

Learning To Be Imperfect: The Ultimatum Game*

JOHN GALE

*Department of Economics, University of Wisconsin, 1180 Observatory Drive,
Madison, Wisconsin 53706*

KENNETH G. BINMORE

*Department of Economics, University College London, Gower Street,
London WC1E 6BT England*

AND

LARRY SAMUELSON

*Department of Economics, University of Wisconsin, 1180 Observatory Drive,
Madison, Wisconsin 53706*

Received August 22, 1994

This paper studies interactive learning processes that are subject to constant perturbations or "noise." We argue that payoffs in the Ultimatum Game are such that responders are more apt to be "noisy" than are proposers and show that as a result the learning process readily leads to outcomes that are Nash equilibria but not subgame-perfect. We conclude that game theorists should not restrict attention to the subgame-perfect equilibrium when predicting laboratory behavior in the Ultimatum Game. *Journal of Economic Literature* Classification Numbers: C70, C72. © 1995 Academic Press, Inc.

* Financial support from National Science Foundation Grant SES-9122176 and the Deutsche Forschungsgemeinschaft, Sonderforschungsbereich 303 at the University of Bonn, is gratefully acknowledged. We thank Drew Fudenberg and Joseph Harrington for helpful discussions, and thank two referees and an associate editor for helpful comments. We are grateful to the Department of Economics at the University of Bonn and the Institute for Advanced Studies at the Hebrew University of Jerusalem, where part of this work was done, for their hospitality. The authors first encountered the possibility of using the replicator dynamics to obtain unusual results in the Ultimatum Game in a manuscript of James Andreoni and John Miller. This paper was prepared for the Nobel Symposium at Björkborn, Sweden, 18-20 June 1993.

1. THE ULTIMATUM GAME

Consider two players with a dollar to divide. The rules of the *Ultimatum Game* specify that player I begins by making an offer of $x \in [0, 1]$ to player II, who then accepts or refuses. If player II accepts, player I gets $1 - x$ and player II gets x . If player II refuses, both get nothing. Traditional game theory predicts that the play of this game will result in the unique subgame-perfect equilibrium in which player II plans to accept whatever she is offered and player I offers player II nothing.¹

In the first of many experiments on this and related games by numerous authors, Güth *et al.* (1982) found that the modal offer was $\frac{1}{2}$ and that player I had roughly half a chance of being rejected if he offered about $\frac{1}{3}$ of the sum of money available. Binmore *et al.* (1989) reported qualitatively similar results in their replication of the Ultimatum Game experiment. There have been many related studies in the interim, surveyed by Bolton and Zwick (1993), Güth and Tietz (1990), Roth (1994), and Thaler (1988).

Critics of traditional game theory have quoted these results (along with the early results on the finitely repeated Prisoners' Dilemma and games involving the private provision of public goods) as demonstrating that the optimizing paradigm on which game theory is based is fundamentally mistaken. Instead, so the story goes, people simply honor whatever social norm is appropriate to the situation. Frank (1988) is particularly eloquent on this subject. In bargaining games, for example, it is popular to assert that people "just play fair."

Many game theorists have responded by dismissing laboratory results as irrelevant to actual behavior. We agree that the results of poorly designed experiments are irrelevant. Binmore (1992, p. 51) stresses that an experimentalist or game theorist should be cautious about making predictions unless the following criteria are satisfied:²

- The game is reasonably simple;
- The incentives are adequate;
- Sufficient opportunity for trial-and-error learning is provided.

¹ If offers must be made in whole numbers of cents, other subgame-perfect equilibria also exist, but player II never gets more than one cent in any of these.

² As experimental techniques in economics have become increasingly sophisticated, the importance of these factors has come to the fore. As advocated by Smith (1991), it is now commonplace to offer experimental subjects large incentives instead of the negligible amounts considered appropriate by many psychologists. At the same time, the introduction of computer technology has made it possible to use interactive demonstrations to teach subjects the rules of the game quickly and efficiently and to give the subjects the experience of large numbers of repetitions of the game. A survey to Ledyard (1992) of recent experiments concerning the private provision of public goods is revealing. In experiment after experiment, subjects are reported to approach the game-theoretic equilibrium as the incentives increase and the subjects' experience with the game becomes extensive.

On the other hand, game theorists cannot ignore experiments that persistently refute their predictions when all three criteria are satisfied. In the case of the Ultimatum Game, the relevant experiments have been replicated too often for doubts about the data to persist. A theory predicting that real people will use the subgame-perfect equilibrium in the Ultimatum Game is therefore open to question.

At first glance, the case for subgame-perfection in the Ultimatum Game seems ironclad. This is a two-player game of perfect information in which each player moves only once.³ Player I need only believe that player II will not play a weakly dominated strategy to arrive at the subgame-perfect offer. But the deletion of weakly dominated strategies is an *eductive* principle (cf. Binmore (1987, 1988)), whereas we believe that the principles to which one must appeal when predicting actual behavior, in the laboratory or elsewhere, are almost always *evolutive* in character. That is to say, the outcomes we observe are not the product of careful reasoning but of trial-and-error learning.

This paper demonstrates that interactive learning processes readily lead to outcomes in the Ultimatum Game that are Nash equilibria but not subgame-perfect.⁴ We argue that game theorists were therefore wrong to put all their eggs in the subgame-perfect basket when predicting laboratory behavior in the Ultimatum Game. A case exists for predicting that interactive learning will result in the selection of one of the other Nash equilibria of the game.

Section 2 begins by showing that if the initial conditions are not too close to the subgame-perfect equilibrium, then the replicator dynamics can converge to Nash equilibria in the Ultimatum Game that are not subgame-perfect. Given the relation between trembles and subgame-perfect equilibria (Selten (1975)), such a result is of interest only if it is robust in the presence of relevant perturbations. We therefore introduce noise into the replicator dynamics. When this noise is small in absolute terms but relatively larger in the population of responders, we find that asymptotic attractors survive which are Nash equilibria but not subgame-perfect. But why do we expect responders to be noisier than proposers? Our reason is to be found in the structure of the Ultimatum Game. When noise levels are allowed to depend on the potential cost of making an error, the system

³ In particular, the criticism that subgame-perfection calls for players to regard their opponents as perfectly rational after having received evidence to the contrary (cf. Binmore (1987/1988)) has no force in the Ultimatum Game.

⁴ This finding is not without precedent. Binmore (1990), Samuelson (1988, 1993, 1994), and Samuelson and Zhang (1992) give simple examples showing that the deletion of weakly dominated strategies is at best a dubious activity in an evolutionary context.

can endogenously produce a situation with more noise in the responding population.

Section 3 explains how we think these calculations should be interpreted in the light of the experimental data. In the process, we comment on a complementary learning-based analysis of the Ultimatum Game due to Roth and Erev (1993). We also explain why we expect to see initial conditions in laboratory experiments that lead the dynamics to equilibria that are not subgame-perfect. In particular, we suggest that initial play reflects decision rules that have evolved in real-life bargaining situations that are superficially similar to the Ultimatum Game. These bargaining games generally feature more symmetric allocations of bargaining power than the Ultimatum Game, yielding initial play in Ultimatum Game experiments that need not be close to the subgame-perfect equilibrium.

Section 4 discusses how we think (out-of-equilibrium) behavior can persist in which people “leave money on the table.” This section also comments on the use of “fairness” explanations for the outcomes of Ultimatum Game experiments.

The results of Section 2 are established by numerically computing trajectories for the replicator dynamics. In order to provide some insight into the forces that drive these results, Section 5 studies a variant of Selten’s (1978) Chain-Store Game, which we reinterpret as a two-offer simplification of the Ultimatum Game. In this simpler setting, an analytic study of the evolutionary dynamics is possible. The same analysis also provides a possible resolution of the well-known chain-store paradox that does not require incomplete information assumptions and applies even when only one potential entrant exists.

We employ the replicator dynamics throughout. Why are such biologically motivated dynamics relevant? First, Section 6 presents an aspiration-level model of learning that leads to the replicator dynamics. Börgers and Sarin (1993), Binmore and Samuelson (1993), Cabrales (1993), and Schlag (1994) similarly present learning models that lead to the replicator dynamics, suggesting that the replicator dynamics are of more than merely biological interest. Second, the analysis of Section 5 isolates the smoothness properties of the learning model that drive our results, revealing that qualitatively similar results will hold in a wide variety of learning models (including variants of Roth and Erev’s (1993) model).

2. NUMERICAL CALCULATIONS

This section studies a version of the Ultimatum Game. The players must split a “pie” of size 40. The set of offers available to the proposer

is $I = \{1, 2, \dots, 40\}$.⁵ (Note that an offer i is the amount that player I proposes that player II should get rather than the amount player I demands for himself.) An action for player II is a choice from the set $\{Y, N\}$. Her strategies are therefore functions $f: \{1, 2, \dots, 40\} \rightarrow \{Y, N\}$. However, we assume that player II is restricted to functions of the form $f(i) = Y$ ($i \geq j$) and $f(i) = N$ ($i < j$) for some $j \in \{1, 2, \dots, 40\}$. We can then identify player II's strategy with the minimum acceptable offer j and the set of pure-strategy pairs can be identified with $I \times I$. The forty pure-strategy Nash equilibria are (i, i) ($i = 1, 2, \dots, 40$). Since $i = 0$ is excluded, the unique subgame-perfect equilibrium is $(1, 1)$.

For an evolutionary analysis, we assume that player I is drawn from an infinite population of proposers and that player II is drawn from an infinite population of responders. The fraction of proposers who make offer i at time t is denoted by $x_i(t)$. The fitness $\pi_i(t)$ of a proposer using offer i at time t is taken to be the expected payoff to a player I who makes offer i when his opponent is drawn at random from the population of responders at time t . Average fitness in the population of proposers at time t is $\bar{\pi}_I(t) = x_1(t)\pi_1(t) + \dots + x_{40}(t)\pi_{40}(t)$. The standard replicator equation for the evolution of $x_i(t)$ is given by

$$\dot{x}_i = x_i(\pi_i - \bar{\pi}_I) \quad i \in \{1, 2, \dots, 40\}. \quad (1)$$

Similarly, we let $y_j(t)$ be the fraction of responders playing strategy j at time t , with $\pi_j(t)$ being the fitness of a responder using strategy j and $\bar{\pi}_{II}(t)$ be the average fitness of responders, so that

$$\dot{y}_j = y_j(\pi_j - \bar{\pi}_{II}) \quad j \in \{1, 2, \dots, 40\}. \quad (2)$$

The evolution of the whole system is determined by the 80 equations given by (1) and (2).

In the terminology of Hofbauer and Sigmund (1988), every Nash equilibrium is a rest point of the replicator dynamics.⁶ It is easy to show that many of these Nash equilibria are local attractors. In addition, the calculation reported in cell $(0, 0)$ of Table I shows that, with uniform initial conditions,⁷

⁵ In what seemed crucial cases, we also computed solutions for games with $I = \{1, 2, \dots, 100\}$ without significantly altering the results.

⁶ A rest point r is a fixed point of the dynamics. A local attractor l has the property that, for each neighborhood V with $l \in V$, there is another neighborhood U with $l \in U \subseteq V$ such that any trajectory that begins in U remains in V . An asymptotic attractor a is a local attractor with the property that all trajectories which begin in a small enough neighborhood of a converge to a .

⁷ That is, each population begins with each strategy being played by $\frac{1}{40}$ of the agents in that population.

the system converges to an equilibrium in which player II receives a little more than 20% of the pie. We therefore have immediate occasion to cast doubt on the subgame-perfect prediction in an evolutive context.

However, it may appear that Nash equilibria which fail to be subgame-perfect equilibria are attractors only because the long-run operation of the replicator dynamics allows some strategies to approach extinction, and hence artificially excludes the evolutionary pressure against weakly dominated strategies that would otherwise eliminate them. We therefore turn our attention to models in which small fractions of all possible strategies are continually injected into the population—including those that test the “rationality” of responders who refuse positive offers. Only if the survival of Nash equilibria that are not subgame-perfect is robust in the presence of such noise can we realistically argue against the subgame-perfect prediction.

It is natural to see the noisy model as an evolutionary gloss on Selten’s (1975) trembling-hand story, which he used to justify subgame-perfect equilibria in games like the Ultimatum Game. However, caution is necessary before pressing the analogy too far. Samuelson and Zhang (1992) show that adding noise to the replicator and other evolutionary dynamics does not necessarily lead to the elimination of weakly dominated strategies. The question of whether only subgame-perfect equilibria can survive in a noisy evolutionary environment therefore remains open.

Noise in an interactive learning system may arise in many ways and cause perturbations of various types. We therefore think it important to be clear on the source of the noise to be studied.⁸ This in turn requires that we take a little more care than is usual in modeling the agents.

We envisage an agent as a stimulus–response mechanism with two modes of operation: a playing mode and a learning mode. Its playing mode operates when it is called upon to choose a strategy in one of a large number of games that it repeatedly plays against different opponents. Its behavior in each game is triggered by a stimulus that is determined by the manner in which the game is framed. (By a “game-frame,” we mean more than the game itself. We include also the context in which the game is encountered and the manner in which its rules are described.)⁹ When it receives such a stimulus s it responds by playing a strategy $D(s)$. If the learning mode were absent, an agent could therefore be identified with a fixed decision rule D that maps a set of stimuli into a set of strategies.

⁸ We depart from that part of the refinement literature which follows Kohlberg and Mertens (1986) in demanding robustness in the face of all conceivable perturbations. There is no reason to suppose that a system will necessarily be adapted to types of noise that it has experienced only rarely if at all.

⁹ For example, it may be relevant whether the interacting agents are a monopoly seller and a buyer, or whether they are the joint winners of a lottery.

However, sometimes an agent will enter its learning mode between games to adjust its current decision rule. When learning, it takes a stimulus s and some information f about the relative success of strategies in the game labeled by s to modify the value of $D(s)$. The learning rule L that it uses for this purpose is assumed to be fixed. We restrict our attention to learning rules that lead to the replicator dynamics largely because this dynamic has been widely discussed in the literature on evolutionary game theory and hence will be familiar (but see Section 6).

Noise may perturb an agent in its decision mode or in its learning mode. Here and in Section 6, we simplify by considering only the second possibility. We then simplify further by assuming that the only source of error lies in the possibility that an agent may mistakenly learn to play a strategy that is adapted to the wrong game.¹⁰ We do not explicitly model the situations that may be confused with the Ultimatum Game. In the case of a misguided proposer, we simply assume that he makes each offer i in the Ultimatum Game with probability θ_i . If the fraction of proposers at time t who misread the game is always δ_I , and the usual arguments leading to the replicator equation apply to the fraction $1 - \delta_I$ of the proposing population who do not misread the game, then we are led to the “noisy replicator equation”

$$\dot{x}_i = (1 - \delta_I)x_i(\pi_i - \bar{\pi}_I) + \delta_I(\theta_i - x_i) \quad (3)$$

for the evolution of the fraction $x_i(t)$ of agents in the proposing population who play strategy i . The corresponding equation for the population of responders is

$$\dot{y}_j = (1 - \delta_{II})y_j(\pi_j - \bar{\pi}_{II}) + \delta_{II}(\psi_j - y_j), \quad (4)$$

where δ_{II} is the fraction of the responding population who misread the game and ψ_j describes the choices of such agents.

Section 6 derives (1)–(2) and (3)–(4) from an explicit choice model in which agents sometimes misread their strategic situation. Other choice models can lead to different versions of the dynamics. For example, if a proportion $\delta_I\tau$ of the agents die (or leave the game, or choose to experiment with new strategies) in each time period of length τ , to be replaced by

¹⁰ Although the English language forces us into speaking of players' misreading the game or learning to play better, it should be emphasized that our agents do not monitor what is going on except insofar as this is modeled by the learning rule with which they are endowed. The fact that they have a learning rule at all makes them more flexible than the stimulus–response machines that are often considered, since their decision rules for playing games evolve over time, but the learning rule that governs how decision rules evolve is fixed.

TABLE I
FIXED NOISE CALCULATIONS

		δ_{II}				
		0.1	0.01	0.001	0.0001	0
δ_I	0.1	7	2	1	1	1
	0.01	9	7	3	1	1
	0.001	9	9	7	3	1
	0.0001	9	9	9	7	1
	0	9	9	9	9	9

novices who play each strategy i with probability θ_i , then we are led to the equation $\dot{x}_i = x_i(\pi_i - \bar{\pi}_1) + \delta_i(\theta_i - x_i)$. This corresponds to noise in agents' playing mode rather than their learning mode. Our theoretical analysis in Section 5 includes this case as well as (3) by examining dynamics of the form $\dot{x}_i = \Delta_i x_i(\pi_i - \bar{\pi}_1) + \delta_i(\theta_i - x_i)$. Alternatively, Binmore *et al.* (1993) examine a choice model that gives rise to the dynamics $\dot{x}_i = x_i(\pi_i - \bar{\pi}_1)/\bar{\pi}_1$. Van Damme (1987) and others work with a discrete version of this dynamic. We have reported numerical calculations using this discrete dynamic in Binmore and Samuelson (1994), and we indicate how the results differ from those reported here as we proceed.

What determines θ_i and ψ_j ? These presumably reflect rules of thumb or behavior learned in other games, and as a result we have little to say about their precise form. For most of our calculations, we will assume that these represent a uniform distribution over strategies. We discuss how the specification of θ_i and ψ_j affects the results at the end of this section.

Table I reports calculations for various values of δ_I and δ_{II} .¹¹ The rows in Table I correspond to different values of δ_I . The columns correspond to different values of δ_{II} . In each case, the system was initialized with each of the 40 possible strategies being played by $\frac{1}{40}$ of each population. The mistake probabilities were also taken to be uniform, so that $\theta_i = \psi_j = \frac{1}{40}$.

The entries in Table I are the model offers made by player I after the system has converged to a point where the proportion of each population playing each strategy is unchanging in its first 15 decimal places. In each case, the frequency with which the model offer is played at this point is 1.00 to at least two decimal places. The equilibrium behavior of responders

¹¹ The difference equation $x_i(t + \tau) - x_i(t) = \tau[(1 - \delta_i)x_i(\pi_i - \bar{\pi}) + \delta_i(\theta_i - x_i)]$ is used to approximate Eq. (3) where we set $\tau = 0.01$. The robustness of the approximation was tested by repeating a sample of the calculations with much smaller values of τ .

is much more diffuse, but is very highly concentrated on strategies less than or equal to the modal offer. Hence, offers are rejected with only a very tiny probability. For example, in the cases when the model equilibrium offer made by player I is 9, a significant fraction of responders would accept each of the offers between 1 and 9 in equilibrium (with virtually no responders insisting on more than 9)—but the fraction of responders who will refuse anything lower than 9 is high enough to make it unprofitable for proposers to reduce their offer.

Table I shows that, if the noise level among responders is sufficiently small relative to that of proposers, then the subgame-perfect equilibrium appears. However, if the noise level in the responding population fails to be small enough compared with the noise level in the proposing population, then outcomes appear that are far from the subgame-perfect equilibrium. If responders are noisy enough compared with proposers, then player II gets a little more than 20% of the pie.¹²

Section 5 provides an analytic explanation of these results for a simple special case, but the intuition is straightforward. It is, for example, weakly dominated for player II to refuse an offer of 10%. There will therefore always be some evolutionary pressure against this strategy because, in a noisy population, the set of proposers who make such low offers is continually renewed. However, if this fraction of the proposing population becomes sufficiently small, the pressure against refusals of 10% will be negligible compared with the drift engendered by the noise in the responding population. Hence, if responders are noisy enough relative to proposers, then sufficiently many responders can reject offers of 10% that it is not a best response for proposers to offer less and we can reach outcomes that are not subgame-perfect.

Why should we anticipate that there will be more noise in the population of responders than in the population of proposers? Recall that we envision the noise arising as a result of an agent misreading the game when learning and hence acquiring an inappropriate behavior. The context is that of a boundedly rational agent without sufficient computational power to devote full attention to all of the many games that compete for its attention. However, the frequency with which learning errors are made is unlikely to be independent of the potential costs. Instead, we expect the likelihood of a learning error to depend on how much it currently matters in payoff

¹² For the dynamic $x_i(t+1) = x_i(t) + x_i(t)(\pi_i - \bar{\pi})/\bar{\pi}$, Binmore and Samuelson (1994) find results that are much the same as reported in this paper, though in (1994) we find that player I's noise level need only be at least as high as player II's in order to give subgame-perfection. The outcome for cases in which player II's noise level is higher is again 9. This difference arises because, near the subgame-perfect equilibrium, the divisor π becomes especially small for responders. This accentuates the learning portion of the noisy replicator dynamic, causing the responding population to seem less noisy.

terms what strategy is played in the game.¹³ In more familiar terms, the assumption will be that the players are more diligent in identifying games correctly when their potential gains and losses are large, and more prone to misread games when their potential gains or losses are small.¹⁴

In the Ultimatum Game, the result of making such an assumption is that responders will tend to be noisier when proposers are making low offers, because the responders will then have less at stake. In particular, such endogenously determined noise will lead responders to be noisier than proposers if the system should get close to the subgame-perfect equilibrium.

To explore this question further, we performed calculations in which the noise levels were endogenized along the lines discussed above. We took

$$\delta_k(t) = \frac{\alpha\beta}{\alpha + \lambda_k(t)} \quad k = \text{I, II}, \quad (5)$$

where α and β are constant and $\lambda_k(t)$ is the difference between the maximum and minimum of the expected payoffs attached to player k 's strategies, given the current distribution of strategies in the opposing population. When this difference is zero, as is nearly the case for responders at the subgame-perfect equilibrium, the noise level takes its highest value of β . If the difference could increase all the way to infinity, $\delta_k(t)$ would decrease to zero.¹⁵

¹³ This formulation is consistent with the spirit of Myerson's (1991) proper equilibrium, which refines the idea of a trembling-hand equilibrium by making more costly mistakes less likely.

¹⁴ Such an assumption adds more complexity to the stimulus-response mechanism used to model an agent. The mechanism must now incorporate a device that responds to changes in its environment by diverting computational capacity between monitoring and other tasks according to the estimated rewards from the different activities. Like the learning rule, this device is assumed to be fixed.

¹⁵ The difference between the maximum and minimum payoff is an arbitrary measure of the payoffs that are at stake in a game. Calculations using alternative measures, such as the variance of the payoffs to player k 's strategies, with each strategy taken to be equally likely in the variance calculation, produced analogous results (Binmore and Samuelson (1994)). A more realistic measure would perhaps use a sample of past payoffs rather than employing all current payoffs. On the other hand, we suspect that people are indeed often able to make educated guesses about their compatriots' current payoffs without necessarily being at all well informed about the strategies that secure the payoffs. Academic economists, for example, are often able to estimate their colleagues' salaries quite closely. Extreme payoffs are especially likely to attract comment.

TABLE II
CALCULATIONS WITH ENDOGENOUS NOISE

α	β	Offer	$\delta_I(\infty)$	$\delta_{II}(\infty)$
10	1	9	0.26	0.52
10	0.1	9	0.024	0.053
10	0.01	9	0.0024	0.0053
10	0.001	9	0.00024	0.00053
1	1	9	0.032	0.1
1	0.1	9	0.0031	0.01
1	0.01	9	0.00031	0.001
1	0.001	9	0.000031	0.0001
0.1	1	9	0.0032	0.011
0.1	0.1	9	0.00032	0.0011
0.1	0.01	9	0.000032	0.00011
0.1	0.001	9	0.0000032	0.000011
0.01	1	9	0.00032	0.0011
0.01	0.1	9	0.000032	0.00011
0.01	0.01	9	0.0000032	0.000011
0.01	0.001	9	0.00000032	0.0000011

Table II summarizes calculations with endogenized noise for various values of the two constants α and β listed in the first and second columns. The third column shows the modal offer made in equilibrium. The frequency with which the modal offer made in equilibrium is again 1.00 to at least two decimal places, and responders' strategies range between the modal offer and zero, with virtually no rejections. The fifth and sixth columns show the noise levels in the two populations after equilibrium is achieved.

Endogenizing the noise leads to an equilibrium in which the responder population is noisier than the proposer population. It is therefore not surprising that the equilibrium outcome is not subgame-perfect. In fact, the equilibrium offer is again close to 20%.

How robust are these results? First, consider the question of initial conditions. The calculations reported in Tables I and II are based on a uniform initial distribution of offers over $I \times I$. We also performed $1600 = 40 \times 40$ other calculations to explore the dependence of the results on the initial conditions. Table III shows the modal equilibrium offers for some of these initial conditions for the case of endogenous noise with $\alpha = 1$ and $\beta = 0.1$. The entry in row i and column j is the modal equilibrium offer when the system is started with all proposers playing i and all responders playing j . The frequency of the modal equilibrium offer remains 1.00 to at least two decimal places.

Space precludes showing the whole table. The table extends downward

TABLE III
CALCULATIONS WITH VARYING INITIAL CONDITIONS

	1	2	3	4	5	6	7	8	9	10	11	12	13	14	...
1	1	2	3	4	5	6	7	8	9	10	10	10	10	10	...
2	1	2	3	4	5	6	7	8	9	10	10	10	10	10	...
3	1	2	3	4	5	6	7	8	9	10	10	10	10	10	...
4	1	2	3	4	5	6	7	8	9	10	10	10	10	10	...
5	1	2	3	4	5	6	7	8	9	10	10	10	10	10	...
6	1	2	3	4	5	6	7	8	9	10	10	10	10	10	...
7	1	2	3	4	5	6	7	8	9	10	10	10	10	10	...
8	1	2	3	4	5	6	7	8	9	10	10	10	10	10	...
9	1	2	3	4	5	6	7	8	9	10	10	10	10	10	...
10	1	2	3	4	5	6	7	8	9	10	10	10	10	10	...
11	1	2	3	4	5	6	7	8	9	10	10	10	10	10	...
12	1	2	3	4	5	6	7	8	9	10	10	10	10	10	...
13	1	2	3	4	5	6	7	8	9	10	10	10	10	10	...
14	1	2	3	4	5	6	7	8	9	10	10	10	10	10	...
15	1	2	3	4	5	6	7	8	9	10	10	10	10	10	...
16	1	2	3	4	5	6	7	8	9	10	10	10	10	10	...
17	1	2	3	4	5	6	7	8	9	10	10	10	10	10	...
18	1	2	3	4	5	6	7	8	9	10	10	10	10	10	...
19	1	2	3	4	5	6	7	8	9	10	10	10	10	10	...
20	1	2	3	4	5	6	7	8	9	10	10	10	10	10	...
21	1	2	3	4	5	6	7	8	9	10	10	10	10	10	...
22	1	2	3	4	5	6	7	8	9	10	10	10	10	10	...
23	1	2	3	4	5	6	7	8	9	10	10	10	10	10	...
24	1	2	3	4	5	6	7	8	9	10	10	10	10	10	...
25	1	2	3	4	5	6	7	8	9	10	10	10	10	10	...
⋮	⋮	⋮	⋮	⋮	⋮	⋮	⋮	⋮	⋮	⋮	⋮	⋮	⋮	⋮	⋮

just as one would anticipate on the basis of its existing pattern. For cases in which proposers initially play at least 10, it extends to the right as one would expect (except that the cells (38, 37), (38, 38), (39, 37), (39, 38), (39, 39), (40, 37), (40, 38), (40, 39), and (40, 40) yield outcomes of 9 rather than 10). For cases in which the initial offer is less than 10, we find the outcome to be 10 as long as the initial response is not too high (generally, up to 25, though higher for lower initial proposals). Higher initial responses yield lower final outcomes. We found the outcomes of initial conditions in which the proposer offers at least 10 to be robust to the values of α and β .¹⁶ The outcomes for cases in which proposers initially offered less

¹⁶ Using the specification of the basic dynamic as $x_k(t + 1) = x_k(t) + x_k(t)(\pi_i - \bar{\pi})/\bar{\pi}$ and taking $\lambda_k(t)$ to be the variance of the payoffs accruing to each of agent k 's strategies (rather than the difference between the maximum and minimum payoff) gives similar results, though the 10's are replaced by 9's.

than 10 were somewhat more sensitive, though the outcome was always at least as large as the minimum of the initial offer and initial proposal.¹⁷

The most striking feature of Table III is the robustness of a modal equilibrium offer above 20%. This offer appears for a large collection of initial conditions, including all those in which the initial proposal and response are at least 20%. The next section explains why we believe these are most likely to be the relevant initial conditions.

How do our results depend upon the specification of noise, i.e., on θ_i and ψ_j ? By changing these specifications, we can obtain different results. The distribution of noise among the responders is especially important. If we alter ψ_j to put relatively more weight on offers and responses near zero, equilibrium outcomes can be achieved in which the responder gets less than 20% of the pie. Causing more weight to be put on somewhat higher offers gives outcomes in which the responders get more than 20% of the pie.

It is interesting to note, however, that changing the values of θ_i and ψ_j that are attached to relatively high offers has virtually no effect on the outcome. For example, we changed the mistake probabilities so that θ_{20} and ψ_{20} , the probabilities of the "fair" offer and response, took various values up to 0.95 (with the remaining values of θ_i and ψ_j remaining equal to one another). We might view this as a case in which the rule of thumb to which most noisy players resort when not paying attention to the game is to split the pie evenly, with a minority of such players adopting completely random rules of thumb that attach equal probability to all strategies. This change had almost no impact on the results of the calculations.

Probability attached to the "fair" response of 20 has little effect because this offer lies above the range of potential modal equilibrium offers (given that the remaining noise is uniformly distributed). As a result, the response earns a low payoff and the learning dynamics ensure that little probability accumulates on this response. The important considerations involve the distribution of noise over those responses that are lower than potential modal equilibrium offers. These responses earn almost identical payoffs and noise plays a major role in determining the equilibrium proportions of responders choosing each of these strategies. The key here is whether the noise in the responder population can amass a sufficient proportion of responders on a strategy y to make it unprofitable for proposers to make offers $x < y$ and hence sustain an outcome in which the modal offer is y . This depends upon the relative noise level of the two populations

¹⁷ Whenever proposers initially make smaller offers than responders will accept, the dynamics begin with a race between proposers and responders, with each adjusting to match the other's strategy. The outcome of this race can be sensitive to parameters of the model when all responders initially get very low payoffs, as is the case for low proposer offers.

and on the specification of the noise. With uniform noise and sufficiently noisy responders, offers less than or equal to 10 can be sustained (cf. Table III) but higher offers cannot. If the responder's noise concentrates more of its probability near (away from) zero, then lower (higher) modal offers will result (provided in the latter case that the increased weight is not directed to offers such as 20 that are too high to be possible equilibrium modal offers).

It is clear that we cannot place too much significance on the particular value of the equilibrium offer of a little more than 20% that repeatedly emerges in the calculations, and we are anxious that our results not be remembered for this number. Different specifications of the model can give different numbers. The important feature of the results is that the equilibrium offer is frequently far from subgame-perfect, even when the noise levels are made very small indeed. This result requires primarily that the responding population be relatively more noisy than the proposing population. But this is the configuration of noise levels that appears if players tend to be less noisy when making decisions that are more important.

3. RELEVANCE TO EXPERIMENTAL DATA?

How do we think the calculations of the previous section might be relevant to the experimental data? We think it useful to distinguish four time spans:¹⁸

1. In the *short run*, one should anticipate that behavior is driven primarily by norms that are triggered by the framing of the problem. The framing may well elicit norms that are ill-adapted to the laboratory situation. If these norms have been strongly reinforced outside the laboratory, they may be hard to shift. We suspect that the "irrational" behavior studied by the school of Kahneman and Tversky (1987) often falls into this category.

2. In the *medium run*, subjects begin to learn—as emphasized by Andreoni and Miller (1993), Crawford (1991, 1992), Miller and Andreoni (1991), Roth and Erev (1993), and numerous other authors.

3. In the *long run*, this interactive learning process may converge on an equilibrium of the game.

4. In the *ultralong run*, there may be jumps between equilibria when

¹⁸ Our general views on the evolution of social norms inside and outside laboratories have been reported elsewhere (Binmore and Samuelson (1994)).

random shocks jolt the system from one basin of attraction to another—as emphasized by Young (1993), Kandori *et al.* (1993), and Samuelson (1994).

In the ultralong run, we expect an evolutionary process to select the subgame-perfect equilibrium. However, our guess is that the ultralong run is *too* long a run to be relevant to the available experimental data. We also do not think that the replicator dynamics provide a useful model of the ultralong run.¹⁹

But, as we have argued, the replicator dynamics do have a role to play as long-run approximations to certain simple learning rules. We therefore believe that the asymptotic properties of the replicator dynamics may be relevant to the long-run outcome of interactive learning in the laboratory. If so, then it is significant that our calculations of the long-run behavior of noisy replicator dynamics in the Ultimatum Game should generate equilibria that are far from subgame-perfect.

We believe that our calculations are relevant to the short run as well as the long run. As Section 6 explains, we think it possible to regard our dynamics as a crude but instructive model of social evolution (as well as of interactive learning in the laboratory).

More importantly, the social norm (or norms) triggered in the short term by laboratory experiments on the Ultimatum Game have presumably evolved to guide behavior in real-life bargaining situations that are superficially similar to the Ultimatum Game in some respects. We must therefore examine long-run behavior in these external situations for the origin of the norms that guide short-run behavior in laboratory experiments on the Ultimatum Game.

The real-life bargaining situations that have shaped the norms which subjects bring to the laboratory will be complicated by informational, reputational and other effects that are controlled away in the laboratory. The pure Ultimatum Game represents an extremal case in the class of real-life bargaining situations, because all the power is on the side of the proposer. If a social norm adapted to the pure Ultimatum Game leads to the proposer offering about 20% to the responder, we should therefore anticipate that bargaining norms adapted to a wider class of bargaining games will assign *more* than 20% to the responder. If this guess is correct,

¹⁹ In Section 6, an implicit appeal is made to the law of large numbers when studying the long run, so that the underlying stochastic learning is smoothed into a deterministic process. Binmore and Samuelson (1993) argue that one must refrain from such appeals when studying the ultralong run, and work directly with the stochastic system instead. If one were to work directly with the stochastic system in the current paper, results would emerge concerning the expected waiting time until reaching the subgame-perfect equilibrium in the ultralong run. The noisier the responding population relative to the proposing population, the longer the system lingers near long-run equilibria that are not subgame-perfect—and hence the longer and less relevant the ultralong run.

TABLE IV
MEDIUM RUN EQUIVALENT OF TABLE I

		δ_{II}				
		0.1	0.01	0.001	0.0001	0
δ_I	0.1	9 (13)	7 (74)	7 (69)	7 (69)	7 (69)
	0.01	9 (15)	9 (12)	9 (12)	9 (12)	9 (12)
	0.001	9 (15)	9 (12)	9 (12)	9 (12)	9 (12)
	0.0001	9 (15)	9 (12)	9 (12)	9 (12)	9 (12)
	0	9 (15)	9 (12)	9 (12)	9 (12)	9 (12)

we should therefore envisage the initial conditions for learning in the laboratory as allocating more than 20% to the responder and hence as lying in the basin of attraction of the 20% equilibrium offer of Table III.

What do our calculations tell us about the medium run? Table IV is a medium-run version of Table I. It differs from Table I in that the modal offer is reported on the first occasion at which no change in consecutive iterations was detected in the first five decimal places of the fractions of proposers making each offer. (Table I does the same, but with fifteen decimal places.) The number in parentheses following each modal offer is a measure of how much learning was necessary before a temporary stability in the first five decimal places was achieved.²⁰ In Table I, the frequency with which modal offers were used was 1.00 to at least two decimal places. The frequency with which the modal offers were used at the time reported in Table IV is at least .98.²¹

Table IV shows that the system *always* (with uniform initial conditions and perturbations) goes quite quickly to a modal offer of about 20%. But Table I shows this to be a medium-run result. In the long run, the system sometimes moves away to the subgame-perfect equilibrium. Only when $\delta_I < \delta_{II}$ is the medium-run behavior a useful guide to the long-run behavior of the system.

These results complement those of Roth and Erev (1993), who report Ultimatum-Game simulations that spend extended periods of time, in the

²⁰ The measure is the number of iterations of the discrete dynamic described in note 11 multiplied by the step size. In the model of Section 6, τ of the population has an opportunity to change strategies in each iteration of the discrete equation. Our measure therefore provides a crude approximation to the aggregate number of times that members of the entire population have assessed their strategies. The measure is intended to serve as a correlate for the number of rounds of an experiment required to reach temporary stability.

²¹ If we ask for stability in only the first three decimal places, the first line of this table would read 9 (7) 9 (6) 9 (6) 9 (6), with the modal offer being played with frequency at least 0.85 in each case.

medium run, near equilibria that are Nash but not subgame-perfect. Roth and Erev suggest explaining the experimental data on the Ultimatum Game as a set of medium-run observations of a learning process.

We agree that much experimental data consists of a series of snapshots of medium-run phenomena. This is especially true of Ultimatum Game experiments, where both dispersion in proposals as well as rejected offers often persist into the later rounds of the experiments, both of which can only be medium-run phenomena in our model. However, we do not think it follows that theories of long-run behavior can be neglected. As our analysis of Section 5 suggests, long-run predictions of theoretical models of interactive learning will often depend only on *qualitative* features of the models. By contrast, medium-run predictions must be expected to depend upon the fine details of the interactive learning process. We therefore think that current theoretical techniques are more likely to be successful when applied to long-run rather than medium-run phenomena. Rather than seeking to explain experimental data in which medium-run behavior has been elicited, we therefore think there is a strong case for designing experiments with a view to eliciting long-run behavior.²² The contribution of this paper is to argue that, in such experiments, there is no compelling reason why predictions should favor subgame-perfect equilibria over other Nash equilibria.

4. LEAVING MONEY ON THE TABLE

The previous sections argue that attention needs to be paid to Nash equilibria in the Ultimatum Game that are not subgame-perfect. Such equilibria require that the responder be prepared to refuse low positive offers. If offered a choice between something and nothing, such a responder would therefore sometimes choose nothing. Such behavior is outlawed in conventional economic modeling. Perhaps for this reason, a common response to our argument is an incredulous "Why would anyone leave money on the table?"

That positive offers should be refused in the short term is easy to understand. Short-term behavior in the Ultimatum Game is likely to be governed by social norms that are triggered by the framing of the laboratory experiment. Rather than being adapted to the pure Ultimatum Game, such social norms will presumably have evolved for use in everyday cousins of the Ultimatum Game. In everyday life, we rarely play pure

²² Binmore *et al.* (1992), for example, obtain very close convergence to equilibrium in less than 40 repetitions in a complicated bargaining game by offering high incentives and helping the subjects with sophisticated computer graphics.

take-it-or-leave-it games. In particular, real-life games are seldom played under conditions of total anonymity. A refusal of something positive may therefore serve to maintain a reputation for toughness. Even when we do play anonymously, outside options are often available. For example, in the take-it-or-leave-it auction used by stores to sell their goods, a refusal of something positive may simply indicate a willingness to search elsewhere for a better deal. Norms that call for refusals in commonly occurring "take-it-or-leave-it" situations therefore make good evolutionary sense. Given that such norms exist, it is unsurprising if they are sometimes inappropriately triggered in laboratory experiments. Short-run refusals of positive offers in the pure Ultimatum Game therefore create no problem for orthodox game theory.

However, we argue that Nash equilibria that are not subgame-perfect should be taken seriously even in the long run. Notice first that such equilibria actually require very few offers to be rejected, because proposers learn not to make such offers. Nevertheless, responders must stand ready to reject some positive offers.

On this subject, it is useful to observe that people clearly *do* sometimes leave money on the table. Frank (1988), for example, reminds us about tipping behavior in restaurants that are never to be visited again. After the waiter makes your change, you can either pocket the entire amount or leave the customary percentage on the table. Nearly everyone chooses the latter option—including economists!

A kibitzer may ask *why* we leave money on the table. Most people are satisfied with the explanation that leaving a tip is a custom that it would be uncomfortable to violate. If pressed, they might attribute the discomfort to the unfairness involved in disappointing the server's expectations.

Such considerations have led a number of authors to downplay strategic explanations of experimental behavior in favor of various theories of "fair play". Sophisticated versions of this approach sometimes build a taste for "fairness" into the utility functions attributed to the subjects. Ochs and Roth (1989) discuss such a utility function in explaining the medium-run results of an alternating offers bargaining experiment. Bolton (1991) explicitly constructs such a utility function for this purpose.

We agree that subjects find their emotions engaged in bargaining situations. They also frequently explain their bargaining behavior in the laboratory in terms of "fairness." But an approach that takes these facts at their face value is in danger of explaining too much and too little. Fairness theories explain too much because, by choosing one's fairness notion with sufficient care, one can justify a very wide range of outcomes. At the same time, such theories explain too little because they provide no insight into the origin of the fairness norms to which appeal is made.

We believe that a more fruitful approach is to ask how the custom of

leaving money on the table can survive. Our answer is quite simple. The amounts involved and the frequency with which the situation arises are too small to provide sufficient evolutionary pressure to eliminate the phenomenon in a noisy environment. What then of the folk explanation in terms of the discomfort felt at violating a fairness norm?

In responding to such questions, it is important to appreciate that the evolutionary approach we advocate reverses the standard *explicans* and *explicandum* of the folk explanation and of economic theory. Our players are not members of the species *Homo economicus*. They do not optimize relative to fixed preferences. They simply have decision rules for playing games. When a player switches from a less profitable to a more profitable strategy, he does not do so because he thinks that the switch is optimal—he is just acting as a stimulus–response mechanism.

This model of *Homo sapiens* raises the question of how it *feels* for one's actions to be programmed as a result of past experience. Here the *post hoc, ergo propter hoc* fallacy awaits the unwary. It is easy to say that I preferred to take this foolish action rather than that wise action *because* I got angry. But we feel angry because adrenalin and other chemicals have been released into our bloodstream by a process which is only very partially under our conscious control. Angry feelings are a conditioned reflex to certain learned stimuli. Such conditioned reflexes survive because the behaviors they induce have evolutionary advantages. Rather than seeking to explain a particular behavior in terms of the angry feelings that accompany it, we therefore do better to explain the angry feelings in terms of the evolutionary advantages of the behavior. In brief, being angry or fearful or amorous is how it feels to be a stimulus–response mechanism.

Of course, none of us like to admit that much of our behavior is little more than a set of conditioned reflexes. We prefer to offer more flattering rationalizations of the behavior. For example, the stimulus of receiving an offer of only 10% in the Ultimatum Game may be sufficiently irritating that we turn the offer down. If asked *why* we refused, we may then rationalize our behavior by arguing that irritation is an entirely appropriate response to an “unfair” offer of 10%. Indeed, such an explanation may become institutionalized and so reinforce the behavior that it “explains.” But we see no more reason to believe that “fairness norms” are fixed and immutable than that economic agents always maximize money. We believe that players usually find their way to a long-run equilibrium of trial-and-error learning without having any clear understanding of the strategic realities of the game they are playing. They simply learn that certain stimulus–response behaviors are effective. After the game, they may rationalize their behavior in various ways. In bargaining experiments, they often say that the long-run equilibrium to which they found their way is “fair.” But, from an evolutionary perspective, how they explain

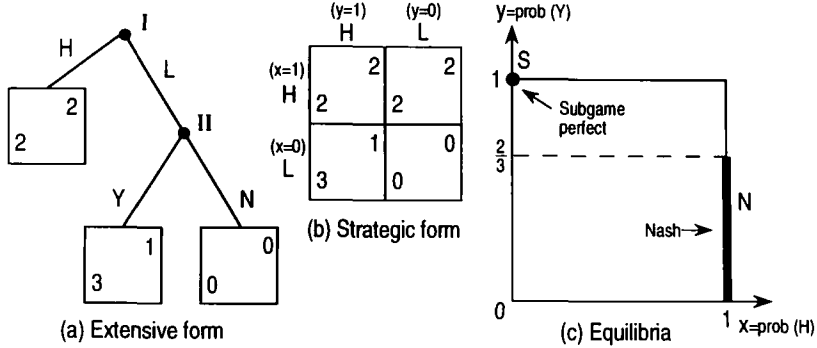


FIG. 1. Ultimatum Minigame.

their own behavior to themselves and others is an epiphenomenon. If they had found their way to another equilibrium, they would be offering some other explanation.²³ Economists who fit utility functions to observed behavior would similarly find themselves proposing a different utility function.

In summary, we believe that attention should be focused on the evolution of *behavior*. If a type of behavior that prompts people to leave money on the table survives, it will be because there is insufficient evolutionary pressure to remove it. Fairness explanations may be offered as rationalizations of the behavior. Such stories may even be incorporated into the workings of the stimulus–response mechanism. But the details of how the mechanism actually works or how we explain its workings to ourselves are secondary. The primary consideration is why some behavior patterns survive in a population while others will necessarily perish. Only after this question has been answered is it worthwhile to ask why some of the stories we tell ourselves are washed away by the evolutionary tide while others remain high and dry.

5. AN ULTIMATUM MINIGAME

To identify the forces that drive our computational results, this section provides an analytical study of the simplified version of the Ultimatum Game shown in Fig. 1. In this Ultimatum Minigame, player I can make

²³ One can observe the fairness norms evolving in the laboratory. In Binmore *et al.* (1992), the median long-run equilibrium claim in a laboratory implementation of the Nash Demand Game turns out to be a very good predictor of the median claim said to be “fair” in a computerized postexperimental debriefing, even though subjects are randomly chosen for an initial conditioning that directs their subsequent play to different long-run equilibrium claims.


S

FIG. 2. Phase diagram, no noise.

a high offer (H) or a low offer (L). If he makes a high offer, it is assumed that player II accepts. If he makes a low offer, player II may accept (Y) or refuse (N).

The Ultimatum Minigame has the same structure as Selten's (1978) Chain-Store Game. Although we do not press the point, our conclusions in this section therefore provide a possible resolution of the chain-store paradox that applies even in the case when there is just one potential entrant.

Figure 1(c) shows the pairs (x, y) that represent equilibria in the Ultimatum Minigame, where x and y are the probabilities with which H and Y are played. There is a unique subgame-perfect equilibrium S at $(0, 1)$ in Fig. 1(c), and a component N of Nash equilibria occupying the closed line segment joining $(1, \frac{2}{3})$ and $(1, 0)$.

Figure 2 shows the trajectories of the standard replicator dynamics in the Ultimatum Minigame:

$$\dot{x} = x(1-x)(2-3y) \quad (6)$$

$$\dot{y} = y(1 - y)(1 - x). \quad (7)$$

We summarize the key properties of these trajectories in Proposition 1 (cf. note 6):

PROPOSITION 1. *The subgame-perfect equilibrium S is the unique asymptotic attractor of the unperturbed replicator dynamics. With the exception of $(1, \frac{2}{3})$, the Nash equilibria in the set N are local attractors.*

The fact that interior points of N are local attractors does not seem to us an adequate reason for regarding them as alternatives to the subgame-perfect equilibrium S . To draw this conclusion, we feel it necessary that the result should survive in the presence of noise that continually replaces strategies that the replicator dynamics drives to extinction. We accordingly require strategies to be at least local attractors in slightly perturbed versions of the dynamics where extinction is not a possibility. We accordingly study the perturbed replicator dynamics, defined by

$$\dot{x} = \Delta_I x(1 - x)(2 - 3y) + \delta_I(\frac{1}{2} - x) \quad (8)$$

$$\dot{y} = \Delta_{II} y(1 - y)(1 - x) + \delta_{II}(\frac{1}{2} - y). \quad (9)$$

In the case when $\Delta_k = 1 - \delta_k$ ($k = I, II$), these equations are analogues to (3)–(4) of Section 3.

Figures 3 and 4 show the trajectories for the perturbed replicator dynamics. None of the points in N is a local attractor in Fig. 3, where responders and proposers are equally noisy; but there exists an asymptotic attractor in Fig. 4, where responders are noisier than proposers.

More formally, we are interested in what happens when the noise in (8)–(9) is small, so that $(\delta_I, \delta_{II}, \Delta_I, \Delta_{II})$ is close to $(0, 0, 1, 1)$. We fix $\phi = \delta_{II}\Delta_I/\delta_I\Delta_{II}$ and consider the limit as $(\delta_I, \delta_{II}, \Delta_I, \Delta_{II}) \rightarrow (0, 0, 1, 1)$ in two cases:

$$\text{Case 1: } 0 < \phi < 3 + 2\sqrt{2}.$$

$$\text{Case 2: } 3 + 2\sqrt{2} < \phi.$$

Since $3 + 2\sqrt{2} \sim 5.8$, responders are appreciably noisier than proposers in the second case.

LEMMA 1. *Let R be the set of rest points of the system (8)–(9) for values of $(\delta_I, \delta_{II}, \Delta_I, \Delta_{II})$ near $(0, 0, 1, 1)$.*

In Case 1, the set R has at most one limit point, which is the subgame-perfect equilibrium $S = (0, 1)$.

S




FIG. 3. Phase diagram, comparable noise ($\delta_I = \delta_{II} = 0.01$).


In Case 2, the set R has at most three limit points $S = (0, 1)$, $(1, \underline{y})$ and $(1, \bar{y})$. The points $(1, \underline{y})$ and $(1, \bar{y})$ lie in the set N of Nash equilibria.

Proof. Writing $\dot{x} = \dot{y} = 0$ and $(\delta_I, \delta_{II}, \Delta_I, \Delta_{II}) = (0, 0, 1, 1)$ in (8)–(9) yields $(0, 0)$, $(0, 1)$ and $(1, y)$ as candidates for the limit points of R . The first point is a source for the unperturbed dynamics, and is easily excluded as a limit point of R . We now consider the possible values of y . To this end, write $\dot{x} = \dot{y} = 0$ in (8)–(9) and then set $x = 1$. We then obtain the equation

$$\phi = \frac{y(1-y)}{(2-3y)(2y-1)}.$$

This equation, illustrated in Fig. 6, has two solutions, \underline{y} and \bar{y} , satisfying $\frac{1}{2} < \underline{y} < 2 - \sqrt{2} < \bar{y} < \frac{2}{3}$ when $\phi > 3 + 2\sqrt{2}$. When $\phi < 3 + \sqrt{2}$, the equation has no solutions y satisfying $0 \leq y \leq 1$. ■

S



N

FIG. 4. Phase diagram, more noise in population II ($\delta_I = 0.01$, $\delta_{II} = 0.1$).

PROPOSITION 2. Let $A(\delta_I, \delta_{II}, \Delta_I, \Delta_{II})$ be the set of asymptotic attractors of the system (8)–(9) given values $(\delta_I, \delta_{II}, \Delta_I, \Delta_{II})$.

In Case 1, the set A has a unique limit point as $(\delta_I, \delta_{II}, \Delta_I, \Delta_{II}) \rightarrow (0, 0, 1, 1)$, which is the subgame-perfect equilibrium $S = (0, 1)$.

In Case 2, the set A has two limit points as $(\delta_I, \delta_{II}, \Delta_I, \Delta_{II}) \rightarrow (0, 0, 1, 1)$, which are $S = (0, 1)$ and $(1, \underline{y})$. (The point $(1, \underline{y})$ is a limit of saddles.)

The first case gives rise to the phase diagram in Fig. 3; the second case to the phase diagram in Fig. 4.

Proof. The proof of the first statement is straightforward, and we consider only the second. The right side of (8)–(9) defines a function $F: \mathbb{R}^2 \rightarrow \mathbb{R}^2$ for which

$$DF(x, y) = \begin{pmatrix} \Delta_I(2 - 3y)(1 - 2x) - \delta_I & -3\Delta_I x(1 - x) \\ -\Delta_{II}y(1 - y) & \Delta_{II}(1 - 2y)(1 - x) - \delta_{II} \end{pmatrix}.$$




FIG. 5. Phase diagram, endogenous noise ($\alpha = \beta = 0.01$).

The trace of this matrix is negative when $\frac{1}{2} < y < \frac{2}{3}$ and $x > \frac{1}{2}$. We therefore consider the limiting value of its determinant.

Multiply the second column of $\det DF(x, y)$ by $2y - 1$ and then make the substitution $\delta_{II}(2y - 1) = 2\Delta_{II}y(1 - y)(1 - x)$, which holds at a rest point by virtue of (8)–(9). Factor out the term $\Delta_I\Delta_{II}(1 - x)$ and write $\delta_I = 0$ and $x = 1$ in what remains. We then have to sign the determinant

$$\begin{vmatrix} 2 - 3y & 3(2y - 1) \\ y(1 - y) & 2y^2 - 2y + 1 \end{vmatrix} = y^2 - 4y + 2.$$

The roots of the quadratic equation $y^2 - 4y + 2 = 0$ are $2 - \sqrt{2}$ and $2 + \sqrt{2}$. It follows that the determinant is positive when $y = \underline{y} \leq 2 - \sqrt{2}$ and negative when $y = \bar{y} < \frac{2}{3} < 2 + \sqrt{2}$. Thus \underline{y} is an asymptotic attractor and \bar{y} is not. ■

We next briefly consider the effect of endogenizing the noise in (8)–(9) by writing $\Delta_I = c_I(1 - \delta_I(t))$ and $\Delta_{II} = c_{II}(1 - \delta_{II}(t))$, where c_I and c_{II} are




FIG. 6. Determining \underline{y} and \bar{y} .

constants and $\delta_I(t)$ and $\delta_{II}(t)$ are given by (5). We study the rest points of this system as β approaches zero, for small values of α . Figure 5 shows the trajectories for this case.²⁴ We have:

PROPOSITION 3. *Let A be the set of asymptotic attractors for (8)–(9) with endogenous noise when $\beta > 0$ is small. For sufficiently small α , A has two limit points $S = (0, 1)$ and $(1, \underline{y})$ as $\beta \rightarrow 0$. As $\alpha \rightarrow 0$, these limit points converger to $(0, 1)$ and $(1, \frac{1}{2})$.*

Proof. The first statement follows from an argument analogous to that of the previous proposition, along with the observation that, as $x \rightarrow 1$, the difference between the minimum and maximum of player II's payoffs approaches zero, so that $\delta_{II}(t)/\delta_I(t)$ gets very large for sufficiently small α . As $\alpha \rightarrow 0$, the limiting ratio of $\delta_{II}(t)/\delta_I(t)$ approaches infinity, in which case we see from Fig. 6 that the limiting values of \underline{y} and \bar{y} approach $\frac{1}{2}$ and $\frac{2}{3}$ respectively. ■

We can identify the forces behind these results. When nearly all proposers are playing H , the pressure for a responder to play Y is weak.

²⁴ An asymptotic attractor exists, though it is difficult to see, at approximately $(1, 0.505)$.

Adding noise to the strategies of the proposer pushes the population of proposers away from H , and increases the pressure for responders to move towards Y . This pushes the system toward the subgame-perfect equilibrium. However, the responders are also noisy. In the absence of other forces, this noise creates a drift that would eventually result in the responding population being split equally between agents who play Y and agents who play N . But the best reply for a proposer against such a mix in the responding population is to play H . If the drift in the responding population is sufficiently strong, it overpowers the countervailing tendency towards the subgame-perfect equilibrium. As a result, the responding population remains close enough to a half-half mix of Y and N that the best reply for proposers continues to be H . The subgame-perfect equilibrium then fails to be selected.

The same intuition can be expressed in a more quantitative form. The dynamic system give by (6)–(7) can be represented as a vector field on the state space $[0, 1]^2$, associating with each point (x, y) a vector

$$(x(1-x)(2-3y), y(1-y)(1-x)) \quad (10)$$

indicating the direction and strength of moment at (x, y) . Along the component N illustrated in Fig. 1c, these vectors are all zero vectors. The perturbed system (8)–(9) is the sum of the vector field given by (10) and a vector field of perturbations given by

$$\begin{aligned} &((\Delta_I - 1)x(1-x)(2-3y) + \delta_I(\frac{1}{2} - x), \\ &(\Delta_{II} - 1)y(1-y)(1-x) + \delta_{II}(\frac{1}{2} - y)). \end{aligned} \quad (11)$$

Because the vector (10) is zero on N , the behavior of (8)–(9) on N is driven by the perturbations given in (11). The key question here is whether we can find a subset $\bar{N} \subset N$ such that the perturbations on \bar{N} point into the basin of attraction of \bar{N} in the *unperturbed* dynamics given by (10). If such an \bar{N} exists, then the perturbations have the effect of continually pushing points near \bar{N} back into the basin of attraction of \bar{N} . The dynamics (8)–(9) will then have an asymptotic attractor that is close to \bar{N} and which converges to \bar{N} as noise levels become small.

What does this have to do with relative noise levels? Let $\bar{N} = [0, \underline{y} + \varepsilon]$ for some small $\varepsilon > 0$. Notice that this set includes the half-half mixture between Y and N , and recall that responder perturbations are pushing responders toward the half-half mixture of Y and N . At the same time, proposer perturbations are pushing proposers towards the strategy L , which in turn creates pressure for responders to switch to Y . The larger

the ratio δ_{II} to δ_I , the stronger the net perturbation pushing responders toward the half-half mixture, and the more likely the resulting perturbations to point into the basin of attraction of \bar{N} . Notice that only relative noise levels matter in determining the *direction* of the perturbation vectors, which (combined with the fact that (10) is zero on N) explains why the argument holds for arbitrarily small absolute noise levels.

Two implications of this argument are immediate. First, if the payoff 3 is replaced by 5 in the Ultimatum Minigame, then the best response to a half-half mixture of Y and N is L , so that the perturbed dynamics can never point into the basin of attraction of a subset of N . The dynamics then lead to the subgame-perfect equilibrium regardless of relative noise levels. Second, even after replacing 3 by 5, we could induce the perturbations to point into the basin of attraction of a subset of N if we altered the specification of the responders' behavior (when noisy) to put a large probability on N rather than a probability of $\frac{1}{2}$ on N . We see here the sensitivity of the results to the specification of noisy behavior.

Second, this argument also provides an idea as to how sensitive our results are to the specification of the unperturbed learning dynamic, which we have taken to be the replicator dynamic. The precise form of this dynamic is not particularly important, as long as the points in N are rest points and the set N has a basin of attraction into which perturbations can point. This excludes pure best-reply dynamics, in which even an arbitrarily small payoff difference between Y and N causes all responders to immediately switch to Y . However, virtually any system in which growth rates of strategy proportions are smooth, increasing functions of expected payoff differences (with growth rates being zero when all strategies have the same expected payoff) has the desired property.²⁵ The existence of an asymptotic attractor close to the component of equilibria that are not subgame perfect therefore holds for a wide class of dynamic processes, although the precise location of this attractor will be sensitive to the specification of the process.

These results address long-run behavior. What about the medium run? Figures 3 and 4 reveal medium-run behavior matching that reported in Table IV for the full Ultimatum Game. The trajectories in Fig. 3 reach the subgame-perfect equilibrium S in the long run, but in the medium run they can first approach the set N of Nash equilibria that are not subgame perfect. In Fig. 4 (where $\delta_I < \delta_{II}$), some trajectories again approach N in the medium run, but in these case these trajectories never leave the vicinity of N .

These results allow us to return to the question of fairness. Experimental results in bargaining games are now seldom explained *purely* in terms of

²⁵ Samuelson (1988) describes such systems as "cardinal" (as opposed to "ordinal").

fairness norms. Prasnikar and Roth (1992) and Roth *et al.* (1991), for example, suggest that some experimental results are best described in terms of a trade-off between strategic and fairness factors. Their most striking example contrasts the Ultimatum Game with the Best Shot Game. The latter is a public-goods-provision game that, like the Ultimatum Game, features a unique subgame-perfect equilibrium in which the first mover's payoff is much higher than the second mover's.

Prasnikar and Roth find experimental outcomes for the Ultimatum Game that are not close to subgame perfection. On the other hand, their Best Shot outcomes are consistent with a subgame-perfect explanation. They suggest that fairness considerations are able to wrestle outcomes away from subgame-perfection in the Ultimatum Game, but are overwhelmed by strategic considerations in the Best Shot game. In adopting this interpretation, they note that the Best Shot Game has only one pure-strategy Nash equilibrium that is not subgame perfect. The payoff pair resulting from this Nash equilibrium is (0.4, 3.7). The payoff pair at the rival subgame-perfect equilibrium is (3.7, 0.4). Prasnikar and Roth argue that there is then very little scope for learning to reinforce movements of behavior away from the subgame-perfect equilibrium, and hence very little scope for fairness considerations to gain a foothold. This contrasts with the Ultimatum Game, where the presence of Nash equilibria that are close to the unique subgame-perfect outcome provides opportunity for movements away from subgame perfection (perhaps induced by fairness considerations) to be reinforced.

Our methodology offers a potential explanation of such results that, in keeping with our discussion in Section 4, does not require treating "fairness" as a primitive concept. To make this point, we contrast the Best Shot Minigame of Figure 7 with the Ultimatum Minigame of Figure 1. In the Best Shot Minigame, player I has the option of making a high (H) or low (L) contribution to the public good. If player I makes a high contribution, the player II is assumed to make a low contribution. If player I makes a low contribution, then player II has the choice of a high or low contribution. The payoffs are such that there is no gain to both players making a high contribution and each player is better off if the other player makes the high contribution.

As Fig. 7 indicates, the Best Shot Minigame has the same qualitative features as the Ultimatum Game, with a subgame-perfect equilibrium S of (L, H) and a component N of Nash equilibria in which player I plays H . However, the strategic incentives associated with these equilibria differ, with this difference making it very much more likely that our model will select a Nash equilibrium that is not subgame-perfect in the Ultimatum Minigame than in the Best Shot Minigame. In particular, notice that N is much smaller than the corresponding Ultimatum Game component. Nash




FIG. 7. Best Shot Minigame.

equilibria that are not subgame perfect in the Best Shot Game require that player II uses *H* with a probability no higher than approximately 0.11. As a result, the basin of attraction of this component relative to the unperturbed replicator dynamics is small. Moreover, if perturbations are introduced in which *H* and *L* receive the same probability from players who misread the game, then the perturbed dynamics cannot point into the basin of attraction of *N*. In contrast to the Ultimatum Minigame, no local attractor can therefore be found close to *N* and the subgame-perfect equilibrium is necessarily selected in the Best Shot Minigame.²⁶

6. THE RELEVANCE OF REPLICATOR DYNAMICS

Why do we think the replicator dynamics, with their origins in biology, are relevant? Börgers and Sarin (1993) have shown that the replicator dynamics can serve as an approximation to simple learning models related to that used by Roth and Erev (1993). In this section we present a simple model of social evolution that also leads to the replicator dynamic. Our purpose is not to argue that the replicator dynamics represent “the” right model, but only to argue that dynamics of their general type are worthy of our attention.²⁷

We interpret our model as one of social evolution because it relies on the ability of players to observe others’ strategies; information that is generally not available in the laboratory. The reinforcement learning model of Börgers and Sarin (1993) is perhaps better suited as a model of learning

²⁶ A local attractor could be created close to *N*, but this requires perturbations that are very heavily weighted toward *L*.

²⁷ See Binmore and Samuelson (1993) and Cabrales (1993) for similar arguments. See Bendor *et al.* (1991) for another aspiration learning model.

in experiments. As explained in Section 3, both types of learning seem relevant to explaining experimental data.

Divide time into discrete periods of length τ . In every period, each agent retains his current strategy with probability $1 - \tau$. With probability τ , the agent compares his payoff with an aspiration level Λ , which is random and uniformly distributed on $[l, L]$. If the agent receives a payoff exceeding Λ , then the agent does not switch strategies. If the agent's payoff falls short of Λ , then the agent randomly chooses a new strategy. The probability that a given strategy is chosen is taken to be the proportion of the population playing that strategy. For example, it may be that the agent chooses a new strategy by randomly selecting another member of the population and imitating his strategy.

Why is Λ random? Our preferred interpretation here is that the aspiration level is actually fixed while the payoffs in the game are random,²⁸ though we find it analytically convenient to work with the equivalent formulation of a random aspiration level. This is consistent with the view we used to motivate our noisy dynamics, namely that players are constantly involved in a multitude of different games and may misperceive the precise nature of the game.

Let $p_i(t)$ be the probability that the aspiration level Λ exceeds the payoff from a proposer's strategy i in period t . We assume that exactly $p_i(t)$ of the proposers playing strategy i at time t are dealt an aspiration level in excess of their payoffs, and that these proposers switch to new strategies in exactly the same proportions as these strategies are used in the population of proposers. Then,

$$x_i(t + \tau) = x_i(t)(1 - \tau p_i(t)) + \sum_{j \in S} \tau p_j(t) x_j(t) x_i(t). \quad (12)$$

Note that $p_i(t) = (L - \pi_i(t))/(L - l)$, where $\pi_i(t)$ is the payoff to strategy i in period t . Hence,

$$\frac{x_i(t + \tau) - x_i(t)}{\tau} = x_i(t) \frac{\pi_i(t) - \bar{\pi}_1(t)}{L - l},$$

where $\bar{\pi}_1(t)$ is the average payoff over all proposers' strategies. Taking the limit as $\tau \rightarrow 0$ leads to a continuous-time version of the dynamic:

$$\dot{x}_i = x_i \frac{\pi_i - \bar{\pi}_1}{L - l}. \quad (13)$$

²⁸ We then implicitly assume that the dispersion of the payoff distribution around its mean does not vary over strategies or players. More realistic assumptions would lead to more complex dynamics.

A similar equation can be derived for responders. The replicator dynamic given by (1)–(2) is the special case in which time has been rescaled so as to eliminate the constant $L - l$.²⁹

Now suppose that each proposer ignores the learning process with probability δ_1 in each period. Given that the learning process is ignored, with probability τ such an agent abandons his strategy regardless of aspiration level considerations and randomly chooses a new strategy, giving strategy i probability θ_i .

The dynamic given by (12) now becomes:³⁰

$$x_i(t + \tau) = (1 - \delta_1)\{x_i(t)(1 - \tau p_i(t)) + \sum_{j \in S} \tau p_j(t)x_j(t)x_i(t)\} + \delta_1[x_i(t) + \tau(\theta_i(t) - x_i(t))].$$

But $p_i(t) = (L - \pi_i(t))/(L - l)$. As $\tau \rightarrow 0$, we obtain that

$$\dot{x}_i = (1 - \delta_1)x_i(t) \frac{\pi_i(t) - \bar{\pi}_1(t)}{L - l} + \delta_1(\theta_i - x_i(t)). \quad (14)$$

The noisy replicator dynamic given by (3)–(4) is the special case in which $L - l = 1$ for both populations.

We have assumed that the two populations are governed by identical learning rules. There are two obvious ways in which they may not be. First, the populations may be characterized by different values of $L - l$. This is equivalent to saying that the unperturbed dynamics presented in (1)–(2) may proceed at different speeds of the two populations.³¹ Consider the Ultimatum Minigame. The faster is the relative rate at which population I learns, the larger is the basin of attraction of the component of Nash equilibria N that are not subgame perfect. This makes it more likely that the perturbations on this component will point into its basin of attraction, and hence more likely that the dynamics do not lead to the subgame-perfect equilibria.

Alternatively, the rates of learning in the perturbed dynamics (3)–(4) may be different. This would correspond to a situation in which learn draws come at different rates for the two populations.³² It is easy to show

²⁹ See Taylor and Jonker (1978) and Hofbauer and Sigmund (1988). This rescaling depends on an assumption that the distribution of the aspiration level is the same for the two players.

³⁰ Samuelson and Zhang (1992) examine an analogous dynamic, with the random choices interpreted as errors in passing strategies (or genes) from one generation to the next.

³¹ One might, for example, make the speed of learning endogenous by linking it to the payoff magnitudes involved, just as we have done with noise levels.

³² Again, these rates might be linked to payoff differences.

that changing the rate at which learning proceeds is equivalent to rescaling payoffs (i.e., multiplying by a constant) and noise levels.³³ As the relative rate at which learning proceeds for population II increases, we are again more likely to observe outcomes that are not subgame-perfect.

7. CONCLUSION

To the question of whether the subgame-perfect equilibrium should be regarded as the one and only game-theoretic prediction for the Ultimatum Game, we hope that we have provided a convincing and firmly negative answer. But what distinguishes our model from other theories of equilibrium selection in the long run, notably fairness theories? Neither our theory nor fairness theories are open to straightforward refutation, since both leave an apologist with ample room for maneuver in explaining the data. In particular, our theory requires tailoring the initial conditions and the noise that perturbs the dynamics to the experimental environment. We hope that our models will not have to be altered radically in moving between environments, as seems to be necessary with fairness models. However, the final word on these questions will have to await further research on a variety of other games. We hope to report such results soon.

REFERENCES

- ANDREONI, J., AND MILLER, J. H. (1993). "Auctions with artificial adaptive agents," mimeo, University of Wisconsin, Madison, WI; Carnegie Mellon University, Pittsburgh, PA.
- BENDOR, J., MOOKHERJEE, D., AND RAY, D. (1991). "Aspiration-based adaptive learning in two person repeated games," mimeo, Indian Statistical Institute, 1991.
- BINMORE, K. (1987/1988). "Modeling Rational Players: I and II," *Economics and Philosophy*, 3, 179–214; 4, 9–55.
- BINMORE, K. (1990). *Essays on the Foundations of Game Theory*, Cambridge, MA: Basil Blackwell.
- BINMORE, K., (1992). *Fun and Games*. Lexington, MA: Heath.
- BINMORE, K., AND SAMUELSON, L. (1993). "Muddling through: Noisy equilibrium selection," mimeo, University College London; University of Wisconsin.
- BINMORE, K., AND SAMUELSON, L. (1994). "An Economist's Perspective on the Evolution of Norms," *J. Institutional Theoretical Economics* 150, 45–63.
- BINMORE, K., SAMUELSON, L., AND VAUGHAN, R. (1993). "Musical chairs: Modelling noisy evolution," mimeo, University College London; University of Wisconsin.

³³ For example, increasing the rate at which time proceeds by replacing \dot{x} with $k\dot{x}$ for $k < 1$ is equivalent to multiplying all payoffs by $1/k > 1$ and multiplying the noise level by $1/k$ (or multiplying the parameters α and β by $1/k$ in our endogenous noise cases).

- BINMORE, K., SHAKED, A., AND SUTTON, J. (1989). "An Outside Option Experiment," *Quart. J. Econ.* **104**, 753–770.
- BINMORE, K., SWIERZBINSKI, J., HSU, S., AND PROULX, C. (1992). "Focal points and Bargaining," Discussion Paper, University of Michigan.
- BOLTON, G. E. (1991). "A Comparative Model of Bargaining: Theory and Evidence," *Amer. Econ. Rev.* **81**, 1096–1136.
- BOLTON, G. E., AND ZWICK, R. (1993). "Anonymity versus punishment in ultimatum bargaining," mimeo, Pennsylvania State University.
- BÖRGERS, T., AND SARIN, R. (1993). "Learning through reinforcement and replicator dynamics," technical report, University College London.
- CABRALES, A. (1993). "Stochastic replicator dynamics," mimeo, University of California, San Diego.
- CRAWFORD, V. P. (1991). "An 'Evolutionary' Interpretation of Van Huyck, Battalio, and Beil's Experimental Results on Coordination," *Games Econ. Behav.* **3**, 25–59.
- CRAWFORD, V. P. (1992). "Adaptive dynamics in coordination games," Department of Economics Working Paper 92-02R, University of California, San Diego.
- FRANK, R. (1988). *Passions within Reason*. New York: Norton.
- GUTH, W., SCHMITTBERGER, R., AND SCHWARZE, B. (1982). "An Experimental Analysis of Ultimatum Bargaining," *J. Econ. Behav. Organ.* **3**, 367–388.
- GÜTH, W., AND TIETZ, R. (1990). "Ultimatum Bargaining Behavior: A Survey and Comparison of Experimental Results," *J. Econ. Psychol.* **11**, 417–449.
- HOFBAUER, J., AND SIGMUND, K. (1988). *The Theory of Evolution and Dynamical Systems*. Cambridge: Cambridge Univ.
- KAHNEMAN, D. AND TVERSKY, A. (1988). "Rational Choice and the Framing of Decisions," in *Decision Making*. Cambridge: Cambridge Univ.
- KANDORI, M., MAILATH, G. J., AND ROB, R. (1993). "Learning, Mutation, and Long Run Equilibria in Games," *Econometrica* **61**, 29–56.
- KOHLBERG, E., AND MERTENS, J.-F. (1986). "On the Strategic Stability of Equilibria," *Econometrica* **54**, 1003–1038.
- LEDYARD, J. (1992). "Public goods: A survey of experimental research," working paper, California Institute of Technology.
- MILLER, J. H., AND ANDREONI, J. (1991). "Can Evolutionary Dynamics Explain Free Riding in Experiments?" *Econ. Lett.* **36**, 9–15.
- MYERSON, R. B. (1991). *Game Theory: Analysis of Conflict*. Cambridge, MA: Harvard University Press.
- OCHS, J., AND ROTH, A. E. (1989). "An Experimental Study of Sequential Bargaining," *Amer. Econ. Rev.* **79**, 355–384.
- PRASNIKAR, V., AND ROTH, A. E. (1992). "Considerations of Fairness and Strategy: Experimental Data from Sequential Games," *Quart. J. Econ.* **106**, 865–888.
- ROTH, A. E. (1994). "Bargaining Experiments," in *Handbook of Experimental Economics* (J. Kagel and A. E. Roth, Eds). Princeton, NJ: Princeton University Press, forthcoming.
- ROTH, A. E., AND EREV, I. (1993). "Learning in Extensive-Form Games: Experimental Data and Simple Dynamic Models in the Intermediate Term," mimeo, University of Pittsburgh; The Technion.
- ROTH, A. E., PRASNIKAR, V., OKUNO-FUJIWARA, M., AND ZAMIR, S. (1991). "Bargaining and Market Power in Jerusalem, Ljubljana, Pittsburgh, and Tokyo: An Experimental Study," *Amer. Econ. Rev.* **81**, 1068–1095.

- SAMUELSON, L. (1988). "Evolutionary Foundations of Solution Concepts for Finite, Two-Player, Normal-Form Games," in *Theoretical Aspects of Reasoning about Knowledge* (M. Y. Vardi, Ed.). San Mateo, CA: Morgan Kaufmann.
- SAMUELSON, L. (1993). "Does Evolution Eliminate Dominated Strategies?" *Frontiers of Game Theory* (K. Binmore, A. Kirman, and P. Tani, Eds). Cambridge, MA: MIT Press.
- SAMUELSON, L., (1994). "Stochastic Stability in Games with Alternative Best Replies," *J. Econ. Theory* **64**, 35–65.
- SAMUELSON, L., AND ZHANG, J. (1992). "Evolutionary Stability in Asymmetric Games," *J. Econ. Theory* **57**, 363–391.
- SCHLAG, K. H. (1994). "Why imitate, and if so, how?" mimeo, University of Bonn.
- SELTEN, R. (1975). "Reexamination of the Perfectness Concept for Equilibrium Points in Extensive-Form Games," *Int. J. Game Theory* **4**, 25–55.
- SELTEN, R. (1978). "The Chain-Store Paradox," *Theory Dec.* **9**, 127–159.
- SMITH, V. (1976). "Experimental Economics: Induced Value Theory," *Amer. Econ. Rev.* **66**, 274–279.
- TAYLOR, P. D., AND JONKER, L. B. (1978). "Evolutionarily Stable Strategies and Game Dynamics," *Math. Biosci.* **40**, 145–156.
- THALER, R. H. (1988). "Anomalies: The Ultimatum Game," *J. Econ. Perspect.* **2**, 185–206.
- VAN DAMME, E. (1987). *Stability and Perfection of Nash Equilibria*. Berlin: Springer-Verlag.
- YOUNG, P. (1993). "The Evolution of Conventions," *Econometrica* **61**, 57–84.