

Institutional Equilibrium Selection by Intentional Idiosyncratic Play

Suresh Naidu † and Samuel Bowles †† *

October 5, 2004

Abstract

We introduce intentional idiosyncratic play in a standard stochastic evolutionary model of equilibrium selection, where the equilibria represent distributional conventions between members of two classes. Intentional idiosyncratic play 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, while the opposite is true in the standard (unintentional) approach. Second, where sub-population sizes and error 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 opposite is true in the standard dynamic. 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. Institutional transitions are less frequent under the intentional dynamic, and its long run average selects the stochastically stable state more reliably than the unintentional dynamic.

Keywords: Stochastic Stability, Nash Bargaining Solution, Multiple Equilibria, Institutions

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. We are grateful to Jorgen Weibull and Peyton Young for comments on earlier versions of this work, especially that appearing in Bowles(2004), and to participants in the working group on inequality in the long run at the Santa Fe Institute.

1 Introduction

Some institutions are much more common than others across the sweep of time and space. Markets, 50-50 crop shares, and monogamy are examples. Do these ways of organizing social interactions have common properties that account for their frequent emergence and long-term persistence? Why do feasible alternatives (direct barter, highly unequal crop shares, and polyandry, for example), emerge infrequently, and typically suffer rapid eclipse when they do? Can we say anything in general about the properties of those very common institutions sometimes called “evolutionary universals?”.

The development of stochastic evolutionary game theory (Foster and Young (1990), Young (1993b), Kandori, Mailath, and Rob (1993), Young (1993a), Binmore, Samuelson, and Young (2003) and its application to the dynamics of economic and other institutions (Young (1995) and Young and Burke (2001)) suggests an affirmative answer. Using these models, institutions may be represented as conventions, and stochastic shocks that induce idiosyncratic individual behaviors occasionally displace a population from the neighborhood of one convention to another. 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 (1998), the resulting equilibrium selection process generates a long term history in which populations tend to spend most of their time at conventions that are Pareto-efficient, risk-dominant (in two by two games), and (in a particular sense to which we will return) egalitarian.

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 forward-lookingness into an otherwise myopic updating process.

As far as we are aware, no models involving state-dependent error distributions have been studied, although Bergin and Lipman(1996) and Van Damme and Weibull(1999) have explored models with state-dependent error rates. Young(1993) shows, however, that the error distributions are irrelevant to the stochastically stable state, provided they have full support. Therefore, we must examine state-dependent error distributions that do not have full support, a condition that is natural given our interpretation of idiosyncratic play as directed and not merely mistakes.

Our modification to the standard dynamic is motivated by our belief that agents who act idiosyncratically in economic conflicts are acting in-

tentionally, 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).

Like Van-Damme and Weibull(2001) and Bergin and Lipman (1996), who conclude that “models or criteria to determine “reasonable” mutation processes. should be a focus of research in this area, 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 violation of an existing norm motivated by dissatisfaction with the status quo. (In our penultimate section we explore the effect of economic polarization on the rate of idiosyncratic play, thereby making it state dependent.)

There has been a large recent literature characterizing 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(1993b) examines the Nash demand game, and shows the Nash bargaining solution is stochastically stable. Young(1998) examines contract games, and shows that the Kalai-Smorodinsky solution is stochastically stable. Troger(2002) studies stochastic stability in a “hold-up” model, where the bargaining follows a first-stage investment decision. Agastya(2004) 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.

However, a common assumption in these papers, and the entire stochastic adjustment literature, is that the idiosyncratic noise takes the form of errors, in that there is no relevant systematic bias in the strategies played when non-best-responses occur. A possible explanation for this omission is that it might not matter; Young(1993a) shows that as long as errors have full support, the stochastically stable state is unchanged. However, as we will see, if 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, while the opposite is true in the standard (unintentional) approach. Second, as one would expect, in the intentional dynamic where

sub-population sizes and error 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 opposite is true in the standard dynamic. The conventions that are selected as stochastically stable under the intentional idiosyncratic play dynamic, not surprisingly, 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.

In order to establish a benchmark for contrast, in the next section we reproduce a version 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 version 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 (and Foster), 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 substantial (non vanishing) error rates and with differing group sizes.

In section 3 we introduce our intentional idiosyncratic play modification and demonstrate the results mentioned above as well as two others: i) 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; and ii) where the intentional and the standard dynamic select the same convention, the selection under intentional is more robust to substantial error rates (in the sense that the population spends more of its time in the selected state).

In the penultimate section we introduce three extensions. First, we show that the results of the intentional dynamic demonstrated above for the contract game apply equally to the Nash demand Game. Second, adapting van Damme and Weibull (2001), we show that we can generate the rate of idiosyncratic play endogenously as the equilibrium of a separate game embedded in the larger population dynamic here modeled (one which would plausibly model the collective action problem facing those seeking to displace the status quo). Finally, we adopt the measure of economic polarization suggested by Esteban and Ray (1994) to capture the insight that the rate of idiosyncratic play may be greater in more polarized societies and hence state-dependent.

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 according to the row player's preferences, and inversely to the column player's preferences. Clearly the diagonal of the game matrix constitutes the set of Nash equilibria, and they are all 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 also consider a matching dynamic with noise. Each period, a finite number of players from each population are matched to play the above contract game. Each time they are matched, agents choose a best response to the distribution of strategies in the opposing population, or, with small probability ϵ they play a non-best-response. However, if an agent is matched in two consecutive rounds, there is a small probability $(1 - \nu)$ that they play the same strategy they played last period. This “inertia” is necessary to ensure convergence.

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

The dynamic described above is a familiar myopic best-response dynamic with inertia, as in Young(1998), Samuelson(1997) and Agastya(2004). Denote the best-response functions $BR_R(q) : \Delta_C \rightarrow \Delta_R = e_{\text{argmax}_j q_j a_j}$ for the row population and $BR_C(p) : \Delta_R \rightarrow \Delta_C = e_{\text{argmax}_j p_j b_j}$ for the column population), where e_i is the i 'th standard basis vector of \mathbb{R}^N (or \mathbb{R}^M). If there are multiple best responses, the agents randomize uniformly over all of them.

With probability ν , each member of each population is given an opportunity to revise their strategy. Without introducing any errors, or reasons to play anything different (e.g. conformity), they will just choose the best-response to the distribution of play in the last period. This defines a Markov process: $P^\nu : X \rightarrow X$, defined by $P^\nu(\theta' | \theta) = \text{Prob}(\theta -$

$(x_1, x_2) + x_1 BR_R(q) + x_2 BR_C(p)$ where $x_1 \sim Bin(N, \nu)$, $x_2 \sim Bin(M, \nu)$ where $Bin(N, x)$ is a binomial distribution with N draws with probability of success given by x.

Following Young(1998) we 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.

Suppose that when agents can revise their strategies, they play a non-best response with probability ϵ . Thus, with probability ϵ they play a uniform distribution U . Later, we will use the following general distribution: $U(i, j)$ is 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 $U(0, K - 1)$ we just gives a Markov process defined by: $P^{\nu, \epsilon}(\theta' | \theta) = Prob(\theta - (x_1, x_2) + (x_1 - \sum_{i=0}^{K-1} y_{1i}) BR_R(\theta) + (x_2 - \sum_{i=0}^{K-1} y_{2i}) BR_C(\theta) + y_1 + y_2)$ where $y_1 \sim MN(K, x_1, U(0, K - 1))$ and $y_2 \sim MN(K, x_2, U(0, K - 1))$, with $MN(N, k, f)$ 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

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(1990) we call these stochastically stable states, with U-stability referring to stability under the perturbation process described in the preceding paragraph.

The proofs in this literature have largely been done using a result from Friedlin and Wentzell(1984), 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(1993), defines the resistance of a transition from state i to state j as the unique R_{ij} that satisfies $0 < \lim_{\epsilon \rightarrow 0} P_{ji}^\epsilon / \epsilon^{R_{ij}} < \infty$. If we build a complete weighed digraph on 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*. See Young(1993) for details.

We shall make use of three propositions proved in Binmore-Samuels-Young(2004):

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*

¹a directed graph where every node save one has only one edge exiting it, The node with no exiting edge is called the root.

an edge that is maximal over all edges, deleting that edge will give the minimal tree.

Proposition 2.3 (Binmore-Samuelson-Young 2004 Proposition 10). . If the set of contracts is given by $0, f(0), \delta, f(\delta)$, if f is strictly concave and δ is sufficiently small, then the stochastically stable state, in the U -dynamic, is the Kalai-Smorodinsky solution.

Our first observation is that the transitions between equilibria in the contract game are instigated by the class that loses from the transition. That is, the resistance of the transition from contract i to contract j is given by

$$R_{ij} = \min(\lceil \frac{Ma_i}{a_i + a_j} \rceil, \lceil \frac{Nb_j}{b_i + b_j} \rceil). \quad (1)$$

If $N=M$ is sufficiently large, then we can use the *reduced resistances* $r_{ij} = \min(\frac{a_i}{a_i+a_j}, \frac{b_j}{b_i+b_j})$. Note that if $b_i > b_j$ and $a_i < a_j$ then $r_{ij} = \frac{a_i}{a_i+a_j} < 1/2 < \frac{b_i}{b_i+b_j} = r_{ji}$. So i is U -stable.

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 monomorphic states 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 exponentially, then the side that mutates faster does worse. We show this in the 2-contract case, and note that it generalizes to the K -contract case easily.

Consider a 2-contract game with payoffs $((a_0, b_0), (a_1, b_1))$ ($a_0 > a_1, b_0 < b_1$). Assume the row population has population size N and the column population is size M with $N > M$.

Contract	0	1
0	a_0, b_0	$0, 0$
1	$0, 0$	a_1, b_1

Then the resistances (see Young(1998)) are given by $R_{ij} = \min(\lceil \frac{Ma_i}{a_i+a_j} \rceil, \lceil \frac{Nb_j}{b_i+b_j} \rceil)$. This is because with asymmetric group sizes, we cannot abstract from the actual population sizes in computing the amount of idiosyncratic play necessary for a transition to occur, so we cannot follow Young in going to reduced resistances, which divide all the resistances by the population sizes.

Thus, if M and N are sufficiently large we divide all the resistances by M , we can see that the resistances will be given by $R_{ij} = \min(\frac{a_i}{a_i+a_j}, \lceil \frac{N}{M} \frac{b_j}{b_i+b_j} \rceil)$

In addition, assume the mutation rate for row is ϵ^T instead of ϵ , then the resistances will be given by $R_{ij} = \min(N \frac{a_i}{a_i+a_j}, TM \frac{b_i}{b_i+b_j})$,

Thus, for T(or N/M) sufficiently large, the resistances will be $R_{01} = \frac{a_0}{a_1+a_0}$, and $R_{10} = \frac{a_1}{a_1+a_0}$. We thus get that the transition from contract 1 to contract 0 occurs when the idiosyncratic play of the columns, who prefer contract 1, that is causing the transition. Similarly for the transition from 0 to 1 being driven by the idiosyncratic play of the row players.

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 $r_{ij} = \min(\frac{a_i}{a_i+a_j}, \frac{Tb_i}{b_j+b_i})$ which is equal to $\frac{a_i}{a_i+a_j}$ for sufficiently large T , which is increasing in a_j and decreasing in a_i . 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 $a_j/a_j + a_1$, while the lowest outgoing edge is $a_1/a_1 + a_j$. 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.

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 some of these problems with the U-dynamic, we now define a new dynamic Γ^I or the I-dynamic, where the error distribution chosen by agents is supported only on the strategies that would be beneficial relative to the current state. This involves a degree of foresight, something that is missing in the standard, purely myopic-with-errors dynamic.

3.1 Intentional Dynamics

This dynamic is somewhat more complicated, involving as it does state-dependent error distributions. First, given $\theta = (p, q)$ define:

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

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 $y_i \sim MN(K, x_i, U(0, K - 1))$ in the above model, we have instead that $y_1 \sim MN(K, x_1, U(0, i^R(\theta)))$ and $y_2 \sim MN(K, x_2, U(0, i^C(\theta)))$

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 $(N, 0, 0, \dots, 0), (0, 0, 0, \dots, M)$ 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 . But this follows from the fact that from θ the column population can mutate to the state $M, 0, \dots, 0$ and the same time that the row population mutates to the state $0, 0, 0, \dots, N$ which leads the row population to respond (with no mutations) with $N, 0, \dots, 0$, and the column population to respond with $0, 0, 0, \dots, M$.

It is clear that as $\epsilon \rightarrow 0$, $P^{\nu, I}(\epsilon) \rightarrow P^\nu$. It is also clear that the process converges exponentially. Let

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

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.

This definition only makes sense in our particular context, where we restrict our attention to contract games, where every Nash-equilibrium is Pareto-optimal, so that there are no mutually beneficial contracts from which both classes gain.

It is obvious that $\infty > \lim_{\epsilon \rightarrow 0} \frac{P^{j,k, \nu, \epsilon, I}(\epsilon)}{\epsilon^{R_{jk}^I}} > 0$, 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.

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 Young(1993), 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.

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

Table 2: U-Resistances for Example 1

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

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 symmetric contract games, the U-stable equilibrium is I-stable.*

Proof. a) Follows from the fact that $R_{01}^I = \frac{a_0}{a_1+a_0} \geq \frac{b_1}{b_0+b_1} = R_{10}^I$ iff $R_{01}^U = \min(\frac{a_0}{a_0+a_1}, \frac{b_0}{b_1+b_0}) < \min(\frac{a_1}{a_0+a_1}, \frac{b_1}{b_1+b_0}) = R_{10}^U$ b) Follows trivially since in symmetric games $R_{jk} = R_{jk}^I$, so the minimal U-tree is also the minimal I-tree. □

After this, one might be tempted to think that there is no substantial difference between U-stability and I-stability, however, the example in Table 1 illustrates otherwise.

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 in Figure 1.





Table 3: I-Resistances for Example 1

Root/Trees			
0	1.583	1.455	1.830
1	1.500	1.628	1.733
2	1.935	1.702	1.689

However, with intentional error distributions (the I-dynamic), the tree resistances are given in Table 3. The full set of trees is given in Appendix B.

So the minimal U-tree has root 0, with resistance 1.455, shown in figure 2. 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 ($a_0b_0 = \max_i a_i b_i$). This is a general difference, as illustrated by the proposition below ².

Proposition 3.3. *Assume equal group sizes, and let the contracts be given by $(a_i, b_i = f(i))$, $a_i = i\delta$, $i \in (1, \dots, 1/\delta - 1)$ where f is a strictly decreasing, strictly convex 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. See Appendix. □

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, we show that the Nash solution is I-stable. The difference stems from the fact that the intentional resistances are always lowest 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

²We have also proved this for the arbitrary 3x3 contract game.

implausible, compared to the “neighboring” transition feature of the I-dynamic.

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 $((a_0, b_0), (a_1, b_1))$ ($a_0 > a_1, b_0 < b_1$). Assume that contract 1, favoured by the column 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 .

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

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 M , and the best contract for them is K-1. Then choosing M, N, T such that $N \frac{a_{K-1}}{a_{K-1} + a_i} = R_{(K-1)i}^I < \min_j R_{j(K-1)}^I = \min_j MT \frac{b_j}{b_j + b_{K-1}} \forall i$. Then, by the local resistance test, the K'th contract is I-stable.

3.3 Convergence

In this subsection we compare the convergence properties of our modified I-dynamic to the standard U-dynamic. We will demonstrate two facts: 1) the maximum waiting time until the system arrives in the stable state is higher in our model than in the standard dynamic. 2) For ϵ that are not arbitrarily small, the I-dynamic spends much more time in its stable state than the U-dynamic.

Both of these observations are straightforward consequences of our definition. 1) follows since the maximum waiting time is a function of the resistance of the minimal tree. So, since the I-dynamic has higher resistances in general, it will take longer for any particular transition to occur, and, for generic games, the maximum waiting time is at least as large as under the U-dynamic. A simple way of seeing this follows from noting that in the contract game, all of the R^I are greater than $1/2$, while

the R^U are all less than 1/2. Since all trees must have the same number of edges, it must be that the resistance of the minimal I-tree is greater than the resistance of the minimal U-tree.

To get a more precise result, we use a result from Ellison(2000), that bounds the maximum waiting time until the first entry into the stable state as a function of the modified coradius of a state x , which is a measure of how difficult to enter x . Note that the modified Coradius (written CR_I^* , see Ellison(2000)) of the I-stable equilibrium is given by the resistance of the transition from either of the adjacent contracts (if δ is small these will be close), and so, without loss of generality, we can write $CR_I^*(i_{Nash}) = \frac{a_{Nash} + \delta}{a_{Nash} + a_{Nash} + \delta}$. But, $CR_U^*(i_{KS}) = \frac{a_{min}}{a_{min} + a_{KS}}$.

If the Nash and KS solutions are both in the contract set, then we can see that, for sufficiently small δ

$$\begin{aligned} CR_U^*(i_{KS}) &= \min\left(\frac{a_0}{a_0 + a_{KS}}, \frac{b_{K-1}}{b_{K-1} + b_{KS}}\right) < \frac{a_0}{a_0 + a_{KS}} \\ &= \frac{1}{1 + a_{KS}/\delta} < \frac{1}{1 + \frac{a_{Nash}}{a_{Nash} + \delta}} = \frac{a_{Nash} + \delta}{a_{Nash} + a_{Nash} + \delta} = CR_I^*(i_{Nash}) \end{aligned}$$

Since the maximum waiting time until first making a transition to the stable state is $O(\epsilon^{CR^*})$, this shows that the I-dynamic is slower, in the sense of taking longer to reach the stable state, than the U-dynamic. However, we are not overly discouraged by these long transition times, as the literature has developed numerous mechanisms for speeding up the waiting time, from local interaction (Ellison 1993) to errors in the payoffs (Binmore, Samuelson, Vaughan 1995).

Another limit, one that has been somewhat neglected in the literature, is the rate of convergence as $\epsilon \rightarrow 0$. Young(1998a) proves that in symmetric 2x2 games, the convergence properties of the model are reasonable, in that for any ϵ , we can find a population size that will guarantee that at least $1 - \epsilon/2$ of the agents are playing the stochastically stable (U-stable, in our language) equilibrium with probability arbitrarily close to 1. However, Young's result relies on letting the population size get arbitrarily large. We are interested in the finite population properties of the model, because we are concerned with the transition process that actually occur between different contracts in groups of historically relevant sizes. Taking either arbitrarily large population sizes or arbitrarily small idiosyncratic play rates means that transitions will almost never occur. While this might be helpful to understand why an observed contract may be stable, it is less helpful in understanding the mechanism that generates transitions between contracts. In addition, Binmore-Samuelson-Young prove in their paper that if the dynamics are continuous, taking the population size to infinity before taking ϵ to 0 selects a different contract than taking $\epsilon \rightarrow 0$ first. It is not clear which of either of these processes best illuminates real historical processes.

Do the results for the limit as ϵ goes to zero indeed hold for non-vanishing error rates? We illustrate the answer to this question for the 2x2 case equal population sizes, and with no inertia ($\nu = 0$), although we think the idea generalizes. Let $\alpha = \frac{a_0}{a_0+a_1}$, and $\beta = \frac{b_0}{b_0+b_1}$. Without loss of generality, assume that 0 is both I and U stable, and that $\alpha < \beta \iff 1 - \beta < 1 - \alpha$. Note also that $\alpha < 1/2 < \beta$. The dynamic is therefore best response, and we ignore the cases where the limit cycle $(1,0) \rightarrow (0,1) \rightarrow (1,0)$ is reached, and assume the population size N is such that the mixed equilibrium is never reached.

This gives us a transition matrix that looks like:

$$\begin{array}{ccc} \text{Contract} & 0 & 1 \\ 0 & 1-u & u \\ 1 & v & 1-v \end{array}$$

The ergodic distribution for this matrix is given by $\mu(0) = v/u + v, \mu(1) = u/v + u$. We will look at different values for d and c under the two dynamics. With the U-dynamic, both populations can engage in sufficient idiosyncratic play to engender a transition. So first define:

$$\begin{aligned} P_R &= \sum_{k=\lceil \alpha N \rceil}^N \binom{N}{k} \epsilon^k (1-\epsilon)^{N-k} \\ P_C &= \sum_{k=\lceil \beta N \rceil}^N \binom{N}{k} \epsilon^k (1-\epsilon)^{N-k} \\ Q_R &= \sum_{k=\lceil (1-\alpha)N \rceil}^N \binom{N}{k} \epsilon^k (1-\epsilon)^{N-k} \\ Q_C &= \sum_{k=\lceil (1-\beta)N \rceil}^N \binom{N}{k} \epsilon^k (1-\epsilon)^{N-k} \end{aligned} \tag{2}$$

Let $u^U(\epsilon) = P_R + P_C$ and $v^U(\epsilon) = Q_R + Q_C$. So the ergodic distribution is defined by $\mu^U(0, \epsilon) = v^U/v^U + u^U$. Let $p^I = P_C$ and $q^I = Q_A$, then $\mu^I(0, \epsilon) = u^I/u^I + v^I$.

Proposition 3.4. $\mu^U(0, \epsilon) \geq \mu^I(0, \epsilon)$

Proof. This follows immediately from establishing that $u^U/v^U - u^I/v^I \geq 0 \iff \frac{P_R+P_C}{Q_R+Q_C} > P_C/Q_R \iff Q_R P_R > Q_C P_C$. Now note that if we write $L(\alpha) = Q_C P_C$, then $L(\beta) = Q_R P_R$ and L is a single peaked function that achieves its unique maximum at $1/2$. Thus $Q_R P_R$ is larger since $1/2 - \alpha > \beta - 1/2 > 0$, which follows from the fact that $\alpha < 1/2 < \beta$ and $\alpha < 1 - \beta$. \square

Figure 3 illustrates the magnitude of the difference as a function of ϵ .

In short, the I-dynamic has attractive features: the people who induce transitions are those who benefit, higher rates of idiosyncratic play and

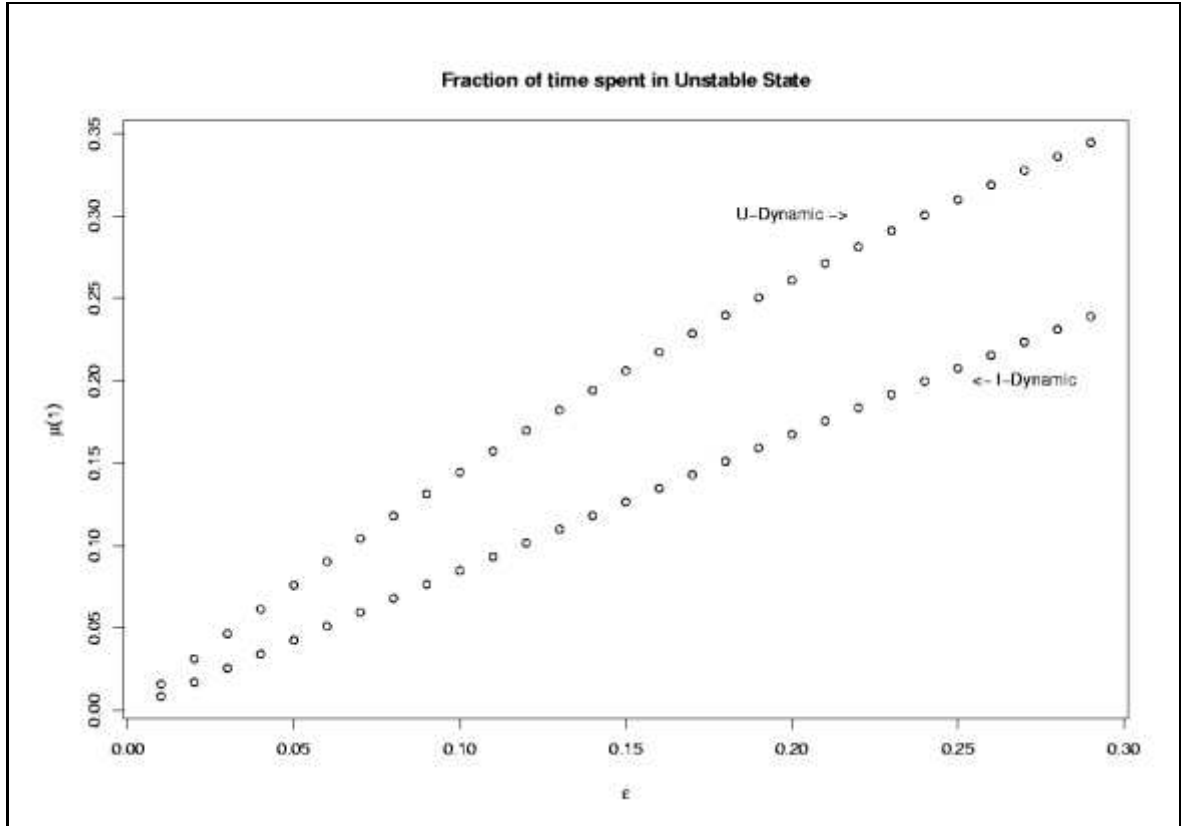


Figure 1: $a_0 = b_0 = 1$ and $a_1 = 1.3$ and $b_1 = .7$ for various values of ϵ

smaller group size are beneficial, and the predictions of of the model are robust under substantial rates of idiosyncratic play. But like the models using the U-dynamic, it is quite abstract, the restrictions on the error-generating process are quite minimal, and has so far it has been applied to a single institutional setting, the contract game. Can the model be extended to encompass other settings and a less minimalist conception of idiosyncratic play?

4 Nash Demand and Double-Sided Auctions

We now consider the effects of intentional idiosyncratic play in alternative specifications of the bargaining game. We first consider the Nash demand game, the stochastically perturbed properties of which were examined by Young(1993b). The difference between the Nash demand Game and the

contract game is that in the former, agents get their offer even if the offers don't agree with their opponents. Contracts that don't exhaust the surplus can be struck. That is, in our previous notation, now 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	0, 0
<i>j</i>	a_j, b_i	a_j, b_j

A recent paper by Agastya(2004) explores the stochastic stability of various 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 . The game matrix is given below.

<i>Contract</i>	<i>i</i>	<i>j</i>
<i>i</i>	a_i, b_i	0, 0
<i>j</i>	$\rho a_j + (1 - \rho)a_i, \rho b_j + (1 - \rho)b_i$	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, the I-stable contract approximates the Nash Bargaining solution*

Proof. Suffices to show this for the Nash demand game, the argument is the same for the Double-Sided Auction. This result follows from the fact that $r_{ij}^I = f(s_i) - (f(s_j)/f(s_i))$, if $s_i < s_j$ and $r_{ij}^I = s_i - s_j/s_i$, if $s_i > s_j$ which is exactly the value of $r_{ij}^U = \min(\frac{f(s_i) - f(s_j)}{f(s_i)})$ U dynamic □

The reason they agree is that the off-diagonal payoffs make the population that stands to lose from a transition more willing to accommodate the idiosyncratic play of the opposing population. The lowest number of idiosyncratic plays required to make a transition from i to j is equal to the resistance of the population that stands to lose, because they do not have to take 0 if they mismatch, but instead get a positive payoff. Thus, since the U-stable is determined by the lowest resistance transitions, and the lowest resistance transitions happen to be those generated by the idiosyncratic play of the population that stands to benefit from the transition, the U-stable and I-stable states are identical.

5 State-Dependent Mutation Rates

In this section we explore alternative, state-dependent idiosyncratic play rates under our assumption of state-dependent idiosyncratic distributions. The first subsection adapts a result from Van-Damme and Weibull(1999) to the case where individuals are not determining their error rates as an

individual choice problem, but instead we look at the error rates as the outcome of a N-player public goods game.

5.1 Collective Action

Our purpose here is to generate the ϵ probability of idiosyncratic play endogenously, as the equilibrium of a separate, within-population model. We can show that our results above are robust to ϵ modelled as the mixed strategy Nash equilibrium of a particular public-good game specification.

Let ϵ_i be the probability that agent i (in a population of size N) plays idiosyncratically. Let θ be a state of the system. $B(\theta)$ is a function that computes the prospective payoff from idiosyncratic play that the population is currently at θ . δ represents a behavioral parameter, for example the “pleasure of agency” (Wood 2003), or the effectiveness of within-population political organization. Many empirical studies (e.g. Wood 2003, Snyder and Tilly 1972, McAdam 1986, Thompson 1959) of insurgent collective action argue that the actual material payoffs to a class are of only secondary importance, they matter only insofar as they facilitate the psychological predisposition to revolt (Moore 1978). However, we maintain the assumption that the δ parameter is very small in relationship to the payoffs involved in the larger population game. Our aim is to show that relatively small individual probabilities of engaging in collective action can still generate population-level historical transitions among institutions.

Suppose each agent adopts a mixed strategy, deviating from the status quo conventions (say, by striking) and with probability ϵ and best responding with probability $1 - \epsilon$. Then we write the payoffs to an agent i in population j choosing to strike with prob ϵ_i as $\pi_{\epsilon_i}^{i,j}(\epsilon_1, \epsilon_2, \dots, \epsilon_N, \delta, B(\theta)) = u(\epsilon_1, \dots, \epsilon_N, \delta, B(\theta)) - v(\epsilon_i)$, with u continuous and increasing in all arguments, and v continuous, increasing, and convex.

We also assume that $\pi_{\epsilon_i}^{i,j}(0, \sigma_{-i}, 0, B(\theta)) = 0 \forall \sigma_{-i}, \forall \theta \in X$ where σ_{-i} is an $N-1$ vector of the other players strategies. This assumption just says that the best-response when there are no gains to collective action is to never participate. We assume that this payoff function is identical for all agents in a given population, so we suppress the i superscript. The first-order conditions characterizing the Nash equilibrium are

$$\pi_{\epsilon_i}^j(\epsilon_1, \epsilon_2, \dots, \epsilon_n, \delta, B(\theta)) = u_{\epsilon_i} - v'(\epsilon_i) = 0$$

We restrict ourselves to symmetric Nash equilibria (which exist since the strategy space is compact and convex and all the payoff functions are symmetric), $\epsilon_i = \epsilon_j, \forall i, j$, and thus just write

$$\pi_i^j(\epsilon, \delta, B(\theta)) = 0 \forall i$$

This implicitly defines a continuous function $\epsilon^j(\delta, B(\theta))$ and $\lim_{\delta \rightarrow 0} \epsilon^j(\delta, B(\theta)) = 0$. We borrow the terms *nice*, *similar*, and *regular* from Van Damme and Weibull, but redefine them so that they apply in our public-goods game setting, rather than the individual choice setting explored in their paper.

Definition 5.1. π^i is *nice* if $\lim_{\delta \rightarrow 0} \epsilon^i(\delta, \lambda B(\theta)) / \epsilon^i(\delta, B(\theta)) > 0 \forall \lambda > 0, \theta \in X$

Definition 5.2. π^1, π^2 are *similar* if $\lim_{\theta \rightarrow 0} \epsilon^1(\delta, B(\theta)) / \epsilon^2(\delta, B(\theta)) > 0 \forall \theta \in X$

Definition 5.3. π^1, π^2 is a *regular* function profile if both are nice and they are similar.

Proposition 5.4. *Given a K-contract game, if π^1, π^2 is a regular function profile, then the I-stable state is the same as that with constant error rates at all states.*

Proof. By proposition 3 in Van-Damme and Weibull, regularity implies that for $\theta, \theta' \in X$ there exist $\alpha, \beta, \bar{\delta}$ such that $\beta > \lim_{\delta \rightarrow 0} \frac{\epsilon^1(\delta, B(\theta))}{\epsilon^2(\delta, B(\theta'))} > \alpha$ since X is finite, we can find $\alpha, \beta, \bar{\delta}$ such that $\forall \theta, \theta' \in X, \beta > \lim_{\delta \rightarrow 0} \frac{\epsilon^1(\delta, B(\theta))}{\epsilon^2(\delta, B(\theta'))} > \alpha$. Therefore all the ϵ go to 0 at the same rate at all states, so the resistances are unchanged, and the I-stable state is unchanged. \square

Examples of CA-games that admit regular function profiles are not hard to find. Consider the game given by the utility functions: $\pi^i(\epsilon_1, \epsilon_2, \dots, \epsilon_N, \delta, B(\theta)) = \delta B(\theta) \prod_{j=0}^N \epsilon_j - \epsilon_i^2 / 2$ for player i in the row population.

This gives first-order conditions $\delta B(\theta) \prod_{j \neq i} \epsilon_j - \epsilon_i = 0 \forall i \in (0, \dots, N)$. If we consider only symmetric Nash equilibria (characterized by a common value of ϵ for all players), we get $\delta B(\theta) \epsilon^{n-1} - \epsilon = 0$, which, restricting ourselves to non-zero solutions, gives us that $\epsilon^R(\delta, B(\theta)) = (\delta B(\theta))^{2-n}$. Then, clearly, $\lim_{\delta \rightarrow 0} \frac{\epsilon^R(\delta, B(\theta))}{\epsilon^R(\delta, B(\theta'))} = (B(\theta) / B(\theta'))^{2-n} > 0$ so the functions π^i are nice.

If the i 'th player in the column population of identical size N , faces a similar within population payoff function given by $\psi^i(\epsilon_1, \epsilon_2, \dots, \epsilon_N, \delta, B(\theta)) = \delta B(\theta) \prod_{j=0}^N \epsilon_j - \epsilon_i^2 / 2$, then the identical calculation shows that the ψ^i are nice. A trivial computation then shows that $\lim_{\delta \rightarrow 0} \frac{\epsilon^R(\delta, B(\theta))}{\epsilon^C(\delta, B(\theta))} > 0 \forall \theta \in X$, which implies regularity.

5.2 Polarization and Conflict

While we have already implicitly introduced a function profile that is not similar when we considered different idiosyncratic rates above, in this subsection we develop an empirically relevant non-nice error function. First notice that *any* error function of the form $\delta^{B(\theta)}$ is not nice if B is non-constant, since there will be states θ, θ' such that $\lim_{\delta \rightarrow 0} \frac{\epsilon(\delta, B(\theta))}{\epsilon(\delta, B(\theta'))} = 0$

But what functions are empirically plausible for a choice of B ? We would like to capture the idea that idiosyncratic play will be greater in highly polarized societies. Esteban and Ray(2004) introduce and axiomatize a measure of polarization. Polarization is related to inequality, but is independent of it, and the Esteban and Ray measure is designed to capture two effects of a trait distribution: identity and alienation. Identity is a measure of how closely one is related to ones nearest neighbors in the distribution, and alienation is a measure of how far away one and ones neighbors are from other groups in the distribution. Esteban and Ray estimate the polarization of various income distributions, and find increased conflict where polarization is high. Formally, given a discrete distribution of a real-valued trait $p(w)$, polarization is given by $\Psi(f) = \sum_x \sum_y p_x^{\alpha+1} p_y |x - y|$, where $\alpha \in [.25, 1]$.

This maps naturally onto our state-space, consisting as it does of a distribution of strategies and concomitant payoffs. We get that, for $\theta = (p, q) \in X$

$$\begin{aligned} \Psi(\theta) = & \frac{1}{(M+N)^2} \left(\sum_{i=0}^{K-1} \sum_{j=0}^{K-1} p_i \frac{p_i}{N+M}^\alpha p_j M^{-1} |a_i q_i - a_j q_j| \right. \\ & + \sum_{i=0}^{K-1} \sum_{j=0}^{K-1} q_i \frac{q_i}{N+M}^\alpha q_j N^{-1} |b_i p_i - b_j p_j| + \\ & \sum_{i=0}^{K-1} \sum_{j=0}^{K-1} q_i \frac{q_i}{N+M}^\alpha p_j |b_i \frac{p_i}{N} - a_j \frac{q_j}{M}| + \\ & \left. \sum_{i=0}^{K-1} \sum_{j=0}^{K-1} p_j \frac{p_j}{N+M}^\alpha q_i |b_i \frac{p_i}{N} - a_j \frac{q_j}{M}| \right) \quad (3) \end{aligned}$$

This expression is the sum of the within-class polarization in both classes(the first two terms), then the second two terms measure the between class polarization. At a recurrent contract i ($\theta = N e_i, M e_i \in X$), the within-class polarization is 0(since all agents play the same contract and get the same payoff), and the between-class polarization is given by $\Psi(e_i, e_i) = (\frac{N}{N+M}^{\alpha+1} \frac{M}{N+M} + \frac{M}{N+M}^{\alpha+1} \frac{N}{N+M}) |a_i - b_i|$, which for any given contract, is maximized when the relative population sizes are the same(1/2) and the contract payoffs are as distant as possible, and minimized when the contracts divide the prize equally .

If we let a state-dependent idiosyncratic play rate be given by $\epsilon(\delta, \theta) = \delta^{1/(1+\Psi(\theta))}$, it is easy to see that if the polarization at a given state is very high, its stability will be reduced. Assume only two contracts and equal population sizes, for example:

Example 2		
Contract	0	1
0	.2, 5	0, 0
1	0, 0	1, 1

Note that both have the same risk-factor, so in the absence of state-dependent idiosyncratic play, both should be I-stable. However, if we use $\epsilon(\delta, \Psi(\theta)) = \delta^{\Psi(\theta)}$, and note that the polarization at 0 is $.5^{\alpha+1}4.8$ (the polarization at 1 is 0). Taking $\alpha = 1$ we get that $R_{01}^I = \frac{1}{1+\Psi(Ne_0, Ne_0)} \frac{5}{6} = 5/(6 \times 2.2) < 1/1.2 = R_{10}^I$ so 1 becomes the only I-stable state. The essence of the calculation is that the unequal, polarized state has a higher level of idiosyncratic play.

6 Conclusion

Transitions among real institutions take place under conditions vastly more complex than the models we have considered. Not without reason did the historian Eric Hobsbawm identify institutional change as the most difficult problem in the study of history and society. 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 surrounding the real historical equilibrium selection process. The second is deviant challenges to the status quo and the occasional concession of best-responding members of the opposing group that occasionally results when the level of deviance is sufficiently great. The end of Communist rule in many countries and the demise of Apartheid in South Africa appear to reflect this pattern.

The stochastic evolutionary approach thus 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.

7 Appendix A

To prove proposition 3.3 we first need a lemma:

Lemma 7.1. *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 $r_{ij} = f(a_i)/(f(a_i) + f(a_i + (j - i)\delta))$ 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 $r_{ij} = a_i/a_i + j\delta$, which will be lowest for $j = i - 1$ \square

Proof of Proposition: $a_{j^*} = \operatorname{argmax}_{a_i} a_i f(a_i)$. 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 $s, f(s)$ $a_1 < s < a_2$ note that $sf(s)$ is a convex function of s , therefore it has a unique maximum at a_{j^*} . Therefore $sf(s)$ is increasing for $s < a_{j^*}$ and decreasing for $s > a_{j^*}$.

Now, consider a node $i < j^*$, we just need to show that $\frac{f(a_i)}{f(a_i) + f(a_i + \delta)} < \frac{a_i}{a_i + a_i - \delta}$, which would show that the transition to the right is cheaper than the transition to the left. This is equivalent to $f(a_i + \delta)/f(a_i) > (a_i - \delta)/a_i$ subtracting 1 from both sides, we get $\frac{f(a_i + \delta) - f(a_i)}{f(a_i)} > -\delta/a_i$ which is approximately equal to (for small δ) $f'(a_i)a_i + f(a_i) > 0$ which reduces to $(a_i f(a_i))' > 0$ which is given by our statement that $sf(s)$ is increasing. The case $i > j^*$ follows symmetrically, with the least edge pointing to the left. However, the actual Nash product maximand j^* can have an edge exiting to the left or to the right, which will give us a cycle of length 2.




We must show that the edge exiting j^* is maximal over the whole tree. Note that the edge exiting j^* will have resistance equal to $r_{j^*j^*+1} = \frac{f(j^*)}{f(j^*) + f(j^*+1)}$ or $r_{j^*j^*-1} = \frac{j^*}{j^*-1 + j^*}$. Since $r_{j,j+1}$ is strictly increasing in j for $j < j^*$ and $r_{j-1,j}$ is strictly decreasing in j for $j > j^*$, it follows that the edge exiting j^* is maximal over the entire tree, since it is greater than all the edges to the right and to the left. QED.

8 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.