $\{a^\tau, a^{\tau+1}, \dots\}$ of the true path from time τ on.

discount factors are sufficiently small. (Note that no matter how small a player's positive discount factor, her current beliefs may be so close to the break-even between two actions that future effects are decisive in her current choice.[12])

It is possible to construct many other examples satisfying the assumptions of the general model. To come full circle, we can use the framework to alter fictitious play so as to improve its performance in Game 1. We make the weights geometric, rather than arithmetic, and simultaneously add some constant positive probability that players play last period best response instead of what is dictated by the re-weighted fictitious play beliefs:

$$\mu_{-i}^{t} = \begin{cases} \beta\mu_{-i}^{t-1} + (1-\beta)a_{-i}^{t-1}, & \text{with probability } \alpha, \\ a_{-i}^{t-1}, & \text{with probability } (1-\alpha), \end{cases}$$

where here "a_{-i}^{t-1}" represents the *belief* putting unit weight on the action a_{-i}^{t-1}. This system of measures clearly satisfies best response entropy and, proceeding as if $\alpha = 1$, one can show that

$$x_n = \begin{cases} 0, & \text{if } n < \ln(1/4)/\ln\beta, \\ 1, & \text{otherwise}, \end{cases}$$

is a uniform (across t) lower bound on the probability that the players think it likely that (Heads, Out) will be repeated again after it has been played n times in a row. ($n < \ln(1/4)/\ln\beta \Leftrightarrow 1-\beta^n \geq (3/4) = \lambda$.)

The bounding sequence here is almost always zero: if (Heads, Out) is ever played $\ln(1/4)/\ln\beta$ times in a row, it is played forever after. Contrast this "lock-in" dynamic with the example above wherein there is always some chance of not playing (Heads, Out) no matter how many times it has been played. Hurkens' (1994) model of learning by forgetful players also satisfies the assumptions of the general model with a bounding sequence that is almost always zero —always zero for all n larger than the bound on memory. Sonsino (1994) generates convergence to patterns of play (see the conclusion for more on patterns); convergence there also operates by a similar "lock-in" dynamic.

Section 2 of the paper sets out the general framework of the model. Section 3 proves the two intermediate results. The main convergence result is proven in Section 4, which also discusses the consistency of the assumptions. Section 5 concerns the size of minimal inclusive sets in special classes of games and Section 6 concludes the paper.

1. GENERAL FRAMEWORK AND ASSUMPTIONS

Fix a stage game $G = (A_1, \ldots, A_m; \pi_1, \ldots, \pi_m)$, where A_i is player i's finite set of *actions* and $\pi_i\colon A_1 \times \cdots \times A_m \to \Re$ is i's *payoff function*. For any

[12] For more details, please see the appendix to Sanchirico (1996a).

subset $E_{-i} \subseteq A_{-i}$ of opponent action profiles,[13] let $\Delta(E_{-i})$ denote the set of all probability measures ψ_{-i} on A_{-i} with $\psi_{-i}(E_{-i}) = 1$. Extend π_i to an *expected payoff function* u_i: $A_i \times \Delta(A_{-i}) \to \Re$ in the usual manner. Denote the set of *(stage game) best responses for player i to the belief* $\psi_{-i} \in \Delta(A_{-i})$ as $b_i(\psi_{-i})$. The set of *(stage game) best responses for player i to beliefs on any subset* $E_{-i} \subseteq A_{-i}$ is $b_i \circ \Delta(E_{-i}) \equiv \bigcup_{\psi_{-i} \in \Delta(E_{-i})} b_i(\psi_{-i})$. Finally, for any subset $E \subseteq A$ of action profiles whether or not rectangular, define $b \circ \Delta(E) = (b_1 \circ \Delta(E_{-1}), \ldots, b_m \circ \Delta(E_{-m}))$.

The model's assumptions concern the manner in which the history of play affects the likelihood of players' belief hierarchies regarding opponents' current actions. A formal statement of these assumptions, then, requires both a definition of such belief hierarchies and a probability space in which to cast statements about likelihoods. For the first task we borrow from Tan and Werlang's (1988) adaptation of "types" to uncertainty regarding strategic intent; let Θ_i denote the topological space of *stage game types* for player i with respect to the set of opponent action profiles A_{-i}, as in their Definition 3.9 (applied to A_{-i} rather than A). For the second task we provide the following definition.

DEFINITION 1: Define the *probability space of action / belief paths* for the game G to be the tuple $([\Theta \times A]^\infty, \Im, P)$, where: (i) $[\Theta \times A]^\infty$ denotes the set of all sequences $\{\theta^t, a^t\}$ of profiles of stage game types and actions, (ii) $\Im$ is the product σ-algebra on $[\Theta \times A]^\infty$ constructed from the Borel sets on each copy of Θ and the power set on each copy of A, and (iii) P: $\Im \to \Re$ is a probability measure.[14]

The object here is to generate common knowledge of *strategic intent* from repeated play; common knowledge of *rationality* is assumed from the onset. Assumption 0, which translates the assumption of common knowledge of rationality into our probability space, borrows more from Tan and Werlang (1988). First, the subset $K_i \subseteq \Theta_i$ (from their Definition 5.2) represents the set of all types for player i consistent with common knowledge of rationality. Second, since by their Theorem 3.1 each θ_i may be regarded as a probability measure on $A_{-i} \times \Theta_{-i}$, we may let $\theta_i(A_{-i})$ and $\theta_i(\Theta_{-i})$ denote the marginal of θ_i on A_{-i} and Θ_{-i}, respectively.

[13] I use the following conventional notation for products. A product set $X_1 \times \cdots \times X_n$ is denoted interchangeably as X. Given any subset S of a product $X_1 \times \cdots \times X_n$ (whether or not the subset is itself a product), S_i denotes the projection of S onto the ith factor and S_{-i} denotes the projection of S onto the product of all factors except the ith.

[14] Two technical notes about this probability space: First, since each pair of profiles of *repeated* game strategies and beliefs (types) induces a unique sequence $\{(\theta^t, a^t)\}$, specifying P is the same as specifying a measure on these repeated game objects (with an appropriately defined θ-algebra). Second, specifying one "big" P over all *sequences* of stage game beliefs and actions is essentially equivalent to specifying a separate probability measure over current beliefs and actions at each "node," $\{\theta^t, a^1, \ldots, \theta^{t-1}, a^{t-1}\}$. See again Ash (1972, Theorem 2.7.2).

ASSUMPTION 0:[15] $\forall t \geq 1, \forall i = 1, \ldots m$, $P([a_i^t \in b_i(\theta_i^t(A_{-i}))] \cap [\theta_i^t \in K_i]) = 1$.

In the introduction's example the history of play affected the probability that players thought a given subset E_{-i} of profiles (there, a singleton) likely to recur. In general, it need only affect the probability that the set is "salient:" every player i *either* believes it likely that his opponents will play in E_{-i} in the incipient stage game, *or* believes it likely that his opponents hold such beliefs, *or*, believes it likely that his opponents believe *their* opponents hold such beliefs, *or*, etc... up to any order. (For simplicity the parameter λ and hence the qualifier "likely" is left out of the formal definition of salience. The generalization is easily conceived, yet tedious to denote.)

For any product $X_1 \times \cdots \times X_m$, and any subset S_k of any factor X_k, let $\langle S_k \rangle$ denote the "slab" of S_k, that is, the subset $\{x \in X_1 \times \cdots \times X_m | x_k \in S_k\}$.

DEFINITION 2: Fix a rectangular subset of action profiles $E = E_1 \times \cdots \times E_m \subseteq A$. Define, for all i, the set $S_i(1)(E) = \{\theta_i \in \Theta_i | \theta_i(A_{-i}) \in \Delta(E_{-i})\}$. Continuing inductively, given $S_j(n-1)(E)$ for each player j, define

$$S_i(n)(E) = \Big\{\theta_i \in \Theta_i | \forall j \neq i, \forall a_j \in \operatorname{supp} \theta_i(A_j),$$
$$\text{either } a_j \in E_j \text{ or } \theta_i\big(\langle \{a_j\} \rangle \cap \langle S_j(n-1)(E) \rangle\big) > 0\Big\}.$$

Then, define for all i, the set $S_i(E) = \bigcup_{n=1}^{\infty} S_i(n)(E)$. Lastly, define $S(E) \equiv S_1(E) \times \cdots \times S_m(E)$. The subset E is said to be *salient* at time t, if $\theta^t \in S(E)$.

Assumption 1 is the source of the uniformly summable feedback discussed in the introduction. Assumption 2 is the source of best response entropy. Both assumptions are parameterized, the former by the class of subsets to which it applies, the latter by the length of recent history, ρ. This parameterization allows for two modes of convergence in the main theorem. For each $\Gamma \subseteq 2^A$ define:

ASSUMPTION $1(\Gamma)$: *For all subsets of action profiles $E \in \Gamma$, there exists a summable sequence $\{x_n\}$ such that*:

(2) $$\forall t \geq 1, \forall 1 \leq n \leq t-2,$$
$$P(\theta^t \in S(E) | a^{t-1}, \ldots, a^{t-n} \in E;\ a^{t-n-1} \notin E) \geq 1 - x_n,$$
if defined.

For each $\rho \geq 1$ define:

ASSUMPTION $2(\rho)$: *There exists $\varepsilon > 0$ such that for all $t \geq 1$ and all $\{a^1, \ldots, a^{t-1}\}$, if $a \in b \circ (\{a^{t-\rho}, \ldots, a^{t-1}\})$, then $P(a^t = a | \{a^1, \ldots, a^{t-1}\}) \geq \varepsilon$, if defined.*

[15]Assumption 0 is less general than it might be. First, all results hold with a sufficiently small, nonzero discount rate. Second, (as with all the assumptions) it need only hold for almost all t. Third, it suffices that players' play a best response to some belief that is "almost" supported on the support of their current beliefs and that this is "almost" common knowledge.

2. INTERMEDIATE RESULTS

2.1 *Inclusive Sets, Feedback, and Eventual Absorption*

Following Basu and Weibull (1991) let us say that a nonempty subset of action profiles $E = E_1 \times \cdots \times E_m \subseteq A$ is *(best response) inclusive*, if $b_i \circ \Delta(E_{-i}) \subseteq E_i$, $\forall i$. In Game 1 the entire set of profiles and (Heads, Out) as a singleton are the only inclusive sets. The importance of inclusive sets in this model of learning lies in the following lemma, which says that if rationality is common knowledge, then whenever an inclusive set is salient it will in fact be played in. The lemma thus establishes that inclusive sets feed back on themselves under Assumption 1. That this implies eventual absorption is the content of Theorem 1.

LEMMA 1: *If P satisfies Assumption* 0 *and I is inclusive, then for all $t \geq 1$, $P(a^t \in I \mid \theta^t \in S(I)) = 1$, if defined.*

The proof, which appears in the Appendix, is inductive on orders of belief. For intuition, note that if i believes her opponents will play in I, then she, being rational, will herself play in I, since I contains all best responses to itself. Similarly, if i believes both that her *opponents* are rational and that they think *their* opponents will play in I, then she must think that her opponents will themselves play in I. Then, again, *she* will play in I.

Let $[a^t \in I$ i.o.$]$ denote the event that play is in I "infinitely often"—in set notation $\bigcap_{t=1}^{\infty} \bigcup_{s=t}^{\infty} [a^s \in I]$. Let $[a^t \in I$ a.a.$]$ be the event that play is in I "almost always," $\bigcup_{t=1}^{\infty} \bigcap_{s=t}^{\infty} [a^s \in I]$.

THEOREM 1 (Eventual Absorption): *For all $\rho \geq 1$ and all $\Gamma \subseteq 2^A$, if Assumptions* 0, 1(Γ), *and* 2(ρ) *hold and I is an inclusive set in Γ, then $P(a^t \in I$* a.a.$\mid a^t \in I$ i.o.$) = 1$, *if defined.*

A sketch of the following proof appears in the introduction.

PROOF: Let I satisfy Assumption 1 with sequence $\{x_n\}$. From Lemma 1 it follows that $\forall t \geq 1$, $\forall 1 \leq n \leq t-2$,

$$(3) \qquad P(a^t, \ldots, a^{t-n} \in I;\ a^{t-n-1} \notin I) \geq P(\theta^t \in S(I);\ a^{t-1}, \ldots, a^{t-n} \in I;\ a^{t-n-1} \notin I).$$

Combining (3) with Assumption 1 yields: $\forall t \geq 1$, $\forall 1 \leq n \leq t-2$,

$$(4) \qquad P(a^t, \ldots, a^{t-n} \in I;\ a^{t-n-1} \notin I) \geq (1 - x_n) P(a^{t-1}, \ldots, a^{t-n} \in I;\ a^{t-n-1} \notin I).$$

Re-indexing, write (4) as: $\forall t \geq 1$, $\forall n \geq 1$,

$$(5) \qquad P(a^{t+n+1}, \ldots, a^{t+1} \in I;\ a^t \notin I) \geq (1 - x_n) P(a^{t+n+1}, \ldots, a^{t+1} \in I;\ a^t \notin I).$$

By Assumption 2 and the fact that I is inclusive, we may take $x_n < 1$, all n. Now for any t, the family of inequalities in (5) indexed by n yields, by iterative substitution:

$$(6) \qquad P\left(\bigcap_{n=1}^{m} I^{t+n+1} \cap I^{t+1} - I^{t}\right) \geq \prod_{n=1}^{m} (1 - x_n) P(I^{t+1} - I^{t})$$

where I have written, and will henceforth write I^t for the event $[a^t \in I]$. The fact that $\{x_n\} \subseteq [0,1)$ and $\Sigma x_n < \infty$ implies that $\lim_{m \to \infty} \Pi_{n=1}^{m}(1 - x_n)$ exists and is a strictly positive number, call it ξ. (See, e.g., Knopp (1971).) Taking the limit of both sides in (6) yields: $\forall t \geq 1$,

$$(7) \qquad P\left(\bigcap_{n=1}^{\infty} I^{t+n+1} \cap I^{t+1} - I^{t}\right) \geq \xi \cdot P(I^{t+1} - I^{t}).$$

Now the sequence of sets $\{\bigcap_{n=1}^{\infty} I^{t+n+1} \cap I^{t+1} - I^{t}\}_{t=1}^{\infty}$ is disjoint. Hence, summing the left side of (7) over all $t \geq 1$ yields a number less than 1. Then since $\xi > 0$, (7) implies $\Sigma_{t=1}^{\infty} P(I^{t+1} - I^{t}) < \infty$. Therefore, by the first Borel-Cantelli lemma $P((I^{t+1} - I^{t})$ i.o.$) = 0$. A standard argument shows $[I^t$ i.o.$] - [I^t$ a.a.$] \subseteq [(I^{t+1} - I^{t})$ i.o.$]$ and the result follows. *Q.E.D.*

2.2 *Plateaus, Entropy, and Eventual Repulsion*

An inclusive set is said to be *minimal* if it does not strictly contain another inclusive set. The *plateau* of an inclusive set is constructed by removing from the inclusive set all smaller inclusive sets nested therein. (By convention the plateau of a minimal inclusive set is the empty set.) Thus, the plateau of the entire set of profiles in Game 1 is the Matching Pennies component plus (Tails, Out). In Game 1, but not in general, the plateau of the entire set of profiles corresponds to the *grand plateau*: what remains of the entire set of profiles when we remove all minimal inclusive sets.

Plateaus, and perhaps also the grand plateau, will be eventually repelling if recent history is sufficiently long. What length suffices depends on certain properties of the game's best response correspondence, which are summarized in the following notion of "size." (All results hold if we take the size of the game to be the number of profiles.)

DEFINITION 3: For all subsets $E \subseteq A$ and all action profiles $a \in E$, a *best response chain* from a to E is a finite sequence of action profiles $\{a(1), \ldots, a(n)\}$ satisfying: (i) $a(1) = a$, (ii) $\forall 1 \leq k \leq n$, $a(k) \in b \circ \Delta(\{a(1), \ldots, a(k-1)\})$, and (iii) $a(n) \in E$. The *span* of the chain $\{a(1), \ldots, a(n)\}$ is defined as $\max_{1 \leq k \leq n}\{\min j\colon a(k) \in b \circ \Delta(\{a(k-j), \ldots, a(k-1)\})\}$. Let a be an element of a nonempty plateau E. The *size of the profile*, a, is the smallest span across all chains from a to $\neg E$. The *size of the game* $s(G)$ is the largest size across all profiles.

Recall that in discussing the application of Theorem 2 to Game 1, we noted the existence of a best response chain (of span 1) from every profile to (Heads, Out). More generally, one can show for all finite games the existence of such a chain (not necessarily of span 1) from every point in every plateau to that plateau's complement. This insures that "size" is well defined and, though implicit in the proof, is central to eventual repulsion.

THEOREM 2 (Eventual Repulsion from Plateaus): *Let P satisfy Assumption* 2(ρ). (i) *If* $\rho \geq s(G)$, *then for all plateaus* $E \subseteq A$, $P(a^t \in E \text{ a.a.}) = 0$. (ii) *If* $\rho \geq |A|$, *then* $P(a^t \in \Pi \text{ a.a.}) = 0$, *where* Π *is the grand plateau.*

A sketch of the following proof appears in the introduction.

PROOF: I prove only part (i) of the theorem. Part (ii) follows in the same manner. Since the plateau $E \subseteq A$ is finite, at least one of its members is played infinitely often, so that $[E^t \text{ a.a.}] = [E^t \text{ a.a.}] \cap \bigcup_{a \in E}[a^t = a \text{ i.o.}]$. Then by the subadditivity of P,

$$P(E^t \text{ a.a.}) \leq \sum_{a \in E} P([a^t = a \text{ i.o.}] \cap [E^t \text{ a.a.}]). \tag{8}$$

Now take any $a \in E$. Let $\{a(1), \ldots, a(n)\}$ be a chain from $a(1) = a$ to $\neg E$ that has the smallest span of all chains from a to $\neg E$ (i.e. a chain whose span is the size of a). We show

$$P(\{a^{t-n}, \ldots, a^{t-1}\} = \{a(1), \ldots a(n)\} \text{ i.o.}) = P(a^t = a \text{ i.o.}). \tag{9}$$

The "$\leq$" direction is obvious since $a(1) = a$. To show "$\geq$" suppose, *contra*, that $r < n$ is the largest index for which (9) does hold with "$\geq$". Now $a(r+1) \in b \circ \Delta(\{a(r-\rho+1), \ldots, a(r)\})$, since $\{a(1), \ldots a(n)\}$ is a chain with smallest span among those from a to $\neg E$ and $s(G) \leq \rho$. Hence, Assumption 2(ρ) insures the existence of $\varepsilon > 0$ such that $\forall t \geq 1$ and each *individual* history $\{a^1, \ldots, a^{t-1}\}$ that has $\{a^{t-r}, \ldots, a^{t-1}\} = \{a(1), \ldots, a(r)\}$, $P(a^t = a(r+1) | \{a^1, \ldots, a^{t-1}\}) \geq \varepsilon$. Therefore, by a standard result (see, e.g., the appendix to Sanchirico (1996a)),

$$\begin{aligned} &P(\{a^{t-r}, \ldots, a^{t-1}\} = \{a(1), \ldots, a(r)\} \text{ i.o.}) \\ &\quad = P(([a^t = a(r+1)] \cap [\{a^{t-r}, \ldots, a^{t-1}\} = \{a(1), \ldots, a(r)\}]) \text{ i.o.}) \\ &\quad = P(\{a^{t-r}, \ldots, a^t\} = \{a(1), \ldots, a(r+1)\} \text{ i.o.}), \end{aligned}$$

contradicting our supposition. This proves equation (9), which in turn implies

$$\begin{aligned} &P([a^t = a \text{ i.o.}] \cap [E^t \text{ a.a.}]) \\ &\quad = P([\{a^{t-n}, \ldots, a^{t-1}\} = \{a(1), \ldots, a(n)\} \text{ i.o.}] \cap [E^t \text{ a.a.}]). \end{aligned}$$

But since $a(n) \notin E$, $P([\{a^{t-n}, \ldots, a^{t-1}\} = \{a(1), \ldots, a(n)\} \text{ i.o.}] \cap [E^t \text{ a.a.}]) = 0$, and so $P([a^t = a \text{ i.o.}] \cap [E^t \text{ a.a.}]) = 0$ also. This holding for all $a \in E$, the result follows from (8). *Q.E.D.*

3. CONVERGENCE AND CONSISTENCY

Neither Assumption 1 nor 2 guarantees convergence on its own. Fictitious play's performance in Game 1—as observed in the opening paragraphs of this paper—is fully consistent with Assumptions 0 and $1(\Gamma)$, even with Γ taken to be the power set of A. Moreover, any probability measure that draws actions from Game 1 in i.i.d. uniform fashion is consistent with Assumptions 0 and $2(\rho)$, for any ρ.

The *interaction* of these assumptions, however, produces a strong form of convergence to minimal inclusive sets, as proven below in Theorem 3. The theorem is perhaps best understood visually. Figure 3 depicts a stage game ((0,0) payoffs are not shown) as a contour map with smaller inclusive sets marked with darker shading. Figure 4 translates this contour map into three dimensions (making clear the choice of the term "plateau"). The second intermediate result, eventual repulsion, guarantees that we do not remain on the highest plateau almost always, implying that we are infinitely often in one of its "holes." In particular, calling the smaller hole S and the larger L, eventual repulsion yields $P([a^t \in S \text{ i.o.}] \cup [a^t \in L \text{ i.o.}]) = 1$. (Note that this is not the same as "$P([a^t \in S \text{ i.o.}]) = 1$ or $P([a^t \in L \text{ i.o.}]) = 1$."). Now split the event $[a^t \in S \text{ i.o.}] \cup [a^t \in L \text{ i.o.}]$ into two (intersecting) sections, $[a^t \in S \text{ i.o.}]$ and $[a^t \in L \text{ i.o.}]$, and consider first $[a^t \in S \text{ i.o.}]$. The *first* intermediate result, eventual absorption says that *conditional* on $[a^t \in S \text{ i.o.}]$, the event $[a^t \in S \text{ a.a.}]$ has probability one. Similarly, conditional on $[a^t \in L \text{ i.o.}]$, $[a^t \in L \text{ a.a.}]$ receives probability one. Together with $P([a^t \in S \text{ i.o.}] \cup [a^t \in L \text{ i.o.}]) = 1$, these conditions imply (with some Boolean manipulation) that $P([a^t \in S \text{ a.a.}] \cup [a^t \in L \text{ a.a.}]) = 1$. In words, probability is divided between those sequences of play that stay in the larger hole always after some point and those sequences that stay in the smaller hole

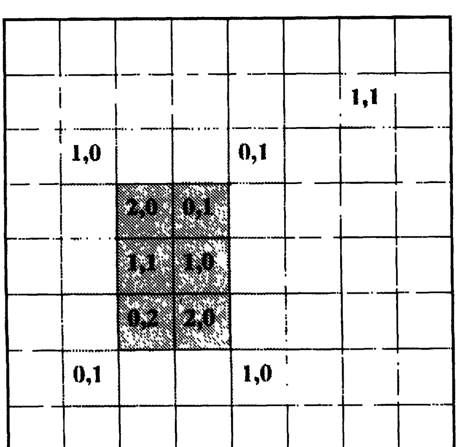

FIGURE 3

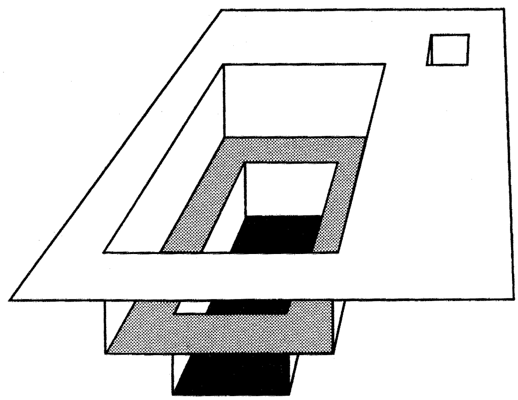

FIGURE 4

always after some point. Now, conditional on either of these (disjoint) "almost always" events we reapply eventual repulsion and absorption in parallel fashion to whatever holes (even smaller inclusive sets) there may be in either S and L, respectively. (Since there are no holes in S, we work again with S itself.) When we eventually reach a level at which none of the inclusive sets has a hole, we have established that there is probability one on those sequences of play that after some point remain in one of the stage game's minimal inclusive sets.

This describes the process of convergence when Assumption 2 holds with recent history of relatively short duration and Assumption 1 applies to all inclusive sets. Only one iteration is necessary if recent history is long enough to guarantee that the *grand* plateau does not absorb, in which case Assumption 1 need only be applied to *minimal* inclusive sets.

THEOREM 3 (Convergence): *Let P satisfy Assumptions* 0, 1(Γ), *and* 2(ρ). *If either* (*i*) *Γ contains all G's inclusive sets and $\rho \geq s(G)$, or* (*ii*) *Γ contains all G's minimal inclusive sets and $\rho \geq |A|$, then* $P(\cup\{[a^t \in I \text{ a.a.}] | I \text{ is minimal inclusive}\}) = 1$.

PROOF: Case (ii) follows directly from Theorems 1 and 2. For case (i) see the appendix to Sanchirico (1996a).

If ρ is smaller than the size of the game, then Assumption 0, 1(Γ), and 2(ρ) do not imply convergence. Game 2, for instance, is of size 2. One can construct a measure P satisfying Assumption 0, 1($\Gamma = A$), and 2($\rho = 1$), that puts probability 1 on the sequence of actions generated by "last-period-best-response" starting from (Heads, Heads): namely, (Heads, Heads), (Heads, Tails), (Tails, Tails), (Tails, Heads), (Heads, Heads), etc....

Convergence is of little interest if no probability measure could possibly satisfy the assumptions from which it has been shown to follow. Mere existence, however, is trivial. For any game G, let $\psi_1, \dots, \psi_m$ be a mixed equilibrium in one of G's minimal inclusive sets. The measure P that, at each t, (i) puts unit weight on the type profile θ such that (a) each player's belief about incipient opponent play is $\psi_1 \times \cdots \times \psi_{i-1} \times \psi_{i+1} \times \cdots \times \psi_m$ and (b) "(a)" is common knowledge, and (ii) draws play according to $\psi_1 \times \cdots \times \psi_m$, is consistent with Assumption 0, Assumption 1 applied to all subsets, and Assumption 2 for any length of recent history. The appendix to Sanchirico (1996a) shows that nontrivial measures exist for all games.

4. MINIMAL INCLUSIVE SETS IN SPECIAL CLASSES OF GAMES

This section concerns the size of minimal inclusive sets in classes of games for which *other* learning processes have been shown to converge (in a manner weaker than proven here). The main result relies on a general theorem whose proof is straightforward and so omitted. (For more details, see Sanchirico (1996b).)

Let E be a rectangular subset of action space A. The *restriction* of stage game G to E, denoted G_E is the finite game with strategy sets E_i and payoffs $\pi_i|E$. Let R be a property defined on the set of all finite games. We say R is *restrictable to inclusive sets*, if for all games G with property R and all inclusive sets I in G, the restriction G_I also has the property. The property: "has no more than one (pure strategy) equilibrium," for instance, is restrictable to inclusive sets, since all equilibria in the restriction of G to an inclusive set are equilibria in G as well. The property is not, however, restrictable to general subsets of profiles. The property "has no less than one (pure) equilibrium" is *not* restrictable, even to inclusive sets.

THEOREM 4: *If property R implies the existence of a pure strategy Nash equilibrium and is restrictable to inclusive sets, then in all games with property R for which all pure equilibria are strict,*[16] *all minimal inclusive sets are singletons consisting of strict equilibria.*

	Heads	**Tails**	**Out**
Heads	1, -1	-1, 1	$-\frac{1}{2}, \frac{1}{2}$
Tails	-1, 1	1, -1	$-\frac{1}{2}, \frac{1}{2}$
Out	$\frac{1}{2}, -\frac{1}{2}$	$\frac{1}{2}, -\frac{1}{2}$	0,0

FIGURE 5.—Game 2.

For lack of space the reader is referred to the papers cited for formal definitions of the classes in the following corollary. However, one new concept is required. Because the subset of a complete lattice may not be a complete lattice in its own right, general supermodular games are not restrictable to inclusive sets and so not subject to Theorem 4. Let us say, then, that an ordinal supermodular game (Milgrom and Shannon (1994)) is *restrictable*, if for all inclusive sets I, each I_i is a complete lattice in its own right. Contained in this subclass are all supermodular games whose strategy sets may be *completely* ordered, including all those analyzed by Krishna (1991).

COROLLARY 1 (Singleton Minimal Inclusive Sets in Special Classes of Games): *Let G be a finite game. All of G's minimal inclusive sets are singletons consisting of strict equilibria, if any of the following hold*:

(i) *G is a restrictable ordinal supermodular game all of whose pure equilibria are strict*,

(ii) *G is an ordinal potential game* (*Monderer and Shapley* (1993*b*)) *all of whose pure equilibria are strict*,

(iii) *G is a game with identical interests* (*Monderer and Shapley* (1993*a*)) *all of whose pure equilibria are strict*,

(iv) *G has the marginal bandwagon property* (*Kandori and Rob* (1992)).

PROOF: By Theorem 4, we need only show that each property (i)–(iv) is restrictable to inclusive sets and implies the existence of a pure equilibrium (with a slight variation in case (iv)). In all cases, restrictability follows directly from the definitions. For existence: (i) Existence of a pure equilibrium is given by Milgrom and Shannon (1991, Theorem 15). (ii) Existence is given by Monderer and Shapley (1993b, Corollary 2.2). (iii) Existence is noted by Monderer and Shapley (1993b, p. 9). (iv) One can show that the definition of the marginal bandwagon property implies the existence of a *strict* equilibrium.

5. CONCLUSION

One direction for future research would be to extend the model to nonsimultaneous stage games, including several plays of a given normal form. This would in turn allow convergence to patterns of equilibria in the simultaneous game, as in Sonsino (1994), but without the lock-in dynamic that operates there. It would also be broad enough to encompass convergence to equilibria that are not simple sequences of equilibria in the simultaneous game. Extending the model in this direction would also force confrontation with the observability problems identified in Fudenberg and Kreps (1988).

[16] Clearly some sort of genericity requirement (such as the requirement here that "all pure equilibria are strict") will be necessary, since these classes of games typically include those with constant payoffs.

Other directions for future research include discovering the precise relationship between the "parameters" in the two assumptions (Γ, $\{x_n\}$, ρ, and ε) and the speed of convergence, and analyzing the issue of which minimal inclusive sets are likely to be selected.

Dept. of Economics, Columbia University in the City of New York, 420 W. 118th St., 10th Floor, New York, New York 10027, U.S.A.

Manuscript received July, 1994; final revision received January, 1996.

APPENDIX: PROOF OF LEMMA 1

Take any inclusive set I. I give the proof in two steps.

Step 1: First I show that $\theta \in K \cap S(I)$ implies $\theta_i(A_{-i}) \in \Delta(I_{-i})$ for all i. For this it suffices to show that for all n and all i, $\theta_i \in K_i \cap S_i(n)(I)$ implies $\theta_i(A_{-i}) \in \Delta(I_{-i})$. The proposition is true for all i and $n=1$ by definition of $S_i(1)(I)$. Continuing inductively, suppose that, for all j, $\theta_j \in K_j \cap S_j(n-1)(I)$ implies $\theta_j(A_{-j}) \in \Delta(I_{-j})$. Take any i and any $\theta_i \in K_i \cap S_i(n)(I)$. I must show $\theta_i(A_{-i}) \in \Delta(I_{-i})$. To this end take any $j \neq i$ and any $a_j \in \operatorname{supp} \theta_i(A_j)$. I claim $a_j \in I_j$. By definition of $S_i(n)(I)$ either (i) $a_j \in I_j$, in which case we are done, or (ii) $\theta_i(\langle\{a_j\}\rangle \cap \langle S_j(n-1)(I)\rangle) > 0$. (Recall that θ_i may be regarded as a measure on $\Theta_{-i} \times A_{-i}$ and that the notation "$\langle\ \rangle$" denotes the inverse image of the projection mapping.) Consider case (ii). First, since $\theta_i \in K_i$, we know that $\theta_i(\langle K_j(m-1)\rangle) = 1$, for all m. Therefore, $\theta_i(\langle K_j\rangle) = 1$ and, in turn, $\theta_i(\langle\{a_j\}\rangle \cap \langle K_j \cap S_j(n-1)(I)\rangle) > 0$. Then by the inductive hypothesis, $\theta_i(\langle\{a_j\}\rangle \cap \langle\{\theta_j \in \Theta_j | \theta_j(A_{-j}) \in \Delta(I_{-j})\}\rangle) > 0$. Second, $\theta_i \in K_i \subseteq K_i(1)$ also means that $\theta_i(\langle\{(\theta_j, a_j) \in \Theta_j \times A_j | a_j \in b_j(\theta_j(A_{-j}))\}\rangle) = 1$, where the inverse projection here is with respect to the factor Θ_j in $\Theta_{-i} \times A_{-i}$. Combining these two implications yields $a_j \in b_j \circ \Delta(I_{-j})$. Then since I is inclusive, $a_j \in I_j$.

Step 2: For all $t \geq 1$,

$$
\begin{aligned}
&P([\theta^t \in S(I)]) \\
&\quad = P([\theta^t \in K \cap S(I)]) && \text{(Assumption 0)} \\
&\quad \leq P\left(\bigcap_{i=1}^{m} [\theta_i(A_{-i}) \in \Delta(I_{-i})] \cap [\theta^t \in S(I)]\right) && \text{(Step 1)} \\
&\quad = P\left(\bigcap_{i=1}^{m} [a_i^t \in b_i(\theta_i^t(A_{-i}))] \cap \bigcap_{i=1}^{m} [\theta_i^t(A_{-i}) \in \Delta(I_{-i})] \cap [\theta^t \in S(I)]\right) \\
&&& \text{(Assumption 0)} \\
&\quad \leq P([a^t \in b \circ \Delta(I)] \cap [\theta^t \in S(I)]) && \text{(Definition of } \Delta(E_{-i})) \\
&\quad \leq P([a^t \in I_i] \cap [\theta^t \in S(I)]) && (I \text{ is inclusive}).
\end{aligned}
$$

Hence, $P([a^t \in I_i] \cap [\theta^t \in S(I)]) = P([\theta^t \in S(I)])$ or $P([a^t \in I_i] | [\theta^t \in S(I)]) = 1$, if defined. *Q.E.D.*

REFERENCES

ASH, R. (1972): *Real Analysis and Probability*. New York: Academic Press.

BASU, K., AND J. W. WEIBULL (1991): "Strategy Subsets Closed Under Rational Behavior," *Economics Letters*, 36, 141–146.

FUDENBERG, D., AND D. KREPS (1988): "A Theory of Learning, Experimentation and Equilibrium in Games," mimeo, M.I.T. and Stanford Graduate School of Business.

——— (1993): "Learning Mixed Equilibria," *Games and Economic Behavior*, 5, 320–367.

GUL, F. (1991): "Rationality and Coherent Theories of Strategic Behavior," mimeo, Stanford Graduate School of Business.

HURKENS, S. (1994): "Learning by Forgetful Players: From Primitive Formations to Persistent Retracts," CentER for Economic Research, Tilburg University, The Netherlands, Discussion Paper No. 9437.

KALAI, E., AND E. LEHRER (1993): "Rational Learning Leads to Nash Equilibrium," *Econometrica*, 61, 1019–1045.

KANDORI, M., AND R. ROB (1992): "Bandwagon Effects and Long Run Technological Choice," mimeo, University of Tokyo and University of Pennsylvania.

KNOPP, K. (1971): *Theory and Application of Infinite Series*. New York: Hafner Publishing.

KRISHNA, V. (1991): "Learning in Games with Strategic Complementarities," mimeo, Harvard Business School.

MILGROM, P., AND J. ROBERTS (1991): "Adaptive and Sophisticated Learning in Normal Form Games," *Games and Economic Behavior*, 3, 82–100.

MILGROM, P., AND C. SHANNON (1994): "Monotone Comparative Statics," *Econometrica*, 62, 157–180.

MONDERER, D., AND L. S. SHAPLEY (1993a): "Fictitious Play Property for Games with Identical Interests," mimeo, The Technion, Haifa, and the University of California, Los Angeles.

——— (1993b): "Potential Games," mimeo, The Technion, Haifa, and the University of California, Los Angeles.

PEARCE, D. (1984): "Rationalizable Strategic Behavior and the Problem of Perfection," *Econometrica*, 52, 1029–1050.

SANCHIRICO, C. W. (1996a): "A Probabilistic Model of Learning in Games," Columbia University, Dept. of Economics, Discussion Paper Series No. 9596-22.

——— (1996b): "Minimal Inclusive Sets in Special Classes of Games," Columbia University, Dept. of Economics, Discussion Paper No. 9596-23.

SONSINO, D. (1994): "Learning to Learn, Pattern Recognition, and Nash Equilibrium," mimeo, Stanford Graduate School of Business.

SHAPLEY, L. S. (1964): "Some Topics in Two Person Games," *Advances in Game Theory*, 26, 1–28.

TAN, T., AND S. WERLANG (1988): "The Bayesian Foundations of Solution Concepts of Games," *Journal of Economic Theory*, 45, 370–391.