Working Paper Series

[redacted]

DEPARTMENT OF AGRICULTURAL ECONOMICS

BERKELEY

CALIFORNIA AGRICULTURAL EXPERIMENT STATION

University of California

[redacted]

Division of Agricultural Sciences
UNIVERSITY OF CALIFORNIA

Working Paper No. 257

ECONOMIC DEVELOPMENT AND INSTITUTIONAL CHANGE: TOWARD A MODELING FRAMEWORK

by

Irma Adelman and Thomas F. Head

California Agricultural Experiment Station
Giannini Foundation of Agricultural Economics
April 1983

ECONOMIC DEVELOPMENT AND INSTITUTIONAL CHANGE: TOWARD A MODELING FRAMEWORK

Irma Adelman and Thomas F. Head*

> It is currently a mark of sophistication in presenting economic models not to mention institutions. But for all that, it is a significant trait of contemporary economics that, despite this omission, it manages somehow to find support for institutional change. It is a neat trick, but it cannot hide the fact that, in thinking about institutions, the analytical cupboard is bare.
>
> [T. W. Schultz, p. 1113]

In the years since Nobel laureate Theodore W. Schultz made these remarks to an annual meeting of the American Agricultural Economics Association, attitudes as well as techniques have undergone change. Yet, despite these shifts, the economics profession is still far from arriving at an accepted and well-tested theoretical framework for the analysis of institutional change. We consider here this state of knowledge and potential future directions with particular reference to the study of developing countries.

Economic development rarely takes place without extensive and important changes in the very fabric of human interaction. The organisms of society go through nothing less than metamorphosis during the process of development. The economic results of such transformations are measured by our conventional indicators of national income, employment, distribution, etc. Demographic and welfare-related measures give us a further picture of the implications of development. But hidden behind these indicators is the fact that the organisms we study are, in many crucial respects, no longer alive at the end of the process and that new forms of life have emerged to take their place.

We suggest here that the institutional changes associated with development too often have been overlooked by economists as important areas of research. Institutional change is frequently mentioned, but it is seldom studied in the same rigorous manner that one examines other economic adjustments. An important problem in the advancement of economic thought is the generation of an adequate framework for the scientific study of institutional change. As Schultz stated: "the analytical cupboard is bare." Much of the writing in this area has never gone beyond the processes of formulating taxonomies, identifying linkages, pinpointing institutional constraints, and brainstorming about institutional innovations. This largely descriptive work has failed to stock the analytical cupboard with the conceptual apparatus needed for a disciplined study of institutions. Unease with this condition has prompted a new wave of institutional research. In Section I of this paper, we review this budding body of work as well as its antecedents; in Section II we outline a few of our own ideas about modeling institutional change; and in Section III we present illustrative applications of the model developed in Section II.

Before proceeding, a few words about definitions are necessary. We employ the term "institution" in its conventional sense, using it to denote the social patterns and arrangements through which and in which economic transactions take place. At least three unique classes or levels of institutions may be identified: (1) the overarching set of cultural values and mores which form a context for economic behavior, (2) the laws and regulations which specify the roles of the game, and (3) the contractual arrangements which are used to effect transactions. Although detailed investigations will call for significant variations in technique at each of these three levels, we argue that work at all levels can draw upon a general theoretical core. Our purpose here is to give an account of this common ground and to report on our progress in formulating an economic model of institutional change.

I. THE LITERATURE ON ECONOMIC INSTITUTIONS

We review here recent contributions to the economic theory of institutions. However, before proceeding to contemporary works, at least brief acknowledgment must be given to the major figures in the history of economic thought who have trod the ground of institutional inquiry. Without any claim of comprehensiveness, we select Karl Marx, John R. Commons, Thorstein Veblen, and Wesley C. Mitchell for comment.

Marx, in a category by himself as the founding and dominant figure of a major school of economic and political thought, obviously concerned himself with the structure of society and with the generation of a perspective which did not take institutions as exogenously given but, instead, attempted to explain class formation, the interaction of classes, and transformations from one economic order to another. Of central concern in the Marxian theory of value and exploitation are the institutions of labor exchange and the related pattern of ownership and control of the means of production. The expropriation of surplus product under various institutions of labor exchange and the mechanisms of transformation from feudalism to capitalism and subsequently to socialism represent essential lines of inquiry to the Marxist. Particular emphasis in recent years has been placed on the theory of the state (A. de Janvry, pp. 183-97; __ Jessop; and __ Holloway and __ Picciotto).

We also should not proceed without acknowledgment of the group of American economists known as the Institutionalists. The dominant personalities in this chapter in the history of economic thought were Veblen, Commons, and Mitchell. These writers shared a preoccupation with institutional variables--Veblen emphasizing cultural patterns and customs; Commons focusing on labor, industrial organization, and the legal foundations of economic transactions; and Mitchell attempting to advance the understanding of economic institutions through gathering voluminous statistical data about them. The Institutionalists reacted to a perceived narrowness of both the content and the methodology of the orthodoxy of their day; however, they failed to establish a lasting major alternative school of thought and are remembered today largely as eccentrics rather than major figures in the development of economic theory.

While the very mention of the term, institution, brings to mind the former Institutionalists, it should be clear that most modern economists who focus upon institutions would not single out the Institutionalists as intellectual ancestors of particular importance to their work nor would these modern economists, in general, seek to separate themselves from the mainstream of economic thought. A typical disclaimer is that of James Roumasset: "In explaining the existence and evolution of institutions the new institutional economics uses conventional economic tools such as benefits, costs and equilibrium. In explaining resource allocation and income distribution, the new approach uses institutions in conjunction with rather than as an alternative to neoclassical theory (pp. 1 and 2). The modern investigators are, in general, attempting to endogenize significant phenomena which, hitherto, have often been assumed as exogenously given; these researchers frequently modify and expand their theoretical and methodological "tool kits" but very rarely find themselves desiring to be set apart from the mainstream of economists. In some cases the treatment of institutions has involved a new championing of neoclassical economics.

This last point is perhaps most boldly illustrated by the work of Gary Becker in which a wide range of behavioral norms from discrimination to marriage are analyzed as market phenomena. Although Becker has come to symbolize an extreme in this regard, he is certainly not alone in his desire to understand social institutions from the perspective of neoclassical economics. For example, major contributions focusing on property rights and public choice have been made by H. Demsetz and A. Alchian (Demsetz, 1967, 1969; and Alchian and Demsetz; see, also, A. Downs and J. Buchanan and G. Tullock). In a similar vein, questions in the economic history of institutional change have been treated by Douglas North and his collaborators (__ Davis and __ North, 1970, 1971; __ North and __ Thomas; and North). Though each of these investigators (from Becker through North) has his own unique outlook and questions, they all proceed from the common point of departure that neoclassical theory--a conceptual tool kit which has been very successful in explaining market behavior--also has great usefulness in understanding nonmarket decision making. Among development economists, one attempt to forge a development theory which would adequately encompass institutional change is that of John P. Powelson. He undertakes the herculean task of explaining the process of growth through a study of institutional selection and institutional effectiveness which draws not only upon economic theory but also the work of sociologists and political scientists. Powelson, like North, covers a wide canvas; however, the theoretical underpinnings of

Powelson's work are less focused and more tentative than North's and thus have not offered a clear target for criticism or a concrete paradigm stimulating a new body of research.

As would be expected, much of the recent successful research on this topic has concentrated on one or another specific institutions in contrast to general theories of institutional change. The benefit of such work is that it is more amenable to precise formulations and empirical validation; the cost is that specialized, fragmented inquiries are seldom easily synthesized in a manner which readily yields usable generalizations about the development process. There is every reason to expect the tension between these two forms of inquiry to persist and to be a distinguishing trait of a lively intellectual process.

Of the specific institutions which have received treatment in the development literature, one of the most thoroughly studied during the past decade or so has been the set of contractual arrangements used in rural land and labor transactions. An important early contribution to this body of research was that of Steven N. S. Cheung who began a line of work in which the terms of sharecropping contracts were treated endogenously, and this field of research has naturally expanded beyond sharecropping alone. In the absence of an institutional setting offering competitive markets in land, labor, credit, and other factors of production, a variety of institutional arrangements can emerge and be used to mediate economic allocations.

We list a representative sampling of the work on contractual arrangements in agriculture. D. Newbery and J. Stiglitz have provided theoretical treatments of sharecropping with emphasis on the perfect information and risk sharing (Newbery, 1975, 1977; Stiglitz; and Newbery and Stiglitz). C. Bell and P. Zusman (1976, 1980) have considered these issues in the context of a model in which the contract between landlord and tenant is the outcome of a simultaneous dyadic bargaining process. The interlinked nature of transactions in land, labor, and credit in rural factor markets has received considerable emphasis in the work of P. Bardhan and of P. Bardhan and __ Rudra. Interlinked contracts have been given particular attention in the modeling efforts of A. Braverman and J. Stiglitz and of __ Mitra. An excellent critical survey of the literature on contractual arrangements and rural labor markets in developing countries has recently been written by H. Binswanger and M. Rosenzweig.

Writing more broadly, the dynamics of institutional change in the development process have been emphasized by Y. Hayami and V. Ruttan (see, also, Binswanger and Ruttan). This line of inquiry examines the economic inducements to technical and institutional change. Particular emphasis has been given to the process of public sector research and to the institutional structure through which change is facilitated. In contrast to the view which sees institutions as constraints to technical change and development, Hayami and Ruttan argue that "institutional reform is appropriately viewed more as a response to the new opportunities for the productive use of human and material resources opened up by advances in technology than as a precondition for agricultural development" (p. 258). In a more general sense, institutional change is portrayed as a significant element in the process of adaptation to changing economic circumstances. As costs and opportunities change, a society is faced with the challenge of altering its behavioral rules and patterns to fit new circumstances.

"Reform" may, of course, not be the word which would always apply to institutional adjustments. This would especially be so from the viewpoint of any particular agent or group who finds itself worse off rather than better off after an institutional change. Such a case is illustrated in recent research by M. Kikuchi and Y. Hayami on the compensation of landless agricultural labor in a Philippine village. They report an institutional innovation which lowered the effective wage rate. Over time, as rice yields improved and labor became more abundant, the marginal product of harvesters declined. In the absence of a fully functioning labor market, wage adjustments were made through institutional change. The shift was from the Hunusan system (in which workers received one-sixth of the harvest) to the Gama system (which is similar to Hunusan but restricts employment by requiring workers to weed fields without compensation in exchange for the right to participate in the harvest). From 1959 to 1976, there was a shift from 0 percent to 83 percent of the farmers using such a system. Another example given by Kikuchi and Hayami illustrates that such economic adjustments are not always in favor of landowners. Because of land reform regulations and other social forces, landlords were limited in their ability to raise rents. As the marginal returns to land rose, many tenants captured the surplus through various forms of subrenting.

Moving away from the development literature, we conclude this section by noting two recent contributions to the general economic literature on institutions. Both works employ a game theory believing that it offers the most fruitful framework for studying the institutional changes which emerge from the interaction of maximizing individuals. E. Thompson and R. Faith build a hierarchical model of strategic behavior which they argue has particular applicability to the choice of political-economic systems. With an eye toward developing a general model suitable for examining the evolution of property rights, systems of coordination, political organizations, and other institutions which solve recurring social and economic problems, Andrew Schotter makes a major contribution to the institutional literature. In particular, Chapter 3, "A Mathematical Theory of Institution Creation," coauthored with S. Berman, presents an n-person, noncooperative supergame in which agents make repeated choices over an infinite time horizon. Schotter wants to build a model in which the history of the play of the game impacts the choices made; this leads him to a mathematical formulation in which selected and surviving social institutions are the absorbing state of a stochastic process. Although Schotter's theory is presented in a provisional manner and "in no way purports to be fully mature," it does make a substantial contribution to its stated intent of being "a first step in an attempt to liberate economics from its fixation on competitive markets as an all-encompassing institutional framework (pages xi and 1). Schotter has developed one of the more significant formal models of institutional change to date, and we will thus have occasion to frequently refer to his work in our next section.

II. THE MODELING APPROACH

We develop here a modeling framework which is intended to be applicable to a wide class of problems encountered in the study of economic institutions. These problems possess the following characteristics in common:

1. Agents are assumed to be rational in the sense that they make choices which they believe to be best for themselves, but there exists no necessary postulate of rationality for the system as a whole.
2. Agents acquire information about their choice problems in the course of repeated decision making over time and alter their decisions accordingly.
3. Institutional changes are endogenously determined by the interaction of agents behaving in the manner described by characteristics 1 and 2.

The first and third characteristics lead us to a game-theoretic approach in which strategy options correspond to various institutional patterns. The second characteristic calls for the incorporation of appropriate adjustment schemes into our model. The result is a modeling framework sufficiently general as to encompass a variety of development phenomena but also precise enough to yield rigorous and detailed analysis.

Our approach owes much to the solution concept for n-person noncooperative games developed by John C. Harsanyi and Reinhard Selter (Harsanyi, 1975, 1977; and Harsanyi and Selten). However, our particular use of game theory departs from the Harsanyi-Selten (H-S) procedure in two significant respects. First, we render the problem in terms of repeated plays of a game in contrast to the single-play version for which the H-S solution concept is defined. Second, we allow for a wide variety of adjustment behavior instead of the unique, precisely defined solution process used by H-S. The first departure from H-S associates our model with the general approach used by Andrew Schotter; however, our second departure introduces a treatment of learning behavior not contained in other models of institutional change.

A brief outline of the H-S solution concept serves as a useful starting point. Their approach makes an important contribution to the study of rational behavior under uncertainty by extending Bayesian decision theory from the one-person to the n-person case. Within the context of noncooperative games, it offers a theoretical process by which agents attain a convergence of expectations and strategy choices. Convergence to a solution involves a "tracing procedure" which ascribes continually decreasing weight to prior expectations (first-order information) and continually increasing weight to other players' likely responses (second-order information).

It is assumed that a prior probability distribution $p_i = (p_i^1, p_i^2, \ldots, p_i^{K_i})$ will be constructed in which p_i^k ($k = 1, 2, \ldots, K_i$) represents the subjective probability that a player other than i will assign to player i using strategy k. It is assumed by H-S that all players share identical expectations of all players other than themselves and, thus, all players other than i formulate the same p_i. One may then summarize all n players' mutual expectations with the vector of prior probability distributions, $p = (p_1, p_2, \ldots, p_n)$. Player i's expectations of all other $n - 1$ players is symbolized by $p_{-1} = (p_1, \ldots, p_{i-1}, p_{i+1}, \ldots, p_n)$.

It seems reasonable to assume that a solution to the game would be chosen with reference to the prior probability distributions. However, which solution will be chosen is not obvious. One possible procedure for identifying a solution would be what H-S refer to as the "naive Bayesian approach." In this approach each player uses a strategy combination $s_i = (s_i^1, s_i^2, \ldots, s_i^{K_i})$ in which player i assigns the probability s_i^k to the <u>k</u>th pure strategy; s_i represents player i's best reply to p_{-i} (i.e., player i formuates s_i in such a way as to obtain a maximal payoff given that all other players will behave according to i's expectations as represented by

p_{-i}). Although this procedure might appear to present a reasonable solution, this is usually not the case. The strategy combination $s = (s_1, s_2, \ldots, s_n)$ generated through the naive Bayesian approach will very likely not correspond to a Nash equilibrium since there is no reason to expect all such s_i's to be best replies to each other. If s is not an equilibrium point of the game, it will not be a candidate for the solution to the game since at least one player will have incentive to move away from this point. However, s is significant in that it provides the starting point of the H-S tracing procedure.

Through use of the tracing procedure strategy combination, s is modified in a systematic and continuous manner until the equilibrium point s* is attained as the solution to the game. A linear tracing procedure defined by H-S is based on a family of auxiliary games, G^t, with $0 \leq t \leq 1$. In this new family of games, all elements of the original game, G, are retained except that the payoff function for player i in G^t now becomes

$$H_i^t(s_i, s_{-i})\ t\ H_i(s_i, s_{-i}) + (1 - t)\ H_i(s_i, p_{-i})$$

under the linear tracing procedure. In this procedure the payoff function is at all times a convex combination of $H_i^0 = H_i(s_i, p_{-i})$ which represents the payoff function which player i would maximize in the naive Bayesian approach and $H_i^1 = H_i(s_i, s_{-i})$ corresponding to the payoff function in the original game, G = G'.

The tracing procedure begins at $(t, s^t) = (0, s)$ and follows a feasible path to an end point (1, s*). In most cases the linear tracing procedure is well defined; however, when it is not, H-S use a logarithmic tracing procedure which adds an additional logarithmic term to the payoff function, H_i^t.

This later procedure is always well defined and selects the same solution point as the linear tracing procedure in those cases for which the linear tracing procedure is well defined.

Of particular interest here is the modification of expectations which is carried out in this solution process. At any point in the linear tracing procedure, player i entertains a subjective probability distribution $\pi_i(t, s)$ which may be expressed as:

$$\pi_i(t, s) = t\, s_{-i} + (1 - t)\, p_{-i}.$$

At the starting point, full weight is placed on the prior probability distribution since there exists no predictive confidence in the original, naive Bayesian s_{-i} being associated with an equilibrium point. Yet, as the process continues, player i will put more and more confidence in predictions of other players' behavior until, at the end, player i will identify s^*_{-i} with complete predictive certainty.

In the H-S solution procedure, no action is taken until the end of the process; however, at each step along the way, player i may be characterized by the strategy combination which would be chosen if an action had to be taken at moment t. That is, at any moment t, player i's strategy choice would be a best reply (in the context of the original game G) to the probability distribution $\pi_i(t, s)$. The process is one of "continually increasing predictive certainty" (Harsenyi, 1975, p. 75). The H-S process represents one method for revising strategic expectations, and it is one which will serve as a useful concept to which we will compare other procedures for revising expectations of the behavior of other agents. Looking ahead briefly, we will see that, although the H-S tracing procedure possesses very desirable mathematical properties which always assure a solution, a good many applications do not permit the assumption of the adjustment behavior upon which the H-S solution concept rests. This in no way detracts from the important role which the H-S solution concept plays in establishing a theoretical approach to a large group of problems; it simply means that particular applications will require departures from the purely conceptual model presented by H-S.

Starting from this point of departure, a first modification is to transform the problem from a static to a dynamic one. Although H-S use the symbol t and refer to stages of the solution process with low values of the parameter t as being "former" to stages with higher t values, this time dimension is purely a mathematical device, and the entire tracing from $t = 0$ to $t = 1$ takes place within a timeless instant in much the same way that a Walrasian auctioneer considers bids and offers in a timeless process before the opening of a market.

Schotter introduces time by instituting repeated plays of the same game so that t now has a different meaning and takes on integer values conventionally indexing each play of the game. As in the H-S procedure, Schotter begins with prior probability distributions of the form $p_i = (p_i^1, p_i^2, \ldots, p_i \ )$ which express the expectations which all other players share of player i's likely choices from the k_i available strategies. However, Schotter imposes the assumption that all such prior probability distributions are initially uniform, and, thus, $p_i = (1/(k_i), 1/(k_i), \ldots, 1/(k_i)$ with all strategies having equal probability. (We will return to this assumption later. Although it may be reasonable in the context of a game where all players have complete mutual ignorance, we regard this as one extreme special case rather than the general case to be used in all models of institutional change.)

At the beginning of Schotter's process, players formulate the equivalent of H-S's naive Bayesian solution; for each player i, a best reply, s_i, is constructed in response to the expectations reflected in the prior probability distributions, p_{-i}. Although the s_i's making up $s = (s_1, s_2, \ldots, s_n)$ will, in general, correspond in mathematical form to mixed strategies, players do not "play" mixed strategies. In the actual play of the game, only a pure strategy can be played. Thus, for each player i, s_i is viewed as a probability distribution governing a random experiment in which strategy $k_{i,t}$ is selected to be executed at time t. Each player so selects a pure strategy, all players simultaneously carry out these strategies, and the first play of the game is then complete.

As was the case in the H-S solution concept, Schotter's model derives its behavioral characteristics from the assumptions which are imposed on the adjustments which players make upon receiving new information. Going into the second round of play, the prior probability distributions are changed according to an updating scheme specified in the model. The particular updating scheme used by Schotter gradually increases the expectation of those strategies which are used and gradually decreases the expectation of strategies not used. Precisely, if k strategies are available to player i and strategy k has been played by player i at time t, all other players will revise their expectations of i in the following manner:

$$p^k_{i,t+1} = p^k_{i,t} + \varepsilon$$

$$p^j_{i,t+1} = p^j_{i,t} - \frac{\varepsilon}{k - 1} \qquad j \neq k$$

where ε is an arbitrarily small, positive, fixed number. Similar to the assumption made in H-S, Schotter simplifies the problem by assuming that all players, other than i, look at i's behavior in the same way. Thus, they all share the same original expectations about i, use the same updating scheme, and consequently entertain precisely the same revised expectations about i's behavior in the next play of the game.

Without presenting in detail the complete structure of Schotter's model, its essence is captured in the updating scheme given above. Through repeated plays of the game and repeated application of the updating scheme, the behavior of the model is worked out. A solution is reached when all players come to expect the same strategy n-tuple to be played with probability one. This is analogous to the H-S solution in that the solution must be an equilibrium point; however, it differs from H-S in that the Schotter model need not always select the same solution. Instead, Schotter attempts to depart from a deterministic model by making the choice of a solution dependent upon a path which is determined by a stochastic process.

Schotter considers this framework to be a way of incorporating historical dynamics into the model. He emphasizes that the same sequence of repeated plays will not likely yield the same result. In his own words, ". . . which stable institutional arrangement emerges to solve a particular problem is a random event that depends crucially . . . on the history of the play of the game and the process of norm [or expectation] creation" (Schotter, p. 53). Although this is technically true throughout Schotter's analysis, it is not difficult to show that in practice the procedure need not perform stochastically and that in many problems it will result in the same deterministic solution point as in the H-S tracing procedure. We illustrate this below for a simple, but very general, class of games.

Although the model presented by Schotter in sections 3.1-3.6 (pp. 54-79) looks at first blush to be quite different from the class of games discussed below, we illustrate in the Appendix that there is a great deal of similarity between the two and that, in fact, Schotter's model in sections 3.1-3.6 is one with a best reply structure of the same form as that analyzed below. Unfortunately, Schotter's specification of the game obfuscates the underlying nature of the problem. Consequently, we do not deal here with his entire model but instead consider only his solution procedure.

The equivalence of an approach, such as the one used by Schotter with the H-S concept for a simple class of two-person nonzero-sum games, can be illustrated by use of a tracing map (Harsanyi, 1975, pp. 92 and 93; Harsanyi and Selten, pp. 3/66-3/70). Such games have a payoff matrix as shown in Figure 1 with α, β, γ, $\delta > 0$. The three equilibrium points for all games of this class are (the first two are pure strategies and the third is a mixed strategy equilibrium):

$$E_1 = (A, X)$$

$$E_2 = (B, Y)$$

$$E_3 = (q_1^*, q_2^*)$$

with

$$q_1^* = \left(\frac{\delta}{\beta + \delta}, \frac{\beta}{\beta + \delta}\right)$$

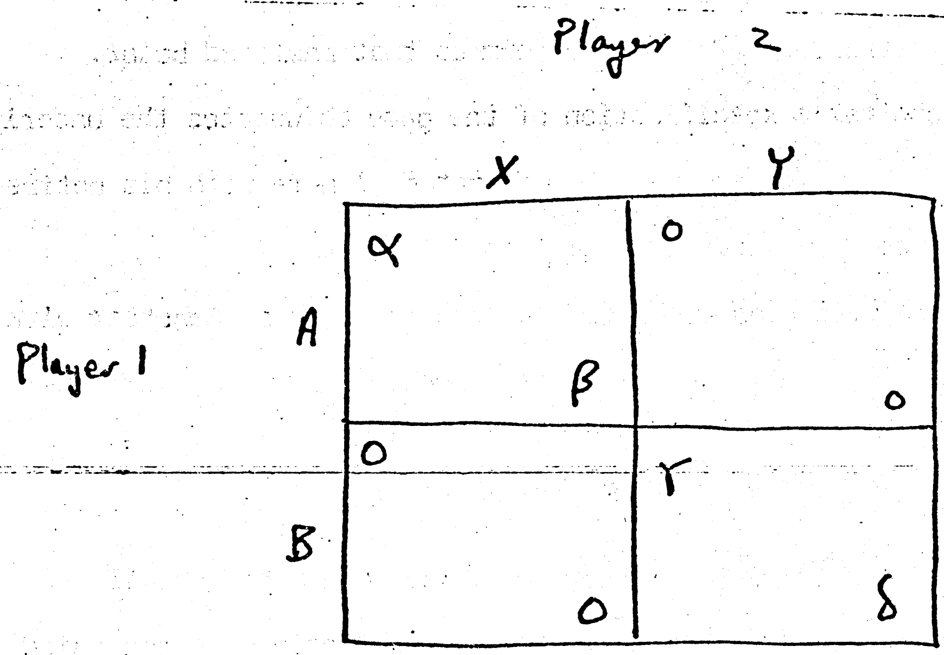

Figure 1

and

$$q_2^* = \left(\frac{\gamma}{\alpha + \gamma}, \frac{\alpha}{\alpha + \gamma}\right).$$

In Figure 2, H-S construct a tracing map for such games.

The broken line, Mq*N, represents a border between the regions which H-S call "source sets." The area below Mq*N is the source set of equilibrium point E_2 in the sense that the H-S tracing procedure will select E_2 as the solution for situations in which the prior vector, p, lies below Mq*N. Likewise, if p lies above Mq*N, it is in the source set of E_1, and this equilibrium point will be selected as the solution.

When p lies on the boundary, Mq*N, the logarithmic tracing procedure must be used, and the results depend upon the location of q*. In brief,

1. If q* and, thus, Mq*N lies above the diagonal BXAY (i.e., $\alpha\beta < \gamma\delta$), E_2 is the solution.
2. If q* and, thus, Mq*N lies below the diagonal BXAY (i.e., $\alpha\beta > \gamma\delta$), E_1 is the solution.
3. If Mq*N coincides with the diagonal BXAY (i.e., $\alpha\beta = \gamma\delta$), E_3 is the solution (Harsanyi and Selten, p. 3168).

Thus, given the tracing map, one can directly identify the solution to the game once the prior vector is known.

A tracing map so constructed provides a vehicle for making illustrative comparisons between the solution procedures used by H-S and Schotter. It will be recalled that Schotter always begins with a prior vector of uniform probabilities. In this case, we then have p being represented by the point in the very center of the tracing map with each player assigning probabilities of

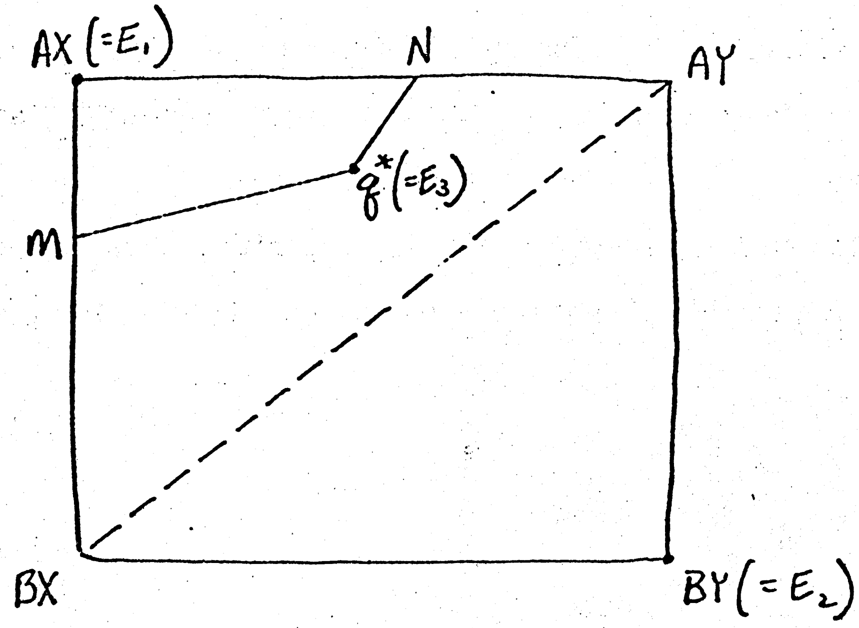

Figure 2.

one-half to the other player using one or the other strategy. To construct an example, let us assume that $\alpha\beta < \gamma\delta$ so that Mq*N will lie above the diagonal BXAY. Starting at p, we can then draw the paths which would be followed in the timeless mental process of the H-S solution concept and in the repeated game using the updating scheme specified by Schotter. We already know from the construction of the tracing map that the H-S solution will be E_2 since p lies in the source set of E_2; and in this case an updating scheme, such as Schotter's, will lead to the same solution.

Figure 3 shows the paths followed by the revision of probabilities in both decision procedures starting from p and following to the equilibrium point E_2. The path following from an updating scheme, such as Schotter's, moves from P to T to V to BY. Probabilities in the H-S scheme also move from P to T but then jump discontinuously to U followed by a continuous move to BY. The small arrows of the Schotter path represent the short, discrete jumps of $\varepsilon\sqrt{2}$ length; as $\varepsilon \to 0$, this path approximates a continuous one. The regularity of the Schotter path, in this case, arises from the fact that best replies are always in pure strategies and, thus, the outcomes of the experiments during each round predictably select the pure strategy best replies.

The above diagram may be used to develop an intuitive sense of the circumstances under which the two adjustment concepts will not lead to the same solution. Furthermore, it will be shown that such circumstances arise only in a degenerate case for which the Schotter path would never reach a stable equilibrium point. Thus, at least for this class of games, when both procedures choose a solution, they choose the identical solution.

An explanation of the equivalence of the solutions obtained by the H-S procedure and by the Schotter procedure will be facilitated by the

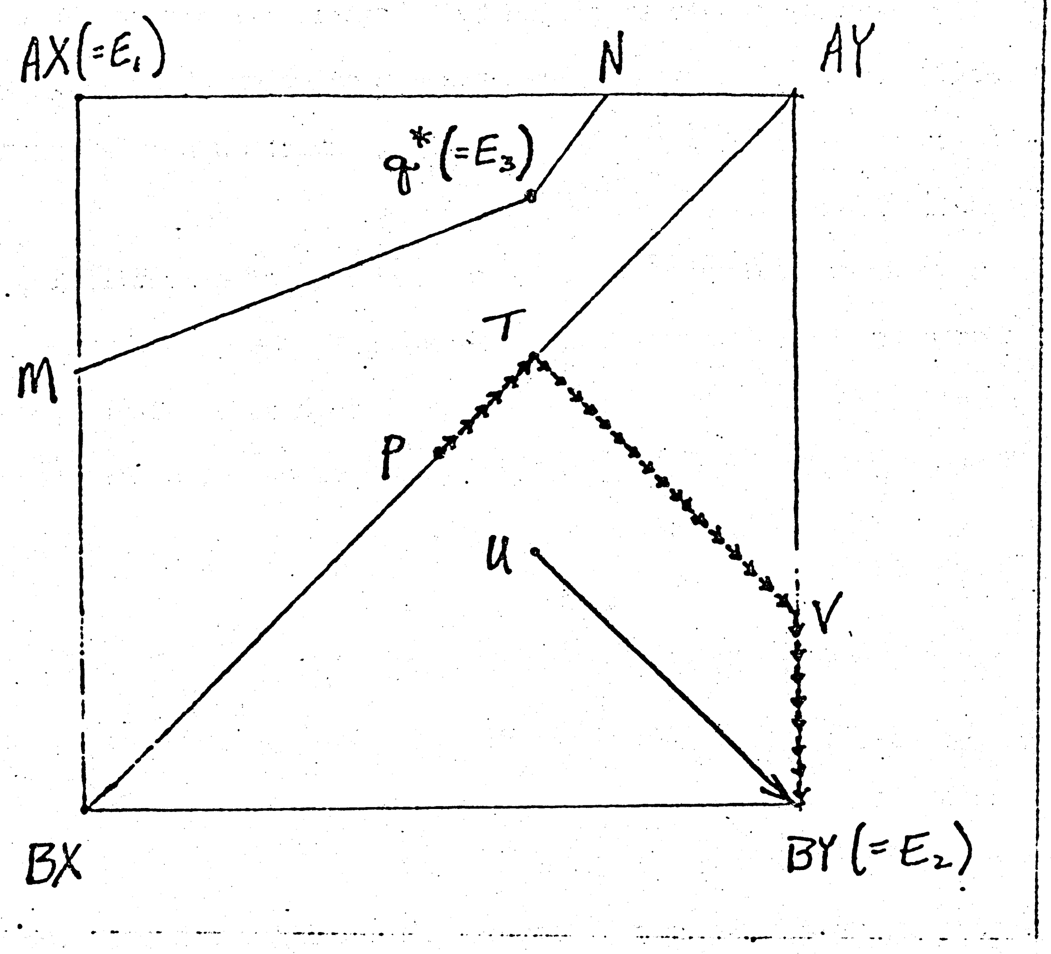

Figure 3

introduction of stability diagrams (Harsanyi and Selten, pp. 2/14-2/18). The stability diagram is constructed in the same space as the tracing map but illustrates a different aspect of the problem. Namely, the stability diagram delineates the best reply structure of the game. The stability diagram for the class of 2 x 2 games under consideration here takes the form shown in Figure 4. Any point in this space may be interpreted as a mixed strategy combination; each rectangular region is labeled with the best reply strategy to any mixed strategy combination falling within the region. For example, AX is the best reply combination to point e; AY (not BX) is the best reply combination to point f. Borderline points represent an indifference between two possible best reply combinations; and the point q* is a best reply to itself. However, the slightest deviation from q* will shift the best reply to one of the four possible pure strategy combinations.

Figure 5 superimposes the tracing map on the stability diagram. A simple heuristic argument will now serve to demonstrate the equivalence of the H-S and Schotter solutions. As long as Mq*N is above the diagonal BXAY, there is no possibility that AX will be a best reply to p or to any p' revised by use of the ε-adjustment scheme. In other words, all movement of expectations will be restricted to the triangle BXAYBY; repeated plays of the game will eventually lead to BY(= E_2), the same unique equilibrium selected by H-S. By a similar argument, if Mq*N lies below diagonal BXAY, the Schotter procedure will select E_1 as in the H-S case.

Finally, if Mq*N coincides with the diagonal BXAY and, thus, the procedure begins with q* = p, the H-S procedure will select E_3 as the equilibrium point. However, the Schotter process, by definition, does not admit a mixed strategy equilibrium point. Instead, on the first play of the game, each of

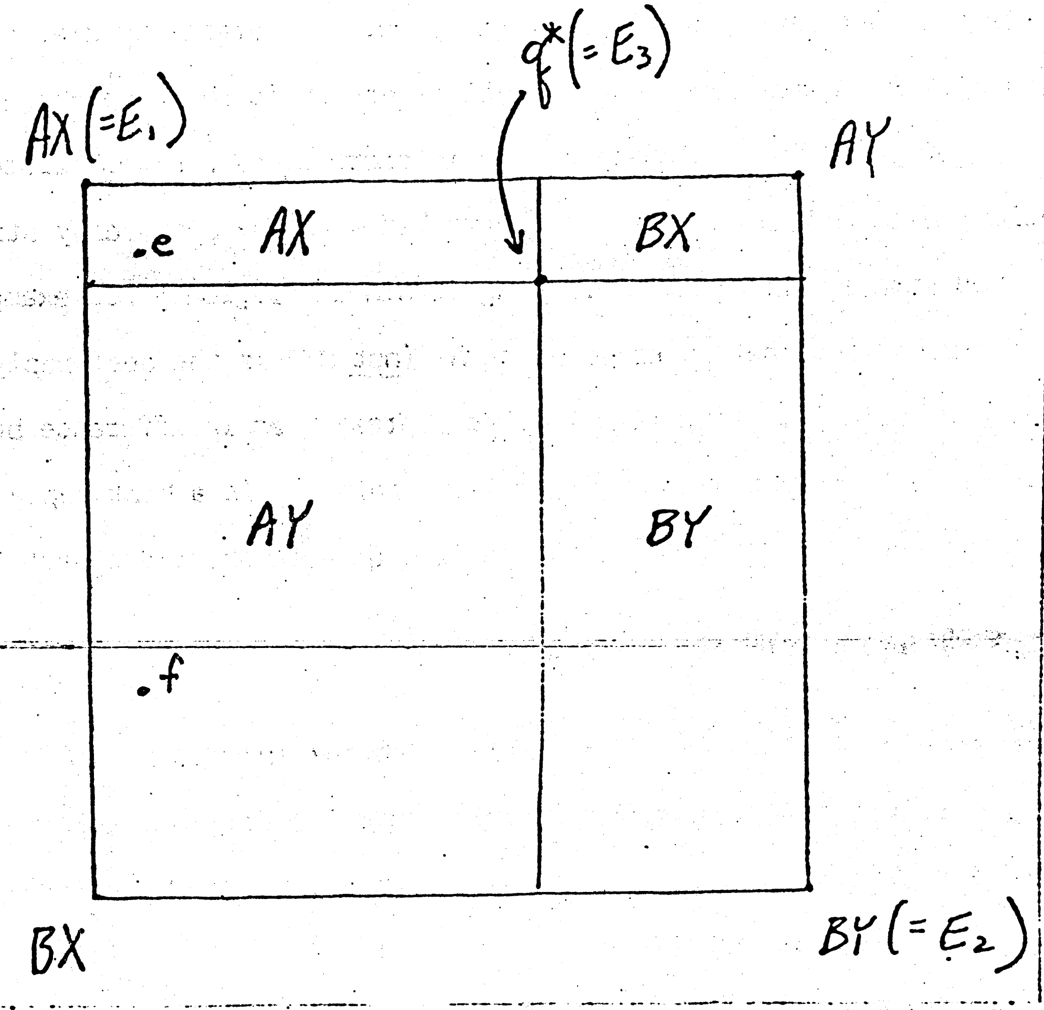

Figure 4.

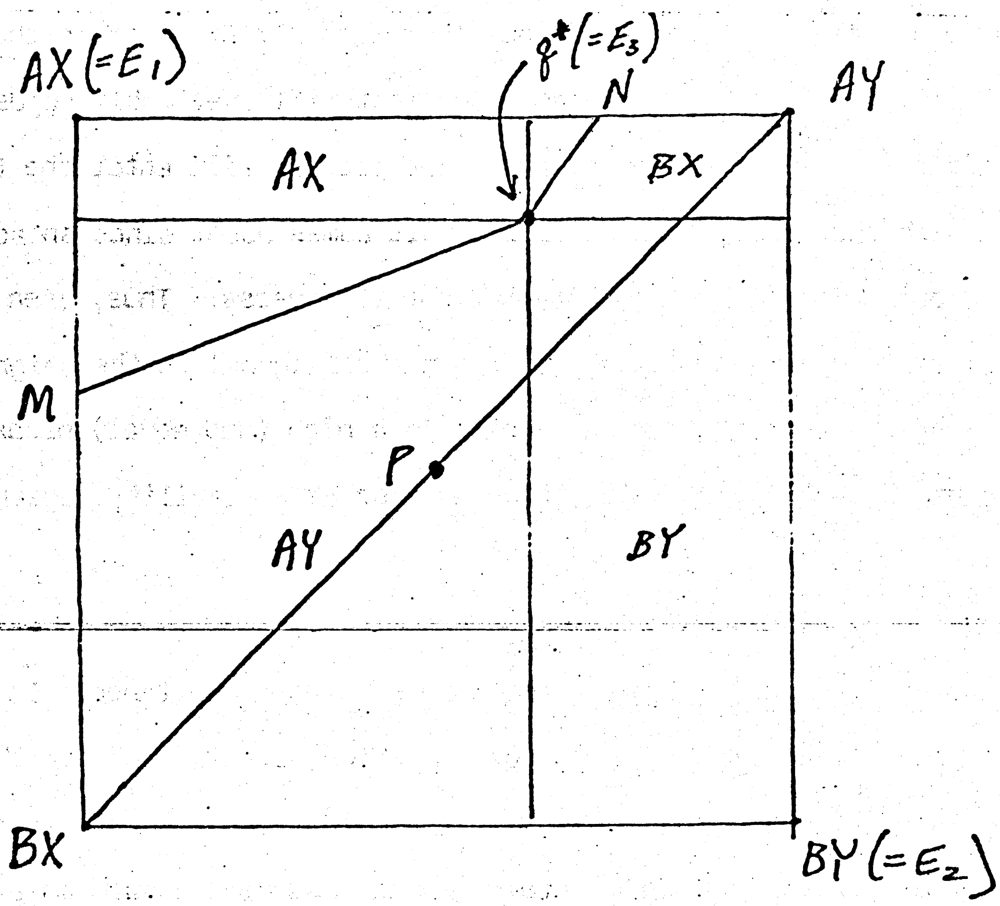

Figure 5.

the four pure strategy combinations will have a probability of one-fourth. If either AX or BY is selected, the process will continue to move in the direction determined on the first move and continue until the pure strategy equilibrium point is reached. On the other hand, if BX or AY is the outcome of the first experiment, the second experiment will inevitably adjust expectations back to the original p and the process will enter the third round as if it were the first; this unusual result comes about since an adjustment toward AY puts p' in BX's stability set and vice versa. Thus, even if there is considerable movement back and forth on the diagonal in the neighborhood of p at the beginning of the process, there is a high (and equal) probability that the process will jump into either AX's or BY's stability region and eventually end up at one of these points.

Although the two processes behave slightly differently when beginning at $p = q^*$, the difference does not seem to be of great importance. In the H-S case, we have a weak equilibrium at $p = q^*$ which, if disturbed in the slightest way, will move toward AX or BY. In the Schotter case, a specific pattern of disturbance is built into the model and will assure an eventual disturbance moving the process toward AX or BY. In sum, the two procedures behave here in essentially the same manner; and, contrary to Schotter's intention, the selection of stable institutional arrangements (i.e., a solution point) is deterministic. The apparent randomness in the Schotter procedure was not obtained from the solution procedure used but from a contrived game formulation dependent upon unfortunate behavioral assumptions.

We present here an alternative solution procedure designed to achieve the desired properties. We introduce new schemes for both the selecting of strategies and the updating of expectations. With regard to strategies, our model assumes that a new best response will not be fully implemented in one period and that beginning strategies will be in the form of some reasonably formulated mixed strategies rather than best response pure strategies. Thus, there will initially be a positive probability of movement in any direction, and these strategies will be gradually altered during play. However, this alone will not be enough to avoid a deterministic outcome in most cases if $\varepsilon \to 0$ since, although "errors" could be made, they would be quite small and the likelihood of the vector of expectations "jumping" into another source set would be quite low. We would still be left with a very accurate, nearly deterministic procedure in terms of the eventual outcome of the choice process. Thus, we also need to allow for more rapid adjustment of expectations.

In contrast to other procedures, we are arguing for more consistent and reasonable assumptions about the behavior of players when they are choosing strategies and when they are formulating expectations about other players. There seems to be little, if any, justification for a scheme which permits adjustments (in a single period) as dramatic as a shift from probability zero to probability one in one's own strategy choices but constrains alterations of expectations about other players' behaviors to an arbitarily small ε. Adjustment speeds need not be the same in both cases, but we suggest that slower adjustment in the first case and faster adjustment in the second case would be a more reasonable assumption in many applications.

In the formulation of what we consider to be more reasonable adjustment schemes, we draw especially upon stochastic learning theory. Here, learning means "the development, either conscious or unconscious, of a standard pattern of responses in particular situations" (J. Cross, 1973, p. 240). Learning behavior has been investigated by psychologists, engineers, mathematicians, and economists. We briefly review this literature here. A more extensive treatment, especially of the engineering literature, is contained in T. Head.

The pioneering work on learning models in mathematical psychology was done by P. Bush and F. Mosteller. Subsequent studies included treatments of the topic by R. Suppes and R. Atkinson and R. Atkinson, G. Bower, and E. Crothers extending the analysis to group interaction and oligopoly. Much of the work by engineers has been done in the Soviet Union (e.g., M. Tsetlin, 1961; Ya. Tsypkin, 1971 and 1973; and M. Tsypkin and A. Poznyak). Tsetlin's work (1973) has included applications to biological systems. In this country, K. Narendra and associates have made many contributions to the literature including a major survey (K. Narendra and M. Thathachar). Also noteworthy is the volume of articles edited by K. Narendra which contains an extensive "Bibliography on Learning Automata," compiled by S. Lakshmivavahan. K. Fu and associates have treated learning automata extensively within the context of adaptive control theory and pattern recognition (e.g., see K. Fu, 1970 and 1971).

The general subject of stochastic learning systems has received rigorous mathematical treatment by M. Iosifescu and R. Theodorescu and M. Norman. A principal investigator among economists has been John G. Cross whose 1973 Quarterly Journal of Economics article presents a general theory of stochastic learning in economic behavior and later works apply that theory to irrigation (Cross, 1978) and to consumer behavior (Cross, 1979). R. Schmalensee evaluates the applicability of stochastic learning models to both firm and household choices under uncertainty, and S. Himmelweit formulates a production model based on stochastic learning.

The essence of stochastic learning models is represented in the reinforcement or adjustment schemes used. In general, when an action is successful in a given period, the probability of selecting that action in the next period is increased, and the probabilities of all other actions are decreased. We illustrate here reinforcement schemes which conform to a simple linear form. "Success" may be measured discretely (e.g., 1 for success, 0 for failure) or continuously (e.g., any real number in the interval [0, 1]). For example, if choice i is successful in period t, next period's action probabilities will be:

$$p_{k,t+1} = (1 - \theta)\, p_{k,t} + \theta(1)$$

$$p_{j,t+1} = (1 - \theta)\, p_{j,t} + \theta(0) \qquad j = 1, \ldots, K \quad j \neq k$$

$$0 < \theta < 1.$$

The learning parameter, θ, is the weight placed on this period's experience; the larger is θ, the more rapid is "learning" or adjustment. Obviously, trade-offs between speed and accuracy can arise when feed back or payoffs vary from period to period.

When feedback or payoffs are continuous, the updating scheme would be:

$$p_{k,t+1} = (1 - \delta)\, p_{k,t} + \delta(\lambda_{k,t})$$

$$p_{j,t+1} = (1 - \delta)\, p_{j,t} + \frac{\delta}{K - 1}(1 - \lambda_{k,t}) \qquad j = 1, \ldots, K \quad j \neq k$$

$$0 < \delta < 1$$

$$0 \leq \lambda_{k,t} \leq 1.$$

If one allows the modification that information may be gathered or formulated on each of the k alterntives in each period, this last updating scheme may be rewritten as:

$$p_{k,t+1} = (1 - \delta)\, p_{k,t} + \delta(\lambda_{k,t}) \qquad k = 1, \ldots, K$$

$$0 < \delta < 1$$

$$0 \leq \lambda_{k,t} \leq 1, \ \sum_{k=1}^{K} \lambda_{k,t} = 1.$$

And the discreet reward scheme given above is simply a special case of this general updating algorithm with $\lambda_{k,t} = 1$ and $\lambda_{j,t} = 0$, $j \neq k$.

With this, we can now move directly to a statement of our solution procedure. It is composed of the following elements.

A.

In general, we consider these to be historically given; that is, most applications begin with information about what players are doing at the start of the time period under consideration. Where this is not the case and there is no other evidence in favor of a particular distribution of prior probabilities, we accept the same specification which Schotter uses (i.e., $p_i^k = 1/K_1$).

B.

Following Schotter's notion of increasing the expectation of observed behaviors, a strategy used by a player in a given period will translate into a feedback of "1" to all other players, and all strategies not choosen will be given a "O." However, rather than his ε-adjustment scheme, we employ the following rule when k is selected by player i in period t:

$$p^k_{i,t+1} = (1 - \theta)\, p^k_{i,t} + \theta$$

$$p^j_{i,t+1} = (1 - \theta)\, p^j_{i,t} \qquad j = 1, \ldots, K^i, \quad j \neq k$$

$$0 < \theta < 1.$$

C.

In contrast to other schemes which permit complete adjustment to new best responses in each period, we envision that strategy choices are also altered gradually according to the following scheme:

$$s^{k\,**}_{i,t+1} = (1 - \delta)\, s^{k\,**}_{i,t+1} + \delta(s^k_{i,t+1}) \qquad k = 1, \ldots, K^i$$

$$0 < \delta < 1.$$

In this scheme, best response strategies are calculated with respect to the expectations formulated by our rule above; however, at each period players cannot make a complete adjustment of s_i.

The procedure outlined here makes plausible the phenomenon of jumping from one source set to another (see Figures 2 and 5) and thus avoids the deterministic behavior of other procedures. The adjustment patterns incorporated in the model are based upon a reasonable behavioral theory and allow for consistent modeling of all adjustment behavior. Summarizing, we have recapitulated in Table 1 the essential elements of the solution procedure discussed in Section II.

TABLE 1

Comparison of Solution Procedures

	Type of game	Initial expectations	Rule used for updating expectations	Rule used for strategy choice in response to expectations
Harsanyi/Selten	static, one time	not generally specified	$\pi_i(t_1\ s) = ts_{-i} + (1 - t)_{p_{-i}}$	best response s_i
Schotter/Berman	dynamic, infinitely repeated game	equiprobable, $p_i^k = \frac{1}{k_1}$	$p^k_{i,t+1} = p^k_{i,t} + \varepsilon$ $p^j_{i,t+1} = p^j_{i,t} - \frac{\varepsilon}{k_i - 1}$, $j \neq k$ $\varepsilon \rightarrow 0$	best response s_i; s_i viewed as a probability distribution governing a random experiment to select pure strategy.
Adelman/Head	dynamic, infinitely repeated game	historically determined; if no basis for specification, $p_i^k = \frac{1}{k_i}$	If player i uses k at t, $p^k_{i,t+1} = (1 - \theta)\ p^k_{i,t} + \theta$ $p^j_{i,t+1} = (1 - \theta)\ p^j_{i,t}$, $j \neq k$ $0 < \theta < 1$	best response s_i is formulated by i, but selection of pure strategy is governed by probabilities calculated by the following: $s^{k}_{i,t+1}{}^{**} = (1 - \delta)\ s^{k**}_{i,t} + \delta\ (s^k_{i,t+1})$ $0 < \delta < 1$

APPENDIX

Schotter's analysis focuses largely on a super game with the following 2 x 2 constituent prisoners' dilemma game (Figure A.1). He uses this constituent game to formulate 4 x 4 super game in which each player chooses a "mode of behavior" or super game strategy. Figure A.2 represents the payoffs (appropriately discounted) in this new game. The strategy $[a_1^1/a_2^1]$, for example, is the mode of behavior which stipulates that player 1 will play the cooperative strategy a_1^1 as long as player 2 plays cooperative strategy a_2^1; otherwise, player 1 will play a_1^2. It should be noted that symmetry is not maintained because the players are assumed to have different discount rates. Schotter then continues his analysis with this new matrix; his choice to do so is regrettable because it needlessly complexifies the presentation.

Two obvious simplifications are in order. First, dominated strategies can be eliminated. That is, if a choice is under all circumstances worse than alternative choices, it can reasonably be eliminated from the strategy set under consideration. In the present case, this would eliminate the second row, player 1's strategy $\sigma^1[a_1^1/a_2^2]$, and the third column, player 2's strategy $\sigma^2[a_2^1/a_1^2]$. This eliminates the obviously inferior behavior (see Figure A.1) of cooperating even when it is obvious that the other player is being noncooperative. Second, in the normal form of the game, there is no difference in the payoffs associated with the two noncooperative modes of behavior available to each player. This occurs because they do not represent different modes of behavior; that is, Agent 2

		Strategy a_2^1 (cooperative)	Strategy a_2^2 (noncooperative)
Agent 1	Strategy a_1^1 (cooperative)	8, 8	0, 9
	Strategy a_1^2 (noncooperative)	9, 0	6, 6

Figure A.1

Player 2

Player 1	$\sigma^2[a_2^1/a_1^1]$	$\sigma^2[a_2^2/a_1^1]$	$\sigma^2[a_2^1/a_1^2]$	$\sigma^2[a_2^2/a_1^2]$
$\sigma^1[a_1^1/a_2^2]$	32, 16	18, 15	21.5, 15.5	18, 15
$\sigma^1[a_1^1/a_2^2]$	28.25, 11	0, 18	26, 14	0, 18
$\sigma^1[a_1^1/a_2^1]$	27, 6	24, 12	32, 0	24, 12
$\sigma^1[a_1^2/a_2^2]$	27, 6	24, 12	27, 6	24, 12

Figure A.2

$$\sigma^1\left[a_1^2/a_2^1\right] \Longleftrightarrow \sigma^1\left[a_1^2/a_2^2\right]$$

$$\sigma^2\left[a_2^2/a_1^1\right] \Longleftrightarrow \sigma^2\left[a_2^2/a_1^2\right].$$

Thus, we can simplify the normal form game to the following 2 x 2 matrix in Figure A.3. And, for purposes of analyzing the best reply structure of the game, we may go one step further and note that the payoff in Figure A.3 may be transformed into the matrix in Figure A.4. This transformation presents the best reply structure of the game.

The matrix in Figure A.4 is a member of the class of games discussed in detail in the illustrations of Section II. It should additionally be noted that this 2 x 2 representation is the only form consistent with Schotter's norm-updating rule which depends upon a one-period observation of the form (a_1^k, a_2^k), k = 1, 2. That is, after each play of the constituent game, players make inferences regarding which modes of behavior are being used. An index which takes on only four possible values obviously cannot be used to distinguish the 16 possible strategy combinations of the 4 x 4 game. This leads to serious confusion in the implementation of Schotter's model and contributes to our reluctance to deal further with the entire model even though we consider the solution procedure itself to be an important step in the development of game-theoretic models of institutional change.

Agent 2

		Cooperative	Noncooperative
Agent 1	Cooperative	9, 1	0, 0
	Noncooperative	0, 0	6, 6

Figure A.4

		Agent 2	
		Strategy $\left[a_2^1/a_1^1\right]$ (cooperative)	Strategy $\left[a_2^2\right]$ (noncooperative)
Agent 1	Strategy $\left[a_1^1/a_2^1\right]$ (cooperative)	36, 16	18, 15
	Strategy $\left[a_1^2\right]$ (noncooperative)	27, 6	24, 12

Figure A.3