ISSN 1471-0498

DEPARTMENT OF ECONOMICS

DISCUSSION PAPER SERIES

WAITING TIMES AND EQUILIBRIUM SELECTION

Alan Beggs

Number 142

January 2003

Manor Road Building, Oxford OX1 3UQ

Waiting Times and Equilibrium Selection

A. W. Beggs
Wadham College
Oxford
OX1 3PN
UK

Current Version: December 2002

Abstract

This paper shows how graphs can be used to calculate expected waiting times in models of equilibrium selection. It also shows how reducing the state space can simplify the calculations both of waiting times and selected equilibria. The results are applied to potential games and games with strategic complementarities.

JEL Classification Nos: C72, C73

Keywords: waiting times, equilibrium selection, graphs

1. Introduction

Foster and Young (1990), Kandori et al. (1993) and Young (1993) introduced the idea that studying evolutionary models with persistent noise may help to select between multiple equilibria in games. They showed that if the noise is small, then the long-run probability with which certain equilibria are played becomes vanishingly small and so may be thought unlikely to be observed. For example in symmetric 2×2 games, under some assumptions, the only equilibrium which survives this test is the risk-dominant equilibrium.

The role of noise in these models is to ensure that it is impossible to become trapped in a bad equilibrium. The noise is taken to be small so that the main force driving behaviour is the payoffs of the underlying game. If noise is small, however, it may take an exceedingly long time to escape from a bad equilibrium. Indeed as the noise tends to zero, escape times tend towards infinity. In order to assess the relevance of the selected equilibrium it is therefore important to estimate how long it is likely to take to reach it.

The most general results on these lines to date are obtained by Ellison (2000). He examines the expected waiting time to reach the selected equilibrium. As the noise becomes small, this becomes larger and larger but it is possible to examine its rate of growth. He obtains useful bounds on this in terms of the so-called 'radius' and '(modified) co-radius'. These are not, however, in general exact.

This paper shows how results of Freidlin and Wentzell (1998) can be used to obtain exact estimates of the growth rates of expecting waiting times. Their 'tree characterisation' of stationary distributions was used by Kandori et al. (1993) and Young (1993). In addition to these results, however, they also show that waiting times can be estimated by using a calculation involving graphs. These results do not appear to have been used before in the economics literature. Ellison (2000)'s results can be obtained as bounds to these estimates.

Models of equilibrium selection are often studied under the assumption that the population size is large, as the assumption of naïve adjustment usually made seems most plausible then. This can make the tree calculations to determine the selected equilibrium quite involved since the state space may then be very large. Kandori et al. (1993) and Young (1993) observed that they can often be simplified by considering a reduced state space consisting only of the limit sets of the process, which may be of much smaller dimension than and indeed often independent of population size.

The current paper extends these observations in two ways. Firstly, it shows that the graph calculations involved in calculating waiting times may also be simplified in the same way. Secondly, it observes that the equilibrium selection, and waiting time, calculations can be even further simplified by iteratively eliminating the limit sets that are easiest to escape from. The calculations required are much simpler than solving the tree problem directly.

The paper applies these results to learning in potential games and games with strategic complementarities. The former have been considered previously by

Blume (1994) but he obtained only numerical estimates of waiting times. The current paper shows that the methods used here lead to a simple analytical expression for waiting times. Kandori and Rob (1995) obtained a partial characterisation of the selected equilibrium in the latter case but the current methods allow a much fuller one. It also provides estimates of waiting times.

The paper proceeds as follows. Section 2 outlines a general framework that encompasses the models of noisy evolution contained in the literature. Section 3 gives the relevant results from Freidlin and Wentzell (1998). They show that the exit time from a given region, when the noise is small, can be determined by finding certain minimum-cost graphs. Section 4 gives examples to illustrate the calculations involved.

Section 5 compares the current results with Ellison (2000)'s measures. It shows how his results can be obtained as bounds to the current ones.

Section 6 applies the results to a model of learning in potential games and shows that the current results lead to a simple characterisation of waiting times. These are tighter than those offered by Ellison (2000)'s approach. Section 7 briefly outlines another equivalent characterisation of waiting times used in the literature on perturbed optimization.

Section 8 shows how the graph characterisations involved in finding waiting times, and also tree calculations, can be simplified by considering spaces of smaller dimension. Kandori et al. (1993) and Young (1993) observed that it was enough to consider the limit sets of the process as these are the only possible long-run stable sets. The section shows that the same observation applies to waiting times. It also observes that the procedure can be extended further and less stable limit sets eliminated iteratively. This enables the selected equilibria to be determined by solving a sequence of shortest path problems without recourse to tree calculations. These results draw on several in the probability literature, particularly those of Catoni et al. (2000).

Section 9 shows that the methods of Section 8 allow a simple derivation of Ellison (2000) 'radius/co-radius' criterion for equilibrium selection. Section 10 applies the results above to learning in games with strategic complementarities. It shows that they offer a fuller characterisation of equilibrium selection than found by Kandori and Rob (1995) and obtains results on waiting times, which they did not.

As noted above the main inspiration for these results is Freidlin and Wentzell (1998). Although their main focus is on processes perturbed by Brownian motion, they provide the key lemmas on perturbed finite Markov chains. It is curious to note that the use of tree calculations appeared much earlier in the economics literature — see Bott and Mayberry (1954). In a general sense they can be traced back to Kirchhoff's papers on electrical networks — see Haken (1977) chapter 4.8 and his references.

Extensive use of Freidlin and Wentzell (1998)'s work has been made in the literature on perturbed optimization. A good survey can be found in Catoni (1999). This literature has made some refinements of Freidlin and Wentzell (1998)'s ideas, particularly in the notion of the 'cycle decomposition' which Freidlin and Wentzell only develop under a general position assumption. Much of this work is geared towards optimization and is of less use in equilibrium selection, although a few results are sketched in Section 7.

As noted above, the paper in the economics literature most closely related to the current one is Ellison (2000). The estimates he gives are simpler to calculate than the current ones. Nevertheless the paper shows that in some economic examples the current methods do lead to sharper and intuitive results. He also provides a simple sufficient condition for an equilibrium to be selected (its 'radius' must be greater than its '(modified) co-radius'). This is often very useful but in the applications studied here it lacks power. The procedure of iterative elimination is straightforward to apply and leads to sharper results.

2. Framework

The models of stochastic evolution considered in the literature can represented abstractly as consisting of the following elements:

1) A finite state space E.
2) A rate function $V : E \times E \to R_+ \cup \{+\infty\}$ where V is irreducible in the sense that the transition matrix $(-\exp V(x,y))$ is irreducible.
3) A family of homogenous Markov chains, $\{X_n\}_{n=0}^{\infty}$, on E indexed by parameter $\beta \in (0,\infty]$ with transition matrices P_β such that,

$$\lim_{\beta\to\infty} \frac{-\ln P_\beta(X_n = y \,|\, X_{n-1} = x)}{\beta} = V(x,y) \tag{1}$$

(with the convention that $\ln 0 = -\infty$).
4) $P_\beta(X_n = y \,|\, X_{n-1} = x)$ is continuous at $\beta = \infty$ for all x, y.

In the literature on equilibrium selection, the state space E summarises the relevant information on how some population has been playing a game. For example, it may represent the number or proportion of players currently playing each strategy. The the chain generated by P_∞ is thought of as the base evolutionary process. The chains P_β represent the process when this is perturbed by mistakes or other noise. The function V measures how likely these perturbations are or, equivalently, their cost. If $V(x,y) = 0$ then direct transitions between x and y are possible in the base process. If $V(x,y) > 0$ then they are impossible and V measures the rate at which their probability tends to zero. If $V(x,y) = \infty$ they are impossible, for large β, even in the unperturbed model.

For example suppose that there are N players in the population and each must choose between playing strategy 1 and strategy 2. Assume that if i players played strategy 1 last period, then strategy 2 is the best response for all players. Suppose that each player plays strategy 1 instead with probability $\epsilon = \exp(-\beta)$. The probability that j players play strategy 1 next period is then

$$\binom{N}{j} \epsilon^j (1-\epsilon)^{N-j} \tag{2}$$

This model satisfies the above assumptions with, for this state, $V(i,j) = j$. Parameterisation by β, rather than ϵ, means that as β becomes large errors become more costly or less likely.

The assumptions imply that, for large β, the chain generated by P_β is irreducible and so has a unique invariant distribution, μ_β. The process generated by P_∞ may have multiple invariant distributions, but the idea in the literature on equilibrium selection is that if $\tilde{\mu} = \lim_{\beta\to\infty} \mu_\beta$ exists, then that is the one selected. It is in a sense the most robust to (these) perturbations. The assumption of continuity at infinity is simply in order to ensure that $\tilde{\mu}$ is indeed an invariant distribution for P_∞.

One may doubt the relevance of these results. For example, $\tilde{\mu}$ may predict that a unique equilibrium is selected. Yet if there are other steady states of P_∞ and the system starts near one of them, it may take an exceedingly long time for the perturbed process to reach it if β is large. It is therefore of interest to estimate waiting times, in particular the expected time to reach one state starting in another.

3. General Results using graphs

It is is well-known that the limiting behaviour of invariant distributions of perturbed Markov chains can be analysed using trees. It seems to be less well known in the economics literature that graphs can also be used to analyse waiting times. This section outlines the relevant results of Freidlin and Wentzell (1998). Some of presentation is guided by Catoni (1999).

Let E be a finite state space, $p : E \times E \to [0,1]$ an irreducible Markov transition matrix and X_n, $n \in \mathcal{N}$, a Markov chain with transition matrix p. Let $W \subset E$ be a subset of E and $\overline{W} = E \backslash W$ its complement. For any oriented graph $g \subset E \times E$, let $g(x) = \{y \,|\, (x,y) \in g\}$ denote the immediate successors of x. More generally $g^n(x) = \bigcup_{y \in g^{n-1}(y)} g(y)$. The set of points that can be reached from x, or its orbit, is $O_g(x) = \bigcup_{n=1}^{\infty} g^n(x)$.

For any set W, one can consider the set of all graphs, without loops, that link each point outside W to it:

Definition 1 *$G(W)$ is the set of oriented graphs $g \subset E \times E$ satisfying*

1) *For any $x \in E$, $|g(x)| = \mathbf{1}_{\overline{W}}$ (no arrow starts from a point of W, exactly one arrow starts from each point outside W),*
2) *g contains no loop.*

The second condition is equivalent to the condition that each point of $\overline{W}$ is linked by g to a point of W.

Definition 2 *For any $x \in E$ and $y \in W$, let*

$$G_{x,y}(W) = \begin{cases} \{g \in G(W) \,|\, y \in O_g(x)\} & \textit{if } x \in \overline{W} \\ G(W) & \textit{if } x = y \\ \emptyset & \textit{if } x \in W \backslash \{y\} \end{cases}$$

That is, if $x \neq y$ $G_{x,y}(W)$ is the set of $G(W)$-graphs which link x to y. Note that by the first condition in Definition 1, x is linked to a unique point in W.

In the case $W = \{x\}$, $G(\{x\})$ is the set of all trees leading to x, familiar from the work of Kandori et al. (1993) and Young (1993).

For a graph g let $p(g) = \prod_{(z,t)\in g} p(z,t)$. That is $p(g)$ is the product of the probabilities of the transitions in g.

The following, now well known, lemma of Freidlin and Wentzell (1998) gives the invariant distribution of p in terms of probabilities of trees:

Lemma 1 *The (unique) invariant distribution of p is given by*

$$\mu(x) = \left(\sum_{g\in G(\{x\})} p(g)\right)\left(\sum_{y\in E}\sum_{g\in G(\{y\})} p(g)\right)^{-1}, \qquad x \in E \tag{3}$$

Its proof is elementary. It forms, however, the basis for much of the work of Kandori et al. (1993) and Young (1993).

Freidlin and Wentzell (1998) also give two lemmas concerning the exit of a Markov chain from a domain. These do not seem to have been used in the economics literature before.

For a domain W, let $\tau(W) = \inf\{n \geq 0 | X_n \in W\}$ be the first time W is reached. Then

Lemma 2 *For any $W \neq \emptyset$, $x \in \overline{W}$, and $y \in W$,*

$$P(X_{\tau(W)} = y | X_0 = x) = \left(\sum_{g\in G_{x,y}(W)} p(g)\right)\left(\sum_{g\in G(W)} p(g)\right)^{-1} \tag{4}$$

and

Lemma 3 *For any $W \neq \emptyset$, any $x \in \overline{W}$,*

$$E(\tau(W)|X_0 = x) = \left(\sum_{y\in\overline{W}}\sum_{g\in G_{x,y}(W\cup\{y\})} p(g)\right)\left(\sum_{g\in G(W)} p(g)\right)^{-1} \tag{5}$$

Lemma 2 states that the probability that when the process hits W the point it hits first is y is given by the ratio of the probability of the $G(W)$-graphs where x is joined to y to the total probability of all $G(W)$-graphs. This is perhaps somewhat intuitive. Lemma 2 is used extensively in the proofs of Section 8.

Lemma 3 governs the expected time to reach W and is the more important for the paper.[1] The set $\cup_{y\in\overline{W}} G_{x,y}(W \cup \{y\})$ is the set of graphs where x is joined to some

[1] Note that E is used to denote expectation in the paper as it should be clear from the context when it refers to this rather than the state space.

point, y, of $\overline{W}$, possibly itself, and not to W. All other points in $\overline{W}$ are joined to the same point or to W. If $x = y$, this means that x has no arrows leaving it. Freidlin and Wentzell (1998) write this set as $G(x \not\to W)$. Roughly speaking, it is the set of graphs where x is not linked to W. Lemmas 2 and 3 appear as Lemma 3.3 and 3.4 in chapter 6 of Freidlin and Wentzell (1998). Freidlin and Wentzell give a proof based on induction on the number of elements in W. Catoni (1999) offers a somewhat smoother proof, similar in spirit to the well known proof of Lemma 1 though not as elementary.

Note that if $\overline{W}$ is a singleton, $G(x \not\to W)$ is empty. In this case, the numerator of (5) is defined to be 1.

The results above apply to any (irreducible) Markov chain. Consider now the framework of Section 2. For a graph g, let its cost be $V(g) = \sum_{(z,t)\in g} V(z,t)$. The empty graph is assumed to have cost zero. Lemma 3 then implies:

Lemma 4 *Under the assumptions of Section 2, for any $x \in E$,*

$$\lim_{\beta\to\infty} -\beta^{-1} \ln \mu_\beta(x) = U(x) \tag{6}$$

where $U : E \to \mathcal{R}$ is defined by

$$U(x) = \min_{g\in G(\{x\})} V(g) - \min_{y\in E} \min_{g\in G(\{y\})} V(g) \tag{7}$$

In particular

$$\lim_{\beta\to\infty} \mu_\beta(x) = 0 \qquad \text{if } x \notin \arg\min_{y\in E} U(y) \tag{8}$$

(8) is the result familiar from the literature on equilibrium selection. $V(g)$ measures the cost of a graph and the result states that any state which is not the root of a minimum cost spanning tree receives no weight. $U(x)$ is sometimes referred to as the quasi-potential function and measures the rate at which the weight on other states goes to zero.

Lemma 2 implies

Lemma 5 *For any $D \subset E$, $D \neq \emptyset$, for any $x \in D$, $y \in \overline{D}$,*

$$\lim_{\beta\to\infty} -\frac{1}{\beta} P_\beta(X_{\tau(\overline{D})} = y | X_0 = x) = \min_{g\in G_{x,y}(\overline{D})} V(g) - \min_{g\in G(\overline{D})} V(g) \tag{9}$$

Denote this quantity by $V_D(x, y)$. In words the probability of hitting y on first exit from the domain D is governed by the cost of the least cost $G(\overline{D})$-graph linking x to y relative to the least cost $G(\overline{D})$-graph. For notational ease, the results here are phrased in terms of exit from a domain, rather than entry, as was the case in Lemmas 2 and 3, but the formulations are equivalent.

More importantly, for current purposes, Lemma 3 implies

Lemma 6 *For any $D \subset E$, $D \neq \emptyset$, for any $x \in D$,*

$$\lim_{\beta \to \infty} \frac{\ln E_\beta(\tau(\overline{D})|X_0 = x)}{\beta} = \min_{g \in G(\overline{D})} V(g) - \min_{y \in D} \min_{g \in G_{x,y}(\overline{D} \cup \{y\})} V(g) \tag{10}$$

Denote the right-hand side of (10) by $H_D(x)$ (H for height), its first term by W_D and the second by $M_D(x)$. Loosely speaking, the expected time to exit the domain D from the point x is governed by how much it costs to link x to $\overline{D}$ rather than to a point in D.

The expected exit time from a domain depends on the initial point. It is convenient to have a notation for the worst case, so let

$$H(D) = \max_{x \in D} H_D(x) \tag{11}$$

To clarify these results, some examples are given in the next Section.

4. Examples

The tree formula for equilibrium distributions is well understood, so the current section gives examples of waiting time calculations. To begin consider a trivial example:

Example 1 Consider a domain consisting of a single point: $D = \{x\}$. For this set. $W_D = \min_y V(x, y)$, that is it is the least cost of connecting x to another point. The only graph on elements of D, without cycles, not connecting it to its complement is the empty graph. By the convention above this has cost zero, hence $M_D(x) = 0$. It follows that in this case

$$H(D) = \min_y V(x, y) \tag{12}$$

Example 2 Consider a chain with the transition costs indicated in Figure 1. In this, and the other figures, transitions not shown have cost infinity or some large positive number, if need be, to be consistent with irreducibility. Let $D = \{x, y\}$. $W_D = 3+2 = 5$. Clearly, the graph $y \to z$ minimises $M_D(x)$. Hence $H_D(x) = 5-2 = 3$. The graph $x \to y$ minimises $M_D(y)$ (note that $y \to z$ connects y to the exterior and so is not admissible). Hence $H_D(y) = 5 - 3 = 2$. $H(D) = 3$.

Example 3 Consider instead Figure 2 and let $D = \{x, y\}$ again. $W_D = 5$ again but now the graph $x \to y$ minimises both $M_D(x)$ and $M_D(y)$. Hence $H(D) = H_D(x) = H_D(y) = 5 - 2 = 3$. At first sight this is puzzling — surely it must take longer to escape D starting from x than y. Indeed it must, but the estimates for waiting times only yield the expectation up to logarithmic order of magnitude as $\beta \to \infty$. Once y is reached return to x is highly unlikely and escape from y takes an order of magnitude longer than escape from x. Escape from y, therefore, dominates the escape time from D even when starting from x. The estimates provided by Ellison (2000) have the same weakness, that only the logarithmic order of magnitude of escape times is estimated.

Example 4 Consider finally the more complicated example shown in Figure 3. Let $D = \{x, y, z\}$. The path $x \to y \to z \to w$ is a least cost graph that links D to w, hence $W_D = 6$. The graph $x \to y \to z$ minimises $M_D(x)$, so $M_D(x) = 3$. Hence $H_D(x) = 6 - 3 = 3$. The same graph also minimises $M_D(y)$ and $M_D(z)$, so $H(D) = 3$.

The next section compares the estimates provided by Ellison (2000) and uses some of these examples to illustrate the differences between those and the current approach.

5. Ellison's Measures

Ellison (2000) provides an upper bound for the exit time from a domain in terms of the 'modified co-radius' and a lower bound in terms of the 'radius'. This section shows how his measures can be obtained from the current approach. The first sub-section considers the radius, the second the modified co-radius.

5.1 The Radius

In order to state Ellison's results, some preliminary definitions are needed. Consider the unperturbed Markov chain of Section 2, P_∞. Recall that a limit set for this chain is a set which is (a) absorbing, (b) every state in it can be reached from every other state in it. An equivalent definition, under the assumptions of Section 2, is

Lemma 7 *Ω is a limit set if and only if (a) every point in Ω can be connected to any other by a path of null cost, (b) any path leaving Ω has positive cost.*

P_∞ has a finite number of limit sets and it is a standard result that the process reaches, and is trapped, in one of these limit sets with probability 1. In applications, often these limit sets correspond to all players playing some Nash equilibrium.

Let $\Omega \subseteq E$ be a union of one or more limit sets of P_∞. The *basin attraction* of Ω, $D(\Omega)$, is defined to be the set of initial states from which the unperturbed process converges to Ω with probability 1:

$$D(\Omega) = \{e \in E | P_\infty (\exists T \text{ s.t } x_t \in \Omega, \forall t > T | x_0 = e) = 1\} \tag{13}$$

A weaker notion is

$$D^w(\Omega) = \{e \in E | P_\infty (\exists T \text{ s.t } x_t \in \Omega, \forall t > T | x_0 = e) > 0\} \tag{14}$$

In other words, the *weak basin of attraction* is the set of all points which have a positive probability of reaching Ω.

If P_∞ is deterministic then (13) and (14) are equivalent and E can be partitioned into the domains of attraction of the different limit sets. If P_∞ is not deterministic, however, neither (13) nor (14) yields a partition in general. There may be transient points which have a positive probability of reaching more than one limit set. For this reason, the union of the ordinary basins of attraction may not cover E, while the weak basins of attraction may not be disjoint. The results below, mutatis mutandis, hold under either definition, so for consistency with Ellison (2000) (13) will be used.

The following Lemma is obvious from the definition and is stated for ease of reference:

Lemma 8 *If Ω is a limit set, any point in $D(\Omega)$ can be connected to any point in Ω by a path of null cost.*

A path p from a set U to a set W is a path such that if $p = (x_1, x_2, \dots, x_T)$ then $x_1 \in U$, $x_T \in W$ and $x_t \notin W$ for $t = 1, 2, \dots, T-1$. Its cost is $V(p)$. Denote the set of paths from U to W by $S(U, W)$. Ellison (2000) defines the radius of the basis of attraction of Ω to be the least cost of any path from Ω out of $D(\Omega)$.

$$R(\Omega) = \min_{p \in S(\Omega, E - D(\Omega))} V(p) \tag{15}$$

One can use the results of the previous section to show that $R(\Omega)$ provides a lower bound to the exit time from $D(\Omega)$.

Proposition 1 *(i) If Ω is a single limit set, then for any $x \in D(\Omega)$, $H_{D(\Omega)}(x) = R(\Omega)$.*

(ii) If Ω is the union of a collection of limit sets, then for any $x \in \Omega$, $H_{D(\Omega)}(x) \geq R(\Omega)$.

(i) follows easily from Lemma 6. The optimal $\overline{D(\Omega)}$-graph must connect all points of $D(\Omega)$ to its complement. Let r be a path that achieves the minimum in the definition of the radius. It is easy to see that all other points can be connected to this path by a sequence of null cost transitions. This graph has cost $R(\Omega)$ and no other graph can have lower cost as all points in Ω must be connected to the exterior of $D(\Omega)$. It follows that $W_{D(\Omega)} = R(\Omega)$. For any point in $D(\Omega)$, $M_{D(\Omega)} = 0$, from Lemmas 7 and 8. Hence, $H_D(\Omega) = R(\Omega)$. The result in part (ii) follows from a similar argument, given in the Appendix. A different proof can be found in Section 9.

It may seem mildly surprising that the exit time from the domain of attraction of a limit set is independent of the point at which the process starts: if it starts near the boundary, surely it is more likely to escape quickly. This is not so, at least on the scale considered here. Although direct escape is possible, the process is much more likely to follow the flow of P_∞, reach Ω rapidly and then be stuck there for a long period of time. This possibility is the dominant contribution to the expected escape time and so the expected escape times from all points in $D(\Omega)$ have the same order of magnitude.

The radius bound is not tight for the sets which are the union of the basins of attraction of several limit sets. Consider Example 2, shown in Figure 1, where each point may be considered as a limit set. Let $D = \{x, y\}$. The radius of D is 2. $H_D(x) = 3$, however.

5.2 The Modified Co-Radius

Ellison (2000) also provides an upper bound on the exit time from a domain. He considers the case of a domain D whose complement $\overline{D}$ is a union of limit sets. His result would in fact go through under the slightly weaker assumption that for all limit sets L, if $D \cap L \neq \emptyset$, then $D(L) \subseteq D$, but for simplicity Ellison (2000)'s formulation will be followed.

Suppose $x \in D$. Let $p = (z_1, \dots, z_T)$ be a path from x to $\overline{D}$ and let $L_1, L_2, \dots, L_r$ be the sequence of limit sets through which the path passes consecutively (repetitions are allowed, but not consecutive ones). x may belong to L_1 and $L_i \cap \overline{D} = \emptyset$ for $i < r$. $L_r \subseteq \overline{D}$. The *modified cost* of this path is defined as

$$V^*(p) = V(p) - \sum_{i=2}^{r-1} R(L_i) \tag{16}$$

The *modified co-radius* is then defined as:

$$CR^*(\overline{D}) = \max_{x \in D} \min_{p \in S(x, \overline{D})} V^*(p) \tag{17}$$

Ellison shows that the modified co-radius provides an upper bound to the exit time from D. This will be deduced from the results of Section 3 below, but first consider some of the examples of Section 4 to see the difference between this measure and the current approach.

The examples illustrate respectively that (a) the co-radius only bounds the worst case for exit times, (b) the modified cost of joining a point to the complement of the domain does not bound exit times for all points (only in the worst cases), and (c) the co-radius bound is not in general tight.

Example 2 Let $D = \{x, y\}$ (see Figure 1). Each point can be regarded as a limit set. The least cost path from y to z is $y \to z$. This passes through no other limit set, so its modified cost is 2. The least cost path joining x to z is $x \to y \to z$. This passes through the intermediate limit set y, which has radius 2. The modified cost of this path is therefore $3 + 2 - 2 = 3$. The modified co-radius of z is therefore 3, which agrees with the calculation that $H(D) = 3$. Note that $H_D(y) = 2$.

Consider instead

Example 3 Again let $D = \{x, y\}$ (see Figure 2). The least cost path from y to z now has cost 3, and this is its modified cost. The path $x \to y \to z$ now has modified cost $2 + 3 - 3 = 2$. The modified co-radius of the domain is still 3, which agrees with the fact that $H(D) = 3$. Note, however, that $H_D(x) = 3$, while the modified cost of the path to z is only 2.

Finally,

Example 4 Let $D = \{x, y, z\}$ (see Figure 3). The radii of y and z are both 1. A path of least modified cost joining x to w is $x \to y \to w$. This has modified cost $2+3-1 = 4$. The paths from y and z have modified cost 3, so the modified co-radius of w is 4. This exceeds $H(D) = 3$.

Of these examples, Example 4 is the most serious if one is interested in equilibrium selection — one may be content with a worst case bound, but it should ideally be tight. Ellison (2000) in fact gives this example (with a few changes in numbers).

The relationship between the modified co-radius and the approach here is, perhaps, not immediately apparent. The former seems to look at a paths, while the latter looks at graphs connecting whole regions of points. The graphs produced, however, must link each point in D to the exterior by a path and a local argument concerning surgery around this path turns out to produce the modified co-radius bound. The argument is given below but the reader who is not interested can skip directly to the statement of Proposition 2. Another proof, where the result emerges as a special case of more general ones, appears in Section 9.

Consider the formula in Lemma 6:

$$H_D(x) = \min_{g \in G(\overline{D})} V(g) - \min_{y \in D} \min_{g \in G_{x,y}(\overline{D} \cup \{y\})} V(g) \tag{18}$$

The worst case for exit from the domain is therefore given by

$$H(D) = \min_{g \in G(\overline{D})} V(g) - \min_{x \in D} \min_{y \in D} \min_{g \in G_{x,y}(\overline{D} \cup \{y\})} V(g) \tag{19}$$

Now $G(\overline{D} \cup \{y\})$ is the set of graphs, without loops, that join every point of D either to $\overline{D}$ or y. $G_{x,y}(\overline{D} \cup \{y\})$ is the set of such graphs that connect x to y. Clearly $G(\overline{D} \cup \{y\}) = \bigcup_{x \in D} G_{x,y}(\overline{D} \cup \{y\})$. Hence (19) can be re-written as

$$H(D) = \min_{g \in G(\overline{D})} V(g) - \min_{y \in D} \min_{g \in G(\overline{D} \cup \{y\})} V(g) \tag{20}$$

In order to bound $H(D)$ it is therefore enough to provide bounds to the first term in (20) and, for each y, the minimand in the second.

Consider therefore a point $y \in D$ and a path linking it to $\overline{D}$, with notation as in the definition of the modified co-radius. Given such a path one can construct a $G(\overline{D})$-graph as follows. Let $\Pi = \{p\} \cup \bigcup_{i=2}^{r-1} D(L_i)$ be the union of the path and the domains of attraction of the intermediate limit sets. For a path of minimum modified cost it can be assume that each L_i occurs only once. All points not belonging to p in $D(L_i)$ $(1 < i < r)$ can be connected to the path at zero cost. Now construct a graph that connects all points in $D \backslash \Pi$ to $\overline{D} \cup \Pi$ at minimum cost (that is a optimal $G(\overline{D} \cup \Pi)$-graph). Call this graph k. It is easy to see that the union of k and the graph within Π is a $G(\overline{D})$-graph. Hence

$$\min_{g \in G(\overline{D})} V(g) \leq V(p) + V(k) \tag{21}$$

Now construct a lower bound to the second term in (20) as follows. Suppose g^* is an optimal $G(\overline{W} \cup \{y\})$-graph. Restricted to the set $D \backslash \Pi$, it is certainly a $G(\overline{W} \cup \Pi)$-graph, since $y \in \Pi$. It may not be optimal, however, so its cost is at least $V(k)$. Consider the graph restricted to the set Π. All points in the set Π must be connected to y or $\overline{D}$. It follows that g^* must connect each limit set L_i, $i = 2, \ldots, r-1$ to the exterior of $D(L_i)$. The cheapest way this can happen is if from each L_i exit occurs along a path with cost equal to the radius, all other transitions following the flow of P_∞ and so having zero cost. All other transitions in Π must have cost at least zero.

It follows that any feasible $G(\overline{D} \cup \{y\})$-graph restricted to Π must have cost at least $\sum_{i=2}^{r-1} R(L_i)$. Hence,

$$\min_{g \in G(\overline{D} \cup \{y\})} V(g) \geq V(k) + \sum_{i=2}^{r-1} R(L_i) \tag{22}$$

Putting (21) and (22) together therefore shows that

$$\min_{g \in G(\overline{D})} V(g) - \min_{g \in G(\overline{D} \cup \{y\})} V(g) \leq \min_{p \in S(y,\overline{D})} V^*(p) \tag{23}$$

Optimizing over y one obtains

Proposition 2 *If $\overline{D}$ is a union of limit sets, then $H(D) \leq CR^*(\overline{D})$.*

As seen in Example 4, the co-radius measure may not be tight. Nevertheless it may often place a useful simple upper bound on waiting times. Indeed one might argue that the path minimisation problems it involves are rather simpler than the graph minimisation problems involved in the more precise bound of Lemma 6. While this may be the case in some examples, the next section provides an economic example of learning in potential games where the approach of Lemma 6 produces simpler and more precise bounds than does the modified co-radius formula. Section 8 shows how the computations involved in Lemma 6 can be simplified.

Ellison (2000) also shows that if D is a union of limit sets for which $R(D) > CR^*(D)$, then all probability mass concentrates on D for large β. Section 9 provides an alternative derivation of this result. This result is only a sufficient condition for equilibrium selection and the next section provides an economic example where it lacks power.

6. Learning in Potential Games

This section examines a model of learning in potential games. These are games in which payoff differences to all players are equal to the first differences of a single function. This notion is due to Monderer and Shapley (1996) and has attracted a fair amount of attention in recent years. Blume (1994) and Blume (1995) analyse equilibrium selection in such games but only give numerical estimates of waiting times. The current section gives an analysis which allows for simple formulae for expected waiting times.

6.1 The Model

The general framework considered is a symmetric game played by N players. The extension to asymmetric games is straightforward. There are M strategies available to each player. Let $x = (x_1, x_2, \ldots, x_M)$ be a vector whose ith entry gives the fraction of players playing the ith strategy. $\pi_i(x)$ denotes the payoff to a player who plays strategy i when the aggregate distribution of play is x.

$\pi_i(x)$ can be interpreted in different ways. For example, in the standard 'random-matching' framework used by Kandori et al. (1993), it would be assumed that each period players are uniformly and randomly matched against one another to play a certain game. $\pi_i(x)$ then denotes the average payoff of strategy i against the average distribution of play. In this case $\pi_i(x)$ is linear. In more general frameworks it need not be. For example, in games representing congestion externalities or search (cf. Diamond (1982)) payoffs pay depend non-linearly on the distribution of strategies.

Suppose that a player changes strategy from i to j. The aggregate distribution of play changes from x to $x^{j\backslash i}$, say. It is assumed that there exists a function U such that

$$\pi_j(x^{j\backslash i}) - \pi_i(x) = U(x) - U(x^{j\backslash i}) \tag{24}$$

for all i, j. A game will be said to be a potential game if this is so.

Note that in (24), the opposite sign is taken to the standard definition. A Nash equilibrium of the game in the current definition corresponds to a point where the potential, U, is a local minimum. Economists prefer to maximise things and so the standard definition takes the opposite sign for U. In other disciplines minimization is the order of the day (for example, find a configuration of minimum energy), so this usage is in this spirit. It has the advantage that, as will be seen shortly, it relates easily to the ideas in Section 3.

6.2 The Learning Scheme

The learning scheme studied is as follows. Each period one agent is selected at random to revise strategy. He picks a strategy at random to compare with his current strategy. It is assumed he observes the aggregate distribution of play and picks the better of the two strategies, but may make errors. These may represent bounded rationality or perturbations to payoffs. These errors are represented by the logit model.

More precisely, it is assumed that the state evolves according to the following transition matrix for $x \neq y$:

$$p_\beta(x,y) = q(x,y)\frac{\exp(\beta\pi_j(y))}{\exp(\beta\pi_i(x)) + \exp(\beta\pi_j(y))} \tag{25}$$

where

$$\begin{cases} q(x,y) > 0 & \text{if } y_i = x_i - \frac{1}{N},\ y_j = x_j + \frac{1}{N}, \text{ some } i,j \\ q(x,y) = 0 & \text{otherwise} \end{cases} \tag{26}$$

q represents the probability that, when the current state is x, a player employing strategy i is chosen to revise strategy and chooses strategy j to consider switching to. The logit term represents the probability that he in fact switches to it. The larger β, the more likely the agent is to choose the better strategy.

Using (24), this implies that, in the framework of Section 2, the chain has cost function (for $x \neq y$):

$$V(x,y) = \begin{cases} \beta\left(U(y) - U(x)\right)^+ & \text{if } y_i = x_i - \frac{1}{N},\ y_j = x_j + \frac{1}{N}, \text{ some } i,j \\ \infty & \text{otherwise} \end{cases} \tag{27}$$

$(x)^+$ denotes $\max\{x, 0\}$.

The assumption that only one player revises strategy at a time is shared by Blume (1994) and Blume (1995). It has the advantage that under it the chain is reversible (see for example Wolff (1989)) and so the equilibrium distribution of the chain is easily determined.

For $M = 2$, this learning scheme is identical to that considered by Blume. If $M > 2$, it differs in that Blume (1995) assumes that players use a logit choice rule defined over all M possible choices, whereas here it is assumed that they only consider a pair of strategies at a time. This choice does not affect the invariant distribution of the chain, but the current formulation makes characterization of waiting times slightly easier. This is discussed further in Section 7. Blume (1995)'s formulation implies that players perform a perturbed global optimization. The current model assumes rather less ability to optimize, in that agents only perform pairwise comparisons at a time.

In the literature on optimization, (27) is referred to as a Metropolis algorithm (see, for example, Catoni (1999)).

6.3 Results

The first result states that in this model the quasi-potential function of Section 3 coincides, up to a constant, with the potential function.

Proposition 3 *If V is given by (24), then the quasi-potential function, $\tilde{U}(x)$, given by Lemma 4 equals $U(x) - \min_{x \in E} U(x)$. It follows that as $\beta \to \infty$ all weight of the invariant distribution concentrates on the set of points of minimum potential.*

This can be proved directly, using reversibility arguments as in Blume (1995), but a simple proof also can be given using the ideas of Section 3. The following property is immediate from (24)

Lemma 9 *For any $x \neq y$, $V(x,y) < \infty$ if and only if $V(y,x) < \infty$ and*

$$U(x) + V(x,y) = U(y) + V(y,x) \tag{28}$$

Suppose that g is an optimal $G(\{x\}$-tree. Any $y \neq x$ must be linked to x by a path, say p. Consider reversing p and linking x to y. This results in a $G(\{y\})$-tree, g'. By Lemma 9, the difference in cost between g and g' is $U(y) - U(x)$. It follows that the minimum-cost $G(\{y\}$-tree has cost at most $U(y) - U(x)$ greater than the optimal $G\{x\}$-tree. Reversing the role of x and y in this argument shows that that the costs differ by exactly $U(y) - U(x)$, which establishes the result.

Note that the only properties of (24) used in this argument are those given by Lemma 9. It follows that Proposition 3 remains true for any other learning dynamic satisfying it. For example, if players used the logit rule on all M choices simultaneously, it is easy to check that Lemma 9 and so Proposition 3 continue to hold. This is Blume (1995)'s result.

To estimate waiting times, consider a set D. For $x \in D$ and $y \notin D$, consider the set of all paths, $\Phi(x,y)$, linking x and y with finite cost: if $p = (z_1, z_2, \ldots, z_T)$, then

$V(z_i, z_{i+1}) < \infty$ all i. For such a path p let

$$H(p) = \sup_{1 \leq t \leq T} U(z_t) \tag{29}$$

be the greatest value of the potential encountered on the way from x to y and let

$$H(x, y) = \inf_{p \in \Phi(x,y)} H(p) \tag{30}$$

be the least possible greatest value of the potential reached on the way from x to y. Then

Proposition 4 *For any set D,*

$$H(D) = \sup_{x \in D} \inf_{y \in \overline{D}} H(x, y) - U(x) \tag{31}$$

Intuitively, $H(x, y) - U(x)$ gives the greatest difference in potential encountered moving between x and y and so measures the energy required to move from x to y (given that the process likes to be at points of low potential). Minimizing over y gives the minimal energy required to escape from D starting at x, and maximising over x finds the worst case.

This result is proved in the Appendix in the case $M = 2$. The general case is not difficult but requires some tedious induction arguments. A proof can be found in Hwang and Sheu (1989) and Hwang and Sheu (1992). The result also follows from the more general ones given in Section 7.

6.4 Discussion

The results above are illustrated by the following example.

Example 5 Consider the transition diagram in Figure 4. This can thought of as deriving from the dynamic in (27) under the potential function shown in Figure 5, with only the 5 values shown feasible (so $N = 5$). Only nearest neighbour transitions are possible. The local minima are at a, c and e. The global minimum is at e, so this is the selected equilibrium. Set $D = \{a, b, c, d\}$. It is straightforward to check that $W_D = 6$ and the worst case for exit from the domain is starting from a, with $M_D(a) = 1$ (realized by the graph $\{a \leftarrow b\} \cup \{c \rightarrow d \rightarrow e\}$). Hence $H_D = 5$. This is equal to the potential difference between a and b, which as can be seen in Figure 5 is the greatest (positive) difference between peak and trough in potential, walking in the direction of e.

For future reference, note that in this example the co-radius of $\{e\}$ is 5, so the co-radius bound on waiting times is exact. Nevertheless, the 'radius/co-radius' bound does not work here: the radius of $\{e\}$ is only 3, which is less than the co-radius. This example is discussed further in Section 8.

For completeness, note that the co-radius bound need not be exact here:

Example 6 Consider the example in Figure 6. The limit sets are a,c,e,g and i. It is straightforward to see that the selected equilibrium is i. The co-radius of i is $10 + (1 - 1) + (2 - 2) + (2 - 1) = 11$. For $D = \{a, b, c, d, e, f, g, h\}$, $W_D = 15$. It is straightforward to check that the worst case for $M_D(x)$ is at $x = a$ and that $M_D(a) = 5$ (realized by either by $b \to c \to \ldots \to i$ or $\{b \to a\} \cup \{c \to d \to \ldots \to i\}$). If follows that $H_D = 15 - 5 = 10$, which is less than the co-radius.

Note that in Example 6, there are 5 limit sets. In a model derived from random matching in a symmetric 2×2 games, one would expect at most 3 limit sets, corresponding to 2 pure and one mixed equilibrium. If one allows non-linear interaction, for example due to congestion or search externalities, there can be many equilibria even in a one-dimensional model.

One feature that may seem surprising is that population size does not explicitly appear in the waiting times found in Proposition 4. In fact, in most games it enters via the potential function. Suppose that the players are randomly matched each period to play a symmetric two-player game. Assume that this underlying game has a potential Q, that is for any player, and strategies i, j, k

$$\pi(i, j) - \pi(k, j) = Q(k, j) - Q(i, j) \tag{32}$$

then it is easy to check that

$$U(x) = \frac{N}{2} \sum_i x_i \sum_j x_j^i Q(i, j) \tag{33}$$

where, if $x_i > 0$, x_j^i gives the proportion of other players playing strategy j from the point of view of a player playing strategy i, is a potential in the sense of (24). Holding the underlying game fixed, therefore, waiting times are indeed dependent on N.

7. An Alternative Characterisation of Waiting Times

The previous section outlined a characterization of waiting times in terms of differences in potential energy. It is in fact possible to develop a general characterization on these lines. This is not so useful in economic models, as opposed to optimization, but it is intuitive and will be briefly sketched out. The reader who is not interested may skip directly to Section 8.

Let U be the quasi-potential function of Lemma 4. In the framework of Proposition 4, replace the definition of $H(p)$ in (29) by

$$H(p) = \sup_{1 \leq t < T} U(z_t) + V(z_t, z_{t+1}) \tag{34}$$

where $p = (z_1, z_2, \ldots, z_T)$ is a path. If $V(x, y) = (U(x) - U(y))^+$, it is easy to see that this reduces to the previous definition. Other definitions remain unchanged. Then

Proposition 5 *For any set D,*

$$H(D) = \sup_{x \in D} \inf_{y \in \overline{D}} H(x, y) - U(x) \tag{35}$$

This result is due to Trouvé (1996). A simpler proof can be found in Catoni (1999).

This is a fairly intuitive result as it relates waiting times to the change in quasi-potential function. Nevertheless, its usefulness is limited as one still needs to determine the quasi-potential.

In optimization applications, the dynamic is often designed so that the quasi-potential function coincides with function to be minimised. In this case, Proposition 5 is quite useful. This has been seen in the previous Section in the case of potential games with the dynamic given in (27). This coincidence also holds for other dynamics. For example, if the logit rule is applied to all choices simultaneously, (27) is replaced by

$$V(x, y) = \begin{cases} \beta \left(U(y) - U^*(x) \right) & \text{if } y_i = x_i - \frac{1}{N},\ y_j = x_j + \frac{1}{N}, \text{ some } i, j \\ \infty & \text{otherwise} \end{cases} \tag{36}$$

where $U^*(x)$ denotes the potential of the best choice given the current distribution of play. It is easy to check that Lemma 9 still holds and so one can also characterise waiting times using Blume (1995)'s formulation, though the final form is not as neat.

Nevertheless, even in potential games with these choice rules, Lemma 9 ceases to hold as soon as one allows simultaneous changes in strategy. While it may be arguable whether simultaneous changes in strategy are plausible, it is a commonly made assumption in models of learning and one would like a method that can deal with these more easily. This is discussed in the next section.

8. Reduced Processes

A difficulty with all the literature on equilibrium selection, including this paper, is that the graph calculations involved are not trivial. When the population size is large, the state space is of a very high dimension. Young (1993) and Kandori and Rob (1995), have noted that the tree calculations involved in finding the equilibrium selected can be simplified by reducing them to ones simply involving the limit sets. The number of limit sets is usually of a much smaller order than, and often independent of, the population size. This section shows that a similar reduction is possible for waiting times.

The Section also shows that one can take the idea of reduction further and eliminate limit sets in a systematic way as well. This allows one to find the equilibrium selected without using tree calculations. Waiting times can also be so calculated. These reductions require the solution of a sequence of shortest path problems, but this kind of problem is relatively straightforward to solve.

The idea of reducing the state space has appeared in one guise or another in varying degrees of generality in different places: in addition to Kandori and Rob (1995) and Young (1993) see, for example, Cerf (1996), Sonin (1999) and, implicitly, Freidlin and Wentzell (1998). The clearest and most general statement seems to be in Catoni et al. (2000) and Catoni (2002). Propositions 6, 7 and 8 are straightforward consequences of results found there. The proofs offered here, however, are more elementary and do not use any results beyond those found in Section 3. The application to limit sets in the Corollary to Proposition 7 and the applications in subsequent sections are new.

8.1 Reduction

Consider the framework of Section 2. Let A be a subset of the state space, E. Let $\tau(n)$ denote the n-th time the process X hits A. The process $X^A_n = X_{\tau(n)}$ is the process X only observed at moments when it lies in A. It is easy to check that

Lemma 10 *X^A is a Markov chain with invariant distribution $\mu^\beta/\mu^\beta(A)$. If $B \subseteq A$ then $(X^A)^B = X^B$.*

The obvious last remark in the Lemma will be used implicitly repeatedly.

If A is the set of all limit sets of P_∞, then the least cost behaviour is to leave $E\backslash A$ rather than cycle around inside it. For, since the process converges to A with probability 1, any point of $x \in E\backslash A$ can be linked to it by a path of zero cost. The following definition generalises this property:

Definition 3 *A set A has the property (MP) if for any $x \in E\backslash A$, there exists a path $z_0, z_1, \ldots, z_T$, with $z_0 = x$, $z_T \in A$, and for all $0 \leq t < T$, $z_t \notin A$ and*

$$V(z_t, z_{t+1}) = \min_{y \neq z_t} V(z_t, y) \tag{37}$$

(37) simply states that any point outside A can be linked to it by a path which takes takes the least cost way to leave every point on it. Note that the hypothesis is trivially satisfied if $E\backslash A$ has only one element.

Proposition 6 *(a) If A has property (MP) then X^A has transitions governed by cost function*

$$\begin{aligned} V^A(x,y) = \min_{z,T} \{ & V(z_0, z_1) + \sum_{t=2}^{T} V(z_{t-1}, z_t) - H(z_{t-1}) : \\ & z_0 = x,\ z_T = y,\ z_t \notin A\ t \neq 0, T,\ z_t \neq z_{t-1}\ t = 2, \ldots, T \} \end{aligned} \tag{38}$$

where $H(z) = \min_{y \neq z} V(z, y)$ and x, y are in A.

(b) If the quasi-potential U is known on A, then for $x \in E\backslash A$

$$\begin{aligned} U(x) = \min_{z,T} \{ & U(z_0) + V(z_0, z_1) + \sum_{t=2}^{T} V(z_{t-1}, z_t) - H(z_{t-1}) : \\ & z_0 \in A,\ z_T = x,\ z_t \notin A,\ z_t \neq z_{t-1}\ t = 1, \ldots, T \} - H(x) \end{aligned} \tag{39}$$

The proof is given in the Appendix. $H(z)$ is nothing other than the the height of z defined in Example 1 of Section 3. The Proposition states roughly that the cost of getting from x to y is the cost of the least cost path through $E\backslash A$, where in $E\backslash A$ costs of transitions are measured relative to the least cost step that could be taken at each stage. $T = 1$ corresponds to moving directly from x to y without passing through $E\backslash A$ and has cost $V(x, y)$.

Part (b) shows how to compute the quasi-potential on $E\backslash A$, if desired. Often interest focusses on A, so this is not necessary. Even if $E\backslash A$ is of interest it may be easier computationally to calculate the quasi-potential on A and then step backwards rather than working directly with the whole of E.

Since any one point set has property (MP) it follows, applying Proposition 6 recursively, that X^A on any domain has a cost function V^A. Hence using Lemma 10, with $U(x)$ the quasi-potential function of Lemma 6, one obtains

Lemma 11 *For any $A \subseteq E$, X^A has transitions governed by some cost function V^A and has quasi-potential function U^A given by $U^A(x) = U(x) - \min_{y\in A} U(x)$.*

The point is that property (MP) allows one to identify V^A simply. U^A may also be rather easier to determine than all of U.

If A is the set of all limit sets, then in addition to satisfying property (MP) it has the property that the process spends much less time in $E\backslash A$ than A. Intuitively, therefore, little is lost in studying the process restricted to A. The next result generalises this idea.

For any sets C, D with $D \subseteq C$ and $x \in C$, let $T_x^C(D)$ be the first time the process reduced to C hits D starting from x.

Proposition 7 *If A has property (MP) and*

$$\max_{x\in E\backslash A} H(\{x\}) \leq \min_{y\in A} H(\{y\}) \tag{40}$$

then

$$\min_{x\in A} U(x) \leq \min_{x\in E\backslash A} U(x) \tag{41}$$

with strict inequality if (40) holds strictly (but not only if).

Moreover, if (40) holds, then for any $B \subseteq A$ and all $x \in A$,

$$\lim_{\beta\to\infty} \frac{1}{\beta} \ln ET_x^A(B) = \lim_{\beta\to\infty} \frac{1}{\beta} \ln ET_x^E(B) \tag{42}$$

and for any $x \in E\backslash A$

$$\lim_{\beta\to\infty} \frac{1}{\beta} \ln ET_x^E(B) \leq \lim_{\beta\to\infty} \frac{1}{\beta} \max_{y\in A\backslash B} \ln ET_y^E(B) \tag{43}$$

The proof is in the Appendix. If (40) holds strictly, then all the minima of U are contained in A, so nothing is lost in confining attention to A. (42) implies that waiting times to hit subsets $B \subseteq A$ can be computed by simply looking at X^A, if the process starts from $x \in A$. If $x \notin A$ this is clearly not so, but (43) implies that the worse cases for hitting B occur when $x \in A$, so to obtain a worst case bound it can be assumed that $x \in A$.

The following is an immediate corollary of Proposition 7:

Corollary *Suppose that P_∞ has limit sets L_i, $i = 1, 2, \ldots, m$ and let $l_i \in L_i$ each i. Let $L = \bigcup L_i$. Then consider the reduced process on $\tilde{L} = \{l_1, \ldots, l_m\}$. It has transitions governed by*

$$V^{\tilde{L}}(l_i, l_j) = \min_{z,T} \{ \sum_{t=0}^{T} V(z_{t-1}, z_t) : z_0 = l_i, z_T = l_j, z_t \notin \tilde{L} \; t \neq 0, T \} \tag{44}$$

Moreover,

$$\arg\min_{x \in E} U(x) = \bigcup \{ L_i : l_i \in \arg\min_{l \in \tilde{L}} U^{\tilde{L}}(l) \} \tag{45}$$

and if $x \in L_i$

$$\lim_{\beta \to \infty} \frac{1}{\beta} \ln ET_x^E(L_j) = \lim_{\beta \to \infty} \frac{1}{\beta} \ln ET_{l_i}^{\tilde{L}}(l_j) \tag{46}$$

(45) is Kandori and Rob (1995) and Young (1993)'s result. (46) implies that waiting times between limit sets can also be calculated in this way.

A proof of the Corollary can be found in the Appendix. Note that the height of every point outside $\tilde{L}$ is zero as every point outside L can be connected to L by path of null cost, and any points within a given limit set can be connected at zero cost (Lemma 7). Hence the absence of H from (44). Note also from Lemma 7 that it costs a positive amount to leave a limit set, so one l_i is retained from each limit set (though the choice is arbitrary).

A simple link to the ideas of Section 5, is

Lemma 12 *In the notation of the Corollary, $H^{\tilde{L}}(\{l_i\}) = R(L_i)$.*

This is clear since the radius of L_i measures the cost of the cheapest path to leave $D(L_i)$. Once it leaves $D(L_i)$ it can reach some limit set for free.

Note that the definition in (44) allows elements of limit sets other than the $\{l_j\}$ to be used in a path. The formula therefore differs slightly from that in Kandori and Rob (1995) and Young (1993) when limit sets contain more than one element. In calculating optimal graphs, however, this is unimportant, as is noted where relevant in the proofs.

The formula in (38) is reminiscent of the modified cost encountered in Section 5. This is not coincidental. This link to the ideas of Section 5 is considered in the next Section.

The procedure above can clearly be applied recursively and this is explored in more detail in the next sub-section. Some examples are considered in the final sub-section.

8.2 Iterative Reduction

As noted, above the reduction procedure can applied recursively. More specifically, let $A_0 = E$ and $V^{A_0} = V$. Let $l(0) = 0$. Then apply the following procedure, for $k = 1, 2, \ldots$:

1. Find a set $A_k \subset A_{k-1}$ satisfying (MP) and (40), with $V^{A_{k-1}}$ as cost function and $H^{A_{k-1}}$ as height function, such that $A_{k-1} \backslash A_k$ is non-empty.
2. Order the elements of $A_{k-1} \backslash A_k$ in increasing order of height, under $V^{A_{k-1}}$, (resolving ties arbitrarily) and label them as $x_{l(k-1)+1}, \ldots, x_{l(k)}$, where $l(k) = l(k-1) + |A_{k-1} \backslash A_k|$.
3. Compute V^{A_k} from $V^{A_{k-1}}$ using Proposition 7.
4. If A_k is a singleton label this $x_{|E|}$ and STOP. Else return to 1.

To understand this note first that step 1 of the algorithm can always be applied: a singleton set of minimum height always satisfies the hypotheses. The requirement that $A_k \backslash A_{k-1}$ be non-empty stops everything being deleted (for example A_k might be single limit set with all points having height zero). Points are successively deleted, and the domain under consideration successively shrunk, until only a single point $x_{|E|}$ remains.

Then

Proposition 8 *(a) $x_{|E|}$ is a point of minimum quasi-potential.*

(b) If for all k, $\min_{x \in A_k} H^{A_{k-1}}(\{x\}) > \max_{y \in A_{k-1} \backslash A_k} H^{A_{k-1}}(\{y\})$, then $x_{|E|}$ is the unique point of minimum quasi-potential.

(c) $H(E \backslash \{x_{|E|}\}) = \max_{y \in A_{k^*}, y \neq x_{|E|}} H^{A_{k^*}}(\{y\})$, where k^* is the last k such that A_k is not a singleton. $H(E \backslash \{z\})$ is the same for all $z \in \arg\min U(y)$.

The proof is in the Appendix.

Part (a) states that $x_{|E|}$ is always a point of minimum quasi-potential. It and part (b) are immediate corollaries of Proposition 7. Note that (b) is required to hold at each step. The result is false in general without this requirement. On the other hand, part (b) is only a sufficient condition for uniqueness. In particular, there is some degree of flexibility in the choice of A_k.

If (b) is not satisfied, part (b) of Proposition 6 can be used to compute the quasi-potential of all points working backwards from $U(x_{|E|}) = 0$ to check if there are any other points of minimum quasi-potential.

Part (c) provides a bound on waiting times. It says that the worst expected time to hit $x_{|E|}$ starting from any point in E can be found by looking at the heights of the last points to be deleted. It is a consequence of (43).

When $x_{|E|}$ is the unique equilibrium selected, or more generally by the Corollary to Proposition 7 part of the unique limit set, this provides a bound on the time for this equilibrium to become relevant. When there are multiple equilibria selected, the second part of (c) states that it measures the time by which *all* of these become relevant. One might be interested in a weaker measure of relevance, the time at which *some* selected equilibrium will have been played (that is $H(E\backslash \arg\min U(x))$). Calculating this requires looking further back in the reduction process.

Note that is not necessary to know all the limit sets of P_∞ to apply this algorithm, although if they are known initial application of the Corollary to Proposition 7 is helpful. This may be advantageous if the long-run behaviour of P_∞ is complex.

8.3 Examples

This section illustrates the results above. First consider again:

Example 4 (See Figure 3.) As discussed in Sections 4 and 5, if $D = \{x, y, z\}$, then $H(D) = 3$ and this is attained at x, y and z. Consider applying reduction to this example. The points of minimal height are y and z, with height 1. One of these may be eliminated, say y. Using (44) this yields Figure 7 (the path x to y to w has reduced cost $2 + (3 - 1) = 4$, hence the new cost of transitions x to w). On this subset, with $D' = \{x, z\}$, the optimal graph that links D' to w is $x \to z \to w$, so $W_{D'} = 5$. $M_{D'}(x) = M_{D'}(z) = 2$ and this cost is realized by the graph $x \to y$. Hence $H_{D'}(x) = H_{D'}(y) = 3$, so waiting times to hit w are as before, in agreement with Proposition 7.

D' itself can be further reduced. x is now the point of minimum height. Eliminating it leaves the domain $D'' = \{z\}$ (see Figure 8). Clearly $H(D'') = 3$, in agreement with the above.

Note one feature of this example. At the first stage either y or z could have been eliminated but not both. The set $\{y, z\}$ does not satisfy property (MP) as for both y and z the least cost path leaving each point leads to each other, so only one of them is eliminated.

The above illustrates the application to waiting times. Now consider equilibrium selection:

Example 5 (See Figure 5.) As noted in Section 6, the selected equilibrium is e, but its radius is less than its modified co-radius. Consider applying reduction to this model. The limit sets are a, c and e. Eliminating b and d yields Figure 9. The element with least height is now c (with height 1). Eliminating c and applying (44) yields Figure 10. a now has least height and eliminating it leaves $\{e\}$. By Proposition 7 or 8, the quasi-potential attains its minimum at e. Moreover since at each stage the points remaining all have strictly greater height than those eliminated, (40) holds strictly. Hence $\{e\}$ is the unique point at which the quasi-potential is minimised. Note that the height of a in Figure 10, the penultimate stage of reduction, is 5, so $H(E\backslash\{e\}) = 5$. This agrees with the direct calculation in Section 6.

The problem for the radius/co-radius formula in this example is that e is not the point with largest radius. a has the largest radius, 5, while e has only radius 3. a is very hard to escape from. It is not the selected equilibrium, however, as once the process escapes it is very unlikely to return: it will rapidly hit c and then e. Once it is at e it is unlikely to escape, but if it ever does so it is highly unlikely to get past c (since $c \to b$ costs 7) but will rather return to e.

One could in fact deduce that e is selected by two applications of the 'radius/co-radius' formula. First $\{a, e\}$ has radius 3 and co-radius 1. This rules out c. Second, $\{c, e\}$ has radius 7 and co-radius 5. This rules out a. Ellison (2000) presents a similar example where grouping works. It is not clear, however, if there is any systematic procedure for finding such groupings.

9. The 'Radius/Co-Radius' Formula

This section applies some of the ideas of the last section to the 'radius/co-radius' criterion for equilibrium selection. It also shows how the waiting time bounds of Section 5 follow from more general results.

The intuition behind much of the equilibrium selection literature, and the 'radius/co-radius' formula, is that a point is stable if it is harder to escape from it than return to it. This sub-section shows how this idea leads to a proof of the formula.

Suppose that one wishes to show only points in the set D can be selected. Let $l \in E \backslash D$ and consider the process reduced to $F = D \cup \{l\}$. By Lemma 10, μ^β is invariant for X^F, so

$$\mu_\beta(l) \sum_{d \in D} p^F_\beta(l, d) = \sum_{k \in D} \mu^\beta(k) p^F_\beta(k, l) \tag{47}$$

Hence

$$U(l) + \min_{d \in D} V^F(l, d) = \min_{k \in D} U(k) + V^F(k, l) \tag{48}$$

It follows that

$$\begin{aligned} &\min_{d \in D} V^F(l, d) < V^F(k, l) \qquad \forall k \in D \\ \implies &\exists k \in D \quad \text{with} \quad U(k) < U(l) \end{aligned} \tag{49}$$

In other words l will not be selected if it is always harder to reach l from D than vice-versa in the process reduced to F. To show this it is enough to bound $V^F(k, l)$ below for $l \in D$ and $\min_{l \in D} V^F(k, l)$ above. Generalised versions of the radius and co-radius bounds respectively do precisely this.

For domains A satisfying (MP), Proposition 6 gives an exact formula for A. (44) does not hold for general domains, but the right-hand side turns out to provide an upper bound:

Lemma 13 *For any domain A,*

$$\begin{aligned} V^A(x, y) \le \min_{z, T} \Big\{ &V(z_0, z_1) + \sum_{t=2}^{T} V(z_{t-1}, z_t) - H(z_{t-1}) : \\ &z_0 = x,\ z_T = y,\ z_t \notin A\ t \neq 0, T,\ z_t \neq z_{t-1}\ t = 2, \ldots, T \Big\} \end{aligned} \tag{50}$$

where $H(z) = \min_{y \neq z} V(z,y)$ and x, y are in A.

Applied to $\tilde{L}$, this upper bound is none other than the modified cost of connecting two limit sets.

Similarly the radius bound comes from the following rather obvious lower bound (see Appendix for proof):

Lemma 14 *If $y \in E \backslash D$ and $F = D \cup \{y\}$, then for any $x \in D$*

$$V^F(x,y) \geq \min_{u \in D, v \in E \backslash D} V(u,v) \tag{51}$$

Applying these Lemmas to the process reduced to $\tilde{L}$ and using the argument of (49) (see Appendix):

Proposition 9 *Let Ω be of a union of limit sets. If $R(\Omega) > CR^*(\Omega)$ then* $\lim_{\beta \to \infty} \mu_\beta(E \backslash \Omega) = 0$.

Similarly, the co-radius and radius bounds on waiting times are examples of general upper and lower bounds applied to the process reduced to $\tilde{L}$.

The following gives an upper bound to waiting times.

Lemma 15 *Let A be any set then*

$$\begin{aligned} H(A) \leq \min_{\substack{z,T \\ x \in A \\ y \in E \backslash A}} \Big\{ & V(z_0, z_1) + \sum_{i=2}^{T} V(z_{i-1}, z_i) - H(z_{i-1}) : \\ & z_0 = x,\ z_T = y,\ z_i \in A\ i < T,\ z_i \neq z_{i-1}\ i = 1, \ldots, T \Big\} \end{aligned} \tag{52}$$

The proof is in the Appendix. The intuition behind it is that any $G(E \backslash A)$-graph must contain a path linking each x to the exterior. A $G(E \backslash A \cup \{x\})$-graph can be constructed by deleting the link leaving x and connecting all other points in the path to $E \backslash A$ or x. The cost of leaving each z_{t-1} is at least $H(z_{t-1})$, so this provides an upper bound on the saving in cost possible in (10) of Lemma 6. The full proof is in the Appendix.

Similarly, the following rather obvious result provides a lower bound on waiting times (see Appendix for proof):

Lemma 16 *For any domain A and $x \in A$,*

$$H_A(x) \geq \min_{\substack{x \in A \\ y \in E \backslash A}} V(x,y) \tag{53}$$

Propositions 1 and 2 follow from applying Lemmas 16 and 15 respectively to the process reduced to $\tilde{L}$.

Corollary to Lemmas 15 and 16 *Propositions 1 and 2 hold.*

The finer details can be found in the Appendix.

The alert reader may be puzzled by the following. Suppose $x \in A$ and $y \in E \backslash A$. Consider the process reduced to $F = \{x, y\}$. On a two point set, the exit time to travel from x to y is (to asymptotic logarithmic order) $V^F(x, y)$. Hence this is bounded by (50). Yet is was observed in Example 2 in Section 5.2 that the modified cost does not always bound waiting times to exit from a domain, only from the worst possible starting point. The resolution is that waiting times for the original process can in general only be calculated from the reduced process when (40) applies.

10. Learning in Games with Strategic Complementarities

This section applies the preceding results to equilibrium selection in games with strategic complementarities. These were studied by Kandori and Rob (1995) but the current approach allows for more detailed results.

The full framework can be found in Kandori and Rob (1995); here only the necessary details are sketched out. Consider a symmetric two-person game with $i = 1, 2, \ldots, n$ strategies and payoff function π. It is assumed that payoffs are supermodular in that $\pi(k, j) - \pi(i, j)$ is strictly increasing in j for $i < k$. For technical reasons a kind of 'continuity' assumption is made: it is assumed that if i is a best response to some mixed strategy α and j to α', then if $i < k < j$, there is $\lambda \in (0, 1)$ such that k is a best response to $\lambda\alpha + (1 - \lambda)\alpha'$.

There is population of M players. Each period players are randomly matched and each pair plays this game. At the end of the period each player may have the opportunity to adjust his strategy, and this probability may depend on the strategy he played as well as his payoff. If he has the chance to adjust he plays a best response to the current distribution of his opponents' play with probability $1 - \exp(-\beta)$, with probability $\exp(-\beta)$ he chooses according some fully mixed probability distribution $m_1, m_2, \ldots, m_n$.

Kandori and Rob (1995) show that the limit sets of this adjustment rule correspond to the pure strategy equilibria of the game (that is all agents play these strategies). Denote these equilibria by N and consider the game reduced to limit sets. Kandori and Rob (1995) show that for large enough M, transition costs between limit sets take the following form

$$V(i, j) = \max_{i \leq k < j} V(k, k+) \qquad i < j \tag{54}$$

$$= \max_{j < k \leq i} V(k, k-) \qquad i > j \tag{55}$$

where $k+$ and $k-$ denote the next equilibrium in N above and below k respectively. If there is no greater or lower equilibrium define the corresponding quantity to be infinite. In other words, adjustment costs depend on the costs of transitions between neighbouring equilibria. These are realised by mutating a large enough number of players to the largest (or smallest) equilibrium strategies and then following best response adjustment.

Consider the following example given by Kandori and Rob (1995) (the underlying payoff matrix can be found there):

Example 7 There are three equilibrium strategies $i = 1, 2, 3$ and for adjacent pairs let $V(i,j) = \beta_{ij}$, with $\beta_{12} = 3/7$, $\beta_{21} = 1/2$, $\beta_{23} = 1/4$ and $\beta_{32} = 1/3$. It follows that $V(1,3) = \max\{\beta_{12}, \beta_{23}\} = \beta_{12}$ and $V(3,1) = \beta_{21}$. Hence $V(3,1) = V(2,1) > V(1,3) = V(1,2) > V(3,2) > V(2,3)$.

Consider applying the reduction procedure of Section 8.2 to this example. Point 2 has the smallest height ($V(2,3)$), so remove it. Transitions through 2 do not reduce cost, so on the set $\{1,3\}$, $V^*(1,3) = V(1,3)$ and $V^*(3,1) = V(3,1)$. 1 now has the smallest radius, so it can be deleted. These leaves 3 and since all deletions were strict, it is the unique equilibrium selected. Moreover the logarithm of the worst expected time to hit 3 tends to $V(1,3)$.

This argument to show that 3 is the unique long-run equilibrium is distinctly shorter than the tree method used by Kandori and Rob (1995). It also yields the waiting time bound, which they do not obtain.

Note that the radius of 3 in the original process $V(3,2)$ is less than its (modified) co-radius $V(2,1)$, so the 'radius/co-radius' formula lacks power here.

The procedure used in the example can always be used. More precisely

1. Let $N^* = N$
2. For each $i \in N^*$ let $V_i^* = \min\{V(i,i+), V(i,i-)\}$, where adjacency is defined relative to N^*.
3. Let $i \in N^*$ be such that V_i^* is least.
4. Let $N^* := N^* \backslash \{i\}$
5. If N^* is a singleton STOP, else go to 2.

The algorithm stops when a single point is left. If all choices are unique at stage 3, then it is the unique equilibrium to be selected and it is always one of the selected equilibria. The height of the penultimate point deleted (measured in the reduced space) gives the worst waiting time to hit it.

On account of the special structure of the problem it is possible to characterise the selected equilibrium without going through all the steps of the algorithm. The form of (54) and (55) implies that direct transitions are always optimal. When a point is deleted, transition costs between all other points are unaffected. Although precisely which equilibrium is selected requires some calculation, it is easy to 'bracket it' by the following procedure. To avoid complications of statement, assume that all finite $V(i,i-)$ and $V(i,i+)$ are distinct.

1. Let $l_1 = 1$ and $r_1 = n$ and $k = 1$.
2. Find the point with largest finite value of $V(i, i+)$ or $V(i, i-)$ in $[l_k, r_k]$.
3. If the maximum is achieved by $V(i, i+)$ some i, set $r_{k+1} = i$ and $l_{k+1} = l_k$, else $l_{k+1} = i$ and $r_{k+1} = r_k$.
4. If $[l_{k+1}, r_{k+1}]$ is a singleton then STOP, else $k := k + 1$ and go to 2.

The meaning of this is as follows. Consider the case $k = 1$ and consider Figure 11. The arrows represent the quantities $V(i, i-)$ and $V(i, i+)$, thought of as pointing left and right respectively, with the size of the arrowheads representing their magnitude. The term 'arrow' will be used with this sense below. The largest arrow is at $i = 3$ and points right. The selected equilibrium lies to the left of $i = 3$. The intuition is that is very hard to escape beyond i. At the next stage, one looks for the largest arrow in the range $[1, i]$ and so on.

Proposition 10 *Assume that all finite $V(i, i+)$ and $V(i, i-)$ are distinct. Then the selected equilibrium is unique and*

(a) For each k, $[l_k, r_k]$ contains the selected equilibrium.

(b) In particular, if i is such that

$$\min\{V(i, i+), V(i, i-)\} > \max_{j \neq i}\{\tilde{V}(j, j+), \tilde{V}(j, j-)\} \tag{56}$$

where $\tilde{V}(j, j+) = V(j, j+)$ if $V(j, j+)$ is finite, 0 otherwise, and similarly for $\tilde{V}(j, j-)$, then i is the selected equilibrium.

(c) If k is the selected equilibrium, V^* is the magnitude of the largest arrow pointing towards k and $T_j(k)$ is the time taken to hit k starting from j then

$$\max_{j \neq k} \lim_{\beta \to \infty} \frac{1}{\beta} \ln ET_j(k) = V^* \tag{57}$$

(b) is Kandori and Rob (1995)'s result. The remaining results are new. (c) is fairly intuitive: the difficulty in reaching k is given by the most difficult movement in its direction. The Theorem is proved in the Appendix. The case without ties is more complicated to state but the generalisation is easy. For uniqueness it is enough that there are never any ties in Step 3, or if there is a tie it occurs at the same point between $V(i, i+)$ and $V(i, i-)$. In particular, (b) holds generally.

11. Conclusion

This paper has shown how Freidlin and Wentzell (1998)'s graphical techniques can be used to obtain estimates of waiting times in equilibrium selection problems and how some of the equilibrium selection calculations can be simplified. These results were applied to potential games and games with strategic complementarities. They should prove useful in other classes of games as well.

Appendix

Proof of Proposition 1

Only part (ii) remains to be proved. As noted in Section 9, the result is immediate if one considers the chain reduced to the limit sets in $D(\Omega)$. The following argument, given for completeness, is essentially an expanded version of this.

Consider the limit sets in Ω. Let $D = D(\Omega)$. From each limit set L_i, pick an element l_i. In a $G(\overline{D})$-graph, g, each l_i must be connected to $\overline{D}$ by a path. This path may pass through other limit sets. Since points within a limit set can be connected at zero cost it may be assumed that if the path from l_i passes through L_j, it passes through l_j. Call the graph consisting of the l_i and these paths the 'skeleton' of g. Conversely, given such a skeleton, without loops, a W_D-graph can be constructed: in every $D(L_i)$, all points can be connected at null cost to this skeleton, by the definition of a limit set. It follows that it can be assumed that every optimal $G(\overline{D})$-graph has such a skeleton, with all other transitions of cost zero.

Consider some point in Ω. It must belong to some L_i, and it may assumed that it equals l_i, since the choice of l_i above is arbitrary. Consider the skeleton of an optimal $G(\overline{D})$-graph. Let the last limit set the path from l_i passes through be L_j. Consider the graph obtained by deleting the path from l_j to the exterior from the skeleton (leaving l_j and any incoming connections). Leave all other paths in the skeleton unchanged. The resulting graph is a skeleton for a $G(\overline{D} \cup \{l_j\})$-graph. It may be completed to make one by joining all other points in $D(\Omega)$ to it by a set of null cost transitions. Since l_i is still connected to l_j, it is not connected to the exterior. The resultant graph therefore belongs to $G(l_i \not\to \overline{D})$. The path from l_j to $\overline{D}$ must cost at least $R(\Omega)$, so deleting it must save at least $R(\Omega)$. It follows that $M_D(l_i) \leq W_D - R(\Omega)$. Hence the result.

Proof of Proposition 4

In the case $M = 2$, $x_1 + x_2 = 1$, so the state can be specified simply by x_1, the fraction of players using strategy 1. The state space can therefore be described as a subset of the unit interval. It is enough to prove the result for the case when D is an interval, since any set is the union of intervals. Since D is finite it may be assumed that $D = (a, b)$, unless D includes 0 or 1 and these cases are similar. Let D have elements $c_1 < c_2 < \ldots < c_S$. Put $c_0 = a$ and $c_{S+1} = b$.

Let g^* be an optimal $G(\overline{D})$-graph. It has the form that elements $c_i \leq c_k$ are connected to a, elements $c_i \geq c_{k+1}$ are connected to b, some k with $0 \leq k \leq S+1$. If $k = 0$, this means all elements are connected to b, similarly if $k = S+1$ all elements are connected to a.

Recall from the proof of Proposition 2, (see (20)) that

$$H(D) = \min_{g \in G(\overline{D})} V(g) - \min_{y \in D} \min_{g \in G(\overline{D} \cup \{y\})} V(g) \tag{58}$$

This is equivalent to

$$H(D) = \max_{y \in D} \max_{g \in G(\overline{D} \cup \{y\})} V(g^*) - V(g) \tag{59}$$

Let c_l be in G. Suppose $l \geq k+1$. Then an optimal $G(\overline{D} \cup \{c_l\})$-graph, g', can be assumed to be identical with g^* to the left of c_l. For it must have the property that for some c_j, all points with $c_i \leq c_j$ are connected to a, with $c_i \geq c_{j+1}$ to c_l. g^* also has this property as $c_l \geq c_{k+1}$. Both graphs solve the problem of choosing c_j to minimise costs, with c_l and the graphs to the right of c_l held fixed. It can therefore be assumed that $c_j = c_k$.

Now to the right of c_l, g' must have the form that for some $c_h \geq c_l$, points with $c_l \leq c_i \leq c_h$ are connected to c_l, with $c_i \geq c_{h+1}$ to b. Holding the first term of (58) fixed at $V(g^*)$, minimising the second term of (58) is equivalent to choosing it to maximise (59). The change in (58) by connecting to h, is

$$\sum_{l}^{h-1} V(c_i, c_{i+1}) - V(c_{i+1}, c_i) + V(c_h, c_{h+1}) \tag{60}$$

If $h-1 \leq l$, the first sum is understood to be empty. That is,

$$\sum_{i=l}^{h-1} U(c_{i+1}) - U(c_i) + (U(c_{h+1}) - U(c_h))^+ \tag{61}$$

In order to minimise costs, or equivalently maximise (59) as g^* can be held fixed, $h \geq l$ must be chosen to maximise this. This is true for all $l \geq k+1$. An inverse characterisation holds for $l \leq k$. Letting l vary freely gives the criterion in the text.

Proof of Proposition 6

Consider first part (a). X^A has transitions governed by (omitting β):

$$p^A(x,y) = p(x,y) + \sum_{w \in E\backslash A} p(x,w)\Pi^{E\backslash A}(w,y) \tag{62}$$

$\Pi^{E\backslash A}(w,y)$ denotes the probability that, starting from w, when the process exits $E\backslash A$ the first point in A it hits is y. $\Pi^{E\backslash A}$ is governed by Lemma 2 and asymptotically by Lemma 5.

Now by property (MP) there is an optimal $G(A)$-graph, g^* say, with only minimum cost transitions. Consider $G(A)$-graphs which connect w to y. Each must contain a path z linking w to y. By property (MP) all other transitions can take place at minimum cost (points off the path are linked either to it or to A). Let such a graph be g'. According to (9), g' is to be chosen to minimise

$$V(g') - V(g^*) \tag{63}$$

Since g' and g^* agree except on z, this equals

$$\sum_{t=1}^{T} V(z_{t-1}, z_t) - H(z_{t-1}) \tag{64}$$

Hence if A satisfies (MP),

$$V_{E\backslash A}(w,y) = \min_{z,T}\Big\{\sum_{t=1}^{T} V(z_{t-1},z_t) - H(z_{t-1}) : \\ z_0 = w,\ z_T = y,\ z_t \notin A\ t \neq T,\ z_t \neq z_{t-1}\ t = 1,\dots,T\Big\} \tag{65}$$

Now from (62),

$$V^A(x,y) = \min\big\{V(x,y),\ \min_{w\in E\backslash A} V(x,w) + V_{E\backslash A}(w,y)\big\} \tag{66}$$

Combining the last two equations yields the result.

To prove part (b), consider the process reduced to $B = A\cup\{x\}$, where $x \in E\backslash A$. Now by Lemma 10, μ_β is invariant for X^B, so

$$\mu_\beta(\{x\})\sum_{y\in A} p^B_\beta(x,y) = \sum_{y\in A}\mu_\beta(\{y\})p^B_\beta(y,x) \tag{67}$$

Hence

$$U(x) + \min_{y\in A} V^B(x,y) = \min_{y\in A} U(y) + V^B(y,x) \tag{68}$$

Note that B satisfies (MP), hence V^B is determined in the same way as in part (a). By assumption, x is linked to some point in A by a path which only uses minimum cost links. The reduced cost of this path, that is the right-hand side of (38), is $H(x)$. Any other path to a point in A costs at least this much, so $\min_{y\in A} V^B(x,y) = H(x)$. This yields the result.

Proof of Proposition 7

To show (41), consider $x \in E\backslash A$. Let g be a tree based at x. By property (MP) x can be linked to some $y \in A$ by a path, p say, which leaves each point at minimum cost. Construct a tree based at g' by replacing all links in g along the path p by the links on p and deleting the out-going link from y. The only change in cost is that incurred along p and at y. x originally had no links, so g' incurs an additional cost of $H(x)$. All other points on p have no higher cost than they did in g. y originally had cost at least $H(y)$, now it has no outgoing link. Hence $V(g) - V(g') \geq H(y) - H(x)$. This proves the result.

To show (42), note that X^A and X agree except when X leaves a point z, enters $E\backslash A$ and then exits it. Only the time to leave z is recorded in X^A. Call such an event an excursion and denote its length in X by $\mathcal{T}_z$ and its length in X^A by $\mathcal{T}^A_z$.

$$E(\mathcal{T}_z) = E(\text{Time to enter } E\backslash A \text{ from } z) + E(\text{Time in } E\backslash A) \tag{69}$$

Note that, implicitly, the above conditions on the fact that such an excursion takes place. The first term is, of course, $E(\mathcal{T}^A_z)$. Now

$$\lim_{\beta\to\infty}\frac{1}{\beta}\ln E(\text{Time to enter } E\backslash A \text{ from } z) \geq H(z) \tag{70}$$

and

$$\lim_{\beta\to\infty}\frac{1}{\beta}\ln E(\text{Time in } E\backslash A) \le H(E\backslash A) \tag{71}$$

while

$$H(E\backslash A) = \max_{y\in E\backslash A} H(y) \tag{72}$$

For, from (20) (see the proof of Proposition 2)

$$H(E\backslash A) = \min_{g\in G(A)} V(g) - \min_{y\in E\backslash A}\min_{g\in G(A\cup\{y\})} V(g) \tag{73}$$

Since A has property (MP), an optimal $G(A)$-graph just follows minimum cost exit paths. An optimal $G(A\cup\{y\})$-graph will therefore agree with this except that the out-link from y is deleted. This saves $H(y)$, so the optimal point to choose is that with largest height.

Hence,

$$H(E\backslash A) = \max_{y\in E\backslash A} H(y) \tag{74}$$

Using (40) and (69), (70),(71) and (74)

$$\lim_{\beta\to\infty}\frac{1}{\beta}\ln\left(E(\mathcal{T}_z) - E(\mathcal{T}_z^A)\right) \le \max_{y\in E\backslash A} H(y) \tag{75}$$

and

$$\lim_{\beta\to\infty}\frac{1}{\beta}\ln E(\mathcal{T}_x^A) \ge \min_{w\in A} H(w) \tag{76}$$

Now

$$ET_x(B) - ET_x^A(B) = \sum_{\text{excursions}} E\mathcal{T}_z - E\mathcal{T}_z^A \tag{77}$$

and

$$ET_x^A(B) \ge \sum_{\text{excursions}} E\mathcal{T}_z^A \tag{78}$$

where to avoid gruesome notation z denotes the (random) starting point for each excursion and implicitly the strong Markov property is used.

It follows that,

$$\frac{ET_x(B) - ET_x^A(B)}{ET_x^A(B)} \le \frac{\max_{z\in A} E\mathcal{T}_z - E\mathcal{T}_z^A}{\min_{z\in A} E\mathcal{T}_z^A} \tag{79}$$

Using (75) and (76), it follows that

$$\lim_{\beta\to\infty}\ln\frac{ET_x(B) - ET_x^A(B)}{ET_x^A(B)} \le 0 \tag{80}$$

The result follows.

The proof of (43) is similar. Suppose the process starts at $x \in E\backslash A$ and consider its position on first exit, after time T_1 say, from $E\backslash A$. Either it has hit B or it is at some point $y \in A\backslash B$. In the second case let the subsequent time to hit B be T_2. Now from (74) the logarithm of the expected value T_1 is bounded above asymptotically by $\max_{z\in E\backslash A} H(z)$. From (40) the time to leave any point of A is at least as great as this. It follows that (a) the logarithm of the expected value of T_1 is bounded asymptotically by $\min_{x\in E\backslash A} H(x)$, and so (b) the logarithm of the expected value of T_1 is dominated asympotically by that of T_2. Hence the result.

Proof of Corollary to Proposition 7

(44) follows immediately from (38) and the fact that (a) every point outside a limit set can be connected to one by a path of null cost, (b) every point within a given limit set can be connect to any other within it by a path of null cost (see Lemma 7). This implies that $\tilde{L}$ satisfies property (MP) and moreover the height of any point outside $\tilde{L}$ is zero. Proposition 7 then implies the remaining results. Recall that all points in a limit set have the same value of U and that the limiting measure must concentrate on limit sets. In (46) note that within a limit set transitions are free, so that (asymptotically) (a) the time to hit any point in L_j from l_i is equal to that to hit l_j, (b) the time to hit L_j is the same starting from any point in L_i. The latter assertions are easily proved using Lemma 6.

Proof of Proposition 8

Parts (a) and (b) follow immediately from Proposition 7 applied repeatedly. The first part of (c) follows by applying (43) applied repeatedly with $B = \{x_{|E|}\}$ and $E = A_k$, $k = 0, 1, 2, \ldots$ and noting that from (74)

$$H^{A_{k*}}(A_{k*}\backslash\{x_{|E|}\}) = \max_{y\in A_{k^*},\, y\neq x_{|E|}} H^{A_{k^*}}(y) \tag{81}$$

To show that $H(E\backslash\{u\})$ is independent of the choice of $u \in \arg\min U(x)$, it is enough to show that whenever some point with $U(y) = 0$ is deleted at stage k there are points z and z' with $U(z) = U(z') = 0$ and $H^{A_k}(A_k\backslash\{z\}) \leq H^{A_k}(A_k\backslash\{y\}) \leq H^{A_k}(A_k\backslash\{z'\})$ which are not deleted. For (a) by Proposition 7 worst expected hitting times of points in A_k can be calculated by looking at A_k, (b) every point with $U(y) = 0$ is deleted eventually except $x_{|E|}$.

The claim follows from (39). To save notation set $E = A_k$ and $A = A_{k+1}$. Let y be the point with $U(y) = 0$ to be removed. Since by assumption the height of any deleted point is at most that of any remaining point, the minimum in (39) can only be attained if there exists a point $z \in A$ with $U(z) = 0$ and a path $p = z_0, \ldots, z_n$ with $z_0 = z$, $z_n = y$, $V(z_0, z_1) = H(y)$ and $V(z_i, z_{i+1}) = H(z_i)$ all $i > 0$. Any $G(\{z\} \cup \{x\})$-graph can be made into a $G(\{y\}\cup\{x\})$-graph with no greater cost by deleting the outlink from y and linking z to y by this path p. The case when x lies on p is trivial. Furthermore, since $U(y) = U(z) = 0$, the optimal $G(\{y\})$-graph has the same cost as the optimal $G(\{z\})$-graph. It follows from (20) that $H(E\backslash\{y\})$ is at least $H(E\backslash\{z\})$. On the other hand, it follows from the definition of y that it is linked by a path q with least-cost transitions to some point z' of A. Using (39), a similar argument to that above shows that $U(z') = U(y) = 0$ and $H(E\backslash\{y\})$ is at most $H(E\backslash\{z'\})$.

Proof of Lemma 13

The argument parallels that of Proposition 6, except that an optimal $G(A)$-graph no longer need have all links of minimum cost. Let g^* be an optimal $G(A)$-graph and let p be a path linking w to y. Modify g^* by replacing out-links from points on p by their corresponding links in p. The resulting graph g' links w to y and agrees with g^* outside p, so (64) holds as before. The class of g' constructed in this way, however, need no longer contain the optimal $G(A)$-graph linking w to y, as it may be desirable to change links off p as well. The right-hand side of (65) is therefore now only an upper bound.

Proof of Lemma 14

Reduction to F can be done iteratively, deleting one element of $E\backslash D$ at a time. Let G be the current state space and let $\{g\}$ be an arbitrary element of $E\backslash D$, $\neq y$, to be removed and let $G' = G\backslash\{g\}$. G' trivially satisfies (MP), so for $x \in D$ and $z \in G'\backslash D$

$$V^{G'}(x,z) = \min\{V^G(x,z), V^G(x,g) + V^G(g,z) - H^G(g)\} \tag{82}$$

Take as inductive hypothesis $V^G(x,z) \geq r^*$ for all $z \in G\backslash D$, where r^* is the right-hand side of (51). By (82) it holds for G'. Since it certainly holds for $G = E$, it holds for all G. Setting $z = y$ yields the result.

Proof of Proposition 9

Since all measure concentrates on limit sets, and all elements within a limit set have the same measure, it is enough (using Lemma 10) to consider the process reduced to representative elements of each limit set. Let $\tilde{L}$ be as in the Corollary to Proposition 7.

Let D be the set of the representatives of the limit sets of $D(\Omega)$. Suppose $l \in \tilde{L}\backslash D$. The formula in Lemma 13 applied to $A = D \cup \{l\}$ for $x = l$ and $y \in D$ is precisely the modified cost of reaching the limit set which y represents from x (see below). Hence $\min_{y\in D} V^A(l,y) \leq CR^*(\Omega)$. On the other hand Lemma 14, applied to $\tilde{L}$ implies $V^A(x,l) \geq R(\Omega)$ for all $x \in D$. The result follows from the argument in the text.

To complete the argument, one final technical point should be dealt with. As was noted after the Corollary to Proposition 7, the formula in (44) allows points in a limit set L_k other than l_k to be used in a path between l_i and l_j, say, when calculating $V^{\tilde{L}}(l_i,l_j)$. $V^{\tilde{L}}(l_i,l_j)$ is therefore not, in general, the same as the minimum cost of a path joining L_i and L_j and touching no other limit set. This does not affect the minimum modified cost of linking two limit sets or their radius, however, as a path passing through l_k can always be found which has the same cost as any given one passing through L_k. Application of Lemmas 13 and 14 does therefore yield the modified cost and radius as claimed.

Proof of Lemma 15

The argument is exactly the same as in the proof of Proposition 2 except simpler. Where it says 'limit set' or 'domain of attraction' replace these by point. For 'radius' read 'height of a point'.

Proof of Lemma 16

Let g be an optimal $G(E\backslash A)$-graph. x must be connected to $E\backslash A$ by a path. Let w be the last point on this path in A and let y be the point in $E\backslash A$ at which it ends. Deleting the link $w \to y$ yields a $G(E\backslash A \cup \{x\})$-graph with cost $W_A - V(w, y)$. This may not be the optimal such graph, so using Lemma 6

$$H_A(x) \geq W_A - (W_A - V(w, y)) = V(w, y) \tag{83}$$

which implies the result.

Proof of Corollary to Lemmas 15 and 16

Proposition 2 follows directly from this result applied to $\tilde{L}$, using the Corollary of Proposition 7 and (43) of Proposition 7. The argument in the last paragraph of the proof of Proposition 9 shows that the radius and (modified) co-radius measures are indeed obtained when Lemmas 15 and 16 are applied to $\tilde{L}$.

Proof of Proposition 10

Consider two points x and y. If the current state space is E^* and z is deleted,

$$V^{E^*\backslash\{z\}}(x, y) = \min\{V^{E^*}(x, y), V^{E^*}(x, z) + V^{E^*}(z, y) - H^{E^*}(z)\} \tag{84}$$

If z is of minimum height

$$V^{E^*}(x, z) + V^{E^*}(z, y) - H^{E^*}(z) \geq \max\{V^{E^*}(x, z), V^{E^*}(z, y)\} \tag{85}$$

Now by (54), (55)

$$V^{E^*}(x, y) = \max\{V^{E^*}(x, w), V^{E^*}(w, y)\} \tag{86}$$

if $E^* = E$ and w lies between x and y. Using this, (84), (85) and induction on the number of elements deleted, it is easy to see that for any sub-domain E^*, $V^{E^*}(x, y)$ is equal to $V^E(x, y)$ if points are deleted in order of height.

The assumption of no ties implies that the choice of point to delete at each stage of the elimination algorithm is unique and so, from Proposition 8, the point selected is unique.

To prove part (a) it is sufficient to consider $k = 1$. Let z be the point with the largest arrow, with value V', and suppose it points to the right. It is enough to show that at least one point of $[1, z]$ always remains. Suppose to the contrary and hence that some point of $(z, n]$ always remains. Let w be the last point of $[1, z]$ to be deleted. Now $V(w, x) = V'$ for any $x \in (z, n]$ from (54). Also w has no adjacent point to the left as all other points in $[1, z]$ have been deleted. It follows that its height at this stage of the algorithm is V'. Since there are no ties, it cannot have minimal height and so cannot be deleted. This is a contradiction.

Part (b) follows immediately from (a). To show (c), suppose z^* is the point that remains after application of the reduction algorithm. Consider the following argument to determine the height of the last point deleted to the left of z^*. Suppose w lies to the left of z^* and has the largest rightwards arrow, with value V'', say. Let x be the last point in $[1, w]$ to be removed. Now $V(x, y) = V''$ for any $y \in (w, z^*]$ by (54) and by assumption no points to the left of x remain. It follows that the height of x at this stage of the algorithm is V''. On the other hand, by assumption $V(y, v) < V''$ if v and y belong to $(w, z^*]$ and $v > y$, again by (54). Hence if $y \in (w, z^*)$ is left at this stage of the algorithm it has height less than x and so is removed first. It follows that the last point to be removed to the left of z^* has height at that stage equal to V''. A similar argument applies to the right of z^* and appealing to Proposition 8 yields the result.

References

Blume, L. (1994). How noise matters. mimeo, Cornell University.

Blume, L. (1995). Population games. mimeo, Cornell University.

Bott, R. and Mayberry, J. (1954). Matrices and trees. In Morgenstern, O., editor, *Economic Activity Analysis*, pages 391–400, New York. John Wiley.

Catoni, O. (1999). Simulated annealing algorithms and markov chains with rare transitions. In *Séminaire de Probabilités XXXIII*, volume 1709 of *Lecture Notes in Mathematics*, pages 69–119. Springer Verlag, Berlin.

Catoni, O. (2002). Statistical learning theory and stochastic optimization. mimeo, Université Paris 6.

Catoni, O., Chen, D., and Xie, J. (2000). The loop erased exit path and the metastability of a biased vote process. *Stochastic Processes and their Applications*, 86:231–261.

Cerf, R. (1996). The dynamics of mutation selection algorithms with large population sizes. *Annales de l'Institut Henri Poincaré*, 32:455–508.

Diamond, P. (1982). Aggregate demand in search equilibrium. *Journal of Political Economy*, 90:881–894.

Ellison, G. (2000). Basins of attraction, long run equilibria and the speed of step-by-step evolution. *Review of Economic Studies*, 67:17–45.

Foster, D. and Young, P. (1990). Stochastic evolutionary game dynamics. *Journal of Theoretical Biology*, 38:219–232.

Freidlin, M. and Wentzell, A. (1998). *Random Perturbations of Dynamical Systems*. Springer Verlag, New York, second edition.

Haken, H. (1977). *Synergetics*. Springer Verlag, Berlin.

Hwang, C.-R. and Sheu, S.-J. (1989). On the weak reversibility condition in simulated annealing. *Soochow Journal of Mathematics*, 15:159–170.

Hwang, C.-R. and Sheu, S.-J. (1992). Singular perturbed markov chains and exact behaviors of simulated annealing processes. *Journal of Theoretical Probability*, 5:223–249.

Kandori, M., Mailath, G., and Rob, R. (1993). Learning, mutation and long-run equilibria in games. *Econometrica*, 61:29–56.

Kandori, M. and Rob, R. (1995). Evolution in the long run: A general theory and applications. *Journal of Economic Theory*, 65:383–414.

Monderer, D. and Shapley, L. (1996). Potential games. *Games and Economic Behavior*, 14:124–143.

Sonin, I. (1999). The state reduction and related algorithms and their applications to the study of markov chains, graph theory, and the optimal stopping problem. *Advances in Mathematics*, 145:159–188.

Trouvé, A. (1996). Cycle decompositions and simulated annealing. *SIAM Journal on Control and Optimization*, 34:966–986.

Wolff, R. (1989). *Stochastic Modeling and the Theory of Queues*. Prentice Hall, Englewood Cliffs, NJ.

Young, P. (1993). The evolution of conventions. *Econometrica*, 61:57–84.

Figures

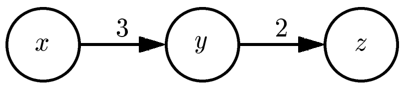

Figure 1

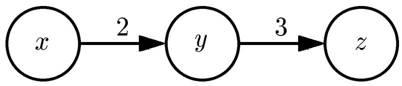

Figure 2

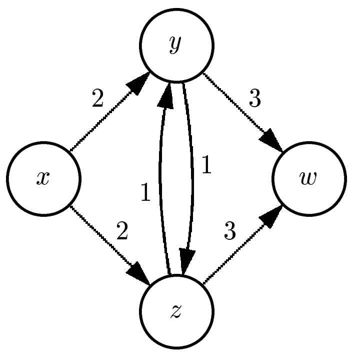

Figure 3

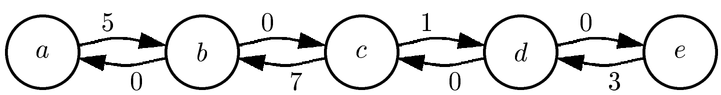

Figure 4

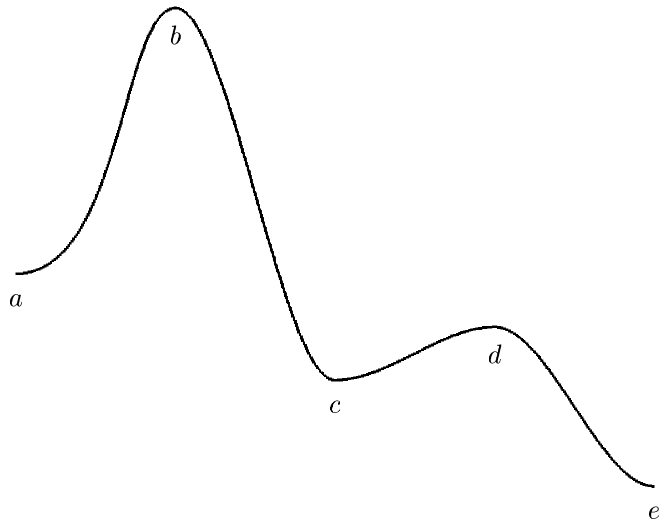

Figure 5

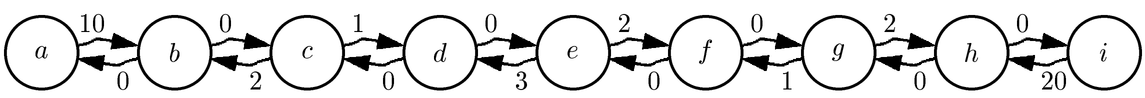

Figure 6

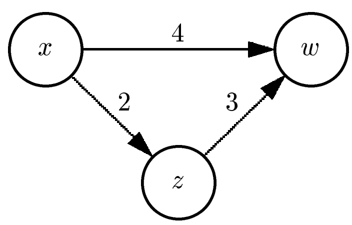

Figure 7

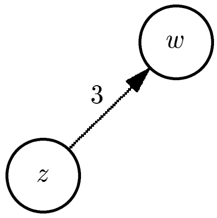

Figure 8

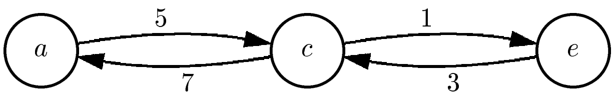

Figure 9

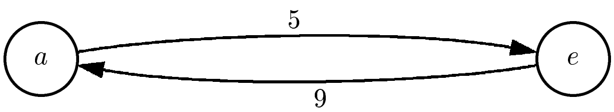

Figure 10

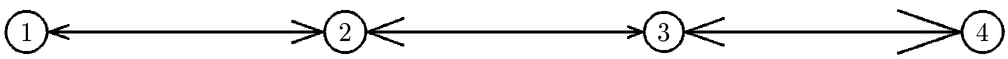

Figure 11