On the Evolution of Pareto Optimal Behavior in Repeated Coordination Problems

Roger Lagunoff*

Revised: April 2, 1997

This paper characterizes the asymptotic behavior of an ongoing society facing a repeated coordination problem. This society has a certain demographic structure: generations of individuals asynchronously supercede their "parents," creating an entry/exit process that allows individuals with possibly different beliefs to enter society. A *self confirming equilibrium (SCE) belief process* describes an evolution of beliefs in this society consistent with a self confirming equilibrium of the repeated game. Due to Fudenberg and Levine (1993), SCE is weaker than Nash as it requires correct forecasts of an individual only along the realized path during the individual's lifetime. Since individuals' beliefs on out-of-equilibrium behavior may vary, an SCE belief process may admit random and heterogeneous forecasts in the form of mutations of beliefs across generations as newborn individuals enter the system.

The main result shows that with belief mutation, for any repeated coordination problem, the Pareto dominant equilibrum is a globally absorbing state of the dynamic process. This result does not involve either of the usual assumptions of myopia or large inertia common in evolutionary models. Nor is this result possible if only Nash rather than self confirming equilibria are considered. *Journal of Economic Literature* Classification Numbers: C7, D8.

*Department of Economics, Georgetown University, Washington DC, 20057, USA. Phone: (202) 687-1510. e-mail: *lagunofr@gunet.georgetown.edu*. I wish to thank seminar participants at Cal Tech, Northwestern, Rutgers, University of Texas, SMU, UCLA, Washington University, and the Midwest Mathematical Economics Meetings at Penn State for helpful discussions. I am also grateful to an associate editor of this journal, two anonymous referees, Larry Blume, Eddie Dekel, David Levine, Akihiko Matsui, Wolfgang Pesendorfer, Bill Zame, and especially George Mailath for extensive comments and suggestions. All errors remain my own.

1 Introduction

A *coordination problem* is described by a static game with multiple Nash equilibria, one of which yields each player his highest possible payoff in the game. Will societies successfully reach and sustain the optimal outcome in repeated coordination problems in the long run? Despite the simple structure of coordination problems, formal theory has generated no consensus on the question.

The most common approach to address coordination issues is an equilibrium approach. By this approach, coordination is readily achieved as a Nash equilibrium of the repeated game. The equilibrium approach requires that all individuals correctly anticipate, not just actual behavior, but behavioral responses to successive layers of counterfactual events as well (i.e, what happens if there is a deviation from the predicted outcome, and a deviation to the predicted outcome of the deviation, and so on). For this reason, equilibria of repeated games are sometimes interpreted as social norms that are enforced by commonly understood "punishments" for deviant behavior. Unfortunately, Nash equilibrium models of repeated games have little predictive content. Even the simplest games have a large multiplicity of outcomes.[1] Furthermore, the source of this common understanding, particularly on counterfactual events, is unclear.

This paper addresses the coordination question by examining the mechanics of an equilibrium process that abandons some arguably less attractive features of Nash equilibrium. The process is assumed to operate in an ongoing society that faces a repeated game of common interest. A common interest game, a term introduced by Aumann and Sorin (1989), is a stage game in which there is a strictly Pareto dominant payoff profile. Our main result is that for any common interest game, the Pareto dominant equilibrium is the unique, globally absorbing state of this process in this sense: starting from any strategy profile, there is a finite random stopping time at which the process settles on the Pareto dominant equilibrium and never departs thereafter.

Specifically, we consider a model of a repeated game. Unlike the standard model, however, there is some inertia in the decisions opportunities of the players. Each period, with positive probability a given player may be unable to change his actions. Moreover, the game is played by an ongoing society of overlapping generations of individuals. Periodically, individuals "die" and are replaced with new entrants who inherit the same physical characteristics as their "parents." However, births and deaths are not observable across different lineages. Similar demographic structures are found in Blume (1994), Matsui and Rob

[1] The clearest expression of this idea is the Folk Theorem. See, for example, Fudenberg and Maskin (1986).

(1991), and Lagunoff and Matsui (1995). All individuals are fully rational in the sense that they maximize discounted expected utility given beliefs about the future path of societal behavior.

Despite the overlapping generations structure, if Nash equilibrium behavior is assumed, then standard Folk Theorem types of results still hold since the environment is stationary and each individual's lifetime is stochastic in length. Whatever process determines the beliefs of new entrants, these beliefs must be perfectly correlated with those of the currently lived individuals whose beliefs, in turn, are always correct. Hence, in a Nash equilibrium, successful coordination in behavior follows from assumed perfect correlation in beliefs.

The present paper emphasizes this "statistical view" of equilibrium as a process on beliefs across generations. In a Nash equilibrium, the belief generating process is a degenerate one. In standard models, "coordination failures" are then rooted in the perfect transmision of inefficient norms across time. However, this need not be the case. We examine an alternative: the *self confirming equilibrium* concept introduced by Fudenberg and Levine (FL) (1993). The self confirming concept weakens the Nash concept by requiring only that all individuals make correct predictions about the realized path of play while they remain in society. Unlike in a Nash equilibrium, beliefs about *out-of-equilibrium* behavior need not be correct.

The self confirming equilibrium alternative admits the possibility of imperfect correlation of beliefs of different individuals. This allows us to introduce noise or mutation in the belief process of new entrants. Behavior then evolves over time due to accumulation of heterogeneous beliefs across generations. With an independence and a full support assumption on mutation, the transmision of existing norms across generations is weak enough to allow society to escape Pareto inferior behavior. Eventually, however, the equilibrium requirement allows sufficient coordination to prevent departure from Pareto superior behavior.

To see how this works, Section 2 gives an example illustrating this result. It shows how the self confirming concept strikes a balance between the coordinating role of equilibrium on the one hand, and the lattitude needed to escape coordination failures on the other. Section 3 describes the formal model. Both the demographic structure and the concept of a self confirming equilibrium belief arrival process are described. Section 4 describes the mutation assumptions and states the main result. We also examine what happens when the crucial assumptions of the model are dropped. Section 5 relates the current paper to existing literature. Finally, Section 6 proves the main result.

2 An Example

The importance of flexibility of beliefs in the equilibrium concept can be illustrated in the following example. The simple 2×2 game below is a coordination game. There are two Pareto ranked Nash equilibria, $\bar{s} = (\bar{s}_N, \bar{s}_B)$ and $s^* = (s^*_N, s^*_B)$. The payoffs are specified so as to emphasize that the following intuition is not specific to a "nongeneric" common interest game.

Bob

Ned	s^*_B	$\bar{s}_B$
s^*_N	5, 6	1, 2
$\bar{s}_N$	1, 0	3, 5

Figure 1

Assume an infinitely repeated game with patient players. However, suppose that there is some inertia so that each period there is some small probability that a player cannot change his behavior. Let Ned and Bob denote the individuals alive at a time t_0 in which the current profile is $\bar{s} = (\bar{s}_N, \bar{s}_B)$. Suppose that both believe that $\bar{s}$ is permanent, i.e., each believes that the other will keep his current action forever. These beliefs are consistent with both Nash and self confirming equilibrium.

Now suppose a stochastic overlapping generations structure, whereby each period there is a small but stationary probability that one of the two players is replaced. However, these replacements are unobserved by the other player. One interpretation is that "Bob" and "Ned" are two firms. The managers of each firm change periodically, but no one in firm i precisely observes the identity of the manager or decision maker in firm j. Let Ned Jr. and Bob Jr. denote both individuals' successors. Suppose that Bob is replaced by Bob Jr. at time t_1 who is born with the belief that the Ned's replacement, Ned Jr., will initiate a switch toward the Pareto superior equilibrium s^*. Bob Jr. believes that if he attempts to preempt this switch by changing his own action to s^*_B, no one will follow him. Bob Jr.'s best response under this belief is to keep his current action until he sees a switch toward s^*. Since Ned and Bob Jr both keep their current action $\bar{s}$ while they face each other, each one's prediction of the other's behavior is correct. For this reason, the scenario is consistent with self confirming equilibrium.

However, this scenario is inconsistent with Nash equilibrium. The reason is that if Ned were to defect to s^*_N, Bob Jr. would then infer that Ned Jr. had replaced Ned and so Bob Jr. will subsequently switch to s^*_N. But if this switch is correctly anticipated by Ned, as required in a Nash equilibrium, then Ned would certainly defect to the dominant equilibrium. This means that Ned's original belief that Bob's lineage remain unconditionally at s_B^* forever is misforecast. The difference between Nash and self confirming equilibrium occurs because Ned's incorrect forecast occurs only on the out-of-equilibrium event that he himself switches to s_N^* which, given his beliefs, he would never do. Self confirming equilibrium beliefs admit this possibility. Nash does not.

To continue, let t_2 be the time at which Ned is replaced by Ned Jr. Since Bob Jr. remains alive at any given future time with positive probability, Ned Jr's beliefs must be consistent with the in-equilibrium prediction of Bob Jr. Specifically, Ned Jr.'s belief must be such that he anticipates that his switch to s_N^* induces an immediate similar switch by Bob Jr. Given this belief, if the probability of inertia is small and individuals are patient enough, then Ned Jr.'s best response is indeed to switch to s_N^*. This entire scenario, sketched in the table below, constitutes a self confirming equilibrium of the repeated game.

Individual		Time	Newborn's belief
Bob	Ned	t_0	"$\bar{s}$ forever"
Bob Jr.	Ned	t_1	"$\{\bar{s} \rightarrow s^*\}$ only after Ned is replaced"
Bob Jr.	Ned Jr.	t_2	"$\{\bar{s} \rightarrow s^*\}$ at the next decision point"

Table 1

So far, the example outlines how new belief arrivals can create a self confirming equlibrium transition to the dominant equilibrium, despite beliefs from the initial cohort which are inconsistent with such a transition. This transition, with the initial beliefs of Bob and Ned, is not possible in any Nash equilibrium.

Now suppose we start with s^*. Substituting the two profiles in Table 1, can the same sequence of beliefs generate an equilibrium transition from s^* to $\bar{s}$? In fact, it cannot.

Since s^* is the Pareto dominant payoff, no individual will depart from s^* unless he anticipates that his failure to do so will lead to a "punishment" phase. In the punishment phase, individuals will move to some inferior profile(s) for some length of time. Failures to follow the punishment strategy is followed by extended punishment phases, and so on. But recall that revision and replacement times of other lineages are not observable. Therefore, any punishment for failing to leave s^* at the appropiate time must begin at some prespecified calendar time. But even if Ned Jr. had planned to switch to $\bar{s}_N$ by some calendar date he may be unable to switch by that specific date due to inertia in the decision process. Therefore, the event that a punishment ensues if Ned Jr. sticks with s^* occurs with positive probability (an "in-equilibrium" event). But given the stochastic replacement, the initial generation of Ned and Bob will remain alive at any future calendar date with positive probability. They both forecast that s^* is permanent. Therefore, reversing the transition in Table 1, when Ned Jr's departs from s^*_N at some date, say, time t_2, this violates the self confirming equilibrium belief of Bob on a set of histories having positive probability. Hence, no departure from s^* occurs with these initial beliefs of Ned and Bob.

What this example demonstrates is a basic asymmetry between a Pareto dominant profile s^* and other profiles such as $\bar{s}$. Starting from $\bar{s}$ there is a self confirming equilibrium transition to s^* even though the initial generation believed in the permanence of $\bar{s}$. However, no such transition is possible from s^*. The reason is that, unlike departing from s^*, Ned Jr. does not require the threat of punishment to depart from $\bar{s}$. For example, in the Table 1 scenario, starting from $\bar{s}$, Ned Jr. strictly prefers to initiate the transition $\{\bar{s} \longrightarrow s^*\}$ knowing that Bob Jr. will follow him. If the same belief sequence had occured starting from s^*, then individual Ned Jr. would prefer to remain in s^*. He can accomplish this simply by keeping his action s^*_N, thereby "pretending" to be the senior Ned. Since Bob Jr. will only switch if he observes a switch from Ned's lineage, Ned Jr. should maintain his current action. Hence, the transition from s^* to $\bar{s}$ does not occur. The difference between Ned Jr.'s incentives in s^* versus $\bar{s}$ comes from the simple fact that s^* Pareto dominates $\bar{s}$, and the knowledge that others cannot condition their behavior on the replacement between Ned and Ned Jr.

To summarize, the "right initial beliefs" make s^*, but not $\bar{s}$ permanent. There is now a clear role for noise in such a process. A "mutation" process over beliefs can hit upon the right sequence of beliefs for departing on $\bar{s}$, and the right sequence of beliefs for remaining in s^*.

3 The Formal Model

We now formalize the informal structure of the preceding example. There are n players who play a standard repeated game. The stage game is described by the tuple $G = (S, u; I)$ where $I = \{1, \ldots, n\}$ is the set of agents. The n-fold product $S = \times_{i \in I} S_i$ is the set of *profiles* $s = (s_1, \ldots, s_n) \in S$ assumed to be finite. Stage game payoffs are defined by $u = (u_i)_{i \in I}$ where $u_i : S \rightarrow \Re$ for each $i \in I$. A *common interest game* is a stage game with the property that there exists a strictly Pareto dominant action profile.[2] This profile will be denoted by

[2] This strictly dominant profile is, of course, a strict Nash equilibrium of G.

s^*. The corresponding payoff to each player is $u_i(s^*) \equiv u_i^*$. The definition is due to Aumann and Sorin (1989).

Now denote the behavior profile at time t by $s(t) = (s_1(t), \ldots, s_n(t)\ )$. This profile is perfectly observable to all players. A period t *behavior history* is a tuple $h^t \equiv (s(1), s(2), \ldots, s(t-1))$ of action profiles observed by time t. The set of period t behavior histories is given by H^t, and let $H = \cup_{t=1}^{\infty} H^t$ denoting the collection of all (finite) behavior histories. An infinite history $h^\infty \in H^\infty$ describes the entire path of play

We assume that in each period there is some probability $\gamma > 0$ that player i is unable to move due to inertia. Let $y_i(t)$ denote the outcome of i's random inertial process, taking value 1 if i has a decision opportunity at time t and value 0 if not. This process is assumed independent and identical across players and across time. Whether or not a player is unable to move is not directly observable to the other players so that certain potential punishments must occur along the equilibrium path. The results will apply to a "neighborhood" of the limiting case of $\gamma \to 0$ in which there is no inertia.

The play at any given time represents a snapshot of a society of overlapping generations. Replacement is assumed to take place stochastically and asynchronously. The process is modelled as follows. Let $\delta > 0$. Each period, with probability $n(1-\delta)$ exactly one of the n players "dies" and is replaced by a successor who inherits the same preferences and information. This replacment is i.i.d. across time. We refer to the index "i" now as the lineage rather than the individual person, Hence, an individual from lineage i remains "alive" in the current period with probability δ. Let $z_i(t)$ denote the accumulated number of replacements to time t; if $j = z_i(t)$ then the current entrant is the jth newborn individual in the ith lineage. The function $z_i(\cdot)$ determines the entire arrival path for lineage i. Let Z denote the set of possible arrival paths $z = (z_1, \ldots, z_n)$.

As with the inertia, an individual does not observe the actual replacement of individuals from other lineages. This assumption plays an important role in establishing the result. The reason for the perfectly asynchronous replacement is to capture the idea of an ongoing society where existing norms of behavior are socially transmitted to newborns as they enter society.[3] The newborn inherits the current norms from the elders.

When a newborn individual j from lineage i comes into the world he chooses a *strategy* f_{ij} which maps histories into the set of actions S_i. Each individual can only change his behavior at decision opportunities. Formally, $f_{ij} : H \times \{0, 1\} \to S_i$. The interpretation is if $y_i(t) = 1$ then this newborn has a decision opportunity in which he chooses action $f_{ij}(h^t, 1)$.

[3]The assumption that individuals enter society "one at a time" is a conceptual and technical simplification. All that required is that *complete* population turnover is not possible in any period.

If $y_i(t) = 0$ then he has no decision opportunity and so $f_{ij}(h^t, 0) = s_i(t-1)$. Notice that only pure behavior strategies are considered. While this entails some loss of generality (even with inertia), it significantly eases the notational burden later on.

Each newborn chooses a strategy which is possibly different than his predecessor. Let $f = (f_{ij})$ denote the entire tuple of strategies, one for each individual who at some point enters the game. The path of behavior generated by f is denoted by $\tilde{s}(f)$.[4] The event that a particular history h is realized in the infinite play paths H^∞ is given by the cylinder set $C(h)$. The strategy tuple f (jointly with γ and δ) determines a measure μ_f on the measurable set of infinite paths $(H^\infty, \mathcal{F})$. Here, $\mathcal{F}$ is the smallest σ-algebra containing all the σ-algebras generated by the cylinder sets $C(h)$. Therefore, $\mu_f(C(h))$ is the probability that the history h is realized under strategy profile f.

Given history h^t, the expected utility for individual ij under strategy tuple f at time t is given by

$$U_i(f|\, h^t) \;=\; E_{\mu_f}\left[\sum_{\tau=t}^{\infty}(1-\delta)\delta^\tau u_i(\tilde{s}(f))\Big|\; h^t\right]. \tag{1}$$

where the expectation operator E_{μ_f} is taken with respect to the measure μ_f given history h^t. Observe that i's discount factor is his retention/replacement rate δ. We could just as easily allow for i to discount separately.

Each individual is assumed to choose a strategy to maximize his expected discounted utility stream given his beliefs about others' strategies. Let F_i denote the set of strategies for individuals in lineage i. A *belief* ϕ_{ij} for individual ij is a probability measure on the space $F_{-i} = \times_{k \neq i} F_k$ of strategy profiles of other lineages (see Section 6 for the formal construction).[5] Individuals within a lineage are identical in every respect other than their beliefs. We will write $\phi_{ij}(D|h^t)$ to denote the regular conditional probability that ij assigns to others' strategies, f_{-i}, being in the set D given the realized history h^t. Let Φ_i be the set of possible beliefs for an individual of lineage i.

Self Confirming Equilibrium

The standard requirements of repeated game Nash equilibria dictate that the prior belief of a newborn coincide with the existing posterior of his predessor, which , in turn, coincides with the objective distribution over, not only the realized path of societal behavior, but also

[4]Formally, $\tilde{s}(f)$ can be recovered by inducting on f: $\tilde{s}(f)(1) = f(h^0, y(1)\ )$, $\tilde{s}(f)(2) = f((h^0, f(h^0, y(1)), y(2)\ )$, ...

[5]We point out an ambiguity in notation. The tuple f is the infinite collection (f_{ij}). However, an individual's belief is defined to be a probability measure on F_{-i} which consists of the $n-1$ tuples of strategies, one from each of the other lineages. Since individual ij cannot distinguish the identity of individuals in other lineages, his belief need not formally distinguish between, say, f_{kj} and $f_{k(j+1)}$.

over successive layers counterfactual possibilities as well. The requirement that individuals, upon entering society, make instantaneously correct forecasts about counterfactual events seems excessive. More importantly, it allows no lattitude for stochastic variation in beliefs of new members of society.

This paper therefore examines an alternative equilibrium concept. We apply the notion of *self confirming equilibrium (SCE)* introduced by Fudenberg and Levine (1993a).[6] In a SCE, individuals need not have mutually consistent and correct forecasts about counterfactual information.

The formal definition is as follows. First, let $\mathbf{\Phi} = \times_{i=1}^{n} \times_{j=1}^{\infty} \Phi_{ij}$, denote the product space of beliefs of all individuals who will, at some point, take part in the game. Let $\boldsymbol{\phi}$ be an element of $\mathbf{\Phi}$. We refer to $\boldsymbol{\phi}$ as a *belief realization.*

Definition 1 A *self confirming equilibrium (SCE)* is a collection of strategies $f^* = (f^*_{ij})$, and a belief and arrival realization $(\boldsymbol{\phi}, z)$ such that

(i) An individual from lineage i who enters at time t with beliefs ϕ chooses a strategy in F_i to maximize

$$E_\phi \left[U_i(f | \, h^t) \right] \tag{2}$$

where E_ϕ is the expectation given probability $\phi(\cdot | h^t)$ over others' strategies.

(ii) Suppose that entrant j from lineage i is alive at time t and let $f \backslash f^*_i = ((f^*_{ij})_{j=1}^{\infty}, \; f_{-i})$. Then for μ_{f^*}-almost every history $h \in H$, and for any event $C(h^t)$,

$$\int_{f_{-i}} \mu_{f \backslash f^*_i}(C(h^t) \Big| \, h) d\phi_{ij} \;=\; \mu_{f^*}(C(h^t) \Big| \, h). \tag{3}$$

In Expression (3), the term $C(h^t)$ is an event that occurs during the entrant's lifetime since h^t is a history which ends when he is still alive. Hence, in a self confirming equilibrium (SCE), each individual's prediction about others' strategies induces a correct prediction about the distribution of observable behavior of society, but only on those paths that are in the support of the realized distribution while that individual remains alive. Definition 1 is almost identical to the definition of Nash equilibrium with an important qualification that in a Nash equilibrium, expression (3) holds for every history h, including out-of-equilibrium paths. Hence, every Nash equilibrium of the repeated game is a SCE. However, the converse is not true.

[6]Similar concepts are found in Kalai and Lehrer (1993b) and Battigalli and Guatoli (1988).

Definition 2 A *self confirming equilibrium (SCE) belief arrival process* denoted by π, is defined as a probability measure on $\mathbf{\Phi} \times Z$ that jointly determines the beliefs and arrival times of each newborn, and satisfies: for π-almost every $(\boldsymbol{\phi}, z)$ pair, there is a collection of strategies (f^*_{ij}) that, together with $(\boldsymbol{\phi}, z)$, constitute a self confirming equilibrium.

Both the self confirming and the Nash belief processes simply describe an evolution of beliefs in this society which is consistent with equilibrium behavior. Both are also endogenous in the sense that they are restricted to satisfy certain equilibrium properties. However, Definition 2 defines a larger class of belief processes than if "SCE" were replaced with "Nash" in the definition. Typically, Nash belief processes are degenerate distributions on belief realizations $\boldsymbol{\phi}$ that are consistent with a Nash equilibrium. For example, in Section 2, a Nash belief process would place probability one on a realization of beliefs $\boldsymbol{\phi}$ where $\phi_{ij}(\{f^*_{-i}\}|h) = 1,\ \forall h \in H$, such that $\tilde{s}(f^*) = \bar{s}$. That is, each individual predicts that all individuals from lineage $i = 1, \ldots, n$, take the action $\bar{s}_i$ consistent with the Pareto inferior Nash equilibrium regardless of past history.

4 The Main Result

Unlike Nash, a self confirming belief process π may be subject to random perturbations or "mutations" in beliefs that generate random fluctuations in behavior. Below, we formally define three "mutations on beliefs" properties.

Some notation is required first. The kth entrant over all lineages is uniquely determined by $k = \sum_i z_i(t)$. Denote by z^k the history of arrivals through entrant k and let ϕ^k denote the n-tuple of realized beliefs of individuals alive at the time of the kth newborn. Now let $\tau(k)$ denote the time that the first k individuals over all lineages have died. Finally, $\kappa = \sum_i z_i(\tau(k))$ denotes the first newborn after the first k entrants have died.

(A1) (Stationary Markov) The belief process π defines a stationary Markov process on the belief realizations $\boldsymbol{\phi}$, with transition probabilities expressed as

$$\pi(\{\phi_k \in A\}\Big|\ \boldsymbol{\phi}^{k-1}, z^{k-1}\ )$$

for any measurable set $A \in \Phi$.

(A2) (Independent Increments) For all z, for each pair k and κ, and each pair of Borel subsets $A_\kappa \in \Phi_\kappa$ and $A_k \in \Phi_k$,

$$\pi\left(\{\phi_\kappa \in A_\kappa\} \cap \{\phi_{\mathbf{k}} \in A_k\}\Big|\ z\right) \ = \ \pi\left(\{\phi_\kappa \in A_\kappa\}\Big|\ z\right) \cdot \pi\left(\{\phi_{\mathbf{k}} \in A_k\}\Big|\ z\right)$$

(A3) (Maximal Support) There is no other SCE belief process that satisfies (A1) and (A2) and has strictly larger support than π.

In words, by Assumption (A1) the process determines the beliefs of the kth individual to enter the world. This belief depends only the previous arrivals and the beliefs of the individuals from each of the n lineages who are alive just prior to the time that individual k enters (this includes k's own "parent") .

Assumption (A2) states that the belief of the entrant κ is drawn independently of the belief of entrant k who dies before κ enters.[7] Assumption (A2) in particular deserves some elaboration. Notice that it requires the beliefs of certain temporally separated individuals to be independently determined while still satisfying an equilibrium condition. The example in Table 1 shows this is possible. In the example, independence is only possible with the beliefs of two individuals: Bob and Ned Jr. What it implies is that nature independently draws among conditional beliefs of Bob Jr over the behavior after the time that Ned is replaced.[8] The realized belief of Bob Jr. "bridges the gap" between the older Ned and the new generation starting with Ned Jr. The SCE belief process therefore allows independent randomness in the determination of a new entrant's beliefs, *only over beliefs about either probability zero events or events that take place after an entire cohort of contemporaries have been replaced.*

Observe that neither (A1) nor (A2) is necesarily at odds with Nash equilibrium. In fact, any Nash equilibrium can be supported by a belief process which satisfies both (A1) and (A2). The reason is simple. In a Nash belief process, each individual correctly predicts all others' strategies. Hence, beliefs of all individuals coincide, satisfying (A1). One can therefore define the support of π to consist of a singleton belief realization $\boldsymbol{\phi}$ that is consistent with some Nash equilibrium. (A2) would therefore be vacuous. It is the *full support* property (A3) that creates an inconsistency with some Nash equilibria of the game. This assumption requires an SCE process to have maximal support with respect to the first two properties. Together with the independence assumption, (A3) rules out most Nash equilibria since they can be represented as degenerate distributions on the belief process. This is critical in establishing the result below.

Given an SCE process π, one may determine the *path process* $\mu_\pi = \int \mu_f d\pi$ by integrating over the realizations $(\boldsymbol{\phi}, z)$ corresponding to a self confirming equilibrium. μ_π is a measure

[7]It may seem strange at first that player can have beliefs that depend on the order of arrival in other lineages whose replacement times he cannot observe. The reason this is not as absurd as from first glance is that individuals do not "choose" their beliefs. Any notion of equilibrium places restrictions on what kinds of beliefs can arise. In particular, the very definition equilibrium requires that replacement times be conditioning events for the (equilibrium) process π, but not for the individuals themselves.

[8]Assumption (A2) could have been rephrased this way.

on sample paths induced by π. Every path process μ_π in which π satisfies assumptions (A1)-(A3) turns out to have the same asymptotic property in common interest games:

Theorem *Given any stage game G, there exists an SCE belief process π satisfying (A1)-(A3). Moreover, if G is any common interest game, there exists a $\hat{\delta} > 0$ and a $\hat{\gamma} > 0$ such that if $1 > \delta \geq \hat{\delta}$ and if $0 < \gamma \leq \hat{\gamma}$, then for any SCE belief process π satisfying (A1)-(A3), for any initial history h^0, and for μ_π-almost every h, there is a finite random time $t(h)$ such that $s(t) = s^*$ for all $t \geq t(h)$.*

Theorem 1 states first that SCE processes satisfying (A1)-(A3) exist. Second, in repeated common interest games with sufficiently patient individuals and with sufficiently little inertia, starting from any profile, there is some stopping time at which the Pareto dominant equilibrium s^* will be reached, and once there, society never departs. The proof is contained in Section 6.

Revisiting Ned and Bob

In obtaining this result, a good deal of structure has been imposed. It is worth revisiting the Example in Section 2 to understand exactly why the result fails when the main assumptions are dropped. There are three crucial properties: (1) unobservable revisions/replacements (2) belief mutation, and (3) self confirming equilibrium (rather than Nash).

1. **Unobservable Revisions/Replacements** Suppose that revisions/replacements were observed by all.[9] Then the transition $\{s^* \to \bar{s}\}$ can be shown to occur infinitely often. The reason is that punishments can now be conditioned on the entry times of new individuals. Hence, Bob Jr. can expect that at the date of Ned Jr.'s entry, Ned Jr. will depart from s^* (under threat of punishment) without contradicting Ned's beliefs that s^* is permanent.

   Naturally the unobservability assumption is restrictive. However, the assumption of complete information about entry times is also restrictive. A reasonable extension is to examine whether imperfect observability can restrict transitions from Pareto superior profiles.

2. **Belief Mutation** Suppose that beliefs do not mutate. Then the transition $\{\bar{s} \to s^*\}$ need not occur. In particular, without either of Assumptions (A2) or (A3) all subsequent individuals may have the same beliefs as Bob and Ned, and so the two

[9]Naturally, if revision dates were not stochastic then this would have the same affect as their being observed.

lineages remain in $\bar{s}$ permanently. More specifically, one might ask why the full support assumption without (A2) is not itself sufficient for the result. One might imagine a process in which after every time t a neighborhood of the Table 1 beliefs occurs with positive probability. However, without (A2), the process may place diminishing probability mass on the neighborhood as time passes so that infinite recurrence is not guaranteed. The Independent Increments property (A2) guarantees that if the neighborhood occurs once, it occurs infinitely often.[10]

3. **SCE** Finally, suppose that the Nash equilibrium concept rather than SCE is considered. Then, again, the transition $\{\bar{s} \to s^*\}$ need not occur. In particular, in a Nash equilibrium Ned correctly anticipates Bob Jr.'s behavior on the out-of-equilibrium event that Ned himself switches to s_N^*. Therefore, Ned should then switch to s_N^* anticipating that Bob Jr. will follow. This means that either Bob and Ned could not have believed that $\bar{s}$ was permanent, or Bob Jr. and Ned Jr. could not have believed that a departure was possible. The result then fails since if the current generation remains unconditionally in $\bar{s}$, all subsequent generations must share the same beliefs and remain there as well.

5 Related Literature

There is an important distinction between this and previous stability results in two models of rational agents where beliefs evolve, at least implicitly, in similar demographic frameworks. Matsui and Rob (1991) consider forward looking behavior in a similar model of population replacement in a repeated, random matching setting. Their requirement that behavior be rationalized by some belief places restrictions on admissible behavioral mutations. Lagunoff and Matsui (1995) explicitly incorporate belief arrivals in the same demographic structure as here. In that paper stability and path properties of differing mechanisms for public goods provision are studied. Both of these papers place few restrictions on the belief process. Global absorbtion is possible in those models only because there is some behavioral inertia relative to discounting, i.e. if $\gamma(1-\delta)$ is large, but not too large. In all three models if $\gamma(1-\delta)$ is large enough, all static Nash equilibria are absorbing states. The difference occurs in the fully rational case, i.e., when $\gamma(1-\delta)$ is small. In those models the process cycles between all profiles. By contrast, global stability of Pareto dominance is obtained here when $\gamma(1-\delta)$ is small. The trade-off is that while the results here do not rely on inertia, this paper requires that beliefs be consistent with a self confirming equilibrium of the repeated game. As a result, this model has very distinct path properties.

[10]The proof utilizes Laws of Large numbers arguments (the Borel Cantelli Lemmas), which requires independence. Full support conditions by themselves are not sufficient.

In addition to the two papers cited above, the dynamics of this model tend to work in a way that is, roughly speaking, reminiscent of some familiar *evolutive approaches*.[11] Both Kandori, Mailath, and Rob (KMR) (1993), and Foster and Young (FY) (1991), for example, draw predictions in repeated coordination games from the interaction of myopic best response dynamics ("adaptation") and random mutation in behavior of randomly matched individuals drawn from a large population. They show that the proportion of time spent by the process on certain equilibrium profiles converges to one as the rate of mutation on behavior converges to zero. Two important differences between this model and these are that, first, mutation in this model operates not directly on behavior itself but on the prior beliefs of new individuals as they enter society; second, because we consider self confirming equilibria of the fully repeated game, individuals here are not assumed to be myopic. A stronger "selection" result is obtained since global convergence to Pareto dominant behavior here takes place almost surely rather than in probability.

The trade-off with the evolutive literature is that by retaining the assumption of equilibrium in some form, a certain degree of coordination of beliefs in this paper is already presumed. It is not clear generally to what extent can the coordination of beliefs assumed in equilibrium be relaxed.[12] Ideally, assumptions that include both foresight and mutation allow both coordinated punishments for social norms, and eventual departure from norms enforced only with excessively sophisticated punishments (i.e. those that operate on high orders of counterfactuals). Presumably, there is a large middle ground of processes that involve foresight and mutation but are less extreme than the SCE and Independent Increments assumptions used here. The hard problem is identifying which kinds of restrictions are most plausible for which environments. Future efforts might be directed toward applications to resolve this issue.

6 Proof of Theorem

Existence

The proof of existence of a process satisfying the assumptions is constructive. Fix an initial history h^0. Recall that F_i denotes the set of strategies $f_i : H \times \{0,1\} \to S_i$ for an individual from lineage i. Let $\mathcal{D}_i$ denote the Borel sets generated by the weak topology on

[11]Binmore (1987) distinguishes between the *evolutive* and *eductive* approaches, the latter defining introspective behavior that may or may not include perfect foresight assumed in equilibrium models. A recent survey of the evolutive approach is found in Mailath (1993).

[12]Using a requirement similar to that found in Kalai and Lehrer (1993a), the Pareto dominant outcome can be shown again to be globally absorbing in a model of rational learning. The conditions, however, are quite strong. See an earlier version of this paper: Lagunoff (1995).

F_i. The open sets of this topology are generated by sets of the form

$$D_\eta(\hat{f}_i) \equiv \left\{\hat{f}_i \in F_i : \left|U_i^m(f|h^0) - U_i^m(\hat{f}_i, f_{-i}|h^0)\right| < \eta, \forall m = 1, \ldots, M\right\}$$

where $\eta > 0$ and U_i^m, $m = 1, \ldots, M$ is any finite collection of bounded continuous functions. $\mathcal{D}_{-i}$ is the product sigma algebra of $F_{-i} = \times_{j \neq i} F_j$.

We define the set of possible beliefs, Φ_i, as the space of probability measures on $(F_{-i}, \mathcal{D}_{-i})$. Now endow Φ with the weak* topology which can be generated by open sets of the form:

$$N_\eta(\phi) \equiv \left\{\hat{\phi} \in \Phi_i : \left|E_\phi[U_i^m(f|h^0)] - E_{\hat{\phi}}[U_i^m(f|h^0)]\right| < \eta, \forall m = 1, \ldots, M\right\} \tag{4}$$

where $\eta > 0$ and U_i^m, $m = 1, \ldots, M$ is any finite collection of bounded continuous functions. A standard Theorem of Alouglu shows that Φ_i is compact in this topology. Let $\mathcal{G}$ denote the Borel σ-algebra generated by the product topology on $\mathbf{\Phi} \equiv \times_i \times_j \Phi_{ij}$.

To construct an SCE belief process π satisfying (A1)-(A3), fix an arrival realization z. A *conditional belief process* $\pi(\cdot|z)$ is a measure on $(\mathbf{\Phi}, \mathcal{G})$. Let $\bar{z}(t) = \sum_i z_i(t)$ denoting the latest individual to be born by time t. Define the subsequence $\{k_m\}$ of $k = 1, 2, \ldots$ recursively by

$$\begin{aligned} \bar{z}(0) &= k_1 \\ \bar{z}(\tau(k_1)) &= k_2 \\ \bar{z}(\tau(k_2)) &= k_3 \\ &\vdots \end{aligned} \tag{5}$$

Each successive k_m denotes a complete turnover in the population of all contemporaries of individual k_{m-1}. Now suppose that entrant k_m is from lineage i. Suppose that k_m remains alive until time $t = \nu_m$.[13] The set of admissible beliefs for this individual is given by

$$\bar{\Phi}_{k_m} = \left\{\phi_{k_m} \in \Phi_i \middle| \exists \text{ SCE } f^*, \ \forall h^{\nu_m}, \int_{f_{-i}} \mu_{f \backslash f_k^*}(C(h^{\nu_m})|\, h) d\phi_{k_m} = \mu_{f^*}(C(h^{\nu_m})|\, h), \ \mu_f - a.e.\ h\right\} \tag{6}$$

The set $\bar{\Phi}_{k_m}$ is the subset of beliefs in Φ for m each of which is consistent with some SCE that continues until history h^{ν_m} at which time individual k_m dies and is replaced.

To satisfy (A1)-(A3) we now construct the process iteratively from the marginals $\pi_k(\cdot|z^k)$ defined on Φ_k. Without loss of generality, reorder the subsequence $\{k_m\}$ so that $k_m = m$. Consider $m = 1$. Since Φ_1 is compact, a measure $\pi_1(\cdot|z^1)$ with full support on $\bar{\Phi}_1$ exists.

[13] Time ν_m is not to be confused with $\tau(k_m)$ which is the time that *all* individuals up to and including entrant k_m have died. Generally, $\tau(k_m) > \nu_m$.

Now suppose that the product measure $\pi_1 \times \cdots \times \pi_m$ has full support on $\bar{\Phi}_1 \times \cdots \times \bar{\Phi}_m$. Inductively, define the product measure $\pi_1 \times \cdots \times \pi_{m+1}$ on $\bar{\Phi}_1 \times \cdots \times \bar{\Phi}_{m+1}$ by

$$\pi_1 \times \cdots \times \pi_{m+1}(A_1 \times \cdots \times A_{m+1} | z^{m+1}) \;=\; \pi_1 \times \cdots \times \pi_m(A_1 \times \cdots \times A_m | z^m) \cdot \pi_{m+1}(A_{m+1} | z^{m+1})$$

where A_m is a Borel subset of Φ_k. This construction satisfies the Independent Increments Assumption (A2). The Kolmogorov Existence Theorem[14] asserts that these finite dimensional distributions extend to a measure, $\pi(\cdot|z)$, on $(\mathbf{\Phi}, \mathcal{G})$.

To show that the measure π is an SCE process, we can construct for almost every realization of π at least one SCE consistent with that realization. For each $\phi_m \in \bar{\Phi}_m$ and each $m = 1, 2, \ldots$, let f^m denote a self confirming equilibrium consistent with ϕ_m. By the definition in (6), one such f^m must exist. Now let f^* satisfy $\tilde{s}(f^*) = \tilde{s}(f^m)$ in the time interval $[\tau(k_{m-1}), \tau(k_m)\ )$. That is, f^* gives rise to the same path as f^m in the time from the death of the first k_{m-1} entrants to the time of death of the first k_m entrants. By construction, each individual's choice f^*_{ij} is a best response to f^*_{-i}. Since the population turnover from $\tau(k_{m-1})$ to $\tau(k_m)$ is complete, there are no forecasting inconsistencies along the path $\tilde{s}(f^*)$. Finally, since the process $\pi(\cdot|z)$ assigns full support to each of the $\bar{\Phi}_k$, assumption (A3) is satisfied. □

Global Absorbtion

Now let G denote a common interest game. To prove global convergence of the path process to the Pareto dominant outcome, it will suffice to prove that for a large enough δ and small enough γ: for a set of histories h having μ_π-probability one there is an unbounded sequence of random stopping times $\{t_m(h)\}_{m=1}^{\infty}$ with $t_1(h) < t_2(h) < \cdots$ such that $s(t_m(h)) = s^*$ for all m, and for some m, $s(t) = s^*$ for all $t \geq t_m(h)$.

For the sequel we will consider the subsequence $\{k_m\}$, which, recall, is defined in (5) as a complete turnover in the population of entrants. We first prove the following Lemma concerning the *Baire* subsets of Φ. These are Borel sets that can be generated by sets of the form $\{\phi \in \Phi : \ H(\phi) > \alpha\}$ where $\alpha > 0$ and H is a real valued, continuous, measurable function defined on Φ.[15]

Lemma 1 *Fix z. Then for any nonempty open Baire set $B \subseteq \bar{\Phi}_{k_m}$,*

$$\pi(\{\phi_{k_m} \in B\},\ i.o.\ m \;\Big|\; z) \;=\; 1 \tag{7}$$

[14]See Billingsley (1986), Theorem 36.2, p.510.

[15]In the compact Hausdorff space Φ, Baire sets are also Borel sets, however, the converse is not necessarily true. See Royden (1988).

Lemma 1 states that with probability one it is the case that for infinitely many m, the belief of entrant k_m lies in the set B.

proof of Lemma 1 Observe that by the Independent Increments and Maximal Support Assumptions, (A2) and (A3), the support of $\pi(\cdot|\phi^{k_m}, z)$ on beliefs $\phi_{k_{m+1}}$ is the entire space $\bar{\Phi}_{k_{m+1}}$. By a theorem in Royden (p. 351, exercise 24a), open Baire sets with nonempty intersection with the support of π have positive measure. It follows that

$$\pi(\{\phi_{k_{m+1}} \in B\}\big| z) \;=\; \alpha \;> 0$$

Let $A_{k_{m+1}} \equiv \{\phi_{k_{m+1}} \in B\}$. By (A2) the $\{A_{k_m}\}$ as defined are mutually independent, and $\sum_m \pi(A_{k_m}|z) = \infty$. Therefore, the second Borel-Cantelli Lemma implies $\pi(A_{k_m},\ i.o.\ m\big|z) = 1$. Hence the result. □

For each i, let $u_i^* \equiv u_i(s^*)$. Now fix some $\epsilon > 0$ and suppose that the kth entrant is from lineage i. Define the open set

$$B_\epsilon^i(h) = \{\phi \in \bar{\Phi}_k : \ \exists\ SCE\ f^*,\ E_\phi[U_i(f^*|\ h)] > u_i^* - \epsilon\} \tag{8}$$

In (8), $B_\epsilon^i(h)$ is the neighborhood of beliefs that puts an individual from lineage i within ϵ of u_i^* under some SCE profile f^*.

Lemma 2 *For each i and each $m = 1, 2, \ldots,$ $B_\epsilon^i(h^{\tau(k_m)})$ is an open Baire set. Moreover, there is a $\bar{\delta}_\epsilon \in (0,1)$ and $\bar{\gamma}_\epsilon \in (0,1)$ such that if $\delta \geq \bar{\delta}_\epsilon$ and $\gamma \leq \bar{\gamma}_\epsilon$, then $B_\epsilon^i(h^{\tau(k_m)})$ is nonempty.*

proof of Lemma 2 Observe that as a function of ϕ, the function $E_\phi[U_i(f^*|h^{\tau(k_m)}]$ is continuous and $\mathcal{G}$-measurable. Therefore, $B_\epsilon^i(h^{\tau(k_m)})$ is an open Baire set. It suffices to prove that $B_\epsilon^i(h^{\tau(k_m)})$ is nonempty for large enough δ and small enough γ.

Fix i. Let $\rho_i[t]$ be a random variable denoting i's first decision point after time t. Now fix m. Without loss of generality, assume that there are q individuals who do not choose s^* at the time of replacement $\tau(k_m)$ of the first k_m entrants. That is, $s_i(\tau(k_m)) \neq s_i^*,\ i = 1\ldots,q$. Now define the sets:

$$\begin{aligned}
D_m^1 \;&=\; \{s_1(\rho_1[\tau(k_m)]) = s_1^* \ \text{ implies } \ s_2(\rho_2[\rho_1[\tau(k_m)]]) = s_2^*\} \\
D_m^2 \;&=\; \{s_2(\rho_2[\rho_1[\tau(k_m)]]) = s_2^* \ \text{ implies } \ s_3(\rho_3[\rho_2[\rho_1[\tau(k_m)]]]) = s_3^*\} \\
&\vdots \\
D_m^q \;&=\; \{s_i(\rho_q[\rho_{q-1}\ldots,\rho_1[\tau(k_m)]\ldots,]]) = s_i^*,\ \forall i\}
\end{aligned}$$

The set D_k^1, for instance, is the event that $i = 1$ choosing s_1^* at the first opportunity after $\tau(k_m)$ implies that $i = 2$ will choose s_2^* at the first opportunity after $\rho_1[\tau(k_m)]$. Now define the sets $\{C_m^i\}_{i=1}^q$ by

$$\begin{aligned} C_m^1 &= \cap_{i=1}^q D_m^i \\ C_m^2 &= \cap_{i=2}^q D_m^i \\ &\vdots \end{aligned} \tag{9}$$

so that C_m^1 is the event that $i = 1$'s switch to s_1^* induces an immediate chain reaction of all $q - 1$ other switches to s^*.

Recall that we have supressed the notational dependence of the path process μ_{f^*} on parameters γ and δ. Given ϵ, we now show that there is a retention probability δ_ϵ and a revision probability γ_ϵ, and a belief ϕ^* such that

$$\int_{f_{-i}} \mu_{f\backslash f^*}(C_m^i \big| \, h) d\phi^* = 1 \text{ implies } \phi^* \in B_\epsilon^i(h), \text{ and } f_i^* \text{ is optimal against belief } \phi^* \tag{10}$$

Observe first that the behavior profile reaches s^* in event C_m^i if i chooses/keeps s_i^*. Moreover, the expected time, $\rho_i[\rho_{i-1}[\ldots\rho_1[\tau(k_m)]\ldots]]$, to reach s^* from time $\tau(k_m)$ monotonically increases in γ. Therefore, if i chooses/keeps s_i^* under f_i^*, then $E_{\phi^*}[U_i(f_i^*, f_{-i}| \, h)] \to u^*$ as $\delta \to 1$ and as $\gamma \to 0$. Property (10) is then satisfied for some $0 < \delta_\epsilon < 1$ and $1 > \gamma_\epsilon > 0$.

Note that if an individual from lineage i has belief ϕ^* he adopts/keeps s_i^* without delay. Since s^* is a Nash equilibrium of the stage game, remaining in s^* is a self confirming equilibrium of the repeated game. Now let $\bar{\delta}_\epsilon$ and $\bar{\gamma}_\epsilon$ denote the infimum and supremum, resp. which satisfy the property in (10). Hence, for each $\delta \geq \bar{\delta}_\epsilon$ and each $\gamma \leq \bar{\gamma}_\epsilon$, $\phi^* \in B_\epsilon^i(h^{\tau(k_m)})$. We conclude that $B_\epsilon^i(h^{\tau(k_m)})$ is nonempty. □

By applying Lemmas 1 and 2, it follows that for each z,

$$\pi(B_\epsilon^i(\, h^{\tau(k_m)}), \ i.o. \ m \, \big| \, z) \ = \ 1$$

That is, within lineage i, there are infinitely many individuals who acquire beliefs in a set $B_\epsilon^i(h)$. By taking ϵ sufficiently small, if $\phi \in B_\epsilon^i(h^{\tau(k_m)})$ then i's expected payoff is approximately u_i^*. It must follow that there is some time $t_m \geq \tau(k_m)$ at which $s(t_m) = s^*$. That is, behavior hits the optimal profile s^* at time t_m. But since $\tau(k_1) < \tau(k_2) < \cdots \infty$ we have established an infinite set of stopping times $\{t_m\}$ with $t_1 < t_2 < t_3 < \cdots \infty$ on a set of μ_π-probability one histories such that $s(t_m) = s^*$ for each m, and for any t there is some m for which $t_m > t$. This proves the first part.

Given the construction of stopping times above, Fix an entrant k from lineage i and a stopping time t_m. Let $A_{t_m} = \{s(t) \neq s^*, \ \ \text{for some} \ \ t \geq t_m\}$ denoting the event of a departure from s^* after t_m. Define the belief $\phi^* \in \bar{\Phi}_k$ to satisfy

$$\int_{f_{-i}} \mu_{f \setminus f_i^*}(A_{t_m} \Big| \, h^{t_m}) d\phi^* = 0$$

where strategy f_i^* is a best response to this belief. This belief places probability zero on a departure from s^* after t_m. Clearly, a best response to this belief is to keep s_i^* at least until one observes a defection from s^*. Assume that f_i^* is such a best response. If individuals from all lineages have this belief, then $f^* = (f_i^*)_{i=1}^n$ is clearly a self confirming equilibrium.

Given some real number $\eta > 0$, define the neighborhood

$$N^i_{\eta,m}(\phi^*) \;\equiv\; \left\{ \phi \in \Phi_i : \int_{f_{-i}} \mu_{f \setminus f_i^*}(A_{t_m} \Big| \, h^{t_m}) d\phi < \eta \right\} \tag{11}$$

Lemma 3 *For any $\eta > 0$, $N^i_{\eta,m}(\phi^*)$ is a nonempty open Baire set.*

proof of Lemma 3 Clearly $N_{\eta,m}(\phi^*)$ is nonempty as it contains ϕ^*. As with Lemma 2 observe that as a function of ϕ, the function $\int_{f_{-i}} \mu_{f \setminus f_i^*}(A_{t_m} \Big| \, h^{t_m}) d\phi$ is continuous and $\mathcal{G}$-measurable. But then the definition in equation (11) conforms with the definition of a Baire set. □

Observe then by Lemma 1, given an arrival realization z the set $N^i_{\eta,m}(\phi^*)$ must re-occur for infinitely many m under $\pi(\cdot|z)$.

Now observe that if for some individual i', $f_{i'}^*(h^t, 1) \neq s_{i'}^*$ for some h^t in a self confirming equilibrium path, then any contemporary of i' must have a belief ϕ satisfying $\int_{f_{-i}} \mu_{f \setminus f_i^*}(A_{t_m} \Big| \, h^{t_m}) d\phi = 1$ since A_{t_m} includes tail events. That is, if there is any departure from s^* at any time after t_m, then one's beliefs must assign probability one to A_{t_m}. Hence, for η small, then for entrant k, $\phi_k \in N^i_{\eta,m}(\phi^*)$ implies that for any history h^t such that k is alive at time t, $\int_{f_{-i}} \mu_{f \setminus f_i^*}(A_{t_m} \cap C(h^t) \Big| \, h^{t_m}) d\phi = 0$. That is, k places probability zero on a departure from s^* while he remains alive. Clearly a best response for k is to keep his action in s_i^* believing with certainty that others do so as well. It follows from the SCE assumption for profile f^* that for all histories h^t in which k remains alive,

$$\mu_{f^*}(A_{t_m} \cap h^t) = 0 \tag{12}$$

That is, the true distribution μ_{f^*} must assign 0 probability of a departure from s^* by time t if k had not been replaced. However, since k's death is unobservable to those in other lineages. the other lineages cannot condition on the replacement time of k.

Now suppose that some entrant ℓ chooses $f_\ell(h,1) \neq s^*_\ell$ after some history h in the support of μ_{f^*}. Let $I(k)$ denote the set of individuals from other lineages who enter when k is alive. Then, by equation (12), any contemporary of k, i.e. any individual k' where $k' \in I(k)$ cannot be the first to depart from s^* without violating the SCE assumption on a set of paths having positive probability. It follows that $\ell \notin I(k)$. For ℓ's departure to be consistent with SCE, it must be the case that ℓ enters the world after all contemporaries of k have died. Since $k' \in I(k)$ never initiates a switch from s^*, no contemporary of k' believes that k' ever initiates a switch. Since this is true for all $k' \in I(k)$ it follows from the SCE requirement that $\ell \notin I(I(k))$ where $I(I(k))$ denotes the contemporaries of the contemporaries of k. Inducting on this definition, it follows that $\ell \notin I(I(\cdots I(k)\cdots))$. That is, since no contemporary of no contemporary of, ... individual k ever initiates a switch, individual ℓ cannot belong to these sets. This implies that $\ell = \infty$, contradicting the fact that some finite ℓ switches from s^* after some finite history h.

We have now established that there is some stopping time that hits beliefs $\phi \in N^i_{\eta,m}(\phi^*)$ for all k alive when s^* is played; given these beliefs, no future individual will ever have a belief that entails that he defect from s^*. This proves the second part, and we conclude the proof. □ □

A Remark about the Proof It is worth re-emphasizing why the transition $\{\bar{s} \to s^*\}$ occurs infinitely often, yet the transition $\{s^* \to \bar{s}\}$ does not occur at all after some stopping time with probability one. First, recall that beliefs from the set $N_{\eta,m}(\phi^*)$ re-occur for infinitely many m. At one such time, no individual departs s^* without the threat of punishment, i.e., subsequent departures from s^* by others. Since revision times and birth/death times are not observed across lineages, departures from s^* must be conditioned on calendar dates. However, due to the stochastic switching times, these departures are in-equilibrium events. Therefore, such departures from s^* must violate beliefs in $N_{\eta,m}(\phi^*)$ of those who may still be alive at the threatened departure date. Hence, SCE is violated. On the other hand, starting from $\bar{s}$ such calendar time departures are not required to induce an individual to unilaterally switch. Hence, once those with beliefs in $N_{\eta,m}(\phi^*)$ exit the population, and the "right" belief sequence occurs in their replacments, one of the replacements may initiate a unilateral departure from $\bar{s}$ which induces others to follow.

References

[1] Aumann, R. and S. Sorin (1989), "Cooperation and Bounded Recall," *Games and Economic Behavior*, 1: 5-39.

[2] Battigalli, P. and P. Guatoli (1988), "Conjectural Equilibrium," mimeo.

[3] Billingsley, P. (1986), *Probability and Measure*, John Wiley and Sons: New York.

[4] Binmore, K. (1987), " Modeling Rational Players: I," *Economics and Philosophy*, 3:179-214.

[5] Blackwell, D. and L. Dubins (1962), "Merging of Opinions with Increasing Information," *Annals of Mathematical Statistics*, 33: 882-86.

[6] Blume, L. (1993), "The Statistical Mechanics of Strategic Interaction," *Games and Economic Behavior*, 5: 387-424.

[7] Blume, L. and D. Easley (1992) "Rational Expectations and Rational Learning," mimeo, Cornell University, July.

[8] Foster, D. and H. P. Young (1990), "Stochastic Evolutionary Game Dynamics," *Theoretical Population Biology*, 38: 219-32.

[9] Fudenberg, D. and D. K. Levine (1993a), "Self-Confirming Equilibrium," *Econometrica*, 61:523-45.

[10] Fudenberg, D. and D. K. Levine (1993b), "Steady State Learning and Nash Equilibrium," *Econometrica*, 61:547-73.

[11] Fudenberg, D. and E. Maskin (1986), "The Folk Theorem in Repeated Games with Discounting or with Incomplete Information," *Econometrica*, 54: 533-56.

[12] Kalai, E. and E. Lehrer (1993a), "Rational Learning Leads to Nash Equlibrium," *Econometrica*, 61: 1019-46.

[13] Kalai, E. and E. Lehrer (1993b), "Subjective Equilibrium in Repeated Games," *Econometrica*, 61: 1231-40.

[14] Kandori, M., G. Mailath, and R. Rob (1993) "Learning, Mutation, and Long Run Equilibria in Games," *Econometrica*, 61: 29-56.

[15] Kurz, M. (1993), "On the Structure and Diversity of Rational Beliefs," mimeo, Stanford University, February.

[16] Lagunoff, R. (1997), "On the Dynamic Selection of Mechanisms for Public Projects," *Journal of Economic Dynamics and Control*, forthcoming.

[17] Lagunoff, R. (1995), "On the Evolution of Pareto Optimal Behavior in Repeated Coordination Problems," mimeo, University of Pennsylvania, April.

[18] Lagunoff, R. and A. Matsui (1995) Evolution in Mechanisms for Public Projects," *Economic Theory*, 6: 195-223.

[19] Mailath, G. (1992), "Introduction: Symposium on Evolutionary Game Theory, *Journal of Economic Theory*, 57: 259-77.

[20] Matsui, A. (1993), "Evolution and Rationalizabilty," mimeo, University of Pennsylvania, March.

[21] Matsui, A. and R. Rob (1991), "Evolution and Rationality in Societal Games," CARESS Working Paper #91-31, University of Pennsylvania.

[22] Matsui, A. and K. Matsuyama (1993), "An Approach to Equilibrium Selection," forthcoming in *Journal of Economic Theory*.

[23] Royden, H.L. (1988), *Real Analysis*, Macmillan: New York.