JOURNAL OF ECONOMIC THEORY 59, 145-168 (1993)

An Evolutionary Model of Bargaining*

H. P. YOUNG

School of Public Affairs/Department of Economics, University of Maryland, College Park, Maryland 20742

Received February 24, 1992; revised September 26, 1992

Individuals from two populations of bargainers are randomly matched to play the Nash demand game. They make their demands by choosing best replies based on an incomplete knowledge of the precedents, and occasionally they choose randomly. There is no common knowledge. Over the long run, typically one division will be observed almost all of the time. This "stochastically stable" division is close to the Nash solution when all agents in the same population are alike. When the populations are heterogeneous, a generalization of the Nash solution results. If there is some mixing between the two populations, the stable division is fifty-fifty. *Journal of Economic Literature* Classification Number: C78.

1. INTRODUCTION

Classical bargaining theory holds that bargains between rational agents depend solely on the agents' *utilities* for the outcomes, not the outcomes themselves. As Nash put it: "What the actual courses of action are among which the individuals must choose is not regarded as essential information... Only the attitudes (like or dislike) of the two individuals towards the ultimate results are considered." [12]. The assumption that only utility matters is implicit in almost all treatments of the bargaining problem. These fall into three broad categories. The axiomatic approach originally introduced by Nash seeks to deduce a solution from plausible first principles [9, 12, 14]. The noncooperative approach (also introduced by Nash) views the outcome as the equilibrium of a one-shot game in which the players make certain demands, which they get provided that the demands do not exceed the amount available [2, 13]. The third approach, introduced by Stahl [17] and Rubinstein [16], views the outcome as the

* This research was sponsored in part by the Santa Fe Institute. The author is indebted to Brian Arthur, David Canning, Ariel Rubinstein, Sangyoung Sonn, and several anonymous referees for comments on an earlier draft. Special thanks go to Dean Foster for providing suggestions and criticism along the way.

0022-0531/93 $5.00

subgame perfect equilibrium of a game in which the players alternate in making offers to one another.

In all three approaches the outcome is highly sensitive to the utility functions of the parties. It is also assumed that these utility functions are common knowledge, or at least that their distribution is common knowledge [2, 8]. Yet it is common knowledge that utility functions are often *not* common knowledge. When the utility functions are not known, bargainers may try to coordinate by appealing to custom, that is, by relying on what is usual and expected in the given situation. A well-known example is fifty-fifty division. Fifty-fifty is customary in many everyday bargains, and it is frequently observed in laboratory settings. Curiously, it is also quite common even when agents are in quite asymmetric positions, as in sharecropping agreements between tenants and landlords [1].

In this paper we propose to explain how such customs arise, and why some (like fifty-fifty) are more likely than others. The essential idea is that bargainers' expectations are shaped by precedent. Each bargainer has limited information about what other bargainers have demanded in past encounters. Each chooses a best reply assuming that these precedents are a reasonable predictor of what other agents are going to demand now. There is no common knowledge. When agents look at different precedents or the precedents conflict, their responses will very likely be uncoordinated. It may happen, however, that a succession of bargainers coordinate *by chance* on the same rule over a period of time. This establishes a set of common precedents and common expectations. People begin to expect that similar bargains will be solved by this rule in the future. Under circumstances that we shall describe below, this positive feedback loop eventually drives society toward a fixed division. It becomes the norm, not because it is inherently "focal," but because early chance events happened to favor it and now everyong expects it.

Of course, a given rule does not remain entrenched forever. Changing economic circumstances may eventually cause the shares to shift in favor of one party or another. Furthermore, there will inevitably be some variability in bargaining outcomes due to differences in agents' behavior and beliefs. When these stochastic influences are included, the long-run behavior of the evolutionary process becomes more complex. In particular, it has no absorbing states. A rule of division will remain in use for a time, after which it will be displaced by another rule. These displacements are irregular and infrequent, but they tend to occur quite rapidly when they do occur. The process then settles down to another period of relative tranquillity in which a fixed rule is in force. We cannot predict which rule will be in use at any given time because of the stochastic nature of the process. Nevertheless we can say that some rules are more likely to be observed than others over the long run. In fact, when the stochastic perturbations are very small, usually one rule will be observed almost all of the time. This *stochastically stable* rule can be calculated explicitly. It depends on the agents' utility functions, but oddly enough it does not depend on the details of the stochastic perturbations so long as they are small.

The results may be summarized as follows. When each class of agents is *homogeneous*—all members of the same class have the same utility function and the same amount of information—the long-run stable division is the one that corresponds to the asymmetric Nash solution. The asymmetry arises from informational differences between the two classes, where the class with more information (the larger sample size) has the advantage. This division does not result from calculations involving other agents' utility functions, but from its inherent stability *given* their utility functions.

When each class is *heterogeneous*, the stable division depends on the types of agents represented in each class, but not on their relative frequency in the population. It therefore generalizes the Nash solution to mixed populations in a substantially different way than Harsanyi and Selten's bargaining solution with incomplete information [8]. When there is even a small amount of mixing between the classes, the heterogeneity within each class no longer matters and the unique stable division is fifty–fifty. This suggests yet another reason (in addition to the usual focal point argument) why fifty–fifty is observed so often in practice, even when bargainers are in asymmetric positions.

2. The Model

Consider two finite, nonempty classes of individuals A and B. To be concrete we shall think of A as the class of landlords and B as the class of tenants. In each period, one landlord meets one tenant and they bargain over their shares of the crop. The share of the landlord will generically be denoted by x and the share of the tenant by y. For technical reasons we are going to assume that there are a finite number of feasible divisions. Let p be a positive integer, and let D be the set of all p-place decimal fractions that are positive and less than or equal to one. D is the set of *feasible demands*, and $\delta = 10^{-p}$ is the *precision* of the demands.

The bargaining process has the following structure. In each period $t = 1, 2, \ldots$, one landlord α is drawn at random from the class A and one tenant β is drawn at random from the class B. They then play the *Nash demand game*: α demands some fraction $x \in D$, β demands some fraction $y \in D$, and they get their demands if $x + y \leqslant 1$; otherwise they get nothing.

Let the demands in period t be denoted by (x_t, y_t). The complete *history* up to and including period t is the sequence $(x_1, y_1), (x_2, y_2), \ldots, (x_t, y_t)$.

Suppose that agents α and β are chosen in period $t+1$. Assume that neither party has prior knowledge or beliefs about the utility function of the other side, or about the distribution of the utility functions in the general population. To decide what to do they rely on the precedents that they happen to have heard about. Formally we may suppose that agent α draws a random sample of k_α items from the last m records $\mathbf{s} = ((x_{t-m+1}, y_{t-m+1}), ..., (x_t, y_t))$. The ratio k_α/m is a measure of α's *information*. We may think of k_α/m as the extent to which α is informed or "gets around." It is not the result of a process in which α sets out to gather an optimal amount of information. Agent α chooses an optimal reply to the observed frequency distribution of the demands made by tenants in the sample that he knows about. More precisely, α forms the cumulative probability distribution $F(y)$ of the demands y_j made by tenants in his sample:

$$\forall y \in D \; F(y) = h/k_\alpha \qquad \text{iff there are exactly } h \text{ demands } y_j \text{ in the sample such that } y_j \leq y.$$

F is a random variable that depends on the particular sample that α happened to draw at time $t+1$.

Agent α is assumed to have a von Neumann Morgenstern utility function $u_\alpha(x)$, where $u_\alpha(x)$ is concave (not necessarily strictly concave) and strictly increasing in α's own share. The utility function is normalized so that $u_\alpha(0) = 0$. We shall assume that $u_\alpha(x)$ is defined for all $x \in [0, 1]$, because later on we shall want to treat the level of precision δ as a variable.

If α demands x and the tenant demands y, then α gets x if $y \leq 1 - x$ and zero otherwise. Thus α believes his *expected payoff* (given the observed cumulative distribution F) is

$$u_\alpha(x)\, F(1-x) + u_\alpha(0)(1 - F(1-x)) = u_\alpha(x)\, F(1-x).$$

We assume that α chooses an amount x_{t+1} in period $t+1$ that maximizes his expected payoff:

$$x_{t+1} = \operatorname{argmax}\, u_\alpha(x)\, F(1-x). \tag{1}$$

If there are several values of x that maximize (1), then α chooses each of them with positive probability.

The tenant behaves similarly: β draws k_β items at random from the last m plays, and chooses an optimal response to the observed distribution of demands made by the landlords in his sample. All best replies are chosen with positive probability. The probability distribution over best replies may depend on the sample, but we assume it is stationary over time. The process is therefore a variation of fictitious play in which each agent reacts to a sample of the opponents' *recent* moves, rather than to all of the opponents' past moves. We shall call this process *adaptive play* [18].

We make no claim that this is how perfectly rational agents would play the game if they were aware of the process in which they are embedded. We are not assuming, however, that the agents in this model are perfectly rational or fully informed. They are merely *sensible*: they possess some information about the world around them, and they choose best replies based on this information. Note that we do not assume that agents learn through repeated plays of the game. We can assume, in fact, that after two individuals play they die and are replaced by agents of the same *type* as their forebears. Each has the same information-gathering propensity and the same utility function as its parent but not the parent's information. Thus each time an agent plays he starts afresh and must ask around to find out what is going on. This assumption is not necessary from a mathematical point of view, but it underscores the fact that "learning" need not occur at the individual level for the process to converge at the social level.

The agents' response rules determine a stationary Markov chain. The *state space* S consists of all sequences $\mathbf{s}$ of length m whose elements are pairs $(x, y) \in D \times D$. Let $p_\alpha(x \mid \mathbf{s})$ be the conditional probability that α demands x given that the state is $\mathbf{s}$, and let $p_\beta(y \mid \mathbf{s})$ be the conditional probability that β demands y given that the state is $\mathbf{s}$. We assume that p_α is a *best reply distribution*, that is, $p_\alpha(x \mid \mathbf{s}) > 0$ if and only if x is a best reply by α to a sample of size k_α drawn from $\mathbf{s}$. Similarly $p_\beta(y \mid \mathbf{s}) > 0$ if and only if y is a best reply by β to a sample of size k_β drawn from $\mathbf{s}$. The samples need not be drawn with uniform probability. For example, recent samples might be more likely than older samples. All that matters is that any appropriate-sized sample *could* be drawn, any best reply to such a sample *could* be made, and the probability of these events (conditional on the state) is stationary.

The process begins at time $t = m$ in some arbitrary initial state $\mathbf{s}^0$, that is, some arbitrarily chosen sequence of m pairs from $D \times D$. (This is clearly more general than assuming that the process begins at time $t = 1$ with a randomly chosen pair (x, y).) In each subsequent period one pair of agents $(\alpha, \beta) \in A \times B$ is drawn at random. Every pair has a positive probability $\pi(\alpha, \beta) > 0$ of being drawn, though it is not necessarily the same probability for all pairs. Given a state $\mathbf{s} = ((x_{t-m+1}, y_{t-m+1}), \ldots, (x_t, y_t))$ we say that $\mathbf{s}'$ is a *successor* of $\mathbf{s}$ if it has the form $\mathbf{s}' = ((x_{t-m+2}, y_{t-m+2}), \ldots, (x_t, y_t), (x_{t+1}, y_{t+1}))$. If the process is in state $\mathbf{s}$ at time t, then it moves to the successor state $\mathbf{s}'$ at time $t+1$ with transition probability

$$P_{\mathbf{ss}'} = \sum_{\alpha \in A} \sum_{\beta \in B} \pi(\alpha, \beta)\, p_\alpha(x_{t+1} \mid \mathbf{s})\, p_\beta(y_{t+1} \mid \mathbf{s}). \tag{2}$$

If $\mathbf{s}'$ is not a successor of $\mathbf{s}$, then $P_{\mathbf{ss}'} = 0$. This Markov process will be called the *evolutionary bargaining process* with precision δ, memory m, information parameters $\{k_\alpha/m, k_\beta/m\}$ and best reply distributions $\{p_\alpha, p_\beta\}$.

3. Convergence of the Evolutionary Bargaining Process

Convention. A state **s** is a *convention* if it consists of some fixed division $(x, 1-x)$ repeated m times in succession, where $x \in D$ and $0 < x < 1$. We shall denote this convention by **x**.

We claim that every convention is an absorbing state of the Markov process P defined by (2). Suppose, indeed, that at time $t \geqslant m$ the process is in the convention **x**. Let α and β be an arbitrary landlord and tenant who bargain in period $t+1$. In every sample drawn by the landlord, the previous tenants always demanded $1-x$, so the landlord's unique optimal reply is $x > 0$. Similarly, in every sample drawn by the tenant, the previous landlords always demanded x, so the tenant's unique optimal reply is $1-x > 0$. (This argument relies on the assumption that $0 < x < 1$.) Thus the optimal demands in period $t+1$ are $(x, 1-x)$, so the state in period $t+1$ is the same as it was in period t.

We now show that, if the information in each class is sufficiently *incomplete*, then the process P converges with probability one to a convention.

Theorem 1. *If at least one agent in each class samples at most half of the surviving records, then from any initial state the evolutionary bargaining process converges almost surely to a convention.*

Proof. We shall prove that there exists a positive integer N and a positive probability p (both independent of t) such that, from any state **s**, the probability of converging to a convention within N steps is at least p. Thus the probability of not reaching a convention within rN steps is at most $(1-p)^r$, which goes to zero as $r \to \infty$. Hence the set of all sample paths that do not land in an absorbing state have probability zero.

Fix two agents α and β, where α has minimum information among all agents in A and β has minimum information among all agents in B. Let α have information $k/m \leqslant 1/2$ and let β have information $k^*/m \leqslant 1/2$. We may assume without loss of generality that $k \geqslant k^*$. Suppose that at time $t \geqslant m$ the process is in state $\mathbf{s} = ((x_{t-m+1}, y_{t-m+1}), ..., (x_t, y_t))$. There is a positive probability that agents α and β (or agents just like them) will be selected to play the Nash demand game in *every* period from $t+1$ to $t+k$ inclusive. There is also a positive probability that α will draw the specific sample $\sigma = ((x_{t-k+1}, y_{t-k+1}), ..., (x_t, y_t))$ in each of these periods, and that β will draw the (possibly smaller) sample $\sigma^* = ((x_{t-k^*+1}, y_{t-k^*+1}), ..., (x_t, y_t))$ in each of these periods. Let x and y be best replies to these samples by α and β, respectively. Then there is a positive probability that α will demand x, and that β will demand y in every period from $t+1$ to $t+k$ inclusive. Hence there is a positive probability of obtaining a *run* $\rho = ((x, y), ..., (x, y))$ from periods $t+1$ to $t+k$ inclusive.

From periods $t+k+1$ to $t+2k$ there is a positive probability that two agents like α and β will be selected to play every time and that both of them will always sample from the run ρ. (Thus α draws the whole sample ρ, and β draws a subsample of k^* items from ρ.) Their unique best replies are $(1-y, 1-x)$, so we obtain a run $\rho' = ((1-y, 1-x), ..., (1-y, 1-x))$ from periods $t+k+1$ to $t+2k$ inclusive.

In period $t+2k+1$, there is a positive probability that a landlord like α will be chosen again and that she draws the sample ρ, and also that a tenant like β will be chosen again and that he draws the sample ρ^* consisting of the last k^* items in ρ'. Their best replies are then $(1-y, y)$. Note that this requires α to look back $2k$ periods, and β to look back k^* periods, both of which are possible because of the assumption that $k, k^* \leqslant m/2$.

In period $t+2k+2$, however, the oldest record in ρ may have disappeared (if $m = 2k$). Nevertheless, α can still draw a sample of size k in which the other side always demanded y, namely, the most recent $k-1$ plays in ρ and the play in the preceding period $t+2k+1$. Meanwhile β can still look at the sample ρ^* because none of these records has disappeared yet. So there is a positive probability that the best replies in period $t+2k+2$ will again by $(1-y, y)$. Continuing in this manner, we see that there is a positive probability of obtaining a run of $(1-y, y)$ for m periods in succession, at which point a convention has been reached.

It follows that, from a given initial state **s**, there is a positive probability (possibly depending on **s**) of reaching a convention within $2k+m$ periods. Since the number of states is finite, there is a positive probability p (independent of the initial state) of reaching a convention within $2k+m$ periods. This concludes the proof of Theorem 1.

The analog of Theorem 1 holds for a more general class of games [18], though a lower bound than $m/2$ may be needed to guarantee convergence (depending on the structure of the game). We do not claim that $m/2$ is the best bound possible in the present case, but without incomplete sampling the result fails. Consider the following example. Let $\delta = 0.1$ and let all agents in A have utility function u and all agents in B have utility function v, where u and v are any concave functions on D such that

$x =$	0	0.3	0.5	0.7	1
$u(x) =$	0	0.35	$0.4\sqrt{2}$	0.75	1
$v(x) =$	0	0.50	$0.5\sqrt{2}$	0.90	1

Suppose that both agents have complete information (they draw complete samples) and suppose that all previous pairs of demands were either

(0.7, 0.5) or (0.5, 0.3). Thus they miscoordinated every time. At time t, let f_t be the relative frequency with which the pair (0.7, 0.5) was played in the preceding m periods, and let $1-f_t$ be the relative frequency with which (0.5, 0.3) was played in the preceding m periods. The landlord's best reply will be either 0.5 or 0.7 depending on the value of f_t. It will be 0.7 if

$$(1-f_t)\, u(0.7) \geqslant u(0.5),$$

and 0.5 if the inequality runs the other way. Similarly, the tenant's best reply is 0.5 if

$$(1-f_t)\, v(0.5) \geqslant v(0.3),$$

and 0.3 if the inequality runs the other way. Substituting for the values of u and v we find that these two inequalities are the same, and hold if and only if $f_t < 1 - 1/\sqrt{2}$. (The inequality is always strict because f_t is rational.) Thus the demands in period $t+1$ are either (0.7, 0.5) or (0.5, 0.3). It follows that if the agents begin in any state where they have always miscoordinated, then they continue to miscoordinate forever.

4. Mistakes and Experimentation

The evolutionary bargaining process defined in (2) is based on the assumption that agents always choose best replies given their information. A more realistic assumption is that agents sometimes makes mistakes or experiment with other choices. These random choices or "trembles" keep the process constantly in motion and test the viability of different regimes, much like mutations in biological models of evolution.

Let $\varepsilon\lambda_\alpha$ be the probability that agent α experiments in any given period, and let $q_\alpha(x \mid \mathbf{s})$ be the conditional probability that α chooses the reply x, given that α is experimenting and that the current state is $\mathbf{s}$. Similarly define $\varepsilon\lambda_\beta$ and $q_\beta(y \mid \mathbf{s})$ for every $\beta \in B$. We assume that $q_\alpha(x \mid \mathbf{s})$ and $q_\beta(y \mid \mathbf{s})$ have full support, that is, all choices are possible in any given state. The parameters $\lambda_\alpha > 0$ and $\lambda_\beta > 0$ define the *relative* probabilities with which particular agents experiment, and the factor ε determines the *absolute* probability with which agents in general experiment. (This model generalizes an approach introduced by Canning [3] and Kandori, Mailath, and Rob [10].)

These random choices yield a stationary Markov chain on S in which the probability of moving from state $\mathbf{s} = ((x_{t-m+1}, y_{t-m+1}), \ldots, (x_t, y_t))$ at time $t \geqslant m$ to the successor state $\mathbf{s}' = ((x_{t-m+2}, y_{t-m+2}), ..., (x_{t+1}, y_{t+1}))$ at time $t+1$ is

$$P^{\varepsilon}_{\mathbf{ss}'} = \sum_{\alpha \in A} \sum_{\beta \in B} \pi(\alpha, \beta) \, [(1-\varepsilon\lambda_\alpha)(1-\varepsilon\lambda_\beta) \, p_\alpha(x \mid \mathbf{s}) \, p_\beta(y \mid \mathbf{s}) + \varepsilon\lambda_\alpha(1-\varepsilon\lambda_\beta) \, q_\alpha(x \mid \mathbf{s}) \, p_\beta(y \mid \mathbf{s}) + \varepsilon\lambda_\beta(1-\varepsilon\lambda_\alpha) \, p_\alpha(x \mid \mathbf{s}) \, q_\beta(y \mid \mathbf{s}) + \varepsilon^2\lambda_\alpha\lambda_\beta q_\alpha(x \mid \mathbf{s}) \, q_\beta(y \mid \mathbf{s})]. \quad (3)$$

$P^{\varepsilon}_{\mathbf{ss}'} = 0$ if $\mathbf{s}'$ is not a successor of $\mathbf{s}$. P^{ε} is the *perturbed* evolutionary bargaining process. Note that the *unperturbed* process P^0 is the same as the process P defined in (2).

Before showing how to compute the stable conventions in particular cases, let us first consider the qualitative behavior of the evolutionary bargaining process when the noise level ε is small. If agents never experiment ($\varepsilon = 0$), then Theorem 1 tells us that the process converges to a convention provided that some agent in each class samples at most half the surviving records. These are the absorbing states of the unperturbed process. If the agents occasionally experiment, then the process has no absorbing states. It still gravitates toward them, but never comes to rest because of the random fluctuations caused by experimentation. Nevertheless, the process is at or near an absorbing state (i.e., a convention) most of the time. Occasionally, a large accumulation of mistakes pushes the process away from a given convention and into the basin of attraction of another convention. The likelihood of such an event is small, but if we wait long enough, it is bound to occur. When it does, the process tends to gravitate toward the new convention. If it reaches it, the process is likely to stay there for awhile before another displacement occurs. This pattern of relative stability punctuated by episodes of instability continues indefinitely. Some conventions, however, are harder to displace than others. The ones that are hardest to displace are the ones that will be observed most often over the long run. When the noise is very small, in fact, the convention(s) that are hardest to displace will be observed *almost all of the time* over the long run.

Let us now make these observations more precise. The perturbed process P^{ε} is irreducible because of the assumption that the experimental distributions q_α, q_β have full support. Hence P^{ε} has a unique stationary distribution μ^{ε}. P^{ε} is also strongly ergodic, so for every $\mathbf{s} \in S$, $\mu^{\varepsilon}_{\mathbf{s}}$ is (with probability one) the relative frequency with which state $\mathbf{s}$ will be observed in the first t periods as $t \to \infty$.

STOCHASTICALLY STABLE CONVENTION. A convention $\mathbf{s}$ is *stochastically stable* if $\lim_{\varepsilon \to 0} \mu^{\varepsilon}_{\mathbf{s}}$ exists and is positive. It is *strongly stable* if $\lim_{\varepsilon \to 0} \mu^{\varepsilon}_{\mathbf{s}} = 1$.

Over the long run, stable conventions will be observed much more frequently than unstable conventions when the probability ε of the perturbations is small. A strongly stable convention will be observed almost *all* of the time when ε is small. This concept was introduced for general evolutionary processes by Foster and Young [6]. Applications of the concept to equilibrium selection may be found in Young [18], Young and Foster [19], Fudenberg and Harris [7], and Kandori, Mailath and Rob [10].

We now describe a general technique for computing the stable conventions, which always exist. Moreover, we show that generically there exists a strongly stable convention, which of course must be unique if it exists.

Mistake. Let $\mathbf{s} = ((x^1, y^1), (x^2, y^2), ..., (x^m, y^m))$ be some state and let $\mathbf{s}' = ((x^2, y^2), ..., (x^m, y^m), (x, y))$ be a successor of $\mathbf{s}$. x is a *mistake* in the transition $\mathbf{s} \to \mathbf{s}'$ if, for every landlord α, x is not a best reply by α to any sample of size k_α drawn from $\mathbf{s}$. Similarly, y is a mistake if, for every tenant β, y is not a best reply by β to any sample of size k_β drawn from $\mathbf{s}$.

Resistance. If $\mathbf{s}'$ is a successor of $\mathbf{s}$, the *resistance* $r(\mathbf{s}, \mathbf{s}')$ of the one-period transition $\mathbf{s} \to \mathbf{s}'$ is the minimum number of mistakes involved in the transition. Clearly $r(\mathbf{s}, \mathbf{s}') = 0$, 1, or 2. For every two states $\mathbf{s}^1$ and $\mathbf{s}^2$, the *resistance* $r(\mathbf{s}^1, \mathbf{s}^2)$ is the least total number of mistakes in any sequence of one-period transitions that leads from $\mathbf{s}^1$ to $\mathbf{s}^2$.

To compute $r(\mathbf{s}^1, \mathbf{s}^2)$ one solves a shortest path problem in a directed graph, namely, the graph in which every state is a vertex, there is an edge directed from $\mathbf{s}$ to $\mathbf{s}'$ if and only if $\mathbf{s}'$ is a successor of $\mathbf{s}$, and its weight is $r(\mathbf{s}, \mathbf{s}')$.

Now define another graph $\mathscr{G}$ as follows. There is one vertex for each convention $\mathbf{x}$, and a directed edge from every vertex to every other. The "weight" or resistance of the directed edge $\mathbf{x} \to \mathbf{x}'$ is the resistance $r(\mathbf{x}, \mathbf{x}')$ of moving from the convention $\mathbf{x}$ to the convention $\mathbf{x}'$.

$\mathbf{x}$-*tree.* An $\mathbf{x}$-*tree* is a collection of edges in $\mathscr{G}$ such that, from every vertex $\mathbf{x}' \neq \mathbf{x}$ there is a unique directed path to $\mathbf{x}$, and there are no cycles. Let $\mathscr{T}_\mathbf{x}$ be the set of all $\mathbf{x}$-trees.

Stochastic Potential. The *stochastic potential* of the convention $\mathbf{x}$ is the least resistance among all $\mathbf{x}$-trees:

$$\gamma(\mathbf{x}) = \min_{T \in \mathscr{T}_\mathbf{x}} \sum_{(\mathbf{x}', \mathbf{x}'') \in T} r(\mathbf{x}', \mathbf{x}'')$$

The following is a special case of Theorem 4 in [18].

THEOREM 2. *The sequence of stationary distributions μ^ε converges to a* *stationary distribution* μ^0 *of* P^0 *as* $\varepsilon \to 0$. *Moreover,* $\mathbf{s}$ *is stochastically stable* ($\mu^0_{\mathbf{s}} > 0$) *if and only if* $\mathbf{s} = \mathbf{x}$ *is a convention and* $\gamma(\mathbf{x})$ *has minimum stochastic potential among all conventions.*

In the remainder of the paper we shall characterize the stable conventions under various assumptions about the composition of the bargaining populations.

5. Homogeneous Populations and the Nash Solution

Let a be a rational fraction, $0 < a < 1$, and let $u(x)$ be a concave utility function defined for all $x \in [0, 1]$. An agent is of *type* (a, u) if he always samples the fraction a of the last m records and his utility function is u. Let A and B be the two classes of agents, and suppose that each class is *homogeneous*: all members of A are of the same type (a, u), and all members of B are of the same type (b, v). Since the classes A and B are treated symmetrically in the model, there is no loss of generality in assuming that $a \geqslant b$. We shall also assume that m is chosen so that ma and mb are integer.

Generic Stability. A division $(x, 1-x)$ is *generically stable* for types (a, u) and (b, v) if the associated convention $\mathbf{x}$ is stochastically stable for all admissible m.

Asymmetric Nash Bargaining Solution. Given types (a, u) and (b, v), the *asymmetric Nash bargaining solution* is the unique division $(x, 1-x)$ that maximizes

$$u^a(x)\, v^b(1-x) \qquad \text{subject to } 0 \leqslant x \leqslant 1. \tag{4}$$

When $a = b$ this reduces to the ordinary Nash bargaining solution.

Theorem 3. *Let A and B be homogeneous populations composed respectively of types* (a, u) *and* (b, v), *where* $a, b \leqslant 1/2$. *For every precision* $\delta > 0$ *there exists at least one and at most two generically stable divisions, and as* $\delta \to 0$ *they converge to the asymmetric Nash bargaining solution.*

Stated informally, this result says that if memory is sufficiently large and the precision is sufficiently small, most people will be using the same rule of division most of the time, and this rule will be close to the asymmetric Nash solution. In particular, agents who are risk averse or poorly informed will fare worse than those who are not, all else being equal. For example, if the landlords are risk neutral and the tenants are risk averse, and if both have the same amount of information ($a = b$), then most of the time the landlords will be getting the larger share. If both sides are risk neutral, but the landlords have more information ($a > b$), then again the landlords will get the larger share most of the time. The reason is that agents with more information are less likely to respond to mistakes by the other side, so they are steadier. Similarly, agents who are less risk averse are more likely to take chances, so they are more demanding.

The proof of Theorem 3 proceeds by a series of lemmas that provide a sharp estimate of the stable conventions for *any* choice of parameters δ and m.

Fix the precision $\delta = 10^{-p}$ and let $D^0 = \{x \in D: \delta \leq x \leq 1 - \delta\}$. Fix $a, b \leq 1/2$ and let m be such that ma and mb are integer. For every real number r let $\lceil r \rceil$ denote the least integer greater than or equal to r. Also, let $r \wedge r'$ denote the minimum of r and r'.

x-*Basin*. For every $x \in D^0$ the **x**-*basin* is the set of all states from which the unperturbed process P^0 converges to the convention **x** with positive probability.

LEMMA 1. *For every $x \in D^0$ the minimum resistance to moving from* **x** *to a state in some other basin is* $\lceil m r_\delta(x) \rceil$, *where*

$$r_\delta(x) = a(1 - u(x - \delta)/u(x)) \wedge b(1 - v(1 - x - \delta)/v(1 - x)) \wedge bv(1 - x)/v(1 - \delta). \quad (5)$$

Before proving Lemma 1, we give the intuition behind it and show how it leads to a proof of Theorem 3. The three terms in expression (5) arise in the following way. To displace an established convention **x** requires one side or the other to demand something other than the conventional amount. Suppose that the tenants demand just a little more than they should, say δ more. The landlords will resist this, and the amount of their resistance is the relative loss in utility that they would suffer by giving up a δ-increment of their current share x, times their sample size. This yields the first term $a(1 - u(x - \delta)/u(x))$. Similarly, if the landlords demand δ more than the conventional amount the tenants will resist, and the extent of their resistance is the relative loss in utility that they would suffer by giving up δ, times their sample size. This yields the second term $b(1 - v(1 - x - \delta)/v(1 - x))$.

A third possibility is that some landlord demands *less* than the conventional amount by mistake. This would be a "silly" mistake, of course, but it could happen. The silliest mistake of all would be to demand only δ. (Recall that we exclude demands of 0, because it makes little sense to "demand" nothing.) If the landlords make enough mistakes of this sort, the tenants will switch to $1 - \delta$, and their resistance to switching is given by the third term, $bv(1-x)/v(1-\delta)$. Of course, the tenants might also make silly mistakes. If they demand δ often enough, the landlords will switch to $1-\delta$. It can be shown, however, that when $a \geqslant b$ the resulting term $au(x)/u(1-\delta)$ is never *strictly* smaller than the term $b(1-v(1-x-\delta)/v(1-x))$ for any concave utility functions u and v. Hence mistakes of the second type have no effect on the long-run stability of various divisions.

The function $r_\delta(x)$ is the minimum of three monotone functions. The first is strictly decreasing in x, the second is strictly increasing in x, and the third is strictly decreasing in x. Hence, as x increases, $r_\delta(x)$ is first strictly increasing and then strictly decreasing. On the subset D it therefore achieves its maximum at one, or at most two, values. We shall show that these are the generically stable divisions, and that, as δ approaches 0, they converge to the asymmetric Nash solution (Lemmas 2 and 3 below).

EXAMPLE 1. Let all agents in A sample 1/3 of the surviving records and have utility function $u(x)=\sqrt{x}$. Let all agents in B sample 1/10 of the surviving records and have utility function $v(y)=y$. The asymmetric Nash solution is (5/8, 3/8). Let $\delta=0.1$. Then $r_\delta(x)$ is the minimum of the following three functions

$$f_1(x)=(1/3)(1-u(x-\delta)/u(x))=(1/3)(1-\sqrt{(1-0.1/x)})$$

$$f_2(x)=(1/10)(1-v(1-x-\delta)/v(1-x))=0.01/(1-x)$$

$$f_3(x)=(1/10)\,v(1-x)/v(1-\delta)=(1/9)(1-x).$$

The function $r_\delta(x)=\min\{f_1(x), f_2(x), f_3(x)\}$ is graphed in Fig. 1. It achieves its maximum on D at the value $x_\delta=0.6$, so by Theorem 3, (0.6, 0.4) is the generically stable division.

Proof of Lemma 1. Let the types (a, u) and (b, v) be given, where $a, b \leqslant 1/2$. Assume without loss of generality that $a \geqslant b$, and let $k=ma$. Suppose that the process is at the convention $\mathbf{x}$, where $x \in D^0=\{x \in D: \delta \leqslant x \leqslant 1-\delta\}$. Let π be a path of least resistance from $\mathbf{x}$ to a state that is in some other basin. Clearly, π must pass through some state $\mathbf{s}$ such that (i) some landlord's best reply to a sample from $\mathbf{s}$ is different from x, and/or (ii) some tenant's best reply to a sample from $\mathbf{s}$ is different from $1-x$. Let $\mathbf{s}$ be the first such state. Without loss of generality we can assume that some landlord (say α) has a best reply $x' \neq x$ to a sample σ of size k drawn from $\mathbf{s}$. By choice of $\mathbf{s}$, every tenant's demand in $\mathbf{s}$ that differs from $1-x$ must have been a mistake. Let $i \leqslant k$ be the total number of mistakes by tenants in the sample σ that the landlord drew.

Now construct another sequence σ' from σ by replacing each mistaken demand of the tenants by the demand $1-x'$. Thus in σ' there are exactly i instances of $1-x'$ on the tenants' side, and $k-i$ instances of $1-x$.

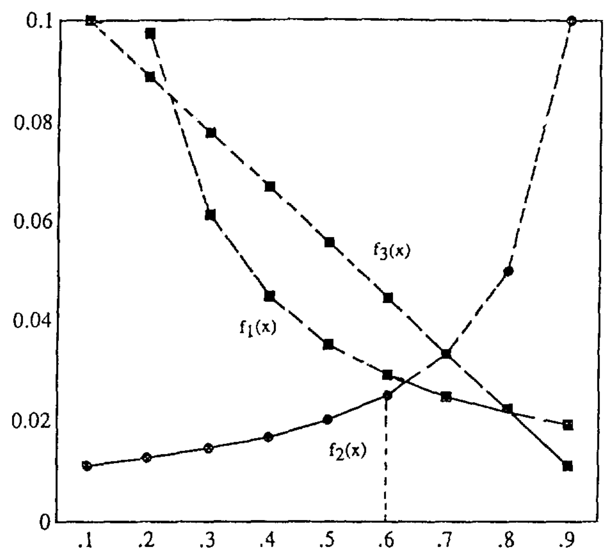

FIG. 1. The function $r_{\delta}(x)=\min\{f_1(x), f_2(x), f_3(x)\}$ for Example 1.

Landlord α's best reply to σ' is still x', because his best reply to σ was x', and the optimality of x' is not compromised when the tenants demand $1-x'$ more often and other things less often.

We now construct an alternative path π' from **x** to the **x**'-basin in which the total number of mistakes is i, from which it follows that π' is also a path of least resistance. Beginning at the convention **x**, let a succession of tenants mistakenly demand $1-x'$ i times in a row. With positive probability the landlord α' drawn in the next period will sample the most recent k plays. This sample σ'' consists of i instances of $1-x'$ and $k-i$ instances of $1-x$. The relative frequency of $1-x'$ is the same in σ' and σ''. Since α' and α have the same utility function, and x' is the best reply by α to σ', it follows that x' is also the best reply by α' to σ''.

With positive probability the landlords will sample σ'' for k periods in succession and reply with x' each time. This establishes a run of x' (by the landlords) without introducing more mistakes. After this there is a positive probability that a succession of k tenants will sample from this run and demand $1-x'$ for k periods in a row. From here it is clear (as in the proof of Theorem 1) that the process converges with positive probability to the convention **x**' with no further mistakes.

Thus we have constructed an alternative least-resistant path from the convention **x** to the convention **x**' in which the only mistakes are an initial succession of i mistaken demands $1-x'$. To compute the least number of mistakes necessary to exit from the **x**-basin, it therefore suffices to consider, for every $x' \neq x$, the least number of initial mistakes $1-x'$ by tenants that would cause a landlord to reply with $x' \neq x$, and the least number of initial mistakes x' by landlords that would cause a tenant to reply with $1-x'$, and to take the smaller of these two numbers.

Choose an arbitrary $x' \neq x$. We distinguish two cases: $x' < x$ and $x' > x$.

Case 1: $x' < x$.

Let the process be in the convention $\mathbf{x}$. Suppose that the landlords make i successive demands of x' that cause some tenant's best reply to switch to $1-x'$ instead of $1-x$. We can assume that $i \leqslant mb$, which is the tenant's sample size. When a tenant samples these i mistaken demands x', together with $mb-i$ of the previous "conventional" demands x, he switches to $1-x'$ provided that $(i/mb)\, v(1-x') \geqslant v(1-x)$, that is

$$i \geqslant mbv(1-x)/v(1-x').$$

Over all feasible $x' < x$ the minimum value of i occurs when $x' = \delta$ and

$$i = [mbv(1-x)/v(1-\delta)]. \tag{6}$$

The other possibility is that the tenants make the mistakes. Suppose that a succession of $j \leqslant k = ma$ tenants demand $1-x'$ by mistake instead of $1-x$. (Note that these are only "mistakes" in the sense that they are not best replies to the existing set of precedents. The tenants would obviously like to get $1-x'$ rather than $1-x$ if they thought they could get away with it.) If some landlord samples the j mistaken demands of $1-x'$, together with $k-j$ of the previous conventional demands $1-x$, then he switches to x' provided that $u(x') \geqslant (1-j/k)\, u(x)$, that is,

$$j \geqslant k(1-u(x')/u(x)) = ma(1-u(x')/u(x)).$$

Over all feasible $x' < x$, the minimum such j occurs when $x' = x-\delta$. Hence

$$j = [ma(1-u(x-\delta)/u(x))]. \tag{7}$$

Case 2: $x' > x$.

An analogous argument shows that, if the landlords make i mistaken demands of x', then some tenant switches provided that

$$i \geqslant mb(1-v(1-x')/v(1-x)).$$

The minimum occurs when $x' = x+\delta$:

$$i = [mb(1-v(1-x-\delta)/v(1-x))]. \tag{8}$$

If the tenants make j mistaken demands of $1-x'$, then some landlord switches provided that

$$j \geqslant mau(x)/u(x').$$

Since all demands are positive, the minimum occurs when $1-x'=\delta$, that is,

$$j=\lceil mau(x)/u(1-\delta)\rceil. \tag{9}$$

Combining (6)–(9) it follows that the least number of mistakes to exit from the **x**-basin is $\lceil mr_\delta(x)\rceil$, where

$$r_\delta(x)=a(1-u(x-\delta)/u(x)) \wedge b(1-v(1-x-\delta)/v(1-x))$$
$$\wedge\ bv(1-x)/v(1-\delta) \wedge au(x)/u(1-\delta).$$

We claim that the last term of this expression is at least as large as the second term for all $x \in D^0$, that is,

$$au(x)/u(1-\delta) \geqslant b(1-v(1-x-\delta)/v(1-x)). \tag{10}$$

Indeed, since v is concave, we have

$$(v(1-x)-v(1-x-\delta))/\delta \leqslant (v(1-x)-v(0))/(1-x),$$

because the rate of loss of utility in going from $1-x$ to $1-x-\delta$ is no more than the rate of loss of utility in going from $1-x$ to 0. Since $v(0)=0$ this implies that

$$1-v(1-x-\delta)/v(1-x) \leqslant \delta/(1-x). \tag{11}$$

Similarly, since u is concave, we have

$$(u(1-\delta)-u(0))/(1-\delta) \leqslant (u(x)-u(0))/x,$$

from which it follows that

$$u(x)/u(1-\delta) \geqslant x/(1-\delta). \tag{12}$$

It may be checked that $x/(1-\delta) \geqslant \delta/(1-x)$ whenever $\delta \leqslant x \leqslant 1-\delta$. Combining this with (11) and (12) we therefore have

$$u(x)/u(1-\delta) \geqslant x/(1-\delta) \geqslant \delta/(1-x) \geqslant 1-v(1-x-\delta)/v(1-x).$$

This, together with the assumption that $a \geqslant b$, proves (10). Hence the right-most term in $r_\delta(x)$ can be omitted, so $r_\delta(x)$ is given by (5). This completes the proof of Lemma 1.

LEMMA 2. *A division $(x, 1-x)$ is generically stable if and only if x maximizes the function $r_\delta(x)$ on D^0.*

Proof. The function $r_\delta(x)$ is first strictly increasing in x and then strictly decreasing in x. It therefore achieves its maximum on D^0 either at a unique value x_δ or at two adjacent values x_δ and $x_\delta + \delta$. Suppose first that the maximum is achieved at a *unique* value x_δ. Construct a directed tree T_δ having root $\mathbf{x}_\delta$ as follows:

(i) for every $x \in D^0$ such that $x < x_\delta$, put the directed edge $(\mathbf{x}, \mathbf{x}+\delta)$ in T_δ and let its weight be $[mb(1 - v(1-x-\delta)/v(1-x))]$;

(ii) for every $x \in D^0$ such that $x > x_\delta$ and $a(1 - u(x-\delta)/u(x)) \leqslant bv(1-x)/v(1-\delta)$, put the directed edge $(\mathbf{x}, \mathbf{x}-\delta)$ in T_δ and let its weight be $[ma(1-u(x-\delta)/u(x))]$;

(iii) for every $x \in D^0$ such that $x > x_\delta$ and $a(1-u(x-\delta)/u(x)) > bv(1-x)/v(1-\delta)$, put the directed edge $(\mathbf{x}, \delta)$ in T_δ and let its weight be $[mbv(1-x)/v(1-\delta)]$.

Then T_δ has the following structure. Every node $\mathbf{x}$ except $\mathbf{x}_\delta$ has a unique outgoing edge. If $\mathbf{x}$ lies to the *left* of $\mathbf{x}_\delta$ the outgoing edge is directed to $\mathbf{x}$'s immediate neighbor on the right (case *i*). If $\mathbf{x}$ lies to the *right* of $\mathbf{x}_\delta$ the outgoing edge is directed *either* to $\mathbf{x}$'s immediate neighbor to the left (in case ii) *or* to the node δ (in case iii). From every vertex other than $\mathbf{x}_\delta$ there exists a unique path to $\mathbf{x}_\delta$. Thus T_δ is an $\mathbf{x}_\delta$-tree. Since T_δ was constructed by choosing at each node the outgoing edge with least resistance, it is clear that T_δ has least resistance among all $\mathbf{x}_\delta$-trees. Fig. 2 illustrates the structure of T_δ for Example 1.

We claim that T_δ has least resistance among *all* $\mathbf{x}$-trees. To prove this, let T be an $\mathbf{x}$-tree for some $x \neq x_\delta$. In T, $\mathbf{x}_\delta$ has a unique outgoing edge e. Let r be the resistance of e. By Lemma 1, $[mr_\delta(x_\delta)]$ is the minimum resistance among *all* possible edges directed away from $\mathbf{x}_\delta$, so $r \geqslant [mr_\delta(x_\delta)]$. By construction, $[mr_\delta(x)]$ is the resistance of the unique edge exiting from $\mathbf{x}$ in the tree T_δ. Further, for every node $\mathbf{x}'$ different from both $\mathbf{x}_\delta$ and $\mathbf{x}$, the resistance of the unique outgoing edge from $\mathbf{x}'$ in T_δ is no greater than the resistance of the unique outgoing edge from $\mathbf{x}'$ in T. Hence

$$r(T) \geqslant r(T_\delta) + r - [mr_\delta(x)].$$

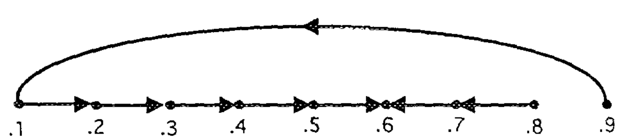

FIG. 2. The x_δ-tree of least resistance for Example 1 ($x_\delta = 0.6$).

Since

$$r \geqslant [mr_\delta(x_\delta)] \geqslant [mr_\delta(x)],$$

it follows that $r(T) \geqslant r(T_\delta)$ as claimed. Moreover, since $\mathbf{x}_\delta$ is the unique maximum of r_δ, we have $[mr_\delta(x_\delta)] > [mr_\delta(x)]$ for all sufficiently large m, so $r(T) > r(T_\delta)$. It follows from Theorem 2 that $\mathbf{x}_\delta$ is a stochastically stable convention for all m, and it is the unique stochastically stable convention for all sufficiently large m.

If r_δ is maximized at two neighboring values $\mathbf{x}_\delta$ and $\mathbf{x}_\delta + \delta$, construct a tree for each of them as above. A similar argument shows that these are the $\mathbf{x}$-trees of least resistance for all sufficiently large m. Hence $\mathbf{x}_\delta$ and $\mathbf{x}_\delta + \delta$ are the unique generically stable conventions. This concludes the proof of Lemma 2.

LEMMA 3. *The maxima of the function $r_\delta(x)$ converge to the asymmetric Nash solution as $\delta \to 0$.*

Proof. For each δ, view $r_\delta(x)$ as a function defined on the whole interval $[0, 1]$. Let x'_δ be the *unique* real value at which the maximum is achieved. It is clear that $|x_\delta - x'_\delta| \leqslant \delta$, where x_δ maximizes $r_\delta(x)$ on D^0. To prove Lemma 3, it suffices to show that the sequence $\{x'_\delta\}$ converges to the asymmetric Nash solution as $\delta \to 0$.

For each precision δ define $f_\delta(x) = r_\delta(x)/\delta$. Clearly x'_δ also maximizes $f_\delta(x)$ on $[0, 1]$. As δ goes to zero, the third term of $f_\delta(x)$ becomes large while the other two terms are bounded. Hence the third term may be ignored in the limit (since $f_\delta(x)$ is the minimum of the three terms). Let $u^{-}(x)$ denote the left-hand derivative of u at x and let $u^{+}(x)$ denote the right-hand derivative of u at x (both of which exist). Then for each x in $(0, 1)$ we have

$$\lim_{\delta \to 0} f_\delta(x) = au\ \ (x)/u(x) \wedge bv^{-}(1-x)/v(1-x) = f(x).$$

Since u and v are concave and bounded, the function $f(x)$ defined in the above expression is strictly quasiconcave and upper semicontinuous. Hence it attains its maximum at a unique value x^*. We claim that x^* is the asymmetric Nash bargaining solution. To see this, suppose first that $u(x)$ and $v(y)$ are differentiable. Then $u'(x) = du(x)/dx$ and $v'(y) = dv(y)/dy$ are continuous and nonincreasing. Thus $au'(x)/u(x)$ is continuous and strictly decreasing, while $bv'(1-x)/v(1-x)$ is continuous and strictly increasing. The maximum of $f(x)$ occurs where the curves $au'(x)/u(x)$ and $bv'(1-x)/v(1-x)$ cross, that is, at the unique value $x^* \in (0, 1)$ such that

$$au'(x^*)/u(x^*) = bv'(1-x^*)/v(1-x^*).$$

This is a necessary and sufficient condition for x^* to be the unique point where the strictly concave function $a \ln u(x) + b \ln v(1-x)$ is maximized. It follows that x^* maximizes $u^a(x)\, v^b(1-x)$, so it is the asymmetric Nash solution.

Suppose now that u and v are not differentiable. Nevertheless they are subdifferentiable. For each $x \in (0, 1)$ consider the set of all lines in R^2 that pass through the point $(x, u(x))$ and are tangent to the curve $u(x)$. The slopes of these lines are the *subgradients* of u at x, and the set of them is the *subdifferential* at x, denoted by $\partial u(x)$ [15]. In general we have $\delta u(x) = \{r : u^-(x) \geqslant r \geqslant u^+(x)\}$. Similarly define $\partial v(y)$ for every $y \in (0, 1)$.

Suppose that f attains its maximum at x^*. Then for all small $\varepsilon > 0$

$$au^-(x^*)/u(x^*) \geqslant f(x^*) \geqslant au^-(x^* + \varepsilon)/u(x^* + \varepsilon).$$

It is a standard result that $u^-(x^* + \varepsilon)/u(x^* + \varepsilon) \to u^+(x^*)/u(x^*)$ as $\varepsilon \to 0$ [15, p. 228]. Hence

$$au^-(x^*)/u(x^*) \geqslant f(x^*) \geqslant au^+(x^*)/u(x^*),$$

and therefore,

$$f(x^*) \in a\, \partial u(x^*)/u(x^*). \tag{13}$$

Similarly, conclude that

$$f(x^*) \in b\, \partial v(1 - x^*)/v(1 - x^*). \tag{14}$$

Consider the strictly concave function $F(x) = a \ln u(x) + b \ln v(1-x)$. Then [15, p. 223]

$$\begin{aligned} \partial F(x) &= \partial(a \ln u(x)) + \partial(b \ln v(1-x)) \\ &= a\, \partial u(x)/u(x) - b\, \partial v(1-x)/v(1-x). \end{aligned} \tag{15}$$

In particular this holds for x^*. From this and (13)–(14) it follows that $0 \in \partial F(x^*)$. Therefore x^* is the unique minimum of F (see [15, p. 264]), so by definition x^* is the asymmetric Nash solution.

The only point that remains to be verified is that, if x_δ^* maximizes $f_\delta(x)$ for every δ, then the sequence $\{x_\delta^*\}$ converges to x^*. Suppose, by way of contradiction, that x' is an accumulation point of the sequence $\{x_\delta^*\}$ that differs from x^*. For notational simplicity we shall assume that $\{x_\delta^*\}$ actually converges to x'. For each δ let

$$C_\delta = \{x \in [0, 1] : f_\delta(x) \geqslant f_\delta(x^*)\}.$$

C_δ is convex, closed, and contains both x^* and x_δ^*. Hence it contains the interval I_δ spanned by x^* and x_δ^*. Since x_δ^* converges to x', the intervals

I_δ converge to the interval I spanned by x' and x^*. Fix some z strictly between x^* and x'. For all sufficiently small δ, $z \in I_\delta$ and hence $z \in C_\delta$. It follows by definition of C_δ that $f_\delta(z) \geqslant f_\delta(x^*)$ for all δ. Hence $f(z) \geqslant f(x^*)$ because f_δ converges pointwise to f. This is a contradiction, however, because x^* uniquely maximizes f. This completes the proof of Lemma 3, and Theorem 3 follows at once.

The above proof shows that Theorem 3 holds under a variety of other assumptions about how agents make mistakes. For example, suppose that agents only make "small" mistakes. In other words, imagine that an agent first samples the past and determines the best reply, but that occasionally he demands a little bit more (or a little bit less) than he should—say δ more or δ less. Then the expression for the resistance function r_δ in (5) is the same except that the third term $bv(1-x)/v(1-\delta)$ is deleted. Since in any event this term is not relevant when δ goes to zero (because it is strictly larger than the first two terms), the proofs of Lemmas 2 and 3 go through much as before. Hence the generically stable division(s) are close to the asymmetric Nash solution when δ is close to zero.

Theorem 3 differs in several important respects from other ways of deriving the Nash solution based on perturbations. In his 1953 paper, for example, Nash suggested that his bargaining solution is the unique limiting outcome of the Nash equilibria of a perturbed version of the Nash demand game in which there is a small uncertainty in the payoff function [13]. A more rigorous argument along these lines was provided by Binmore [2]. Carlsson [4] investigated variations of the Nash demand game in which agents' demands are modified by a small error term, and showed that the Pareto optimal Nash equilibria of this noncooperative game converge to the Nash solution as the error term becomes vanishingly small.

These models of equilibrium stability differ from the present one in that they presuppose that the outcome of the perturbed game *is* a Nash equilibrium. To justify this approach, it must be assumed (among other things) that the properties of the game are common knowledge among the agents. In our model, there is no common knowledge and no assumption that a Nash equilibrium will be played. Instead, equilibrium emerges (in an asymptotic sense) as a consequence of the long-run dynamics, without the players being aware of it.

6. Heterogeneous Populations and a Generalization of the Nash Solution

Theorem 3 shows how the Nash bargaining solution can be selected by a dynamic process that does not require any common knowledge on the part of the agents. This conclusion is based, however, on the special assumption that all agents in the same class have the same utility function. Moreover, the result is very sensitive to this assumption. If just *one* agent has a utility function that differs from the norm in his or her class, the long-run outcome of the process may be completely different from the one described in Theorem 3. The reason is that the long-run behavior of the process is determined by those who would suffer the *greatest proportional loss in utility* from a given reduction in share. If there is just one landlord, for example, who is more risk averse than the others in this sense, then his choices will determine the long-run behavior of the process despite there being many other landlords who are less risk averse.

Let us therefore consider the more general situation in which each population of agents contains various types of individuals. Let T_A be the finite set of types represented in class A, and T_B the set of types represented in class B. Let a^* be the least a among all (a, u) in T_A and let b^* be the least b among all (b, v) in T_B. We may assume without loss of generality that $a^* \geqslant b^*$. We shall also assume that $a^*, b^* \leqslant 1/2$. Then, just as in the proof of Lemma 1, it follows that the minimum resistance to moving from the convention $\mathbf{x}$ into the basin of attraction of some other convention is $[mR_\delta(x)]$, where

$$R_\delta(x) = \min_{(a,u)\in T_A} a(1-u(x-\delta)/u(x)) \wedge \min_{(b,v)\in T_B} b(1-v(1-x-\delta)/v(1-x)) \wedge \min_{(b,v)\in T_B} bv(1-x)/v(1-\delta).$$

As before we are able to omit the terms of form $au(x)/u(1-\delta)$, because for *every* type $(a, u) \in T_A$ we have

$$b^*(1-v^*(1-x-\delta)/v^*(1-x)) \leqslant au(x)/u(1-\delta),$$

where (b^*, v^*) is a type with minimum information in T_B.

It is easy to see that $R_\delta(x)$ is unimodal. An argument analogous to the proof of Lemma 2 shows that $(x, 1-x)$ is generically stable if and only if x maximizes $R_\delta(x)$. Hence there are at most two generically stable conventions and they differ by δ. An argument similar to the proof of Lemma 3 shows that these stable conventions converge to the unique maximum of the strictly quasiconcave function:

$$R(x) = \min_{(a,u)\in T_A} au'(x)/u(x) \wedge \min_{(b,v)\in T_B} bv'(1-x)/v(1-x). \qquad (16)$$

THEOREM 4. *Let A and B be two finite populations of agents and let T_A and T_B be the types represented in each class, where at least one agent in each class samples at most half the records. For every precision $\delta > 0$ there* *exists at least one and at most two generically stable divisions, and as* $\delta \to 0$ *they converge to the unique division* $(x, 1-x)$ *such that* x *maximizes* $R(x)$.

The solution described in Theorem 4 is a generalization of the Nash solution to heterogeneous populations of bargainers that we shall call the *heterogeneous bargaining solution.* It differs from Harsanyi and Selten's solution [8], which depends on the relative frequency with which the utility functions are represented in each class. Specifically, the Harsanyi Selten solution maximizes the product

$$\prod_{u_i \in A} u_i^{\lambda_i}(x) \prod_{v_j \in B} v_j^{\mu_j}(1-x),$$

subject to $0 \leqslant x \leqslant 1$, where λ_i is the relative frequency of u_i in population A and μ_j is the relative frequency of v_j in population B. The heterogeneous bargaining solution, by contrast, depends only on the types of bargainers represented in each population, not on their relative frequency.

EXAMPLE 2. Let population A consist of two types of agents: ($a_1 = 1/4$, $u_1(x) = x$) and ($a_2 = 1/3$, $u_2(x) = x^{1/3}$). Population B consists of two other types of agents: ($b_1 = 1/5$, $v_1(y) = y$) and ($b_2 = 1/2$, $v_2(y) = y^{1/2}$). In population A the criterion $au'(x)/u(x)$ is minimized for every x by the second type: $a_2 u_2'(x)/u_2(x) = 1/(9x)$. In population B the criterion $bv'(y)/v(y)$ is minimized for every y by the first type: $b_1 v_1'(y)/v_1(y) = 1/(5y)$. Hence $R(x) = 1/(9x) \wedge 1/(5(1-x))$. It achieves its unique maximum when $1/(9x) = 1/(5(1-x))$, that is, when $x = 5/14$. In other words, the stable division is $(5/14, 9/14)$ for any two populations A and B consisting of the types mentioned above, irrespective of the distribution of these types within each population.

7. MIXING BETWEEN CLASSES AND THE FIFTY-FIFTY SOLUTION

Under some circumstances it is reasonable to expect that the types represented in each class will be the same, or at least the same with positive probability. For example, suppose that there is some mobility between classes. Every so often a landlord (or the child of a landlord) loses his land and becomes a tenant, and every so often a tenant climbs into the ranks of the landlords. Then at any given time there is a positive probability that the two bargainers could be *any* pair of types drawn from *either* of the two original classes.

This idea can be made more general as follows. Let T be the set of all possible types that either the landlords or tenants could be. Let $\pi(\tau, \tau')$ be the probability that the type-pair $(\tau, \tau') \in T \times T$ is chosen to bargain in any given period. The first entry τ denotes the landlord's type and the second entry τ' denotes the tenant's type. Up to now we have been assuming that T is the union of two classes T_A and T_B and that $\pi(\tau, \tau') > 0$ if and only if $\tau \in T_A$ and $\tau' \in T_B$. The above arguments and theorems generalize, however, to other probability distributions π defined on $T \times T$.

In particular, we say that the distribution π *mixes roles* if every pair in $T \times T$ occurs with positive probability. Mixing is a natural consequence of mobility between classes. It also holds if the classes are rigid but tenants and landlords are drawn from the same "gene pool": every individual has a positive probability of being any type τ, though the probability of being a τ may differ for landlords and tenants. When π mixes roles, it is straightforward to verify that Theorem 4 still holds with $T_A = T_B = T$ in the expression for $R(x)$. This implies that $R(x)$ is symmetric about one-half.

COROLLARY 4.1. *Let T be a common set of types and let $\pi(\tau, \tau') > 0$ be the probability that any given landlord and tenant are of types τ and τ', respectively. Then fifty-fifty is the unique generically stable division.*

This result does not say that society ultimately locks into fifty–fifty and stays there forever. Rather, the evolutionary process favors a solution near fifty–fifty when there is mixing between the classes of bargainers. In other words, if memory is large, the noise is small, and there is some mixing between classes, the chances are good that the customary division will be about fifty–fifty.

REFERENCES

1. P. BARDHAN, "Land, Labor, and Rural Poverty," Oxford Univ. Press, Delhi, 1984.
2. K. G. BINMORE, Nash bargaining theory II, and Nash bargaining theory and incomplete information, *in* "The Economics of Bargaining" (K. G. Binmore and P. Dasgupta, Eds.), Basil Blackwell, Oxford, 1987.
3. D. CANNING, "Average Behavior in Learning Models," Working paper, Pembroke College, Cambridge, 1991.
4. H. CARLSSON, A bargaining model where parties make errors, *Econometrica* **59** (1991), 1487–1496.
5. M. FREIDLIN AND A. WENTZELL, "Random Perturbations of Dynamical Systems," Springer-Verlag, New York, 1984.
6. D. FOSTER AND H. P. YOUNG, Stochastic evolutionary game dynamics, *Theoret. Population Biol.* **38** (1990), 219–232.
7. D. FUDENBERG AND C. HARRIS, Evolutionary dynamics in games with aggregate shocks, *J. Econ. Theory* **57** (1992), 420–441.
8. J. HARSANYI AND R. SELTEN, A generalized Nash solution for two person bargaining games with incomplete information, *Manage. Sci.* **18** (1972), 80–106.
9. E. KALAI AND M. SMORODINSKY, Other solutions to Nash's bargaining problem, *Econometrica* **45** (1975), 513–518.

10. M. KANDORI, G. MAILATH, AND R. ROB, Learning, mutation, and long-run equilibria in games, *Econometrica* **61** (1993), 29–56.
11. T. LENSBERG, Stability and collective rationality, *Econometrica* **55** (1987), 935–962.
12. J. NASH, The bargaining problem, *Econometrica* **18** (1950), 155–162.
13. J. NASH, Two-person cooperative games, *Econometrica* **21** (1953), 128–140.
14. H. RAIFFA, Arbitration schemes for generalized two-person games, *in* "Contributions to the Theory of Games" (H. Kuhn and A. W. Tucker, Eds.), Princeton Univ. Press, Princeton, NJ, 1953.
15. R. T. ROCKAFELLAR, "Convex Analysis," Princeton Univ. Press, Princeton, NJ, 1970.
16. A. RUBINSTEIN, Perfect equilibrium in a bargaining model, *Econometrica* **50** (1982), 97–109.
17. I. STAHL, "Bargaining Theory," Stockholm School of Economics, Stockholm, 1972.
18. H. P. YOUNG, The evolution of conventions, *Econometrica* **61** (1993), 57–84.
19. H. P. YOUNG AND D. FOSTER, Cooperation in the short and in the long run, *Games Econ. Behav.* **3** (1991), 145–156.