Step-by-Step Evolution
with
State-Dependent Mutations[*]

Thomas Norman
Nuffield College, Oxford

April 22, 2003

Abstract

This paper considers the application of Ellison's (2000) "Radius-Modified Coradius" Theorem to models of evolution with state-dependent mutations. A reformulated theorem is presented, with a crucial role played by the most probable evolutionary paths between states. The form of such paths is liable to change outside of the uniform mutations case, with concomitant effects on both long-run selection and expected waiting times. An algorithm for finding these paths is offered, and used to confirm the optimality of "step-by-step" evolution.

1

1 Introduction

Since its inception in the works of Foster and Young (1990), Kandori, Mailath, and Rob (1993), and Young (1993), the field of stochastic adjustment dynamics has been a lively and controversial one. Perturbing the deterministic population dynamics of evolutionary game theory with vanishingly small noise appeared initially to resolve the equilibrium selection problem from a general boundedly rational foundation. Moreover, the standard uniform mutation rate models of Kandori, Mailath, and Rob (1993), and Young (1993) (henceforth KMRY) - in which a player errs (relative to the unperturbed model) with a fixed probability ε each period - reduced the complexity of Freidlin and Wentzell's (1984) Markovian graph-theoretic techniques to simple "mutation-counting".

However, such models were quickly criticised on the grounds that the vanishing noise required for selection results implied unacceptably long transition times to "long-run equilibrium" (Ellison 1993). It was not until Ellison (2000) though that general analytical results on transition times were available; Ellison's (2000) "Radius-Modified Coradius" Theorem not only provided a new technique for characterising the long-run stochastically stable set of an evolutionary model, but also bounded the speed with which evolutionary change occurs. The intuition behind the theorem is that "if a social convention tends to *persist* for a long time after it is established and is sufficiently *attractive* in the sense of being likely to emerge relatively soon after play begins in any other state, then in the long run that convention will prevail most of the time"[1].

Two new measures are employed to exploit this intuition. First, the "radius" of the basin of attraction of a limit set (or a union of limit sets) Ω, $R(\Omega)$, is defined as the minimum number of "mutations" (ε-probability events) necessary to escape the basin of attraction of Ω. This radius provides a bound on the *persistence* of the set Ω. Second, the "coradius" of the basin of attraction of Ω, $CR(\Omega)$, is defined as the maximum over all other states of the minimum number of mutations necessary to reach Ω. This coradius can be shortened by incorporating the effect of "step-by-step" evolution: "large evolutionary changes will occur more rapidly if it is possible for the change to be effected via a series of more gradual steps between nearly stable states"[2]. To capture the increased speed of step-by-step evolution, a new measure is computed by subtracting from the coradius a correction term which depends on the number of intermediate steady states along the evolutionary path and the sizes of their basins of attraction. This "modified coradius", $CR^*(\Omega)$, provides a bound on the *attractiveness* of Ω.

Using these two measures, the "Radius-Modified Coradius" Theorem

---

[1] Ellison (2000), p.18, emphasis added.
[2] Ellison (2000), p.19.

shows that $R(\Omega) > CR^*(\Omega)$ is a sufficient condition for the long-run stochastically stable set to be contained in Ω, and that the expected wait until Ω is reached in this case is $O(\varepsilon^{-CR^*(\Omega)})$.

The aim of this paper is to consider the power of the "Radius-Modified Coradius" Theorem in the face of another line of criticism of stochastic evolutionary game theory, namely the arbitrariness of mutations occurring at a rate independent of the current state of the system. Such "state-independent mutations", embodied in the fixed mutation rate ε of KMRY and others, imply that players make mistakes (or experiment, etc.) with the same probability irrespective of the current strategy frequencies, and thus of the expected payoffs at stake. The effect of relaxing this assumption is dramatic: Bergin and Lipman (1996) demonstrate that, given any model of the effect of mutations, any invariant distribution of the "mutationless" process is close to an invariant distribution of the process with appropriately chosen small mutation rates. This implies that any strict Nash equilibrium of a strategic form game is selected under some suitably chosen mutation model. Intuitively, when the mutation rates vary across states, the size of the relevant basins of attraction is no longer enough to determine the long-run equilibrium; in particular, even though one basin of attraction may be smaller, it may be "deeper" in the sense that mutations out of this basin are less likely.

Bergin and Lipman's findings seemed to cast stochastic evolutionary game theory into a wilderness of indeterminacy far worse than the one which it had sought to escape. However, economically justified models of "state-dependent mutations" still offer the prospect of insight from the Markovian selection tools, albeit at the price of greater complexity. And indeed, many of the early results of the stochastic adjustment dynamics literature - most notably the pre-eminence of the risk-dominant equilibrium in 2×2 coordination games - can be confirmed in this new context (under certain conditions). Particular models of state-dependent mutations doing just this include Myatt and Wallace (1998), van Damme and Weibull (1998), Lee, Szeidl, and Valentinyi (2001), and Norman (2003). The most general analysis of the role of noise in stochastic adjustment dynamics is that of Blume (1999), who finds that the known stochastic stability results are preserved for the (large) class of noise processes satisfying a certain symmetry condition.

Ellison's (2000) "Radius-Modified Coradius" Theorem is framed using the "ε-cost" language of the uniform mutation rate model which Bergin and Lipman (1996) so forcefully criticised. However, as Ellison notes, his model "can easily accommodate state-dependent mutation rates with unbounded likelihood ratios (as in Bergin and Lipman (1996))"[3]. In this context though, the "cost" of a transition has a less clear interpretation; it still measures the order of probability of the transition, but this can no longer be ascertained

---

[3]Ellison (2000), p. 21.

by simply "counting mutations". As a consequence, the application of the theorem in models with state-dependent mutations is not immediately obvious.

The present paper makes this application explicit by modifying the specification of noise in Ellison's model, and hence reformulating the theorem in terms of the underlying transition probabilities. The simplicity of Ellison's theorem is sacrificed somewhat in making this step, but the same fundamental lessons emerge. However, the reformulated theorem highlights the crucial role of the most probable (or "optimal") evolutionary paths between states, and clouds Ellison's "step-by-step" effect on the speed of evolution outside of the uniform mutations case. Nonetheless, an algorithm for finding the optimal evolutionary paths is offered, and used to demonstrate the continued optimality of "step-by-step" evolution. This serves to clarify the precise sense in which Ellison's intermediate "steps" must be "intermediate".

The next section presents the essentials of Ellison's model, modified to emphasise the presence of state-dependent mutations. Section 3 then presents the reformulated theorem and compares it with the original. Ellison's exposition is followed very closely in order to facilitate comparison. Section 4 presents the optimal evolutionary path algorithm, and section 5 applies it to "step-by-step" evolution.

2   Preliminaries

Ellison's (2000) model is unchanged, except that the parameter ε and its associated "cost" function are replaced by a general noise model $g : \mathbb{R} \to \mathbb{R}$, in the sense of Blume (1999), which assigns choice probabilities to payoff differences. This is to emphasise the presence of state-dependent mutations. Ellison's (2000) Definition 1 thus becomes

Definition 1 *A* model of evolution with noise *is a triple* $(Z, P, g(\varpi; \sigma^2))$ *consisting of:*

1. *A finite set Z referred to as the* state space *of the model;*

2. *A Markov transition matrix P on Z;*

3. *A* noise model *$g(\varpi; \sigma^2)$ mapping payoff differences $\varpi(z)$, $z \in Z$, into choice probabilities, given a* noise variance *of σ^2. The noise model defines a family of Markov transition matrices $P(\sigma^2)$ on Z indexed by the parameter $\sigma^2 \in [0, \bar{\sigma}^2)$ such that:*

   (a) *$P(\sigma^2)$ is ergodic for each $\sigma^2 > 0$;*
   (b) *$P(\sigma^2)$ is continuous in σ^2 with $P(0) = P$.*

The thinking behind the noise model here is that trembles from strategies' payoffs (i.e., noise) are generated by a random variable with cumulative distribution function F, mean ν and variance σ^2. This *random utility*-style framework is intuitive for modelling evolution with state-dependent mutations, but it is not entirely general. It cannot, for instance, generate the standard uniform mutation rate model for every game (though it can for any given game). Blume (1999) instead parameterises noise models by a parameter β such that the variance around the best response decreases with $1/\beta$. This too is not general, but it does include the most popular parameterisations of noise and noise reduction; the uniform mutation rate model and the "log-linear model" of Blume (1993) and Brock and Durlauf (1995), for instance, both fit into Blume's scheme. Nonetheless, the results presented in this paper hold for either choice of parameterisation.

Now, property 3(b) in Definition 1 implies that the Markov process $(Z, g(\varpi; \sigma^2))$ converges to (Z, P) as noise vanishes ($\sigma^2 \to 0$). (Z, P) is simply the deterministic dynamic defined by the underlying game, the matching mechanism, the rules for strategy revision, and the default behavioural assumptions made of the players (most frequently best-response to some statistical frequency of play). The "recurrent classes"[4] or "limit sets" of this dynamic represent short- to medium-run equilibria of the system, and constitute the candidates for its long-run stochastically stable set (Young 1993). A given limit set of the unperturbed process (Z, P) is denoted L, whilst Ω describes a union of one or more such sets. $\mathcal{L}$ denotes the union of all of (Z, P)'s limit sets. The *basin of attraction* of Ω is denoted $D(\Omega)$, and is given by

$$D(\Omega) = \{z \in Z \mid \text{Prob}\ \{\exists T \text{ s.t. } z_t \in \Omega\ \forall t > T \mid z_0 = z\} = 1\}$$

This is the set of initial states from which the unperturbed Markov process converges to Ω with probability one.

Following Ellison, $W(x, Y, g(\varpi; \sigma^2))$ will denote the expected wait until a state belonging to the set Y is first reached given that play in the σ^2-perturbed model begins in state x. By examining $\max_{x \in Z} W(x, \Omega, g(\varpi; \sigma^2))$ when σ^2 is small one can address the issue of how quickly a system converges to its long-run stochastically stable set Ω. Of course, $W(\cdot)$ will in general tend to infinity as σ^2 goes to zero, but the speed of convergence can still be judged according to how quickly the waiting times increase as σ^2 vanishes.

Heavy use will also be made of the following related notation. $N(A, B, x)$ will denote the expected number of times states in A occur (counting the initial period if $x \in A$) before the process reaches B (not counting the process as having immediately reached B if $x \in B$) when the process starts at x. Meanwhile, $Q(A, B, x)$ will be the probability that A is reached before

---

[4]$\Omega \subset Z$ is a recurrent class of (Z, P) if $\forall w \in \Omega$, Prob $\{z_{t+1} \in \Omega \mid z_t = w\} = 1$, and if for all $w, w' \in \Omega$ there exists $s > 0$ such that Prob $\{z_{t+s} = w' \mid z_t = w\} > 0$.

B when the process starts at x (not counting what happens in the initial period if $x \in A$ or $x \in B$).

3   The Theorem

The previous section's seemingly minor modification to the noise mechanism employed by Ellison complicates the resulting theorem considerably (though not the underlying analysis, which is substantially unchanged). However, it delivers a reformulation of the "Radius-Modified Coradius" Theorem in terms of the evolutionary model's transition probabilities, rather than its "cost" function. This facilitates the application of the theorem in the presence of state-dependent mutations. The structure of Ellison's presentation is followed very closely in order to facilitate direct comparison. Hence, the bulk of the analysis is relegated to the Appendices, where the Lemmas are presented in the same order as in Ellison (2000).

Theorem 1 *Let $(Z, P, g(\varpi; \sigma^2))$ be a model of evolution with noise, $\mathcal{L}$ be the union of the limit sets of (Z, P), L a single limit set, Ω a union of limit sets, and $\{L_j\}_{j=i}^r$ the limit sets through which the most probable path from a given limit set L_i to Ω passes. Then*
    (a) the long-run stochastically stable set of the model is contained in Ω if

$$\max_{\omega \in \Omega} Q(Z - D(\Omega), \Omega, \omega) = o\left(\min_{L_i \in \mathcal{L} - \Omega} \left(\max_{z \in L_i} Q(Z - D(L_i), L_i, z) \right.\right.$$

$$\left.\left. \prod_{j=i}^{r-1} \frac{\max_{z'' \in L_j} Q(L_{j+1}, \mathcal{L} - (L_j \cup L_{j+1}), z'')}{\max_{z' \in L_j} Q(Z - D(L_j), L_j, z')} \right) \right)$$

(b) for any $y \notin \Omega$,

$$W(y, \Omega, g(\varpi; \sigma^2)) = O\left(\left[\min_{L_i \in \mathcal{L} - \Omega} \left(\max_{z \in L_i} Q(Z - D(L_i), L_i, z) \right.\right.\right.$$

$$\left.\left.\left. \prod_{j=i}^{r-1} \frac{\max_{z'' \in L_j} Q(L_{j+1}, \mathcal{L} - (L_j \cup L_{j+1}), z'')}{\max_{z' \in L_j} Q(Z - D(L_j), L_j, z')} \right) \right]^{-1} \right)$$

as $\sigma^2 \to 0$.

    Proof. Following Ellison (2000), Lemma 1 in the Appendix presents the following characterisation of the steady-state distribution of the Markov process $(Z, g(\varpi; \sigma^2))$:

$$\frac{\mu^{\sigma^2}(y)}{\mu^{\sigma^2}(\Omega)} = \frac{N(y, \Omega, y)}{\sum_{\omega \in \Omega} Q(\omega, \Omega - \omega, y) N(\Omega, y, \omega)}$$

As Ellison notes, the numerator is bounded above by $W(y, \Omega, g(\varpi; \sigma^2))$, so it will suffice for parts (a) and (b) of the theorem to show that

$$W(y, \Omega, g(\varpi; \sigma^2)) = O\left(\left[\min_{L_i \in \mathcal{L} - \Omega}\left(\max_{z \in L_i} Q(Z - D(L_i), L_i, z)\right.\right.\right.$$
$$\left.\left.\left.\prod_{j=i}^{r-1} \frac{\max_{z'' \in L_j} Q(L_{j+1}, \mathcal{L} - (L_j \cup L_{j+1}), z'')}{\max_{z' \in L_j} Q(Z - D(L_j), L_j, z')}\right)\right]^{-1}\right) \qquad \forall y \notin \Omega \qquad (1)$$

and

$$1/N(\Omega, y, \omega) = O\left(\max_{\omega \in \Omega} Q(Z - D(\Omega), \Omega, \omega)\right) \qquad \forall \omega \in \Omega \qquad (2)$$

Lemmas 2 and 6 in the Appendix contain these two results.

It is worth noting that it follows from Lemmas 2 and 3 that the time necessary to leave the basin of attraction of a single limit set is $W(l, Z - D(L), g(\varpi; \sigma^2)) \sim O(1/\max_{z \in L} Q(Z - D(L), L, z))$ for any state l belonging to a limit set L. ∎

Despite its notational complexity, this theorem has a clean interpretation as the natural formulation of Ellison's "Radius-Modified Coradius" Theorem in a state-dependent mutations setting. To see this, note first that $\max_{\omega \in \Omega} Q(Z - D(\Omega), \Omega, \omega)$ is just the probability of the most probable way of escaping Ω's basin of attraction. Hence it is the analog of Ellison's radius,

$$R(\Omega) = \min_{(z_1, \ldots, z_T) \in S(\Omega, Z - D(\Omega))} c(z_1, \ldots, z_T)$$

where $(z_1, \ldots, z_T)$ is a path[5], $S(X, Y)$ is the set of all paths from X to Y, and $c(z_1, \ldots, z_T)$ is the "cost"[6] of the given path. The radius $R(\Omega)$ is just the exponent of ε that gives the order of the most probable ("least cost") way of escaping Ω's basin of attraction.

Meanwhile,

$$\min_{L_i \in \mathcal{L} - \Omega}\left(\max_{z \in L_i} Q(Z - D(L_i), L_i, z)\right.$$
$$\left.\prod_{j=i}^{r-1} \frac{\max_{z'' \in L_j} Q(L_{j+1}, \mathcal{L} - (L_j \cup L_{j+1}), z'')}{\max_{z' \in L_j} Q(Z - D(L_j), L_j, z')}\right) \qquad (3)$$

---

[5] A "path" from a set X to a set Y is defined to be a finite sequence of distinct states $(z_1, \ldots, z_T)$ with $z_1 \in X$, $z_t \notin Y$ for $2 \le t \le T - 1$, and $z_T \in Y$.

[6] The "cost" function measures any given transition's order of probability. In the uniform mutations model, this involves simply counting the number of "mutations" (ε-probability events) required to effect the given transition. Hence, a lower cost $c(z_1, \ldots, z_T)$ implies a higher probability $\varepsilon^{c(z_1, \ldots, z_T)}$ for the required mutations.

is the analog of Ellison's modified coradius,

$$CR^*(\Omega) = \max_{z \notin \Omega} \left(\min_{(z_1,\dots,z_T) \in S(z,\Omega)} \left(c(z_1,\dots,z_T) - \sum_{j=2}^{r-1} R(L_j) \right) \right)$$

In particular, $\max_{z \in L_i} Q(Z - D(L_i), L_i, z)$ in (3) is the most probable way of escaping L_i's basin of attraction, whilst each term in the product is simply the maximum probability of transition to L_{j+1} divided by the maximum probability of leaving L_j's basin of attraction. Note that $\max_{z'' \in L_j} Q(L_{j+1}, \mathcal{L} - (L_j \cup L_{j+1}), z'')$ is the analog of Ellison's "minimum cost" of transition from L_j to L_{j+1}, $C(L_j, L_{j+1}) = \min_{(z_1,\dots,z_T) \in S(L_j,L_{j+1})} c(z_1,\dots,z_T)$. Since $\max_{z \in L_i} Q(Z - D(L_i), L_i, z)$ is the analog of $R(L_i)$ whilst $\max_{z' \in L_j} Q(Z - D(L_j), L_j, z')$ is the analog of $R(L_j)$, the analogy between (3) and Ellison's modified coradius becomes clear.

Recognising these relationships, the parallel between the two theorems is evident. Specialising to the uniform mutations model, the sufficient condition for Ω's long-run stochastic stability in part (a) of the above theorem becomes $\varepsilon^{R(\Omega)} = o(\varepsilon^{CR^*(\Omega)})$. This of course holds precisely when Ellison's condition, $R(\Omega) > CR^*(\Omega)$, is true. For the waiting time in part (b), meanwhile, the expression in (3) is given by $\varepsilon^{CR^*(\Omega)}$ under uniform mutations, so that the parallel is even more readily apparent.

Thus, unsurprisingly, the same basic lessons emerge from the reformulated theorem as from the original. However, two novelties also emerge. First, the effect of "step-by-step" evolution is no longer immediately obvious from the reformulated theorem; it is not clear whether any given intermediate limit set should be passed through in the "step-by-step" fashion. This points up the need to determine the most probable evolutionary path before applying Theorem 1.

Second, and relatedly, the most probable way of escaping a given basin of attraction emerges as crucial for both long-run selection and expected waiting times. In the uniform mutation rate model, "direct jumps"[7] out of basins of attraction are generally most probable; indeed, it is this result which delivers the simple "mutation-counting" approach of models with state-independent mutations. In the state-dependent mutations setting, by contrast, it is no longer clear that direct jumps are "optimal" in this sense, further complicating the results of Theorem 1. This underlines the need to investigate optimal evolutionary paths in general models of state-dependent mutations.

---

[7] By "direct jumps" is meant just enough simultaneous mutations to move between the two states concerned in one period.

4   An Algorithm

Finding the most probable evolutionary path from state 0 to a given state z is a problem of combinatorial optimization.[8] Let $G = (V, P)$ be a directed graph with set $V = Z = \{0, 1, 2, \ldots, N\}$ of vertices, where each arc ij is weighted by the transition probability p_{ij} from the perturbed Markov matrix $P(\sigma^2)$. This digraph is *connected*[9] by virtue of the irreducibility of the perturbed Markov process. Unfortunately, the longest path problem for a cyclic graph is *NP-complete*.[10] However, the optimal evolutionary path problem can be turned into a *shortest* path problem by relabelling the weights

$$a_{ij} = -\log p_{ij}$$

This transformation allows the direct use of Dijkstra's (1959) shortest path algorithm on the transformed graph $G' = (V, A)$.

Algorithm 1 (Dijkstra (1959))
set $u_0 = 0$ and set $T = \{1, \ldots, N\}$
for $k = 1$ to N
   set $u_k = a_{0k}$ (and set $P[k] = 0$)
repeat $(N-1)$ **times**
   let l be a node k in T minimizing u_k
   delete l from T
   for each $k \in T$
      if $u_l + a_{lk} < u_k$ **then** set $u_k = u_l + a_{lk}$ (and set $P[k] = l$)
    return $u_0, u_1, \ldots, u_N$ (and $P(0), P(1), \ldots, P(N)$)

Dijkstra's Algorithm partitions the graph's vertices into two sets, F (fixed) and T (temporary). Initially $F = \{0\}$ and all other vertices belong to T, and at each stage the nearest vertex in T is moved into F. Once the algorithm has been run, $u_0, u_1, \ldots, u_N$ records the shortest *distances* from vertex 0 to each other vertex. The "predecessor" array $P(0), P(1), \ldots, P(N)$, meanwhile, records where the various minima were obtained, and hence allows the recovery of the shortest *paths*. The algorithm's running time is $O(N^2)$.

Proposition 1 *Applying Dijkstra's Algorithm to the transformed graph G' delivers the most probable paths between vertex 0 and all other vertices in the original graph G.*

---

[8]Introductory texts on combinatorial optimization include McDiarmid (1997), Wilson (1996) and Papadimitriou and Steiglitz (1982).

[9]A digraph is "connected" if it cannot be expressed as the union of two digraphs.

[10]An *NP-complete* problem is one that is as hard as any reasonable problem, in a precise sense, and cannot be solved by any known polynomial algorithm (see Papadimitriou and Steiglitz (1982), ch.15).

Proof. The application of Dijkstra's Algorithm to the transformed graph G' gives the path U_z of minimum "length" $\sum_{(i,j)\in U_z} -\log p_{ij}$, for any given z. But this path will clearly maximise $\log\{\prod_{(i,j)\in U_z} p_{ij}\}$, and hence also $\prod_{(i,j)\in U_z} p_{ij}$, the probability of the path from 0 to z. ■

In fact, as a consequence, the following simpler algorithm can be employed on the original graph G.

Algorithm 2
set $u_0 = 0$ and set $T = \{2, \ldots, N\}$
for $k = 1$ to N
   set $u_k = p_{0k}$ (and set $P[k] = 0$)
repeat $(N - 1)$ **times**
   let m be a node k in T maximizing u_k
   delete m from T
   for each $k \in T$
      if $u_m \times p_{mk} > u_k$ **then** set $u_k = u_m \times p_{mk}$ (and set $P[k] = m$)
   return $u_0, u_1, \ldots, u_N$ (and $P(0), P(1), \ldots, P(N)$)

Proposition 2 *Applying Algorithm 2 to the original graph G delivers the most probable paths between vertex 0 and all other vertices in G.*

Proof. The vertex m selected on each pass when maximizing u_k in Algorithm 2 is the same as the vertex l selected when minimizing u_k in Dijkstra's Algorithm, since

$$m = \arg\max_k \prod_{(i,j)\in U_k} p_{ij}$$

$$\iff m = \arg\max_k \sum_{(i,j)\in U_k} \log p_{ij}$$

$$= \arg\min_k \sum_{(i,j)\in U_k} -\log p_{ij}$$

$$= l$$

where U_k is the path associated with u_k for a given k.

Furthermore, if $u_m \times p_{mk} > u_k$ in Algorithm 2, then

$$\left(\prod_{(i,j)\in U_m} p_{ij}\right) p_{mk} > \prod_{(i,j)\in U_k} p_{ij}$$

$$\iff -\log\left\{\left(\prod_{(i,j)\in U_m} p_{ij}\right) p_{mk}\right\} < -\log\left\{\prod_{(i,j)\in U_k} p_{ij}\right\}$$

$$\iff \sum_{(i,j)\in U_m} \left\{-\log p_{ij}\right\} - \log p_{mk} < \sum_{(i,j)\in U_k} -\log p_{ij}$$

which, given that $l = m$, is precisely the condition $u_l + a_{lk} < u_k$ in Dijkstra's Algorithm.

Consequently, applying Algorithm 2 to G is equivalent to applying Dijkstra's Algorithm to G', and the result follows from Proposition 2. ∎

5   Step-by-Step Evolution

As was seen earlier, when the "Radius-Modified Coradius" Theorem is reformulated for a state-dependent mutations setting, it is no longer immediately clear that Ellison's "step-by-step" evolution is optimal. To see that it is, one must identify the limit sets through which a given evolutionary path should pass.

This question can be addressed by applying Algorithm 2 to the digraph $G_{\mathcal{L}} = (V_{\mathcal{L}}, Q)$ with vertices $V_{\mathcal{L}} = \mathcal{L}$ corresponding to each limit set of the model, where each arc ij is weighted by the probability q_{ij} of the most probable path from the basin of attraction of limit set i to that of j.

Definition 2 *A limit set L' is said to be* intermediate *between limit set L and a union of limit sets Ω if*

$$q_{LL'} \cdot q_{L'\Omega} > q_{L\Omega}$$

Proposition 3 *The optimal evolutionary path (as $\sigma^2 \to 0$) from a given limit set L_i to a union of limit sets Ω will pass through all "intermediate" limit sets $\{L_j\}$ in a "step-by-step" fashion.*

Proof. Applying Algorithm 2 to $G_{\mathcal{L}}$ delivers a predecessor array with $P[L_k] = L_{k-1}$ for all "intermediate" limit sets L_j in the sense of Definition 2. ∎

This result confirms that, if there exist intermediate recurrent classes (i.e. Ellison's (2000) "steps" are in place), an optimal evolutionary path will pass through each in a "step-by-step" fashion. Definition 2 also clarifies the precise sense in which a recurrent class must be "intermediate" if it is to constitute such an evolution-facilitating "step".

Appendix

Lemma 1 *Suppose $(Z, P, g(\varpi; \sigma^2))$ is a model of evolution with noise. If $y \in Z$ and $\Omega \subset Z$ with $y \notin \Omega$, then*

$$\frac{\mu^{\sigma^2}(y)}{\mu^{\sigma^2}(\Omega)} = \frac{N(y, \Omega, y)}{\sum_{\omega \in \Omega} Q(\omega, \Omega - \omega, y) N(\Omega, y, \omega)}$$

Proof. See Ellison's (2000) Lemma 1 proof (with σ^2's replacing ε's). ■

Lemma 2 *Suppose $(Z, P, g(\varpi; \sigma^2))$ is a model of evolution with noise and that Ω is a union of limit sets of (Z, P). Then, for any $\omega' \in \Omega$ and any $y \notin D(\Omega)$*

$$\frac{1}{N(\Omega, y, \omega')} \leq \frac{1}{N(\Omega, Z - D(\Omega), \omega')} = O\big(\max_{\omega \in \Omega} Q(Z - D(\Omega), \Omega, \omega)\big)$$

Proof. This is established in the first part of Ellison's (2000) Lemma 2 proof. ■

Lemma 3 *Suppose $(Z, P, g(\varpi; \sigma^2))$ is a model of evolution with noise and that L is a limit set of (Z, P). Then,*

$$W(l, Z - D(L), g(\varpi; \sigma^2)) = O\big(1/\max_{z \in L} Q(Z - D(L), L, z)\big)$$

for all $l \in L$.

Proof. Following Ellison's (2000) Lemma 3 proof, given L we can find a T and a $k > 0$ such that for any $z \in D(L)$ there exists a path $z = z_1, z_2, \ldots, z_T$ with $z_T \in Z - D(L)$ such that the product of the transition probabilities along the path is at least $k \max_{z' \in L} Q(Z - D(L), L, z')$. Conditioning on the outcome of the first T periods we have for any $z \in D(L)$ that

$$W(z, Z - D(L), g(\varpi; \sigma^2)) \leq T +$$
$$(1 - k \max_{z' \in L} Q(Z - D(L), L, z')) \max_{z'' \in D(L)} W(z'', Z - D(L), g(\varpi, \sigma^2))$$

Taking the maximum of the LHS over $z'' \in D(L)$ gives

$$\max_{z'' \in D(L)} W(z, Z - D(L), g(\varpi; \sigma^2)) \leq (T/k)\big(1/\max_{z' \in L} Q(Z - D(L), L, z'))\big)$$

as desired. ■

Lemma 4 *Suppose $(Z, P, g(\varpi; \sigma^2))$ is a model of evolution with noise. Let $\mathcal{L}$ be the union of the limit sets of (Z, P) and suppose L is a single limit set. Then for any $l \in L$,*

$$W(l, \mathcal{L} - L, g(\varpi; \sigma^2)) = O\big(1/\max_{z \in L} Q(Z - D(L), L, z)\big)$$

Proof. For each $l \in L$ let S_l be the set of values of σ^2 for which $W(l, \mathcal{L} - L, g(\varpi; \sigma^2)) = \max_{l' \in L} W(l', \mathcal{L} - L, g(\varpi; \sigma^2))$. Ellison (2000, proof of Lemma 4) establishes that, if S_l is not empty, then for all $\sigma^2 \in S_l$ we have

$$\Big(1 - \sum_{z \in Z - D(L), z \notin \mathcal{L}} Q(z, Z - D(L) - z, l)Q(L, \mathcal{L} - L, z)\Big) W(l, \mathcal{L} - L, g(\varpi; \sigma^2))$$
$$\leq W(l, Z - D(L), g(\varpi; \sigma^2)) + \max_{z \in Z} W(z, \mathcal{L}, g(\varpi; \sigma^2))$$

As Ellison notes, the first term on the LHS of the expression is bounded away from zero because $Q(L, \mathcal{L} - L, z)$ is uniformly bounded away from one for any $z \notin D(L)$. The first term on the RHS is $O(1/\max_{z \in L} Q(Z - D(L), L, z))$ by Lemma 3. The second term on the RHS is finite. It follows from these observations that the desired result holds when $\sigma^2 \in S_l$. Taking the union of these sets over all $l \in L$, it holds for all σ^2. ∎

Lemma 5 *Suppose $(Z, P, g(\varpi; \sigma^2))$ is a model of evolution with noise. Let $\mathcal{L}$ be the union of the limit sets of (Z, P) and suppose that L and L' are two given limit sets. Then for any $l \in L$,*

$$\frac{1}{Q(L', \mathcal{L} - (L \cup L'), l)} = O\left(\frac{\max_{z \in L} Q(Z - D(L), L, z)}{\max_{z' \in L} Q(L', \mathcal{L} - (L \cup L'), z')}\right)$$

Proof. In the proof of his Lemma 5, Ellison (2000) shows that for $l \in L$ we have

$$Q(L', \mathcal{L} - (L \cup L'), l) \geq \frac{1}{|L|} \sum_{t = 1, \ldots, \infty, l' \in L} \text{Prob}\big\{\{z_1, \ldots, z_t\} \cap (\mathcal{L} - L) = \emptyset,$$
$$z_t = l' \mid z_1 = l\big\} \cdot \text{Prob}\big\{L' \text{ is reached before } \mathcal{L} - (L \cup L') \text{ and}$$
$$\text{at most } |L| \text{ periods are spent in } L \mid z_1 = l\big\}$$

where $|L|$ is the number of elements of L.

As Ellison notes, the summation over t and l' of the first terms on the RHS above is bounded below by $N(L, \mathcal{L} - L, l)$. Meanwhile, the second terms on the RHS are uniformly bounded below by $k \max_{z' \in L} Q(L', \mathcal{L} - (L \cup L'), z')$ for σ^2 small for some $k > 0$. Hence we have

$$Q(L', \mathcal{L} - (L \cup L'), l) \geq \frac{1}{|L|} N(L, \mathcal{L} - L, l) k \max_{z' \in L} Q(L', \mathcal{L} - (L \cup L'), z')$$
$$\geq \frac{1}{|L|} N(L, Z - D(L), l) k \max_{z' \in L} Q(L', \mathcal{L} - (L \cup L'), z')$$
$$\geq k' \frac{\max_{z' \in L} Q(L', \mathcal{L} - (L \cup L'), z')}{\max_{z \in L} Q(Z - D(L), L, z)}$$

for some $k' > 0$ using the result of Lemma 2. ∎

Lemma 6 *Let $(Z, P, g(\varpi; \sigma^2)$ be a model of evolution with noise and suppose that Ω is a union of limit sets of (Z, P), whilst $\{L_j\}_{j=i}^r$ are the limit sets through which the most probable path from a given limit set L_i to Ω passes. Then*

$$\max_{y \notin \Omega} W(y, \Omega, g(\varpi; \sigma^2)) = O\left(\left[\min_{L_i \in \mathcal{L} - \Omega}\left(\max_{z \in L_i} Q(Z - D(L_i), L_i, z)\right.\right.\right.$$
$$\left.\left.\left.\prod_{j=i}^{r-1} \frac{\max_{z'' \in L_j} Q(L_{j+1}, \mathcal{L} - (L_j \cup L_{j+1}), z'')}{\max_{z' \in L_j} Q(Z - D(L_j), L_j, z')}\right)\right]^{-1}\right)$$

Proof. Following Ellison's (2000) Lemma 6, it will suffice to establish that

$$\overline{W}(\Omega, g(\varpi; \sigma^2)) \equiv \max_{y \in \mathcal{L} - \Omega} W(y, \Omega, g(\varpi; \sigma^2))$$

$$= O\left(\left[\min_{L_i \in \mathcal{L} - \Omega}\left(\max_{z \in L_i} Q(Z - D(L_i), L_i, z)\right.\right.\right.$$
$$\left.\left.\left.\prod_{j=i}^{r-1} \frac{\max_{z'' \in L_j} Q(L_{j+1}, \mathcal{L} - (L_j \cup L_{j+1}), z'')}{\max_{z' \in L_j} Q(Z - D(L_j), L_j, z')}\right)\right]^{-1}\right)$$

Given $y \in \mathcal{L} - \Omega$ let $z_1, z_2, \ldots, z_T$ be the most probable path from y to Ω, passing through distinct limit sets $L_1, L_2, \ldots, L_r$.

Writing q_{12} for $\min_{y' \in L_1} Q(L_2, \mathcal{L} - (L_1 \cup L_2), y')$ and $W(A, B, g(\varpi; \sigma^2))$ for $\max_{x \in A} W(x, B, g(\varpi; \sigma^2))$, Ellison shows that

$$\overline{W}(\Omega, g(\varpi; \sigma^2)) \leq \frac{W(L_1, \mathcal{L} - L_1, g(\varpi; \sigma^2))}{q_{12} q_{23} \cdots q_{(r-1)r}} + \cdots$$
$$+ \frac{W(L_{r-1}, \mathcal{L} - L_{r-1}, g(\varpi; \sigma^2))}{q_{(r-1)r}}$$

It now suffices to show that the summation on the RHS of the above equation is

$$O\left(\left[\min_{L_i \in \mathcal{L} - \Omega}\left(\max_{z \in L_i} Q(Z - D(L_i), L_i, z)\right.\right.\right.$$
$$\left.\left.\left.\prod_{j=i}^{r-1} \frac{\max_{z'' \in L_j} Q(L_{j+1}, \mathcal{L} - (L_j \cup L_{j+1}), z'')}{\max_{z' \in L_j} Q(Z - D(L_j), L_j, z')}\right)\right]^{-1}\right)$$

Now, by Lemma 5 we know that

$$\frac{1}{q_{j(j+1)}} = O\left(\frac{\max_{z' \in L_j} Q(Z - D(L_j), L_j, z')}{\max_{z \in L_j} Q(L_{j+1}, \mathcal{L} - (L_j \cup L_{j+1}), z)}\right)$$

14

Using this and the result of Lemma 4 we have

$$
\frac{W(L_i, \mathcal{L} - L_i, g(\varpi; \sigma^2))}{q_{i(i+1)} q_{(i+1)(i+2)} \cdots q_{(r-1)r}} = \left(\frac{1}{\max_{z \in L_i} Q(Z - D(L_i), L_i, z)} \right.
$$

$$
\left. \prod_{j=i}^{r-1} \frac{\max_{z' \in L_j} Q(Z - D(L_j), L_j, z')}{\max_{z'' \in L_j} Q(L_{j+1}, \mathcal{L} - (L_j \cup L_{j+1}), z'')} \right)
$$

Hence we have

$$
\overline{W}(\Omega, g(\varpi; \sigma^2)) = O\left(\max_{L_i \in \mathcal{L} - \Omega} \left(\frac{1}{\max_{z \in L_i} Q(Z - D(L_i), L_i, z)} \right. \right.
$$

$$
\left. \left. \prod_{j=i}^{r-1} \frac{\max_{z' \in L_j} Q(Z - D(L_j), L_j, z')}{\max_{z'' \in L_j} Q(L_{j+1}, \mathcal{L} - (L_j \cup L_{j+1}), z'')} \right) \right)
$$

$$
= O\left(\left[\min_{L_i \in \mathcal{L} - \Omega} \left(\max_{z \in L_i} Q(Z - D(L_i), L_i, z) \right. \right. \right.
$$

$$
\left. \left. \left. \prod_{j=i}^{r-1} \frac{\max_{z'' \in L_j} Q(L_{j+1}, \mathcal{L} - (L_j \cup L_{j+1}), z'')}{\max_{z' \in L_j} Q(Z - D(L_j), L_j, z')} \right) \right]^{-1} \right)
$$

as required. ∎

References

BERGIN, J., AND B. LIPMAN (1996): "Evolution with State-Dependent Mutations," *Econometrica*, 64, 943–956.

BLUME, L. (1993): "The Statistical Mechanics of Strategic Interaction," *Games and Economic Behavior*, 4, 387–424.

——— (1999): "How Noise Matters," Working Paper.

BROCK, W., AND S. DURLAUF (1995): "Discrete Choice with Social Interactions I: Theory," SSRC Working Paper No. 9910, University of Wisconsin.

DIJKSTRA, E. W. (1959): "A Note on Two Problems in Connexion with Graphs," *Numerische Mathematik*, 1, 269–271.

ELLISON, G. (1993): "Learning, Local Interaction and Coordination," *Econometrica*, 61, 1047–1071.

——— (2000): "Basins of Attraction, Long-Run Stochastic Stability, and the Speed of Step-by-Step Evolution," *Review of Economic Studies*, 67, 17–45.

FOSTER, D. P., AND H. P. YOUNG (1990): "Stochastic Evolutionary Game Dynamics," *Theoretical Population Biology*, 38, 219–232.

FREIDLIN, M. I., AND A. D. WENTZELL (1984): *Random Perturbations of Dynamical Systems*. Springer-Verlag, Berlin.

KANDORI, M., G. J. MAILATH, AND R. ROB (1993): "Learning, Mutation and Long-Run Equilibria in Games," *Econometrica*, 61, 29–56.

LEE, I. H., A. SZEIDL, AND A. VALENTINYI (2001): "Contagion and State-Dependent Mutations," Working Paper.

McDIARMID, C. J. H. (1997): "Lecture Notes on Combinatorial Optimisation," Mathematical Institute Notes, Oxford University.

MYATT, D. P., AND C. WALLACE (1998): "Adaptive Play by Idiosyncratic Agents," Discussion Paper, Nuffield College, Oxford.

NORMAN, T. (2003): "The Evolution of Coordination under Inertia," Nuffield College Discussion Paper.

PAPADIMITRIOU, C. H., AND K. STEIGLITZ (1982): *Combinatorial Optimization: Algorithms and Complexity*. Prentice-Hall, Englewood Cliffs, NJ.

VAN DAMME, E., AND J. W. WEIBULL (1998): "Evolution with Mutations Driven by Control Costs," Tilburg University: Mimeo.

WILSON, R. J. (1996): *Introduction to Graph Theory*. Pearson Education Limited, Harlow, fourth edn.

YOUNG, H. P. (1993): "The Evolution of Conventions," *Econometrica*, 61, 57–84.