

Equilibrium Selection by Intentional Idiosyncratic Play

Suresh Naidu † and Samuel Bowles †† *

August 11, 2005

Abstract

We introduce intentional idiosyncratic play in a standard stochastic evolutionary model of equilibrium selection in bargaining games. We define intentional mutations as rare play of mixed strategies that are supported only on the set of strategies that would give the idiosyncratic player a higher payoff were sufficiently many others to do the same. In contract games intentional idiosyncratic play alters the standard perturbed evolutionary dynamic, where idiosyncratic play is drawn from the entire strategy set, in ways that are plausible in light of historical studies of institutional transitions. First, the most probable transitions between institutions are induced only by the idiosyncratic play of those who stand to benefit from the switch. Second, where sub-population sizes and idiosyncratic play rates differ cross groups, the group whose interests are favored are those who engage in more frequent idiosyncratic play and who are less numerous. The intentional idiosyncratic play dynamic selects the equilibrium that implements the Nash bargain as the stochastically stable state, while the standard dynamic selects the Kalai-Smorodinsky bargain.

Keywords: Evolutionary Game Theory, Stochastic Stability, Nash Bargaining Solution, Multiple Equilibria, Institutional Transitions, Intentionality

JEL CLASSIFICATION: C73, C78

*†Corresponding Author: snaidu@econ.berkeley.edu. Santa Fe Institute and University of California-Berkeley † † Santa Fe Institute and University of Siena. We would like to thank the MacArthur Foundation, the Russell Sage Foundation and the Behavioral Science Program of the Santa Fe Institute for financial support. Sung-Ha Hwang provided excellent comments on the entire draft. We are also grateful for comments on this paper to Philippe Aghion, Abhijit Banerjee, Dilip Mookherjee, Adam Szeidl, Rajiv Sethi, Chris Udry, Erik Wright, Peyton Young and to participants in the working group on inequality in the long run at the Santa Fe Institute. All remaining errors are our own.

1 Introduction

The development of stochastic evolutionary game theory [14, 27, 26, 17, 6, 5] provides a mechanism for equilibrium selection in models with multiple strict Nash equilibria[19]. Using these models, institutions may be represented as conventions, and idiosyncratic individual behaviors occasionally displace a population from the neighborhood of one convention to another. Models from this literature have been used to explain the evolution of conventions[28], the diffusion of innovations[31], the emergence of classes [13, 2] and crop-shares[32]. When applied to a contract game or other interactions governing distribution between economic classes, the approach allows remarkably strong conclusions about the nature of evolutionarily successful institutions. For example in Young [29], the equilibrium selection process generates a long term history in which populations tend to spend most of their time at conventions that are Pareto-efficient and egalitarian. In addition, high-rationality results in cooperative bargaining theory, such as the Nash and Kalai-Smorodinsky solutions, emerge as unique stochastically stable states in non-cooperative, low-rationality processes of adaptation and mutation.

Here we extend this approach by imposing empirically plausible restrictions on the process generating idiosyncratic play. The works above use a standard adaptive learning dynamic, in which idiosyncratically playing agents randomly draw strategies from their entire strategy set, effectively replicating a mutation-like process. The approach we develop here is that the error distribution is state dependent: when playing idiosyncratically, agents draw from strategies that offer them a better payoff, should sufficiently many others do the same, by comparison with their current payoff. We thus introduce a minimal amount of foresight into an otherwise myopic updating process.

Our modification to the standard dynamic is motivated by our belief that agents who act idiosyncratically in economic conflicts are behaving intentionally, and thus do not “accidentally” experiment with contracts under which they would do worse, should the contract be generally adopted. We have in mind such idiosyncratic play as walking to work rather than riding in the racially segregated (“Negro”) section of the bus or refusing to exchange under the terms of a contract that awards most of the joint surplus to the other party (for example locking out overly demanding employees). We thus seek to synthesize the evolutionary game theory approach with the literature on institutional transitions, providing formal models for processes that sociologists [21, 20] and historians [8, 16] have long found to be empirically important. Axtell and Chakravarty [2] use a similar approach, but they a) limit themselves to agent-based simulations, and b) mandate that a fraction of players always plays a non-best-response, while we provide analytical results and assume that players only play a non-best-response when they play idiosyncratically.

Like Bergin and Lipman [3], who conclude that “models or criteria to determine ‘reasonable’ mutation processes should be a focus of research in this area”, and Van Damme and Weibull[9], our idiosyncratic play is state-dependent. But while these authors make error *rates* state dependent, we make the *distribution* of idiosyncratic play across the strategy space state-dependent. We do this in order to impose a particular structure on the process generating idiosyncratic play, one that we think captures an essential aspect of the process of institutional transitions, namely the intentional vio-

Table 1: Example 1

Contract	0	1	2
0	5,60	0,0	0,0
1	0,0	12,20	0,0
2	0,0	0,0	36,1

lation of an existing norm motivated by dissatisfaction with the status quo.

A large recent literature characterizes the stochastically stable equilibria of various classes of games. A strand of this literature has looked at bargaining games, where the set of strict Nash equilibria are symmetric in strategy and are Pareto-optimal. Young[27] examines the Nash demand game, and shows the Nash bargaining solution is stochastically stable. Young[29] studies contract games, and shows that the Kalai-Smorodinsky solution is stochastically stable. Tröger[22] studies stochastic stability in a “hold-up” model, where the bargaining follows a first-stage investment decision. Ellingsen [10] examines the related but distinct notion of evolutionary stability in bargaining games. Agastya [1] investigates stochastic stability in double-sided auctions, which can be represented as bargaining games where matches that do not exhaust the surplus are decided by randomizing between the contracts that do fully divide the surplus.

A common assumption in these papers and the stochastic adjustment literature is that the idiosyncratic noise takes the form of errors, in that there is no systematic bias in the strategies played when non-best-responses occur. A possible explanation for this omission is that it might not matter; Young [26] shows that as long as errors have full support, the stochastically stable state is invariant to the choice of error distribution. If on empirical plausibility grounds, however, we weaken the assumption of full support, we select different equilibria for some important classes of games.

Intentional idiosyncratic play, we will show, alters the standard evolutionary dynamic in ways that are plausible in light of historical studies of institutional transitions. First, transitions between institutions are induced only by the idiosyncratic play of those who stand to benefit from the switch, in contrast to the standard (unintentional) approach. Second, as one would expect, in the intentional dynamic where sub-population sizes and error rates differ across groups, the group whose interests are favored are those who engage in more frequent idiosyncratic play and who are less numerous. In contract games, the conventions that are selected as stochastically stable under the intentional idiosyncratic play dynamic, differ from those selected under the standard dynamic. Our dynamic selects the convention that implements the Nash bargain, while the standard dynamic selects the Kalai-Smorodinsky bargain. The difference is illustrated in the example in Table 1.

In example 1, the Kalai-Smorodinsky contract is $(1, 1)$, as $\frac{12}{36} = \frac{1}{3} = \frac{20}{60}$, and the Kalai-Smorodinsky solution equates the ratio of the payoffs to the ratio of the players best possible payoff. In contrast, the Nash bargaining solution is $(0, 0)$, since

$5 \times 60 > 12 \times 20 > 36 \times 1$, and the Nash solution is that which maximizes the product of the payoffs. Our primary contribution in this paper is to show that when errors are intentional rather than random, although still far from fully rational, the Nash bargain is generally selected as the ultra-long-run equilibrium.

In order to establish a benchmark for contrast, in the next section we reproduce a version of the standard adaptive stochastic dynamic model and point out some counter-intuitive implications of the institutional transition process that it supports. The main difference between our construction of the standard model and those mentioned above is that, in contrast to Kandori, Mailath and Rob, we have two sub-populations (classes) and in contrast to Young, our agents have but a single period memory and best respond to the (known) distribution of play in the previous period. These modeling differences do not alter the basic results of the unintentional idiosyncratic play model under investigation here. In contrast to both we are interested in the dynamics given by heterogeneous error rates and group sizes.

In section 3 we introduce our intentional idiosyncratic play modification and demonstrate the results mentioned above as well as showing that institutional transitions under the intentional dynamic are among adjacent contracts (those that are the “neighbors” of the status quo contract along the contract frontier in a finite contract space) while the standard dynamic moves between extreme contracts, leapfrogging, as it were, across large segments of the contract space. Section 4 extends the model to double-sided auction and Nash Demand games, and section 5 concludes.

2 Adaptive Stochastic Dynamics

We consider two large sets of agents, called classes (denoted R and C for row and column), playing an asymmetric K-contract game. This has K strategies, with payoff functions given by $\pi^R(i, j) = \pi^C(i, j) = 0$ if $i \neq j$, and $\pi^R(i, i) = a_i, \pi^C(i, i) = b_i$ $i \in (0, 1, \dots, K - 1)$ otherwise. We order the strategies such that if $i < j$ then $a_j > a_i$ and $b_j < b_i$, so the contracts are ordered such that the row player favors higher indices, and the column player favors lower indices. Clearly the diagonal of the game matrix constitutes the set of pure Nash equilibria, and they are all strict and Pareto-optimal. For example, a simple 2 contract game, with both contracts specifying the division of a unit good, is given by (with $a_0 < a_1, b_i = 1 - a_i$):

<i>Contract</i>	0	1
0	$a_0, 1 - a_0$	$0, 0$
1	$0, 0$	$a_1, 1 - a_1$

We can represent this dynamic by a stochastic dynamical system, where the states represent the number of each population playing each strategy. The state space is given by $X = \Delta_R \times \Delta_C$, where $\Delta_R = \{n_0, n_1, n_2, \dots, n_{K-1} \mid \sum_i n_i = N\}$ and $\Delta_C = \{m_0, m_1, m_2, \dots, m_{K-1} \mid \sum_i m_i = M\}$ where N is the size of the row population and M is the size of the column population, and each n_i and m_i is the number of the row and

column population, respectively, that is playing strategy i . Let $p \in \Delta_R$ and $q \in \Delta_C$ be vectors denoting the number of agents playing each strategy in the row and column population, respectively. We will often denote a state as $\theta = (p, q) \in X$.

Denote the best-response functions for the row and column populations respectively,

$$BR_R(q) : \Delta_C \rightarrow \Delta_R = Ne_{\text{argmax}_j q_j a_j} \quad (1)$$

$$BR_C(p) : \Delta_R \rightarrow \Delta_C = Me_{\text{argmax}_j p_j b_j} \quad (2)$$

where e_i is the i 'th standard basis vector of \mathfrak{R}^K . If there are multiple best responses, the agents choose the highest index if they are row players, and the lowest index if they are column players.

We now consider a matching dynamic with noise. The dynamic is a familiar myopic best-response dynamic with inertia. Each period, all players are matched to play the contract game. Each time they are matched, agents play the strategy they played last with probability ν or revise their strategy with probability $1 - \nu$. This ‘‘inertia’’, provided by ν , is necessary to ensure convergence. This defines a Markov process: $P^\nu : X \rightarrow X$, defined by $P^\nu(\theta'|\theta) = \text{Prob}(\theta - (\frac{x_1}{N}p, \frac{x_2}{M}q) + (\frac{x_1}{N}BR_R(q), \frac{x_2}{M}BR_C(p)) = \theta')$ where $x_1 \sim \text{Bin}(N, \nu)$, $x_2 \sim \text{Bin}(M, \nu)$ where $\text{Bin}(N, \nu)$ is a binomial distribution with N draws with probability of success given by ν . Note that, for generic contracting games and sufficiently large population sizes, the only recurrent classes of this Markov process are the strict pure Nash equilibria, where both players coordinate on the same contract[29].

Suppose that when agents can revise their strategies, they play a non-best response with probability ϵ^{TR} for the row players and probability ϵ^{TC} for the column players. Thus, with probability ϵ^{TR} the row players play a strategy drawn from uniform distribution U on the strategy space. As we will see, it will be useful to use the following general distribution: $U(i, j)$ as the density function for the uniform distribution on the strategies $i, i + 1, \dots, j$, with 0 weight on the other strategies. For the standard dynamic, the error distribution is just $U(0, K - 1)$ which gives a Markov process defined by: $P^{\nu, \epsilon}(\theta'|\theta) = \text{Prob}(\theta - (\frac{x_1}{N}p, \frac{x_2}{M}q) + (\frac{1}{N}(x_1 - \sum_{i=0}^{K-1} y_{1i})BR_R(q), \frac{1}{M}(x_2 - \sum_{i=0}^{K-1} y_{2i})BR_C(p)) + (y_1, y_2) = \theta')$ where x_1, x_2 are binomial draws as above, $y_1 \sim \mathcal{MN}(K, x_1, U(0, K - 1))$ and $y_2 \sim \mathcal{MN}(K, x_2, U(0, K - 1))$, with $\mathcal{MN}(N, k, f)$ being the multinomial distribution with N bins, k draws, and distribution f over the bins. Owing to the unintentional nature of the errors, where mistakes that are potentially beneficial are as likely as those that are potentially unfavorable, we call this the U-dynamic, Γ^U .

As is well-known, this Markovian dynamic can be represented by a transition matrix given by (abusing notation somewhat) $P_{\theta\theta'}^{\nu, \epsilon} = P^{\nu, \epsilon}(\theta'|\theta)$. The long-run steady state of the dynamic is then given by the unique vector $\mu(\nu, \epsilon) \in \mathbf{R}^{2K}$ that satisfies $\sum_i \mu_i(\nu, \epsilon) = 1$ and $\mu(\nu, \epsilon)P^{\nu, \epsilon} = \mu(\nu, \epsilon)$. We are interested in the states that have positive weight in the distribution $\mu^*(\nu) = \lim_{\epsilon \rightarrow 0} \mu(\nu, \epsilon)$, following Foster and Young[14] we call these stochastically stable states, with U-stability referring to stability under the perturbation process described in the preceding paragraph.

In general, this ergodic distribution will be the solution to an intractably large system of linear equations. However, the proofs in this literature have largely been done

using a result from Friedlin and Wentzell[15], that expresses the ergodic distribution of a finite irreducible Markov process as the sums of “tree potentials”. This provides a useful method for characterizing the stochastically stable states. Young[26] defines the resistance of a transition from state i to state j as the unique R_{ij} that satisfies $0 < \lim_{\epsilon \rightarrow 0} P_{ij}^{\nu, \epsilon} / \epsilon^{R_{ij}} < \infty$. If we build a complete weighed directed graph, with vertices corresponding to the states of the Markov process, with each edge i, j having weight R_{ij} , the recurrent class will be the root of the in-branching¹ with the least sum of edge weights. We will also call this the *minimal tree*.

Equilibrium selection in this model is determined by the likelihood for each convention that the stochastic process generating idiosyncratic play will realize a number of deviants from the convention sufficient to induce the best responding players on the other side to adopt some other convention. The resistances defined in the previous paragraph are the minimum number of players who must idiosyncratically deviate from the convention i to induce a transition to some other convention j is termed the resistance for this transition or R_{ij} . (For non-vanishing ϵ transitions in both directions will be induced by both classes, but as ϵ goes to zero, the influence on the ergodic distribution $\mu(\nu, \epsilon)$ of the least likely path from i to j can be ignored).

We shall make use of three propositions proved in Binmore, Samuelson and Young[5]:

Proposition 2.1 (Local Resistance Test). *If $\max_j R_{ji} < \min_j R_{ij}$ then i is the root of the minimal tree.*

Proposition 2.2 (Naive Minimization Test). *Take the least edge exiting each node. If the resulting graph has a unique cycle containing an edge that is maximal over all edges, deleting that edge will give the minimal tree.*

Proposition 2.3 (Contract Theorem). *Let $0 < \delta < 1$. If the contracts are equally distant in the row player’s payoff, and given by $(a_i, b_i = f(a_i)), a_i = i\delta, i \in (1, \dots, \frac{1}{\delta} - 1)$ where f is a continuous, strictly decreasing, strictly concave function. Then, for δ sufficiently small, the stochastically stable contract in the U-dynamic is the Kalai-Smorodinsky solution, given by i such that $a_i/a_{max} = b_i/b_{max}$, assuming it lies in the set of contracts.*

Note that the resistance of the transition from contract i to contract j is given by

$$R_{ij} = \min\left\{\left\lceil \frac{MT_C a_i}{a_i + a_j} \right\rceil, \left\lceil \frac{NT_R b_i}{b_i + b_j} \right\rceil\right\}. \quad (3)$$

If $NT_R = MT_C$ is sufficiently large, then we can use the *reduced resistances* $r_{ij} = \min\left\{\frac{a_i}{a_i + a_j}, \frac{b_i}{b_i + b_j}\right\}$. Note that if $b_i > b_j$ and $a_i < a_j$ then $r_{ij} = \frac{a_i}{a_i + a_j} < 1/2 < \frac{b_i}{b_i + b_j} = 1 - r_{ji}$. Simple manipulation then shows that i is stochastically stable if and only if $a_i b_i > a_j b_j$.

¹An in-branching is a directed connected graph where every node save one has only one edge exiting it. The node with no exiting edge is called the root.

Our first observation is that the transitions between equilibria in the contract game are instigated by the class that loses from the transition. For the transition from i to j , it is the more likely of the two paths that is the resistance, transitions are always induced by those who lose as a result. The intuition behind this result is that it always takes more idiosyncratic play by those disadvantaged at the status quo convention to dislodge the best responding members of a class from its preferred contract than to dislodge the best responders in the disadvantaged class to try a contract that is preferable to the status quo.

The fact that the unintentional dynamic takes the minimum of both populations' resistances means that the agents who are inducing the change are those who stand to lose from the tip. The resistance of the transition from i to j is the number of idiosyncratic plays made by the population facing payoffs b_i and b_j . Similarly, the transition from j to i is driven by the idiosyncratic play of the population facing the payoffs a_i and a_j . The most likely path between the two conventions occurs when the losers from the transition make enough mistakes.

This observation gives rise to a number of corollary observations. a) Having a larger group benefits you, and b) if the rates of idiosyncratic play differ, then the side with the faster rate does worse. We show this in the 2-contract case to illustrate the intuition, and then generalize to the K-contract case.

Consider the 2-contract game below with payoffs $(a_0, b_0), (a_1, b_1)$ with $a_0 < a_1$ and $b_0 > b_1$ as usual. Assume the row population has population size N and the column population is size M . Further, assume that the row population has idiosyncratic play probability ϵ^{TR} and the column population has idiosyncratic play probability ϵ^{TC} .

<i>Contract</i>	0	1
0	a_0, b_0	0, 0
1	0, 0	a_1, b_1

Then the resistances are given by $R_{ij} = \min\{\lceil MT_C \frac{a_i}{a_i+a_j} \rceil, \lceil NT_R \frac{b_i}{b_i+b_j} \rceil\}$. If NT_R/MT_C is sufficiently large (the row population has a very slow idiosyncratic play rate and/or a very large population relative to the column population), the resistances will be $R_{01} = \lceil \frac{MT_C a_0}{a_1+a_0} \rceil$, and $R_{10} = \lceil \frac{MT_C a_1}{a_1+a_0} \rceil$. If $a_0 < a_1$ then $R_{01} < R_{10}$ and 1 is U-stable. Note that 1 is the preferred contract for the row players, while it is the column players with relatively faster idiosyncratic play rates and smaller population. This is because both transitions occur as a result of the idiosyncratic play of the column population, who prefer contract 0. The errors of the column players bring about their own reduction in payoffs.

When we have K contracts, $a_i, b_i, i \in 0 \dots K-1, i < j \rightarrow a_i < a_j, b_i > b_j$, this is also easy to see. Note that $R_{ij} = \min\{MT_C \frac{a_i}{a_i+a_j}, NT_R \frac{b_i}{b_j+b_i}\}$ which is equal to $\frac{MT_C a_i}{a_i+a_j}$ for sufficiently large NT_R/MT_C , which is increasing in a_i and decreasing in a_j . Suppose 1 is the contract with the highest row payoff a_1 . Let j denote the second highest a_k since the highest incoming edge to vertex 1 is $\frac{MT_C a_j}{a_j+a_1}$, while the lowest outgoing edge is $\frac{MT_C a_1}{a_1+a_j}$. Since $a_1 > a_j$, this means that the minimum outgoing edge from 1 is greater

than the maximum incoming edge, so by the local resistance test, a_1 is stochastically stable. So the population that is largest, or idiosyncratically plays the least, does best. We summarize these findings in the next proposition:

Proposition 2.4. *Under the U-dynamic, if population sizes are given by N and M , and idiosyncratic play rates are given by T_R and T_C for row and column respectively, then if NT_R/MT_C is sufficiently large(small) then the contract most beneficial, with payoff $a_{max}(b_{max})$, to the row(column) population will be U-stable.*

3 Intentional Idiosyncratic Play

In sum, under the U-dynamic the individuals who induce transitions from one contract to another always lose as a result. Two additional odd results follow from this: those who play idiosyncratically at a lower rate and those who come from larger groups are favored in this dynamic. The reason is that idiosyncratic play by members of a more numerous group with lower rates of idiosyncratic play are less likely to induce a transition (from which they would necessarily lose, if it occurred).

In order to overcome these problems with the U-dynamic, we now define a new dynamic Γ^I or the I-dynamic, where the error distribution for idiosyncratic play is supported only on the strategies that would be beneficial relative to the current state. Instead of idiosyncratic play being a random draw over the entire strategy space, we require that the agent randomly choose only from those strategies that would give a higher payoff, were they to be played by both populations. That is, when the agent's opposing population is at a recurrent class $\theta = (j, j)$, the agent in the row population randomly plays a strategy drawn from the set $S_j = \{i | \pi_R(i, i) > \pi_R(j, j)\}$. We make this precise below. This involves a degree of foresight without imposing full rationality, something that is missing in the standard, purely myopic-with-errors dynamic.

An important caveat is that agents do not take account of the costs of failing to coordinate. While this may be plausible in some circumstances, it is certainly not general. The off-diagonal payoffs should play a role in the error distribution, as in Van-Damme and Weibull[9]. In order to reduce notational clutter, however, we focus on the case when both populations face 0 off-diagonal payoffs.

3.1 Intentional Dynamics

This dynamic is somewhat more complicated, involving as it does state-dependent error distributions. First, we need to denote the domain from which idiosyncratic strategies are drawn. If the system is at a contract i , then the row players will choose from all strategies with index greater than i , while the column players will choose from among all the strategies less than i , since row payoffs are increasing in the index. More generally, for row(column) players at a state $\theta = (p, q)$, this will be the strategy with the least(greatest) index supported by the distribution of play of the column(row) players.

Formally:

$$\begin{aligned} i^R(\theta) &= \max\{i | q_i > 0\} \\ i^C(\theta) &= \min\{i | p_i > 0\} \end{aligned}$$

The error distribution in our case is population- and state-dependent; at state p, q the error distribution is $U(i^R(\theta), K - 1)$ for the row population, and $U(0, i^C(\theta))$ for the column population. Therefore our transition probabilities are now given by equation 1, with the following modification. Instead of $y_i \sim \mathcal{MN}(K, x_i, U(0, K - 1))$ in the above model, we have instead that $y_1 \sim \mathcal{MN}(K, x_1, U(i^R(\theta), K - 1))$ and $y_2 \sim \mathcal{MN}(K, x_2, U(0, i^C(\theta)))$. We denote the transition matrix by $P^{\nu, \epsilon, I}$.

Showing this process is ergodic is straightforward, albeit not trivial, since our errors are not always supported on the entire strategy space. Given state $\theta \in X$, how can we get to state θ' in a finite number of periods?. It suffices to show that we can get to the state $(p_0, q_K) = (N, 0, 0, \dots, 0), (0, 0, 0, \dots, M)$, which is where all members of both populations are at the contract that would be worst for them were it to become an equilibrium, from an arbitrary state $\theta = (p, q)$, since then the errors are supported on the entire strategy space, and therefore any state is accessible from θ .

This follows from the fact that from θ there is a positive probability that the column population will mutate to the state $(M, 0, \dots, 0)$, which is the contract that is most favorable to it, and the same time that the row population mutates to the state $(0, 0, 0, \dots, N)$. Then, there is a positive probability that all agents play a best-response to which leads the row population to respond (with no mutations) with $(N, 0, \dots, 0)$, and the column population to respond with $(0, 0, 0, \dots, 0, M)$.

It is clear that as $\epsilon \rightarrow 0$, $P^{\nu, I}(\epsilon) \rightarrow P^\nu$. It is also clear that the process converges exponentially, since the transition probabilities are polynomials in ϵ . Let

$$R_{jk}^I = \begin{cases} \lceil \frac{NT_R b_j}{b_j + b_k} \rceil & \text{if } a_j < a_k \\ \lceil \frac{MT_C a_j}{a_j + a_k} \rceil & \text{if } b_j < b_k \end{cases}$$

Note that this is well-defined since all of the i are Pareto-optimal. This reflects the fact that if one class loses from a transition to a particular equilibrium, it will never idiosyncratically play the strategy corresponding to that equilibrium. Thus the transition will be generated by the idiosyncratic play of the other, opposing population. The class that stands to benefit from the transition must overcome the resistance of the would-be loser by generating so much idiosyncratic play that the best-response of the losing population is to play a strategy that, when the strategy is an equilibrium, gives them a lower payoff than the current equilibrium.

It is obvious that $\infty > \lim_{\epsilon \rightarrow 0} \frac{P_{jk}^{\nu, \epsilon, I}(\epsilon)}{\epsilon^{R_{jk}^I}} > 0$, since, when we only allow one population to idiosyncratically play, that this is the smallest number of non-best-responses required to make a transition.

Table 2: U-Resistances for Example 1

Root/Trees				
0	0.266	0.297	0.266	
1	0.341	0.310	0.169	
2	0.544	0.371	0.371	

Table 3: I-Resistances for Example 1

Root/Trees				
0	1.583	1.455	1.830	
1	1.500	1.628	1.733	
2	1.702	1.689	1.932	

Definition 3.1. We call a contract *I-stable* if it is the stochastically stable state when transition resistances are defined as above.

We call trees with $R^I(R^U)$ edge weights *I-trees(U-trees)*. From theorem 1 in [26], we know that the I-stable state is contained in the root of the minimal I-tree.

For the rest of this section, we will omit the *I* superscript from the resistances unless there is some ambiguity.

Proposition 3.2. *Assume equal class sizes and idiosyncratic play rates, then:*

- a) *In 2-contract games the risk-dominant equilibrium is I-stable.*
- b) *In 3-contract games, if a contract is both maximin and risk-dominant, then it is I-stable.*

Proof. a) The 2 contract case follows from the fact that $R_{01}^I = \frac{b_0}{b_1+b_0} \geq \frac{a_1}{a_0+a_1} = R_{10}^I$ if and only if $R_{01}^U = \min\{\frac{a_0}{a_0+a_1}, \frac{b_0}{b_1+b_0}\} \geq \min\{\frac{a_1}{a_0+a_1}, \frac{b_1}{b_1+b_0}\} = R_{10}^U$. Thus, 0 is U-stable if and only if 0 is I-stable, and similarly for contract 1. If the inequality is an equality, then both equilibria are I-stable and U-stable.


b) See Appendix.

□


After this result, one might be tempted to think that there is no substantial difference between U-stability and I-stability. However, the example in the introduction (Table 1) illustrated otherwise, and we now turn to investigating the differences in the two dynamics.

In Table 1, the I-stable contact is 0, while the U-stable contract is 1. Table 2, consisting of tree resistances, illustrates the calculations for the U-dynamic(3 trees for each root).

Thus the lowest tree, with resistance 0.169 has root 1. The actual tree is given below.



However, with intentional error distributions (the I-dynamic), the tree resistances are given in Table 3. The minimal I-tree has root 0, with resistance 1.455, shown in the tree below. The full set of trees is given in Appendix B.



Note that the U-stable state in example 1 is the Kalai-Smorodinsky solution ($a_1/b_1 = a_{max}/b_{max}$), while the Intentionally Stable State is the Nash Solution ($a_0 b_0 = \max_i a_i b_i$). This is a general difference, as illustrated by the next proposition.

Proposition 3.3. *Assume equal group sizes and idiosyncratic population sizes, and let the contracts be given by $(a_i, b_i = f(a_i))$, $a_i = i\delta, i \in (1, \dots, 1/\delta - 1)$ where f is a positive, continuous, strictly decreasing, strictly concave function. Then, for δ sufficiently small, the I-stable contract is the one that maximizes the Nash product $sf(s)$, assuming it lies in the set of contracts.*

Proof. This follows as a special case of the next proposition. □

When agents' errors are restricted to be a uniform distribution supported only on strategies that could give them a higher payoff, the results can be different from the unintentional errors case. Counter to the Binmore-Samuelson-Young result presented above, the Nash bargaining solution is I-stable. The difference stems from the fact that the intentional resistances are always lowest in transitions to adjacent contracts, while the Binmore-Samuelson-Young result depends on the fact that with unintentional errors, the lowest resistances are for the transitions to the extreme contracts, i.e. those that are best for one side. This is illustrated by the comparison between the minimal I-tree and U-tree in the figures above. We find this "leapfrogging" feature of the U-dynamic to be historically implausible, compared to the "neighboring" transition feature of the I-dynamic.

The intuition of the result is as follows. The Kalai-Smorodinsky solution is the midpoint of the secant connecting the endpoints of the bargaining frontier, while the Nash solution is the point on the bargaining frontier that is the midpoint of the tangent line at that point. This characterization illustrates why the two dynamics give different solutions. Intuitively, the stochastically stable state maximizes the chance of entering the state. Given the transitions are driven by opposing populations, the stable state will equate the highest probability of entering the state from the left with the highest probability of entering it from the right. The least cost transition into the I-stable

state under the I-dynamic are from the immediately adjacent states, while the least-cost transition into the U-stable state is from the extreme contracts $x = 0$ and $x = 1$. Hence the midpoint of the tangent line in the first case and the midpoint of the secant line in the latter case.

A key property of the Kalai-Smorodinsky solution is that it is invariant to the slope of f , the bargaining frontier, while the Nash solution is sensitive to the slope of f (a steeper curve benefits the column population more). They agree generically only if the bargaining frontier is symmetric around the 45-degree line or linear.

3.2 Population and Mutation Rates

We first note that under the I-dynamic, the relative population and idiosyncratic play rates operate in exactly the opposite direction than the analogous variables in the U-dynamic. Smaller groups with higher rates of idiosyncratic play are favored. To illustrate this, consider a 2-contract game with payoffs as in the previous example. Assume also that contract 1, favored by the row players is risk-dominant, so that $a_1 b_1 > a_0 b_0$. Assume the row population has population size N and the column population is size M with $N > M$. Assume also that the rate of mutations differ by a power T , so that the row population makes mistakes at a rate ϵ^T , while the column population plays idiosyncratically at the rate ϵ (i.e. $T_C = 1$).

The resistances in the intentional dynamic will be $R_{01}^I = NT \frac{b_0}{b_0 + b_1}$ and $R_{10}^I = M \frac{a_1}{a_0 + a_1}$. Contract 0 will be I-stable if and only if $R_{01}^I > R_{10}^I$ which is the case if $TN/M > \frac{b_1 a_1 + b_0 a_1}{b_0 a_1 + b_0 a_0} > 0$. Thus, if the row population, N , is sufficiently large, and/or the row mutation rate, is sufficiently slow (high T), then the risk-dominant contract (contract 1) may not be selected.

For the K contract case, we can easily make a simpler statement, without tight specifications of the precise difference in idiosyncratic play rates or population size necessary to secure the best contract for a given size. Again suppose that the row population has population N , and the best contract for them is K-1. Then choose $MT_C > NT_R$ such that $\max_i NT_R \frac{b_i}{b_{K-1} + b_i} = R_{i(K-1)}^I < \min_j R_{(K-1)j}^I = \min_j MT_C \frac{a_{K-1}}{a_j + a_{K-1}}$. Then, by the local resistance test, contract $K - 1$ is I-stable. The next proposition gives a more precise result.

Proposition 3.4. *Assume unequal group sizes N and M and idiosyncratic play rates T_R and T_C for the row and column populations respectively. Let the contracts be given by $(a_i, b_i = f(a_i))$, $a_i = i\delta$, $i \in (0, 1, \dots, \frac{1-\delta}{\delta}, \frac{1}{\delta})$ where f is a continuous, strictly decreasing, strictly concave function, with $f(0) = 1$ and $f(1) = 0$. Then, for δ sufficiently small, the I-stable contract is the one that maximizes the product $G(a_i) = (a_i e^{2a_i/\delta})^{NT_R} (f(a_i) e^{-2a_i/\delta})^{MT_C}$, assuming it lies in the set of contracts.*

Proof. See Appendix A. □

Note that if $NT_R = MT_C$, the I-stable contract is the symmetric Nash Bargaining solution, which proves proposition 3.3. Note also that since δ is small, the stochasti-

cally stable contract will be close to the best contract for the population with lower population-size and higher idiosyncratic play rate. That is, row players will be favored when $NT_R < MT_C$. We can summarize this in a corollary:

Corollary 3.5. *As δ goes to 0, there are three candidates for equilibrium payoffs a^*, b^* :
if $MT_C > NT_R$ then $a^* = 1$ and $b^* = 0$,
if $MT_C = NT_C$, then $a^* = a_{Nash}$ and $b^* = b_{Nash}$
if $MT_C < NT_R$ then $a^* = 0$ and $b^* = 1$.*

4 Nash Demand and Double-Sided Auctions

We now consider the effects of intentional idiosyncratic play in alternative specifications of the bargaining game, where the off-diagonal payoffs are not necessarily 0. However, the games still have the feature that all the strict Nash equilibria are Pareto-optimal, so our idiosyncratic distribution of play is still well-defined. We first consider the Nash demand game, the stochastically stable(U-stable) equilibrium of which is the Nash Bargaining solution[27]. The difference between the Nash demand game and the contract game is that in the former, all contracts that do not exhaust the surplus can be struck, and so agents get their offer even if the offers do not agree with that of their matched opponent. Formally, in our previous notation, if $i < j$ then $a_i < a_j$, $b_i > b_j$, just as before, but now the off-diagonal payoffs are not all 0. In particular, if $i < j$ then the payoff matrix is given by:

<i>Contract</i>	<i>i</i>	<i>j</i>
<i>i</i>	a_i, b_i	a_i, b_j
<i>j</i>	0, 0	a_j, b_j

A recent paper by Agastya[1] explores the stochastic stability of strict Nash equilibria in a two-sided auction game. In the double-auction game, the payoffs are similar to the Nash demand game, except if the agents fail to exhaust the surplus, with probability ρ the agents' payoff is what the agents would have received in contract i and with probability $1 - \rho$ the payoff is what they would have received at contract j . Agastya considers agents with varying risk-aversion, while we limit ourselves to the risk-neutral case here. The game matrix is given below.

<i>Contract</i>	<i>i</i>	<i>j</i>
<i>i</i>	a_i, b_i	$\rho a_i + (1 - \rho)a_j, \rho b_i + (1 - \rho)b_j$
<i>j</i>	0, 0	a_j, b_j

Proposition 4.1. *If the strategy space is as in Proposition 3.3, but the payoff structure is either (1) Nash demand or (2) double-sided auction games with risk-neutral utility functions, the I-stable contract approaches the Nash Bargaining solution as δ goes to 0.*

Proof. It suffices to show this for the Double-Sided Auction game, as the proof for the Nash Demand game is virtually identical. The resistances are calculated, as usual, by computing the fraction of idiosyncratic players necessary to make their opponents indifferent between two strategies. The result follows from the fact that

$$r_{ij}^I = \begin{cases} \frac{f(s_i) - \rho f(s_j) - (1-\rho)f(s_i)}{f(s_i) + f(s_j) - \rho f(s_j) - (1-\rho)f(s_i)} & \text{if } i < j \\ \frac{s_i - \rho s_j - (1-\rho)s_i}{s_i + s_j - \rho s_j - (1-\rho)s_i} & \text{if } i > j \end{cases}$$

Without loss of generality, we can restrict our attention to only local transitions. Note that $r_{ii+\delta}^I = \frac{\rho(f(i) - f(i+\delta))}{\rho f(i) + (1-\rho)f(i+\delta)}$ and $r_{ii-\delta}^I = \frac{(1-\rho)(i - (i-\delta))}{(1-\rho)i + \rho(i-\delta)}$. Note that if δ is sufficiently small, $r_{ii+\delta}^I$ is equal to $r_{ii+\delta}^U = \min\left\{\frac{\rho(f(i) - f(i+\delta))}{(1-\rho)f(i) + \rho f(i+\delta)}, \frac{i+\delta}{(1-\rho)(i+\delta) + \rho i}\right\}$ which is the resistance generated by the U dynamic. Similarly, $r_{ii-\delta}^I = r_{ii-\delta}^U$. Under these resistances, Agastya[1] shows that the U-stable state is $\operatorname{argmax}_s s^\rho f(s)^{1-\rho}$. Since the resistances are the same for the I-dynamic, the proof works here as well²

□

The reason the U-stable and I-stable state are the same in these games that the off-diagonal payoffs increase the resistance of the winners and decrease the resistance of the losers from a transition. For small δ , this makes the resistance of the losers less than the resistance of the winners in the U-dynamic. Since the resistance of a transition is given by the lower resistance of the two populations, the U-dynamic generates the same resistance for a transition as the I-dynamic.

5 Conclusion

Notwithstanding its abstract nature, the strength of the stochastic approach is that formalizes a dynamic of institutional emergence and demise that highlights two critical aspects of real historical processes. The first is the structure of payoffs given by the different conventions and the resulting conflicts of interest that often drive the real historical equilibrium selection process. The second is the central role of deviants from the status quo and the occasional concession of best-responding members of the opposing group that results when the level of deviance is sufficiently great. The recognition of trade unions in the United States during the first third of the last century, the end of Communist rule in many countries and the demise of apartheid in South Africa appear to reflect this pattern.

However, in addition to the not entirely plausible aspects of the transition process in the U-dynamic that we pointed out earlier, a less attractive feature of this class of models is that as presently formulated, the transition times from one convention to another are implausibly long except for very small populations with substantial rates of idiosyncratic play. The long waiting times characteristic of these models are partly the result of special and empirically unrealistic simplifying assumptions, which could be relaxed to provide a more plausible pace of transition. First, we have abstracted from perturbations of the payoff structure of the game. If these were introduced, they

²Young(1993b) has an analogous proof for the Nash Demand game.

would periodically substantially reduce the critical number of deviants necessary to induce a transition, and could displace a convention even in the absence of idiosyncratic play. Second if there are many similar contracts (δ is small in proposition 2.3 and 3.3), then the transition times are greatly reduced in the I-dynamic, which moves exclusively among adjacent contracts (but not in the U-dynamic). Third, if most interactions are local as is often the case, at least initially, in political movements challenging the status quo, transition times are reduced [11, 6, 31]. Fourth, in models with more than one period of memory, if idiosyncratic players are persistent, in the sense that an agent who plays idiosyncratically continues to play idiosyncratically for several more periods then the resistances are substantially reduced, with a corresponding decrease in the transition time [30]. Fifth, if agents have behaviorally plausible best-response rules, where they may systematically overpredict the idiosyncratic play of the opposing population [23, 18], transitions will occur more frequently.

While we have explored intentional error structures in a relatively narrow class of games, the principle underlying equilibrium selection here is quite general, applying to any model with multiple Pareto-optimal equilibria. Further extensions of the model are possible. For example, one can also generalize the dynamic so that it is defined for games where there are equilibria that are not Pareto-optimal, just by mandating that both parties engage in idiosyncratic play when they are at a non-Pareto-optimal equilibria, reducing it to the standard case. Another extension would make the error distributions sensitive to the off-diagonal payoffs, as currently the agents do not account for the costs of a mismatch when they play idiosyncratically.

The stochastic evolutionary approach also provides a framework open to further steps towards historical realism. Among these are an account of the way in which technical change alters the shape of the contract set, in some periods making highly unequal bargains stochastically stable, and others favoring more egalitarian outcomes. Another is an explicit modeling of non-conformism with the terms of the status quo and particularly its behavioral foundations and its realization through various forms of state-dependent collective action. For example, the rate of idiosyncratic play may then depend the degree of collective action, which in turn may depend on the amount of inequality or group polarization in a particular contract. Exploring these extensions seem like promising directions for future research.

6 Appendix A

Proof of Proposition 3.2(b)

Proof. First normalize the payoffs by dividing each player's payoff by the maximum attainable by that player. This gives us contracts $(a_{min}, 1), (a, b), (1, b_{min})$. So the middle contract (a, b) is maximin. The only candidate for a minimal tree rooted at the middle is $\frac{1}{1+a} + \frac{1}{1+b}$. Note that the resistances of the trees rooted at the first contract are $\frac{a}{a+a_{min}} + \frac{1}{1+a}$ and $\frac{b}{b+b_{min}} + \frac{1}{1+a_{min}}$

The resistances of the trees rooted at the last contract are $\frac{b}{b+b_{min}} + \frac{1}{1+b}$ and $\frac{a}{a+a_{min}} + \frac{1}{1+b_{min}}$

However, risk-dominance implies that $ab > a_{min}$ and $ab > b_{min}$. So, since payoffs are bounded between 0 and 1, we get that $b > a_{min}$ and $a > b_{min}$ so $\frac{a}{a+a_{min}} + \frac{1}{1+b_{min}} > \frac{a}{a+a_{min}} + \frac{1}{1+a}$ and $\frac{b}{b+b_{min}} + \frac{1}{1+a_{min}} > \frac{b}{b+b_{min}} + \frac{1}{1+b}$ so, since the middle contract is risk-dominant, we have $b > a_{min}/a$ and $a > b_{min}/b$, which implies $\frac{1}{1+a} + \frac{1}{1+b} < \frac{a}{a+a_{min}} + \frac{1}{1+a}$ and $\frac{1}{1+a} + \frac{1}{1+b} < \frac{b}{b+b_{min}} + \frac{1}{1+b}$ so the tree rooted at (a, b) is least resistant. \square

To prove proposition 3.4 we first need a lemma:

Lemma 6.1. *The least cost edge exiting a given node i is to an adjacent node.*

Proof. Any edge going to the right ($j > i$) has resistance $R_{ij}^I = NT_R \frac{f(a_i)}{(f(a_i) + f(a_i + (j-i)\delta))}$ since f is decreasing, the lowest edge exiting to the right will have $j = i + 1$. Similarly, any edge going to the left ($j < i$) has resistance $R_{ij}^I = MT_C \frac{a_i}{a_i + j\delta}$, which will be lowest for $j = i - 1$. \square

This lemma says that we only need to consider transitions to adjacent contracts.

Proof of Proposition 3.4: Let $a_{j^*} = \operatorname{argmax}_{a_i} (a_i e^{2a_i/\delta})^{NT_R} (f(a_i) e^{-2a_i/\delta})^{MT_C} = \operatorname{argmax}_{a_i} \frac{1}{NT_R} (\log(a_i) + \frac{2a_i}{\delta}) + \frac{1}{MT_C} (\log(f(a_i)) - \frac{2a_i}{\delta})$.

We will use the Naive Minimization test. Take the least edge from each node, I claim that all the nodes less than j^* point to the immediate right, and all the nodes greater than j^* point to the immediate left.

Given $s, f(s)$ $a_1 < s < a_2$ we first note that the log of $(s e^{2s/\delta})^{NT_R} (f(s) e^{-2s/\delta})^{MT_C}$ is a concave function of s , therefore it has a unique maximum at a_{j^*} . This can be seen by noting that $\log G(s) = NT_R \log s + MT_C \log f(s) + \frac{2}{\delta} (NT_R - MT_C) s$ is strictly concave, as it is the sum of a strictly concave function and a linear function of s . Thus, the maximum can be obtained by using first order conditions. Also this implies that $NT_R (\log(a_i) + \frac{2a_i}{\delta}) + MT_C (\log(f(a_i)) - \frac{2a_i}{\delta})$ is increasing for $s < a_{j^*}$ and decreasing for $s > a_{j^*}$.

Now, consider a node $i < j^*$, we just need to show that $NT_R \frac{f(a_i)}{f(a_i) + f(a_i + \delta)} < MT_C \frac{a_i}{a_i + a_i - \delta}$, which would show that the transition to the right is cheaper than the transition to the left.

This is equivalent to

$$MT_C\left(\frac{f(a_i + \delta)}{f(a_i)} + 1\right) > NT_R\left(\frac{a_i - \delta}{a_i} + 1\right) \iff$$

$$MT_C\left(\frac{f(a_i + \delta)}{f(a_i)} + 1\right) - 2MT_C + 2MT_C > NT_R\left(\frac{a_i - \delta}{a_i} + 1\right) - 2NT_R + 2NT_R$$

Now multiply throughout by $\frac{1}{\delta}$ to get:

$$MT_C\left(\frac{f(a_i + \delta) - f(a_i)}{\delta f(a_i)}\right) + \frac{2MT_C}{\delta} > NT_R\left(\frac{a_i - \delta - a_i}{\delta a_i}\right) + \frac{2NT_R}{\delta} \iff$$




$$MT_C\left(\frac{f(a_i + \delta) - f(a_i)}{\delta f(a_i)}\right) + \frac{2MT_C}{\delta} - NT_R\left(\frac{a_i - \delta - a_i}{\delta a_i}\right) - \frac{2NT_R}{\delta} > 0$$

For small δ , this is close to $MT_C\left(\frac{f'(a_i)}{f(a_i)} + 2/\delta\right) + NT_R\left(\frac{1}{a_i} - 2/\delta\right) > 0$ which reduces to $MT_C\left(\frac{d(\log(f(a_i)) + \frac{2a_i}{\delta})}{da_i}\right) + NT_R\left(\frac{d(\log a_i - \frac{2a_i}{\delta})}{da_i}\right) > 0$ which is given by our statement that $MT_C\left(\log(a_i) + \frac{2a_i}{\delta}\right) + NT_R\left(\log(f(a_i)) - \frac{2x}{\delta}\right)$ is increasing. The case $i > j^*$ follows symmetrically, with the least edge pointing to the left. However, the actual maximand j^* can have an edge exiting to the left or to the right, which will give us a cycle of length 2.




Now we must show that the edge exiting j^* is the largest over the entire tree. This follows from the fact that at j^* , the first-order condition implies that $MT_C\left(\frac{f(a_{j^*+\delta}) - f(a_{j^*})}{\delta f(a_{j^*})}\right) + \frac{2MT_C}{\delta} - NT_R\left(\frac{a_{j^*} - \delta - a_{j^*}}{\delta a_{j^*}}\right) - \frac{2NT_R}{\delta}$ is close to 0. Thus, the resistances $r_{j^*j^*-1}$ and $r_{j^*j^*+1}$ are approximately equal. Since $R_{j^*j^*-1} > R_{jj-1}$ for all $j > j^*$ and $R_{j^*j^*+1} > R_{jj+1}$ for all $j < j^*$, it follows that the least edge exiting j^* is maximal over the entire tree. Our result then follows from the Naive Minimization test.

7 Appendix B: I-resistance Trees




Trees rooted at contract 0.



Trees rooted at contract 1.



Trees rooted at contract 2.



For each root, the least I-resistance tree is indicated by *. The minimum of these, indicated by **, identifies contract 0 as the root of the minimal I-tree.

References

- [1] Murali Agastya. Stochastic stability in a double auction. *Games and Economic Behavior*, 70:281–297, 2004.
- [2] Robert Axtell and Shubha Chakravarty. Activists, radicals, and reactionaries in a multi-agent model of class norms. Brookings Institute Working Paper, 2001.
- [3] Jim Bergin and Bart Lipman. Evolution with state dependent mutations. *Econometrica*, 70:281–297, 1996.
- [4] Ken Binmore, Larry Samuelson, and Richard Vaughan. Musical chairs: Modeling noisy evolution. *Games and Economic Behavior*, 11:1–35, 1995.
- [5] Ken Binmore, Larry Samuelson, and H. Peyton Young. Equilibrium selection in bargaining models. *Games and Economic Behavior*, 46:296–328, 2003.
- [6] Lawrence E. Blume. The statistical mechanics of strategic interaction. *Games and Economic Behavior*, 5:387–424, 1995.
- [7] Samuel Bowles. *Microeconomics: Behavior, Institutions, and Evolution*. Princeton University Press, Princeton, NJ, 2004.
- [8] Robert Brenner. Agrarian class structure and economic development in pre-industrial europe. *Past and Present*, 70:30–75, 1976.
- [9] Eric Van Damme and Jörgen Weibull. Evolution in games with endogenous mistake probabilities. *Journal of Economic Theory*, 106:296–315, 2002.
- [10] Tore Ellingsen. The evolution of bargaining behavior. *Quarterly Journal of Economics*, 112(2):581–602, May 1997.
- [11] Glenn Ellison. Learning, local interaction, and coordination. *Econometrica*, 61(5):1047–1071, September 1993.
- [12] Glenn Ellison. Basins of attraction, long-run stochastic stability, and the speed of step-by-step evolution. *Review of Economic Studies*, 61:567–588, 2000.
- [13] Robert Axtell Josua M. Epstein and H. Peyton Young. The emergence of classes in a multi-agent bargaining model. CSED Working Paper No. 9, 2000.
- [14] Dean Foster and H. Peyton Young. Stochastic evolutionary game dynamics. *Theoretical Population Biology*, 38:219–232, 1990.
- [15] Friedlin and Wentzel. *Random Perturbations of Dynamical Systems*. Springer-Verlag, 1984.
- [16] Eric Hobsbawm. Peasant land occupations. *Past and Present*, 62:120–152, 1974.
- [17] M. G. Kandori, G. Mailath, and R. Rob. Learning, mutation, and long run equilibria in games. *Econometrica*, 61:29–56, 1993.
- [18] Matthew Rabin. Inference by believers in the law of small numbers. *Quarterly Journal of Economics*, 117(3):775–816, 2002.
- [19] Larry Samuelson. *Evolutionary Games and Equilibrium Selection*. MIT Press, Cambridge, MA, 1997.
- [20] Theda Skocpol. *States and Social Revolutions*. Cambridge University Press, New York, NY, 1979.

- [21] Charles Tilly. *European Revolutions 1492-1990*. Cambridge University Press, New York, NY, 1994.
- [22] Thomas Tröger. Why sunk costs matter for bargaining outcomes: An evolutionary approach. *Journal of Economic Theory*, 111(2):375–402, 2002.
- [23] Amos Tversky and Daniel Kahneman. Judgment under uncertainty: Heuristics and biases. *Science*, 185:1124–1131, September 1974.
- [24] Jörgen W. Weibull. *Evolutionary Game Theory*. MIT Press, Cambridge, MA, 1995.
- [25] E. Wood. *Civil War and Insurgent Collective Action in El Salvador*. Cambridge University Press, 2003.
- [26] H. Peyton Young. The evolution of conventions. *Econometrica*, 61(1):57–84, January 1993.
- [27] H. Peyton Young. An evolutionary model of bargaining. *Journal of Economic Theory*, 59(1):145–168, February 1993.
- [28] H. Peyton Young. Conventions. *Journal of Economic Perspectives*, 10:105–122, 1995.
- [29] H. Peyton Young. Conventional contracts. *Review of Economic Studies*, 65:776–792, 1998.
- [30] H. Peyton Young. *Individual Strategy and Social Structure: An Evolutionary Theory of Institutions*. Princeton University Press, Princeton, NJ, 1998.
- [31] H. Peyton Young. The diffusion of innovations. Santa Fe Institute Working Paper, 2003.
- [32] H. Peyton Young and Mary Burke. Competition and custom in economic contracts: A case study of illinois agriculture. *American Economic Review*, 91:3:559–73, 2001.
- [33] F. Zeuthen. *Problems of Monopoly and Economic Welfare*. George Routledge and Sons, London, 1930.