Evolution Through Imitation in a Single Population[1]

David K. Levine and Wolfgang Pesendorfer[2]

First version: September 29, 1999

This version: May 10, 2000

Abstract: Kandori, Mailath and Rob [1993] and Young [1993] showed how introducing random innovations into a model of evolutionary adjustment enables selection among Nash equilibria. Key to this result is that poorly performing strategies may be introduced in sufficient numbers that they begin to perform well. We examine imitation as an alternative and more plausible propagation mechanism. If imitation is much more likely than innovation, it is significantly easier to compute long-run equilibrium. The long-run limit contains only pure strategies. Calculations can be made by comparing pairs of pure strategies to see how well they do against one another. A sufficient condition for a profile to be the unique long-run equilibrium is that it beat all others in pairwise contests. A number of examples are considered.

Note: this version is preliminary and incomplete.

---

[1] We are grateful to support from National Science Foundation Grants.

[2] Departments of Economics UCLA and Princeton.

1. Introduction

In two influential papers Kandori, Mailath and Rob [1993] and Young [1993] showed how introducing random innovations (mutations) into a model of evolutionary adjustment enables predictions about which of several strict Nash equilibria will occur in the very long run. Key to this result is the possibility that strategies that perform poorly may be introduced into the population in sufficient numbers through innovation that they begin to perform well. Here we examine imitation as an alternative propagation mechanism. A striking fact is that if imitation is much more likely than innovation, it is significantly easier to find the long-run equilibrium. First, the long-run limit contains only pure strategies. Second, calculations can be made by comparing pairs of pure strategies to see how well they do against one another. One useful result is that it is sufficient that a strategy profile beat all others in pairwise contests. As we illustrate through examples, this is implied by, but more likely to be satisfied than, the criterion of ½-dominance proposed by Morris, Rob and Shin [1993].

This work stems from our admiration of the existing theory, a desire to apply it to interesting games, and our dissatisfaction with innovation as a propagation mechanism. To us, the theory of evolution with persistent randomness is a theory of the propagation of ideas through a population. Key to analyzing long-run dynamics is the possibility that bad ideas may spread, changing which ideas are good and which are bad. To think of this propagation taking place through random innovations seems highly unsatisfactory. It is hard to think of any significant changes in institutions that have occurred in human history because large numbers of people simply happen to have tried the same thing at more or less the same time. Rather, we believe that ideas are spread through imitation. One example is the change of institutions through civil disobedience, as occurred, for example, in East Germany. Initially a small number of people protested the existing government. It seems likely that rational calculation would show that this was a bad idea – that the probability of being punished severely was substantial. But others imitated, and once the idea spread sufficiently widely, the probability of punishment dropped, and civil disobedience became a best response.

In addition to the work mentioned above, there are several other papers that have a connection to our results. Bergin and Lipman [1994] show that the relative probabilities of different types of noise can make an enormous difference in long-run equilibrium; here we explore on particular theory of how those relative probabilities are determined. Van Damme and Weibull [1998] study a model in which it is costly to reduce errors, and show that the basic 2x2 results on risk dominance go through. By way of contrast, our focus is on larger more complex games. Our result about winning pairwise contests is connected to a result of Kandori and Rob [1993]. We explain this connection in conjunction with introducing our own result about winning pairwise contests.

2. The Model

We study a symmetric normal form game with a single population of players. There is a finite number S of pure strategies, and we write $s \in S$ for a typical pure strategy. Notice that we use the same symbol for the number of pure strategies and the set of pure strategies. Mixed strategies are vectors of probabilities denoted by $\sigma \in \Sigma$. A mixed strategy is called *pure* if it puts unit weight on a single pure strategy; we denote the mixed strategy corresponding to the pure strategy s also by s. The utility of a player depends on his own pure strategy and the mixed strategy played by the population. It is written $u(s,\sigma)$. We assume that u is continuous in σ. A prototypical example is a game in which players from different populations are randomly matched to play particular player roles; we discuss this along with other examples below.

There are m players in the population, each of whom plays a pure strategy. At time t we denote the distribution of strategies in the population by $\sigma_t \in \Sigma$. Initially at time $t = 0$ there is a given initial condition σ_0. In subsequent periods σ_t is determined from σ_{t-1} according to the following "imitative" process.

1) One player i is chosen at random. Only this player changes his strategy.
2) With probability $C\varepsilon$ player i chooses from S randomly using the probabilities σ_{t-1}. This is called *imitation*: strategies are chosen in proportion to how frequently they were played in the population in the previous period.
3) With probability ε^n player i chooses from S randomly with equal probabilities of $1/S$ of choosing each strategy. This is called *innovation*: strategies are picked regardless of how widely used they are, or how successful they are.
4) With probability $1 - C\varepsilon - \varepsilon^n$ player i chooses randomly with equal probabilities among the set of strategies that solve the problem used by the most players

$$\max_{\tilde{s}|\sigma_{t-1}(\tilde{s})>0} u(\tilde{s}, \sigma_{t-1}) .$$

This is called a *relative best response*: it is the best response among those strategies that are actually used by the particular population.

Observe that this process gives rise to a Markov process M on the state space $\Sigma^m \subset \Sigma$ consisting of all mixed strategies consistent with the grid induced by each player playing a pure strategy. Note that all pure strategies are in Σ^m. The process M is positively recurrent, and so has a unique invariant distribution μ^ε. Our goal is to characterize $\mu \equiv \lim_{\varepsilon \to 0} \mu^\varepsilon$.

Our main assumption is that imitation is much more likely than innovation. Specifically

Unlikely Innovation: $n > m$.

In particular this means that as $\varepsilon \to 0$ it is far more likely that every player in the population will change strategy by imitation than even a single player will innovate.
By way of contrast standard evolutionary models of persistent randomness assume

Standard Model: $C = 0$.

Actually the standard model does not use relative best response for players who are not innovating (often called mutating), but typically some variation on the best response dynamic. As we shall see below, this does not make that much difference.

Discussion and Examples

We should first indicate the strong connection between the case of unlikely imitation $C = 0$, and the standard case of innovation (or mutation) and a best-response like dynamic. In matching games, generally speaking, the basins of steady states are determined by fractions of the population. Consequently as the population size goes to infinity, the difference in the number of innovations needed to move from one steady state to another, versus the number required to move back, typically goes to infinity. It is this difference that determines which are the stochastically steady states. On the other hand, while the relative best-response dynamic is different than the best-response dynamic, it requires only a number of innovations equal to the total number of strategies in the game to make sure that every strategy is actually in use. In this interior case, the best-response and relative best-response dynamics are identical. Since this is a fixed number of innovations, when the population size is large enough, it is swamped by the difference in innovations required to move between steady states, and the calculations made for stochastic stability in the best-response case coincide with those in the relative best-response case. As a result, except in knife-edge cases, we do not expect unlikely imitation to yield different results than Kandori, Mailath and Rob [1993], or Young [1993], or subsequent research using the standard model.

We should also note that the assumption of a single-population is significant. In the existing literature, this has been the primary focus of research, although Hahn [1995] does have some results in the multiple population case. Here the single-population assumption not only means that all players are *a priori* identical, but that there is only one population within which ideas spread. The case in which ideas are more likely to spread within particular exogenously or endogenously identified populations of "people like me" is of great interest. The model of Friedman [1998] in which players are sometimes matched with opponents from the same population and sometimes with opponents from a different population provides a natural setting for this type of study, but it is beyond the scope of this paper.

A prototypical example of the type of environment we are studying is a *matching game*. A matching game is defined by a utility function $\tilde{u}^j(a^1, a^2, \dots, a^J)$ where $j = 1, 2, \dots, J$ are player roles, and $a^j \in A$, a finite set, is called an action. Players are randomly assigned to different roles. Strategies are maps from roles to actions $s : \{1, 2, \dots, J\} \to A$. The function $u(s, \sigma)$ is computed by calculating the probability of playing different roles, and meeting opponents playing different roles.

$$u(s,\sigma) = (1/J)\sum_{j=1}^{J}\sum_{\tilde{s}\in S^J|\tilde{s}^j=s} \tilde{u}^j(\tilde{s}^1(1), \tilde{s}^2(2), \dots, \tilde{s}^J(J))\prod_{k\neq j}\sigma(\tilde{s}^k)$$

We make the fairly standard, and in a large population relatively innocuous, simplification that a player does not take into account the fact that he cannot meet himself. Notice that only in the case of two player games is $u(s,\sigma)$ linear in σ, which is the most familiar case. Also of interest are *anonymous matching games* which are matching games in which strategies are restricted to be independent of the player role, so that $S = A$.

3. Basic Results

We begin by establishing some basic results. First, we examine the relative best-response dynamic in the unperturbed case $\varepsilon = 0$. This dynamic is similar in some respects to the replicator: for example, all pure profiles are steady states, but points that are not Nash equilibria are not locally stable. Second, we establish two basic results for the perturbed case $\varepsilon > 0$. When imitation is much more likely than innovation, mixed strategies should be less stable than pure strategies. A mixed strategy can evolve into a pure strategy purely through imitation, while a pure strategy cannot evolve at all without at least one innovation. We confirm this intuition by showing that the limit invariant distribution μ places weight only on pure profiles in Σ^m. We then further study the connection between the limit invariant distribution μ and Nash equilibrium, showing that if the support of μ is a singleton, it must be a Nash equilibrium.

Let μ^0 be an irreducible invariant distribution of the Markov process in which $\varepsilon = 0$. Let ω be the set of mixed strategies in the state space Σ^m that this invariant distribution gives positive weight to. We call such an ω an *ergodic set*. Let Ω be the set of all such ω. Note that this is a set of sets. Let $S(\sigma)$ denote the set of pure strategies used with positive probability in σ. First we establish some basic facts about Ω.

Lemma 3.1: *The sets ω are disjoint. Each set consisting of a singleton pure profile $\{s\} \in \Omega$. If $\sigma, \sigma' \in \omega \in \Omega$ then $S(\sigma) = S(\sigma')$.*

Proof: When $\varepsilon = 0$ we have the *relative best-response dynamic* in which in each player one player switches with equal probability to one of the relative best-responses to the current state. The sets ω are by definition minimal invariant sets under the relative best-response dynamic. That these sets are disjoint is immediate from the definition. Pure profiles are absorbing since no strategy can be used unless it is already in use. This means that every set ω consisting of a single pure strategy is in Ω. To see that have $S(\sigma) = S(\sigma')$, observe that the relative best-response dynamic cannot ever increase the set of strategies in use. If there is a point $s \in S(\sigma), s \notin S(\sigma')$ then the probability that the best-response dynamic goes from σ to σ' is zero, which is inconsistent with the two strategies lying in the same ergodic set.

☑

The third part of the Lemma means that for each $\omega \in \Omega$ we may assign a unique set of pure strategies $S(\omega)$ corresponding to $S(\sigma), \sigma \in \omega$.

The relative best-response dynamic is similar to the better-known replicator dynamic in several respects. Like the replicator dynamic, the relative best-response dynamic is absorbed by pure strategies and so has many steady states. However, as is the case with the replicator, this is offset somewhat by the fact that points that do not correspond to Nash equilibria are locally unstable. By locally unstable, we mean that there is a neighborhood of the state and a change in strategy by a single player that leads to a positive probability of exiting that neighborhood.

Theorem 3.2: *Suppose* σ *is such that for some* $\tilde{s}$ *with* $\sigma(\tilde{s}) > 0$ *we have* $u(s,\sigma) > u(\tilde{s},\sigma)$. *Then for all sufficiently large* n, *if* $\sigma \in \omega$ *it is locally unstable.*

Proof: Since u is continuous, there is a neighborhood of σ in which $u(s,\sigma) > u(\tilde{s},\sigma)$. Also, since $\sigma(\tilde{s}) > 0$ we may assume that this is also the case in this neighborhood. For n sufficiently large the neighborhood must contain points in which one player is playing s. At each such point, the relative best-response assigns positive probability to the number of players playing $\tilde{s}$ decreasing by one, so there is positive probability of exiting the neighborhood.

☑

Turning to the case $\varepsilon > 0$, from a theorem of Young [1993] μ may be described as a probability distribution over Ω. Our intuition that pure strategies are more important in a setting of innovation is confirmed by our first theorem.

Theorem 3.3: *With unlikely innovation the limit invariant distribution* μ *puts weight only on the sets* $\omega \in \Omega$ *that consist of a single pure strategy*.[3]

To prove this theorem, and our additional results, we will use the characterization of μ given by Young [1993].[4] Let τ be a tree whose nodes are the set Ω. We denote by $\tau(\omega)$ the unique predecessor of ω. An ω-tree is a tree whose root is ω. For any two points $\omega, \tilde{\omega} \in \Omega$ we define the resistance $r(\omega, \tilde{\omega})$ as follows. First, a path from ω to $\tilde{\omega}$

---

[3] A similar result has been shown for imitation in the form of genetic algorithms. See Dawid [1999] Theorem 4.2.1, for example.

[4] Although the standard convention in game theory is that a tree begins at the root, Young [1993] followed the mathematical convention that it ends there. We have used the usual game-theoretic convention, so our trees go the opposite direction of Young's.

is a sequence of points $(\sigma_0,\dots,\sigma_K) \subset \Sigma^m$ with $\sigma_0 \in \omega$, $\sigma_K \in \tilde{\omega}$ and σ_{k+1} reachable from σ_k by a single player changing strategy. If the change from σ_k to σ_{k+1} is a relative best-response, the resistance of σ_k is 0; if the change is an imitation the resistance is 1; if the change is an innovation the resistance is m. The resistance of a path is the sum of the resistance of each point in the sequence. The resistance $r(\omega,\tilde{\omega})$ is the least resistance of any path from ω to $\tilde{\omega}$. The resistance $r(\tau)$ of the ω-tree τ is the sum over non-root nodes of $r(\tilde{\omega},\tau(\tilde{\omega}))$. The resistance of ω, $r(\omega)$ is the least resistance of any ω-tree. The following Theorem is proven in Young [1993].

Young's Theorem: $\mu(\omega) > 0$ *if and only if*

$$r(\omega) = \min_{\tilde{\omega}\in\Omega} r(\tilde{\omega})$$

Remark: The set of ω for which $\mu(\omega) > 0$ is called the *stochastically stable set*.

The basic tool for analyzing μ is tree surgery, by which we transform one tree into another and compare the resistances of the two trees. Suppose that τ is an ω-tree. For any nodes $\tilde{\omega} \neq \omega$ we *cut* the $\tilde{\omega}$-subtree separating the original tree into two trees; one the $\tilde{\omega}$-subtree and the other what is left over. This reduces the resistance by $r(\tilde{\omega},\tau(\tilde{\omega}))$. If $\widehat{\omega}$ is a node in either of the two trees, and $\hat{\omega}$ is the root of the other tree, we may *paste* $\hat{\omega}$ to $\widehat{\omega}$ by defining $\tau(\hat{\omega}) = \widehat{\omega}$. This tree has the root of the tree containing $\widehat{\omega}$. The paste operation increases the resistance by $r(\hat{\omega},\widehat{\omega})$, so the new tree has resistance $r(\tau) + r(\hat{\omega},\widehat{\omega}) - r(\tilde{\omega},\tau(\tilde{\omega}))$. These operations can be used to characterize classes of least resistance trees, by showing certain operation do not increase the resistance. They can also be used as below in proof by contradiction, showing that certain trees cannot be least resistance because it is possible to cut and paste in such a way that the resistance is reduced.

Proof of Theorem 3.2: Suppose that $\mu(\omega) > 0$ and that ω is not a singleton pure profile. Let τ be a least resistance ω-tree. Let $\tilde{\omega} = s$ be a singleton pure strategy that is played with positive probability by some $\sigma \in \omega$, that is, $s \in S(\omega)$. Cutting $\tilde{\omega}$ and pasting the root ω to it. Since $\tilde{\omega}$ is a singleton pure profile, it requires at least one innovation to go anywhere, so cutting reduces the resistance by at least n. On the other hand, since $\sigma \in \omega$ and $\sigma(\tilde{\omega}) > 0$, we can go from ω to $\tilde{\omega}$ by no more than m imitations, pasting the root to $\tilde{\omega}$ increases the resistance by at most m. By the assumption of unlikely innovation, this implies that the new tree has strictly less resistance than the old contradicting Young's Theorem.

☑

4. Winning Pairwise Contests is Sufficient

We now establish our first main result: we show that if a pure strategy beats all others in pairwise contests, then it is the unique stochastically stable state. We begin by explaining what it means to win pairwise contests.

Definition 4.1: *Suppose that* $s,\tilde{s} \in S$ *. For* $0 \le x \le 1$ *define a family of mixed strategies* $\sigma(x)$ *by* $\sigma(x)[s] = x, \sigma(x)[\tilde{s}] = 1 - x$ *. If for all* $x \ge 1/2$

$$u(s,\sigma(x)) - u(\tilde{s},\sigma(x)) > 0$$

we say that s **beats** $\tilde{s}$ *.*

Theorem 4.1: *Suppose unlikely innovation, sufficiently large* n *and that* s *beats all* $\tilde{s} \ne s$ *. Then* $\mu(\{s\}) = 1$ *.*

Proof: Suppose that there is other some ω with $\mu(\omega) > 0$. By Theorem 3.2 $\omega = \{\hat{s}\}$ for some pure strategy $\hat{s}$. Let τ be the least resistance ω -tree. Since it is not the root, we may suppose that $\{s\}$ is attached to some $\tilde{\omega}$, and consider cutting it and pasting the root to it. It took at least one innovation plus, since s beats any point in $\tilde{\omega}$, more than $m/2$ imitations to get to $\tilde{\omega}$, so the resistance is reduced by strictly more than $n + m/2$. However, since s beats $\hat{s}$ we can get from $\omega = \{\hat{s}\}$ to $\{s\}$ with one innovation and no more than $m/2$ imitations. So resistance is strictly reduced contradicting Young's Theorem.

☑

The hypothesis, that when half or more of the population is playing s against any other pure profile, all players prefer to play s is connected to the idea of ½-dominance introduced by Morris, Rob and Shin [1993]. The concept of ½-dominance is that when half or more of the population is playing s against any other combination of strategies, it is a best response to play s. The concept here is weaker in two respects: first, s must only beat pure profiles, not arbitrary combinations of strategies. Second, s must win only in the sense of being a relative best-response, it need not actually be a best-response;

a third strategy may actually do better than s, and this is significant as we will see in examples below. On the other hand, ½-dominance clearly implies winning all pairwise contests, so if there is a ½-dominant strategy, from Morris, Rob and Shin [1993] it is stochastically stable with respect to the usual evolutionary dynamic, and it is also stochastically stable when innovation is unlikely. Examples below will show clearly how the usual notion of stochastic stability and the case of unlikely innovation diverge when ½-dominance fails, as well illustrating that in interesting games there can be a strategy that wins all pairwise contests, even though there is no ½-dominant strategy.

Interestingly, Kandori and Rob [1993] study a class of games in which winning all pairwise contests implies ½-dominance. They study single population matching games satisfying the "total bandwagon property" meaning that the best response to any mixed strategy is contained in the support of that mixed strategy. In particular, this means that any pure profile is a Nash equilibrium. They make several other assumptions as well, but as Fudenberg and Levine [1998] point out, these other assumptions are redundant. In a game satisfying the "total bandwagon property" if s wins all pairwise contests, it must be actually be a best-response against every other pure strategy when ½ the population is playing s. Since utility is linear in the population distribution of strategies, this means that s is actually best response against any combination of strategies when ½ the population is playing s, so in fact s is in fact ½-dominant.

Example 4.1: A Specialization Game

We now study a simple game with a unique equilibrium that is mixed. This illustrates that unlikely innovation can be quite different than unlikely imitation. Consider a simple 2×2 symmetric game of specialization: players may specialize in being hunters or gatherers. If both choose the same specialization they consume only one product, resulting in a utility of zero. If they choose different specializations they trade, consume both products, and get a utility of one. The payoff matrix is

	Hunt	**Gather**
Hunt	0,0	1,1
Gather	1,1	0,0

We first assume that this is played as an anonymous matching game. This means the only pure strategies are Hunt and Gather. From symmetry it is obvious both must have equal weight (of ½) in the limit distribution with unlikely innovation. This is very different than the case the standard case: the mixed equilibrium is the unique Nash equilibrium and it is the unique point in Ω since players prefer to do the opposite of what everyone else is doing. This means that it takes one innovation to get to the basin of the mixed equilibrium, while it takes half the population to innovate to get out of the basin. Notice that even with unlikely imitation, we continue to assume the relative best-response dynamic, but this makes little difference as it takes only a single innovation to get out of the non-Nash points in Ω.

We should digress briefly to discuss how mixed strategy Nash equilibria appear in Ω. If the mixed strategy is actually on the grid, then it will be in ω, but this is unlikely. If the mixed strategy is not on the grid, there are two possibilities depending on whether it is stable in the relative best-response dynamic or not. If it is not, then there will not be any point in Ω corresponding the mixed equilibrium. If it is stable, then there will be a small cycle around the mixed equilibrium that will be in Ω – it is this we refer to as the mixed strategy, although strictly speaking it is not.

The hunter-gatherer example reflects an aspect of unlikely innovation that should be disturbing. There can be no stochastically stable mixed equilibria. In this example the only Nash equilibrium is mixed, and the result is that the stochastically stable set is not Nash and gives players less than the minmax as well. However, implicit in this formulation is that mixing takes places by accident, through half of the population doing one thing and half something else. It is fairly well known in the learning literature, for example from Fudenberg and Kreps [1993], that this can be problematic. However, we can also consider mixing through explicit randomization: that is, introduce a mixed strategy as an explicit pure strategy. Suppose that we do this in the example: we add a strategy of randomizing 50-50 between Hunt and Gather. It is apparent that if half the population is following this strategy and half is playing a pure strategy (Hunt, for example) it is better to mix. So the new mixing strategy wins all pairwise contests and is the unique stochastically stable state.

Once we admit the possibility of explicit mixing, however, there is an even better idea than can make its way into the population: using a correlating, or identifying device.

That is, an anonymous game restricts players to actions that are independent of their player roles (which might correspond, for example, to man and woman). The fact is that while in the laboratory it is possible to create anonymous matching games, it is not terribly likely to happen in the field. If the game is played as a non-anonymous matching game, then there are two additional strategies of Hunt when player 1 and Gather when player 2, and vice versa. Either of these strategies is Pareto efficient, and seem to reflect historical patterns of specialization (men hunt, women gather).

Because neither of the two new strategies beats the other, it is useful to consider an extension of the main theorem that gives a sufficient condition for a set of pure strategies to be the only ones receiving positive weight in the limit distribution.

Definition 4.2: *A set of pure strategies* $\hat{S}$ **beats the field** *if each strategy* $s \in \hat{S}$ *beats all* $\tilde{s} \notin \hat{S}$.

The next theorem is specialization of Corollary C.2 in Appendix C.

Theorem 4.2: *Suppose unlikely innovation, sufficiently large* m *and that* $\hat{S}$ *beats the field. Then* $\sum_{s \in \hat{S}} \mu(\{s\}) = 1$.

Proof: See Appendix C, Lemma C.2 (i).

☑

Turning back to the example, we see that Hunt/Gather and Gather/Hunt are the stochastically stable set; it is clear from symmetry that they are equally likely.

Comparison to the standard case is simplified if we break part of the symmetry by supposing that Gather/Gather is safer, and therefore a little better than Hunt/Hunt, so that the payoffs are

	Hunt	**Gather**
Hunt	0,0	1,1
Gather	1,1	1/4,1/4

This does not change the analysis in the case of unlikely innovation. However, in the standard analysis ½ dominance fails. To see this observe that if half the population is playing Hunter/Gather and half is playing Gather/Hunt the best response is Hunt/Hunt, that is neither of the two. Despite this, we can use Ellison [1995]'s methods to show that

Hunt/Gather and Gather/Hunt are the unique stochastically stable set in the standard model. There are many mixed equilibria of this model, but adding one Hunt/Gather innovation to one of these equilibria causes Hunt/Gather to become the best response and vice versa, while it takes quite a few innovations to get out of Hunt/Gather or Gather/Hunt.

5. A Gift Giving Game

We now consider a more extended example. The setting is a gift-giving game introduced by Johnson, Levine and Pesendorfer [1999] to study the evolution of the social norm of cooperation. In this game players live two periods in overlapping generations. Young players are randomly matched against an equal number of old players and must choose whether to give or withhold a gift from the old opponent. Old players are passive and do not have an action. However, the behavior when young is reported on by *information systems* so they may be reward or punished based on what they did when they were young. This is a variation on the model of Kandori [1992], and following Kandori, it is possible to prove a folk theorem for this model if there is "enough" information about old players.

Specifically, we assume that it costs the young player 1 unit of utility for giving a gift, and provides a benefit of $\alpha > 1$ to the older recipient. Payoffs are additive between the two periods of life so that gift-giving is efficient. Note the resemblance of the model to Prisoners' Dilemma. The myopic optimum for the young player is to withhold the gift, just as defection is dominant when the Prisoners' Dilemma is played once. However, the overlapping generations environment allows for a connection between actions when young and payoffs when old, just as repetitions would allow for consequences in later periods to influence earlier actions in a repeated Prisoners' Dilemma.

They key assumption is that enables cooperative play is that young players are (partially) informed about the history of their older opponent. Following Kandori [1992] we model this through *information systems*. An information system provides a signal about past play. We examine the simplest case in which this signal can take on two values, which we describe as a "red flag" or a "green flag." Let $\{r, g\}$ be the set of flags. Formally, an *information system* is a map $i : \{0,1\} \times \{r, g\} \to \{r, g\}$ that assigns an old player a flag based on his own action (the size of the gift – 0 or 1) and the flag of the opponent he met when young. It is easily checked that there are 16 information systems. We denote the set of information systems by I and f_i for the flag corresponding to information service *i*.

Flag vectors f corresponding to the different information systems are observed by young opponents as follows. With probability $1-\eta$ the flag vector observed is equal to the vector assigned by the different information systems. With probability $\eta > 0$ the flag vector observed is chosen randomly according to a uniform distribution on F. We assume that the chance a player is assigned a random flag vector is small: specifically that $\beta := \alpha(1-\eta) > 1$.

Finally, we assume that a young player may consult only one information system. This means that the only feasible strategies are those of the form $s = (a,i)$ consists of the choice of *one* information system $i \in I$ and a map $a:\{r,g\} \to \{0,1\}$ that assigns an action to a flag.

Finally, we must specify how expected utility $u(s,\sigma)$ is determined. Utility to a young player depends on expectations of future play by next period young players σ and on the distribution of flags among current old players. As usual, we take the previous period distribution of strategies σ_{t-1} as a proxy for beliefs about expectations of next period play. We further assume that the distribution of current old player flags is believed to be the steady state distribution[5,6] of flags corresponding to σ_{t-1} being played repeatedly. Since all flag vectors have positive probability, this steady state is unique and we denote the corresponding marginal probability distribution over F by $\phi(\sigma_{t-1})$. Expected utility is then calculated with respect to ϕ.

We are going to focus on two key types of strategies. One is the *always selfish* strategy of making a low transfer regardless of the old opponent's flag vector. The second type of strategy we consider is exemplified by the *green-team* strategy. This uses an information system that assigns a green flag to high transfer against a green flag and low

[5] This long run view of the flag distribution may seem inconsistent, but the overlapping generations structure does not mean that each player only plays once, merely that information about play only persists for a single period.

[6] If beliefs about flag distributions are noisy this will effect results. However, as theorists we are looking for simple cases to build qualitative intuition – numerical analysis requires simulations. Our focus here is on role of imitation in propagation so we do not choose to let it compete with other sources of randomness. The assumption that flag distribution are steady state are only one of many such simplifying assumptions standard in theoretical evolutionary analysis: we do not allow for noise in computing relative best responses, in sampling strategies from the population or in payoffs.

transfer against a red flag, and a red flag to low transfer against green flag and high transfer against a red flag. The strategy itself is to give high transfer on green flag, low transfer on red flag. That is a green flag means the player is a member of the team; team members are supposed to give high transfers to team members and low transfers to non-team members. Behaving as a team member is the ticket for admission to the team; the penalty for failing to behave as a team member is expulsion from the team. The other strategy in this class is *red-team* which uses the same information system, but uses the convention of high transfer on red and low transfer on green.

Theorem 5.1: *If $\beta < 2$ and η sufficiently small the unique stochastically stable state is always selfish; if $\beta > 2$ the unique stochastically stable state places weight ½ on each green-team and red-team.*

Proof: See Appendix A.

☑

Notice that this is very sharp: for $\beta < 2$ we get always selfish; for $\beta > 2$ we get the two team strategies as the stochastically stable set. This can be compared to Johnson, Levine and Pesendorfer [1999] who examine the standard model under slightly different assumptions and with elaborate calculations are able to get the team strategies emerging as stochastically stable only for $\beta > 4$.

6. Characterization of Stochastically Stable States

We first describe Ω in greater detail as a preliminary to defining the modified resistance, which can be used to give an exact characterization of stochastically stable state.

Denote by Ω_k the subset of $\omega \in \Omega$ for which $\# S(\omega) \le k$. Most of the "action" takes place in Ω_2 in the sense that it is unlikely that there are more than two pure strategies being used in the population at any moment of time. The space Ω_2 has a relatively simple topological structure, consisting of pure strategies and the segments connecting them. This has important consequences:

Lemma 6.2: *If $\omega \in \Omega_2$, $\sigma, \sigma'' \in \omega$, $S(\sigma') = S(\sigma)$ with $\sigma(s) > \sigma(s') > \sigma(s'') > 0$ then $\sigma'' \in \omega$.*

Proof: If $\sigma, \sigma'' \in \omega$ there must be a relative best response path, or equivalently a path of zero resistance between the two. Such a path cannot pass through Ω_1, so it must pass through σ'. This means that $\sigma' \in \omega$ as well.

☑

We illustrate Ω in the figure below. Here there are 4 points $\{s^A, s^B, s^C, s^D\}$ in $\Omega_1 = S$. The dashes connecting those pure strategies represent points on the grid Σ^m; that is, they are mixtures between exactly two pure strategies. (Note that the drawing is not exact, since there should be the same number of points $m-1$ between each pair of pure strategies.) There are three points $\{\omega^x, \omega^y, \omega^z\}$ in $\Omega_2 \setminus \Omega_1$. Each one is a set of points in Σ_2^m. The point ω^z has $S(\omega^z) = \{s^B, s^D\}$, and consists of two points in Σ_2^m. The points ω^x, ω^y each have $S(\omega^x, \omega^y) = \{s^A, s^D\}$. The point ω^x consists of two points in Σ_2^m; the point ω^y a singleton in Σ_2^m.

The case of ω^x and ω^z are in a sense typical of points in $\Omega_2 \setminus \Omega_1$ in consisting of two points in Σ_2^m. Two see this, let us refer to a mixed strategy equilibrium in a game using a subset consisting of two pure strategies only as a *relative mixed equilibrium*. Recall that any point in Ω consists of points in Σ_2^m that are minimal invariant under the relative best-response dynamic. Consider $\omega^z = \{\sigma^z, \sigma^Z\}$. Ruling out ties, and remembering that only one player moves at a time, this means that σ^Z is a relative best response to σ^z and vice versa. Consequently, ω^z is a two-cycle. Because only one player moves at a time, there cannot be cycles of length longer than 2 in $\Omega_2 \setminus \Omega_1$. Notice that such a two-cycle arises when σ^z, σ^Z are separated by a non-grid point that is a relative mixed strategy equilibrium stable in the continuous time best response dynamic. Of course if such a relative mixed equilibrium was actually on a grid point, then that point would as a singleton set be in $\Omega_2 \setminus \Omega_1$ such as ω . However, it is unlikely that a relative mixed equilibrium will actually lie on a grid point. If a relative mixed equilibrium is unstable in the continuous time best response dynamic, then there will be no corresponding point in Ω_2.

Suppose that $\omega, \tilde{\omega} \in \Omega_2$. We say that they are *adjacent* if it possible to go from one to the other without passing through another point in Ω_2. More precisely, $\omega, \tilde{\omega} \in \Omega_2$ are adjacent if there is a path $(\sigma_0, \ldots, \sigma_K) \subset \Sigma_2^m$ from ω to $\tilde{\omega}$ such that for all $\hat{\omega} \in \Omega_2$ and $0 < k < K$, $\sigma_k \notin \hat{\omega}$. For example, in the figure, ω , ω are adjacent because $\sigma_0, \ldots, \sigma_4$ is a path between them that does not pass through any $\omega \in \Omega$.

Notice that every point in $\Omega \quad \setminus \Omega$ is adjacent to exactly two other points in Ω , while points in Ω_1 are adjacent to $\# S - 1$ points (3 in the example).

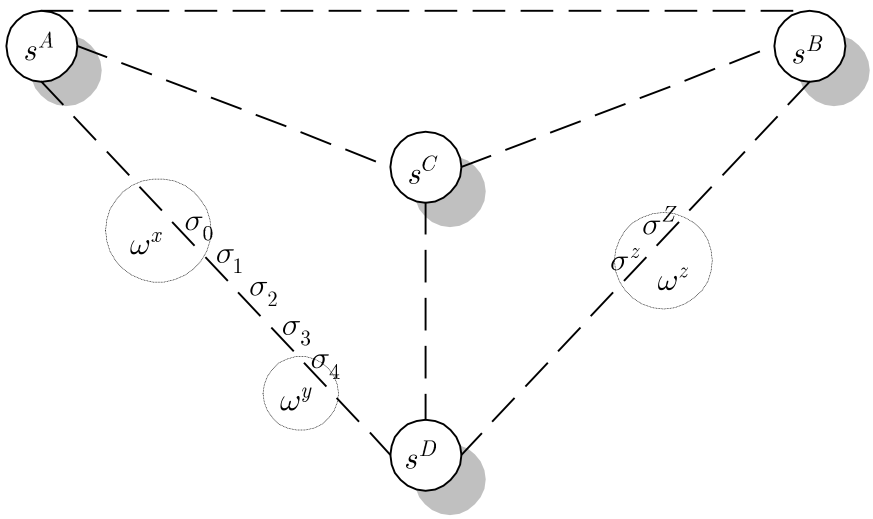

For any two points $\omega, \tilde{\omega} \in \Omega_1$, there is a unique sequence $\omega_0, \ldots, \omega_K \in \Omega_2$ such that $\omega_0 = \omega$, $\omega_K \in \tilde{\omega}$ and ω_k, ω_{k+1} are adjacent. For each k define

$$\underline{r}(\omega, \tilde{\omega}, k) = \sum_{j=1}^{k} r(\omega_j, \omega_{j-1}) + \sum_{j=k+1}^{K-1} r(\omega_j, \omega_{j+1})$$

and define $\underline{r}(\omega, \tilde{\omega}) = \min_k \underline{r}(\omega, \tilde{\omega}, k)$. Note that $\underline{r}(\omega, \tilde{\omega}) = \underline{r}(\tilde{\omega}, \omega)$. The *modified resistance* is $\hat{r}(\omega, \tilde{\omega}) = r(\omega, \tilde{\omega}) - \underline{r}(\omega, \tilde{\omega})$

Theorem 6.2: *With unlikely innovation* $\mu(\omega) > 0$ *if and only if* ω *is the root of a tree with least modified resistance in* Ω_1 .

Proof: In Appendix B.

☑

7. The Continuous Approximation

In many applications it is useful to think of m as being large, as round-off error on the grid is not usually of much economic interest. This leads us to define cost as the limit of per capita modified resistance. Most results about modified resistance then have an analog in terms of cost.

Suppose that $s, \tilde{s} \in S$. As we did earlier, for $0 \le x \le 1$ define a family of mixed strategies $\sigma(x)$ by $\sigma(x)[s] = x, \sigma(x)[\tilde{s}] = 1 - x$. Consider the utility difference

$$U(x) = u(s, \sigma(x)) - u(\tilde{s}, (\sigma(x)) .$$

Define

$$c(s, \tilde{s}) \equiv \int_0^1 1(U(x) < 0)dx - \min_y \left\{ \int_0^y 1(U(x) < 0)dx + \int_y^1 1(U(x) > 0)dx \right\}.$$

For any pure profile s let $T(s)$ be all trees on S with root s .

$$\hat{c}(s) = \min_{\tau \in T(s)} \sum\nolimits_{\tilde{s}} c(\tilde{s}, \tau(\tilde{s})) .$$

We first establish the relationship between modified resistance and cost.

Theorem 7.1: $\lim_{m \to \infty} \underline{r}(s, \tilde{s}) / m = c(s, \tilde{s})$.

Proof: Follows from integration theory and the definitions.

☑

The most useful application of Theorem 7.1 is through an obvious corollary.

Corollary 7.2: *If* $c(\tau) > c(\tilde{\tau})$ *then for all* m *sufficiently large,* $\hat{r}(\tau) > \hat{r}(\tilde{\tau})$.

This corollary means that for m large the least modified resistance trees are a subset of the least cost trees. In particular, least cost is a necessary (although no longer sufficient) condition to be stochastically stable. From Theorem 6.1, we have

Corollary 7.3: *Suppose unlikely innovation and sufficiently large* m *. Then* $\mu(\{s\}) > 0$ *only if*

$$\hat{c}(s) = \min_{\tilde{s} \in S} \hat{c}(\tilde{s}) .$$

Finally, because zero cost connection play an important role in the theory, it is important that zero cost implies zero modified resistance for all sufficiently large m .

Theorem 7.4: *If* $c(s, \tilde{s}) = 0$ *then for all* m *sufficiently large* $\underline{r}(s, \tilde{s}) = 0$.

Proof: To be shown.

☑

8. Further Results

We now give a more detailed characterization of stochastically stable sets.

Definition 8.1: s' *is a* **strongly better response** *to* s *if* $c(s, s') = 0$

If a profile admits no strongly better response, we call it a generalized strict Nash equilibrium.

Definition 8.2: s *is a* **generalized strict Nash equilibrium** *if* $c(s,s') > 0$ *for all* $s' \neq s$.

Since a strategy that does strictly worse cannot be a strongly better response, every strict Nash equilibrium is a generalized strict Nash equilibrium.

Definition 8.3: Σ *is a* **strongly better response cycle** *if*

1) $s \in \Sigma$ *and* s' *a strongly better response to* s *implies* $s' \in \Sigma$

2) $s, s' \in \Sigma$ *implies a sequence* $(s_0, s_1, \ldots, s_K)$ *such that* $s_0 = s, s_K = s'$ *and* s_{k+1} *is a strongly better response to* s_k

Theorem 8.1: *If* s *is the unique long-run stationary state then* s *is a generalized strict Nash equilibrium*

Proof: If not and s is the root of a least cost tree, find $c(s,s') = 0$ and paste it above the root.

☑

Remark: The sum of costs in both directions is less than or equal one.

Theorem 8.2: *In a least cost tree* $c(s, \tau(s)) \leq 1/2$

Proof: If not cut s. Pasting either above or below the root costs less than or equal ½.

☑

Denote the set of pairs with costs strictly between zero and one

$$P = \{(s,s') \mid 1 > c(s,s') > 0\}.$$

Assumption 8.1 [Generic Assumption]: *If* $P_1, P_2 \subseteq P, P_1 \neq P_2$ *then*

$$\sum\nolimits_{p \in P_1} c(p) \neq \sum\nolimits_{p \in P_2} c(p)$$

Theorem 8.3: *Under the assumption 8.1 if* $\Omega^* = \{s \mid \mu(s) > 0\}$ *is not a singleton, it is a strongly better response cycle.*

Proof: If s is the root of a least cost tree, then so are all strongly better responses, since we may cut them and paste above the root without increasing the cost. Suppose that s' is the root of a least cost tree and that there is no zero-cost path from s' to s. Let s'' be the first predecessor of s' in the lowest cost s-tree that is connected with a non-zero cost. Then there is a least cost tree with root s''. But this tree does not have the cost $c(s'',\tau_s(s''))$, so the tree has a different cost then the s-tree by the generic assumption.

☑

Remark: It should be possible without the generic assumption to show that we get at most a collection of generalized Nash equilibria and strongly better response cycles.

Remark: An implication of Theorem 8.3 is that under the generic Assumption 8.1 either there is a generalized strict Nash equilibrium or a strongly better response cycle. We can show directly that this is true in every game. For any point s define the successor set $\vec{S}(s)$ to be the set of points reachable by a sequence of zero cost moves. Suppose there is no generalized strict Nash equilibrium. Then every strategy has a strongly better response. So there exists a point $s \in \vec{S}(s)$. We say that $\vec{S}(s)$ is minimal if there does not exist a point $s' \in \vec{S}(s')$ with $\vec{S}(s') \subset \vec{S}(s)$. Since S is finite there exists a minimal set. We claim that such a minimal set is a strongly better response cycle. Already $\vec{S}(s)$ is closed under the strongly better response operation since by construction it contains all successor points. Let $s' \in \vec{S}(s)$. We must show $s \in \vec{S}(s')$. Suppose not. It may not be the case that $s' \in \vec{S}(s')$, but because there is no strong better response cycle, there is a point $s'' \in \vec{S}(s')$ with $s'' \in \vec{S}(s'')$. Since $s'' \in \vec{S}(s')$ it must be that $\vec{S}(s'') \subseteq \vec{S}(s')$ and by hypothesis $s \notin \vec{S}(s')$ so $s \in \vec{S}(s'')$. But $s' \in \vec{S}(s), s'' \in \vec{S}(s')$ so $s'' \in \vec{S}(s)$ implying $\vec{S}(s'') \subseteq \vec{S}(s)$ and we have contradicted the assume minimality of $\vec{S}(s)$.

Three Strategy Games

We now characterize generic pairwise matching games in which $\#S = 3$. We make the generic Assumption 8.1, and the additional assumption

Assumption 8.2: *For all* $s \neq s'$, $c(s,s') \neq 1/2$. *If* $c(s,s') = 0$ *then* $c(s',s) > 0$.

Case 1: There is no generalized strict Nash equilibrium. Then there is a unique strongly better response cycle consisting of all three strategies, which is the stochastically stable set.

Case 2: There is one generalized strict Nash equilibrium. There is no strongly better response cycle, so generalized strict Nash equilibrium is the stochastically stable set.

Case 3: There are two generalized strict Nash equilibria. Attach the third strategy with zero cost to one of the two equilibria, then we need only worry about how to attach the equilibria; the equilibrium with least cost of getting to the other is the unique point in the stochastically stable set.

Case 4: There are three generalized strict Nash equilibria.
Case 4a: One generalized strict Nash equilibrium beats the field. It is the unique point in the stochastically stable set.
Case 4b: Two generalized strict Nash equilibria beat the field. This reduces to a contest between the two; the one with the least cost of getting to the other is the unique point in the stochastically stable set.
Case 4c: No generalized strict Nash equilibrium beats the field. Consequently each has a unique successor with cost less than ½. This means the only possible least cost trees are linear. The least cost linear tree has at the top the strategy with the greatest cost less than ½ , and this is the unique point in the stochastically stable set.

Notice that the only case where the stochastically stable set is not a singleton is case 1.

Appendix A: Gift-Giving Game

Theorem 5.1: *If* $\beta < 2$ *and* η *sufficiently small the unique stochastically stable state is always selfish; if* $\beta > 2$ *the unique stochastically stable state places weight ½ on each green-team and red-team.*

Proof: If half the population is playing always selfish, high transfer when young can yield an expected gain when old of at most $\beta/2$, while costing 1 today. So if $\beta < 2$, always selfish wins all pairwise contests and is the unique stochastically stable state. This is despite the fact that high transfer by all players Pareto dominates low transfer by all players.

Clearly if green-team is stochastically stable, so is red-team and vice versa. We will show that when $\beta > 2$ the two strategies together beat the field, so, according to Theorem 4.2, they constitute the unique stochastically stable set. Because of symmetry, it suffices to consider just green-team.

Assume $\beta > 2$. Notice first, that green-team beats always selfish. To see this, suppose that half or more of the population is playing green-team, and the rest always selfish. If the old opponent has a red flag, both strategies give a low transfer, and the same expected utility. However, since there is flag noise, the steady state flag distribution has some green flags, and in this case green-team and always selfish behave differently. Always selfish today gives a benefit of 1 today, but green-team yields a benefit of at least $\beta/2$ tomorrow. This follows because with probability $1-\eta$ there is no flag noise, and green-team will receive a high transfer from green team members, while always selfish will not receive any transfer. If there is flag noise, the two strategies get the same expected utility tomorrow. So under the assumption that $\beta > 2$ the payoff to green-team is strictly higher than that of always selfish. By a similar argument, green-team beats never selfish, the strategy of high transfer regardless of flag vectors.

Consequently, we can restrict attention to opposing strategies that depend on the flag from some information system. Without loss of generality, since the actual names of the flags do not matter, we may examine the strategies that give high transfer on green flag, low transfer on red flag.

There are four flag combinations of the form (green-team-flag, non-green-team-flag). Let $\phi \le 1/2$ be fraction of population playing non-green-team. In all cases we can ignore the second period payoffs following a randomly assigned flag. We have the following cases

(g,g): green-team and non-green-team play the same way, so give the same payoff.

(r,r): green-team and non-green-team play the same way, so give the same payoff.

(r,g): green-team gets at least $1+(1-\phi)\beta$; non-green-team gets either $\phi\beta$ or 0 depending on whether it rewards itself for high transfer. Since $\phi \le 1/2$, the advantage to green-team is at least 1.

(g,r) if non-green-team is to beat green-team it has to win this one, since it ties or loses all other combinations, and combinations have positive probability.

Focusing on the (g,r) combination, green-team gets at least $(1-\phi)\beta$. Non-green-team gets either $1+\phi\beta$ or 1 depending on whether it rewards itself for high transfer. If it gets zero, it loses since $\beta > 2$. So it must assign low transfer on red a green flag. Moreover, it cannot assign high transfer on red a green flag, since then it gives green-team an additional $\phi\beta$, and again loses to green-team. So we may assume that non-green-team on a red flag gives green for selfish, red for altruistic.

Suppose in fact non-green-team on a red flag gives green for low transfer, red for high transfer. So green-team is winning at (r,g) and non-green-team at (g,r). Moreover, non-green-team cannot reliably reward at (r,g) (that is, give green to high transfer on green and red to low transfer on green) since then it actually is the green-team strategy. So at (r,g) green-team actually has an advantage of $1+(1-\phi)\beta$. To summarize the advantage to green-team is:

(r,g) $1+(1-\phi)\beta$

(g,r) $(1-\phi)\beta - 1 - \phi\beta$

Now we need the steady state probabilities. Let θ be steady state probability of (g,r). Then ϕ of the population is non-green team, and a fraction $\phi\theta$ of them meet (g,r), so they give a low transfer and have probability $(1-\eta)$ of winding up in (r,g). So the steady state probability of (r,g) is at least $\phi\theta(1-\eta)+\eta/4$. Consequently, the expected advantage of green-team is at least

$$(\phi\theta(1-\eta)+\eta/4)(1+(1-\phi)\beta)+\theta((1-\phi)\beta-1-\phi\beta) =$$
$$\theta((\phi(1-\eta))-1+\eta/4\theta+[(\phi(1-\eta)+\eta/4\theta)(1-\phi)+(1-2\phi)]\beta)$$

Since $(\phi(1-\eta)+\eta/4\theta)(1-\phi)+(1-2\phi)$ is strictly positive for $0 \le \phi \le 1/2$, this is strictly increasing in β. So it suffices to show that this expression is non-negative for $\beta = 2$. Substituting $\beta = 2$ yields

$$\theta(\phi(1-\eta)(3-2\phi)) + 1 - 4\phi) + (\eta/4)(3-2\phi)$$

This is quadratic in ϕ with second derivative $-2\theta(1-\eta) < 0$, so it suffices to show the expression non-negative for $\phi = 0{,}1/2$. At $\phi = 0$ we have $\theta + 3\eta/4 > 0$. At $\phi = 1/2$ we have $\eta(1/2-\theta)$, so it suffices that at the steady state $\theta \le 1/2$.

It is easy to compute the steady state distribution when $\eta = 0$. Conjecture that $\theta = 1/2$. The only way to get to (r,g) is for a non-green-team to meet (g,r), $pr(r,g) = 1/4$. The only way to get to (r,r) is for a non-green-team to meet (r,g), so $pr(r,r) = 1/8$. The only way to get to (g,g) is for anyone to meet (r,r) so $pr(g,g) = 1/8$. Since that adds up to 1, we are done. Now it is pretty clear that adding noise cannot make θ bigger than ½. We will do the local computation for η small: Let $x(\eta)$ be the steady state probabilities if they are unique. Let e be the vector corresponding to a uniform distribution. Then Taylor's theorem tell us that to first order $x(\eta) = (1-2\eta)x(0) + 2\eta e$; so that

$$\theta = (1-2\eta)(1/2) + 2\eta(1/4) =$$
$$(1-\eta)/2$$

☑

Appendix B: Main Theorem

Theorem 6.2: *With unlikely innovation* $\mu(\omega) > 0$ *if and only if* ω *is the root of a tree with least modified resistance in* Ω_1.

The proof requires several steps. In Lemma B.3, we show that we can characterize $\mu(\{s\}) > 0$ by looking at least resistance trees on Ω_2. We then show in Lemma B.4 that there exist least resistance trees in Ω_2 with $\tau(\omega)$ is adjacent to ω. This enables us to characterize $\mu(\{s\}) > 0$ by looking only at trees on Ω. First we establish a useful technical result.

Lemma B.1: *Suppose that* $\mu(\omega) > 0$, *and that* τ *is a least resistance* ω *-tree. Suppose* $\tilde{\omega} \neq \omega$. *If* $\# S(\tilde{\omega}) = 1$ *then* $r(\tilde{\omega}, \tau(\tilde{\omega})) \leq n + m/2$. *If* $\# S(\tilde{\omega}) = 2$ *then* $r(\tilde{\omega}, \tau(\tilde{\omega})) \leq m$.

Proof: By Theorem 3.3 ω is a singleton pure strategy. Let τ be a least resistance ω -tree.

First, suppose that $\# S(\tilde{\omega}) = 1$. Then if we cut $\tilde{\omega}$ and paste it to the root ω we save $r(\tilde{\omega}, \tau(\tilde{\omega}))$ in the cut. Moreover, either $r(\tilde{\omega}, \omega) \leq n + m/2$ or $r(\omega, \tilde{\omega}) \leq n + m/2$. In the former case paste $\tilde{\omega}$ to the root, in the latter paste the root to it. So the new tree has resistance increased by at most $n + m/2 - r(\tilde{\omega}, \tau(\tilde{\omega}))$. From Young's Theorem, this is non-negative, giving the first result.

Now suppose that $\# S(\tilde{\omega}) = 2$. Suppose in particular that $S(\tilde{\omega}) = \{\tilde{s}, \hat{s}\}$. Observe that either $r(\tilde{\omega}, \tilde{s}) \leq m/2$ or $r(\tilde{\omega}, \hat{s}) \leq m/2$. Without loss of generality suppose the former. As in the previous case, we may cut $\tilde{s}$ and paste either above or below the root. The new tree has resistance increased by at most $m/2$. Now, cut $\tilde{\omega}$ and paste it to $\tilde{s}$. The cutting saves $r(\tilde{\omega}, \tau(\tilde{\omega}))$, while the paste adds at most $m/2$ by the way we chose $\tilde{s}$. So we have increased resistance by at most $m/2 - r(\tilde{\omega}, \tau(\tilde{\omega})) + m/2 = r(\tilde{\omega}, \tau(\tilde{\omega})) - m$. From Young's Theorem, this is non-negative, giving the second result.

☑

In Appendix C we prove

Lemma C.1 [Tree Trimming]: *Let* $\Omega^* \subset \tilde{\Omega} \subseteq \Omega$. *Suppose that for every least resistance tree in* $\tilde{\Omega}$*, there is another least resistance tree* τ *with the same root and* $\tau(\omega^*) \in \Omega^*$ *whenever* $\omega^* \in \Omega^*$. *Then least resistance trees in* $\tilde{\Omega}$ *have roots in* Ω^**, and* ω *is the root of a least resistance tree in* $\tilde{\Omega}$ *if and only if it is the root of a least resistance tree in* Ω^*.

Lemma B.2: *With unlikely innovation* $\mu(\omega) > 0$ *if and only if*

$$r_2(\omega) = \min_{\tilde{\omega} \in \Omega_2} r_2(\tilde{\omega})$$

Proof: If $\# S(\omega) = 1$, by Lemma B.1, $r(\omega, \tau(\omega)) \le n + m/2$. On the other hand, if $\# S(\tau(\omega)) > 2$, then $r(\omega, \tau(\omega)) \ge 2n$. With unlikely innovation, this is a contradiction.

If $\# S(\omega) = 2$, by Lemma B.1, $r(\omega, \tau(\omega)) \le m$. On the other hand, if $\# S(\tau(\omega)) > 2$, then $r(\omega, \tau(\omega)) \ge n$. With unlikely innovation, this is again a contradiction.

We conclude that if $\omega \in \Omega_2$ then $\tau(\omega) \in \Omega_2$. The result then follows directly from Lemma C.1.

☑

Lemma B.3*: With unlikely innovation if* $\mu(\omega) > 0$ *there is a least resistance* ω*-tree* τ *in* Ω_2 *such that* $\tilde{\omega}$ *and* $\tau(\tilde{\omega})$ *are adjacent for all* $\tilde{\omega} \ne \omega$

Proof: Let τ be a least resistance ω-tree. Suppose that $\tilde{\omega}$ and $\tau(\tilde{\omega})$ are not adjacent. Then every for every path from $\tilde{\omega}$ and $\tau(\tilde{\omega})$ there must be some $\hat{\omega}$ and $\sigma \in \hat{\omega}$ that lies on that path. In particular, this must be true for any least resistance path from $\tilde{\omega}$ to $\tau(\tilde{\omega})$. In this case, it is obvious from the definition of resistance that

(*) $$r(\tilde{\omega}, \hat{\omega}) + r(\hat{\omega}, \tau(\tilde{\omega})) \le r(\tilde{\omega}, \tau(\tilde{\omega})).$$

If $\tilde{\omega}$ is in the subtree with root $\hat{\omega}$ cut it and paste it to $\hat{\omega}$. Cutting reduces the resistance by $r(\tilde{\omega}, \tau(\tilde{\omega}))$, while pasting adds $r(\tilde{\omega}, \hat{\omega})$. By (*) the resistance of the tree is not increased.

Suppose that $\tilde{\omega}$ is not in the subtree with root $\hat{\omega}$. First we cut $\hat{\omega}$ and paste it to $\tau(\tilde{\omega})$. Then we cut $\tilde{\omega}$ and past it to $\hat{\omega}$. In effect, we place $\hat{\omega}$ in between $\tilde{\omega}$ and $\tau(\tilde{\omega})$. Cutting $\hat{\omega}$ does not increase the resistance. Pasting to $\tau(\tilde{\omega})$ increases the resistance by $r(\hat{\omega}, \tau(\tilde{\omega}))$. Cutting $\tilde{\omega}$ reduces the resistance by $r(\tilde{\omega}, \tau(\tilde{\omega}))$. Pasting it to $\hat{\omega}$ increases the resistance by $r(\tilde{\omega},\widehat{\omega})$. So the resistance is increased by $r(\widehat{\omega},\tau(\tilde{\omega})) - r(\tilde{\omega},\tau(\tilde{\omega})) + r(\tilde{\omega},\widehat{\omega})$. By (*) the resistance of the tree has not been increased.

Now repeat the process, noting that once $\tilde{\omega},\tau(\tilde{\omega})$ are adjacent, the procedure never breaks this link, so the process does not cycle. At the end we have a least resistance ω -tree with the required property.

☑

Proof of Theorem 6.2: Given any Ω_2 tree τ with root in Ω_1, we may define a tree $\Omega_1(\tau)$ in Ω_1 by defining $\Omega_1(\tau)(\omega)$ to be the closest predecessor of ω in Ω_1. If τ has a root in Ω_1 and if $\tilde{\omega}$ and $\tau(\tilde{\omega})$ are adjacent for all $\tilde{\omega} \neq \omega$ we say that τ is an adjacent- Ω tree. By Lemma B.4 and Young's Theorem $\mu(\omega) > 0$ if and only if ω is the root of a least-resistance adjacent- Ω tree. We will show that within the adjacent- Ω trees there is a class of proto- Ω_1 trees for which

$$r(\tau) = \hat{r}(\Omega_1(\tau)) + \sum_{s \neq \tilde{s}} \underline{r}(s,\tilde{s})/2 ,$$

and that there exist least resistance ω proto- Ω_1 trees. Since $\Sigma_{s\neq\tilde{s}}\underline{r}(s,\tilde{s})/2$ is a constant that does not depend on τ, and since $\Omega_1(\cdot)$ preserves the root and maps the space of proto- Ω_1 trees **onto** the space of Ω_1 trees, this implies the desired result.

Consider a node $\omega \in \Omega \setminus \Omega_1$ in an adjacent- Ω tree τ. We refer to this as a binary point. Observe that a binary point can be adjacent to only two points. Since a binary point is not the root, is has an immediate predecessor that is adjacent to it, and it follows that it can have only one immediate successor. If ω has a successor in Ω_1 we call it a *trapped node* if not we call it a *dangling node*.

For two points $s,\tilde{s} \in \Omega_1$, let $\omega_0,\ldots,\omega_K \in \Omega_2$ be the unique sequence such that $\omega_0 = \omega$, $\omega_K \in \tilde{\omega}$ and ω_k,ω_{k+1} are adjacent. Observe the all of these nodes are trapped, or all are dangling. Note that the resistance is

$$r(s,\tilde{s}) = \sum_{k=0}^{K-1} r(\omega_k,\omega_{k+1}) .$$

If $\Omega(\tau)(s) = \tilde{s}$, because τ is adjacent- Ω, in τ the nodes ω_k must be trapped in between $s,\tilde{s}$ in the correct order. This implies that the resistance of τ is the resistance of τ_1 plus the resistance of the dangling nodes.

If neither $\Omega_1(\tau)(s)=\tilde{s}$ nor $\Omega_1(\tau)(\tilde{s})=s$ we say that $s,\tilde{s}$ is a *dangling pair*. In this case, for some k ω_k is a successor of s and ω_{k+1} is a successor of $\tilde{s}$. In the least resistance tree, we clearly must choose k to minimize

$$\underline{r}(\omega,\tilde{\omega},k)=\sum\nolimits_{j=1}^{k} r(\omega_j,\omega_{j-1})+\sum\nolimits_{j=k+1}^{K-1} r(\omega_j,\omega_{j+1}) .$$

Any tree which minimizes this expression for all dangling pairs is called a *proto*- Ω_1 tree. Notice that for any $\tau_1\in\Omega_1$ there is a proto- Ω_1 tree τ with $\Omega_1(\tau)=\tau_1$, so the map $\Omega_1(\cdot)$ is onto as asserted.

To compute the resistance of the dangling nodes in a proto- Ω_1 tree, we sum $\underline{r}(s,\tilde{s})$ over all dangling pairs $s,\tilde{s}$. Alternatively, instead of adding $r_1(\Omega_1(\tau))$ to $\underline{r}$ for all dangling pairs, we may add it to $\underline{r}$ for all pairs, and subtract $\underline{r}(s,\tilde{s})$ for those pairs that are not dangling. In other words,

$$r(\tau)=r_1(\Omega_1(\tau))+\sum\nolimits_{s\neq\tilde{s}}\underline{r}(s,\tilde{s})/2-\sum\nolimits_{s}\underline{r}(s,\tau(s)) .$$

Substituting the definition of the modifed resistance, this is the desired result.

☑

Appendix C: Tree Trimming

Lemma C.1 [Tree Trimming]: *Let* $\Omega^* \subset \tilde{\Omega} \subseteq \Omega$. *Suppose that for every least resistance tree in* $\tilde{\Omega}$, *there is another least resistance tree* τ *with the same root and* $\tau(\omega^*) \in \Omega^*$ *whenever* $\omega^* \in \Omega^*$. *Then least resistance trees in* $\tilde{\Omega}$ *have roots in* Ω^*, *and* ω *is the root of a least resistance tree in* $\tilde{\Omega}$ *if and only if it is the root of a least resistance tree in* Ω^*.

Remark: Notice that this can be applied recursively. That is, after we have applied the lemma once to looking only at trees in Ω^*, we can then apply the lemma to a subset $\Omega^{**} \subset \Omega^*$.

Proof: We refer to a tree for which $\omega^* \in \Omega^*$ implies $\tau(\omega^*) \in \Omega^*$ as a proto-Ω^* tree. It is apparent the proto-Ω^* trees must have roots in Ω^*, so by the hypothesis, least resistance trees in $\tilde{\Omega}$ must have roots in Ω^*.

Next, we map proto-Ω^* trees to Ω^* trees, and show how to compute the resistance of the proto-Ω^* trees from that of the corresponding Ω^* tree. This will enable us to show that it is sufficient to minimize resistance over Ω^*. Specifically, a proto-Ω^* tree $\tilde{\tau}$ gives rise to a unique tree $\Omega^*(\tilde{\tau})$ on Ω^* defined by $\Omega^*(\tilde{\tau})(\omega) = \tilde{\tau}(\omega)$. Define

$$r^*(\Omega^*(\tilde{\tau})) = \sum\nolimits_{\omega \in \Omega^* \setminus root(\tilde{\tau})} r(\omega, \tilde{\tau}(\omega))$$

$$r_-(\tilde{\tau}) = \sum\nolimits_{\omega \in \tilde{\Omega} \setminus \Omega^*} r(\omega, \tilde{\tau}(\omega))$$

Then

$$r(\tilde{\tau}) = r^*(\Omega^*(\tilde{\tau})) + r_-(\tilde{\tau}) .$$

Finally, consider a tree τ^* in Ω^* and a proto-Ω^* tree $\tilde{\tau}$. We may construct another proto-Ω^* tree $\tau^* \wedge \tilde{\tau}$ by

$$\tau^* \wedge \tilde{\tau}(\omega) = \begin{cases} \tau(\omega) & \omega \in \Omega^* \\ \tilde{\tau}(\omega) & \omega \in \tilde{\Omega} \setminus \Omega^* \end{cases}$$

This tree has resistance $r_*(\tau_*) + r_-(\tilde{\tau})$.

Now suppose that τ is a least resistance tree. Let $\tilde{\tau}$ be the corresponding least resistance proto-Ω^* tree. Suppose that $r^*(\tau^*) < r^*(\Omega^*(\tilde{\tau}))$. Then

$$r(\tau^* \wedge \hat{\tau}) = r^*(\tau^*) + r_-(\tilde{\tau}) < r^*(\Omega^*(\tau)) + r_-(\tilde{\tau}) = r(\tilde{\tau}) = r(\tau)$$

contradicting the assumption that τ has least resistance. This proves the "only if" part of the Lemma.

Suppose conversely, that τ^* has least resistance among all Ω^* trees. Let $\tilde{\tau}$ be a minimizer of $r_-(\tilde{\tau})$. Then clearly $\tau^* \wedge \tilde{\tau}$ is least resistance in $\tilde{\Omega}$.

☑

Corollary C.2: *Suppose that* $\Omega^* \subset \tilde{\Omega} \subseteq \Omega$ *, and that for all* $\omega^*, \omega^{**}\Omega^*, \tilde{\omega}, \tilde{\tilde{\omega}} \in \tilde{\Omega} \setminus \Omega^*$.

(i) *If* $r(\omega^*, \tilde{\omega}) > r(\tilde{\tilde{\omega}}, \omega^*)$ *then roots of least resistance* $\tilde{\Omega}$ *trees are in* Ω^*.

(ii) *If in addition* $r(\omega^*, \tilde{\omega}) \geq \min\{r(\omega^*, \omega^{**}), r(\omega^{**}, \omega^*)\}$ *then* ω *is the root of a least resistance tree in* $\tilde{\Omega}$ *if and only if it is the root of a least resistance tree in* Ω^*.

Proof: First, we establish that least resistance trees in $\tilde{\Omega}$ must have roots in Ω^*. If not, then there is some $\tilde{\tilde{\omega}} \in \tilde{\Omega} \setminus \Omega^*$ that is the root of a least resistance tree τ. Moreover, there must be some $\omega^* \in \Omega^*$ with $\tau(\omega^*) = \tilde{\omega} \in \tilde{\Omega} \setminus \Omega^*$. Cutting ω^* saves $r(\omega^*, \tilde{\omega})$. Pasting the root to ω^* costs $r(\tilde{\tilde{\omega}}, \omega^*)$. So by assumption the resistance of the tree has been strictly reduced, contradicting the assumed minimality of τ.

Now suppose that $\omega^{**} \in \Omega^*$ is the root of a least resistance tree τ. If $\tau(\omega^*) = \tilde{\omega} \in \tilde{\Omega} \setminus \Omega^*$, then cutting ω^* saves $r(\omega^*, \tilde{\omega})$. If $r(\omega^*, \tilde{\omega}) \geq r(\omega^{**}, \omega^*)$ then pasting the root ω^* does not increase the resistance; while if $r(\omega^*, \tilde{\omega}) \geq r(\omega^*, \omega^{**})$ then pasting ω^* to the root does not increase the resistance. By assumption, one of these must be the case. Proceeding iteratively, we see that there exists a least resistance tree $\hat{\tau}$ in which $\omega^* \in \Omega^*$ implies $\hat{\tau}(\omega^*) \in \Omega^*$. The result now follows from Lemma C.1.

☑

References

Bergin, J. and B. L. Lipman [1994]: "Evolution with state-dependent Mutations," UCLA Theory Workshop, 4/14/94.

Dawid, H. [1999]: *Adaptive Learning by Genetic Algorithms,* (Berlin: Springer).

Ellison, G. [1995]: "Basins of Attraction and Long Run Equilibria," Universidad Carlos III de Madrid, Second International Conference on Economic Theory: Learining in Games.

Friedlin, M. and A. Wentzell [1982]: *Random Perturbations of Dynamical Systems*, (New York: Springer Verlag).

Friedman, D. [1998]: "On Economic Applications of Evolutionary Game Theory," *Journal of Evolutionary Economics,* 8: 15-43.

Fudenberg, D. and D. K. Levine [1998]: *The Theory of Learning in Games,* (Cambridge: MIT Press).

Fudenberg, D. and D. Kreps [1993]: "Learning Mixed Equilibria," *Games and Economic Behavior*, 5: 320-367.

Hahn, S. [1995]: "The Long Run Equilibrium in an asymmetric Coordination Game," Harvard.

Johnson, P., D. K. Levine and W. Pesendorfer [1999]: "Evolution and Information in a Gift Giving Game," UCLA.

Kandori, M. [1992]: "Social Norms and Community Enforcement," *Review of Economic Studies*, 59: 61-80.

Kandori, M. and R. Rob [1993]: "Bandwagon effects and long run technology choice," University of tokio Discussion Paper 93-f-2.

Kandori, M., G. Mailath and R. Rob [1993]: "Learning, Mutation and Long Run Equilibria in Games," *Econometrica*, 61: 27-56.

Morris, S., R. Rob and H. Shin [1993]: "p-dominance and Belief Potential," *Econometrica*, 63: 145-158.

Van Damme, E. and J. W. Weibull [1998]: "Evolution with Mutations Driven by Control Costs," Tilburg.

Young, P. [1993]: "The Evolution of Conventions," *Econometrica*, 61: 57-83.