Dynamics in Stochastic Evolutionary Models[☆]

David K. Levine[1], Salvatore Modica[2]

Abstract

We characterize transitions between stochastically stable states and relative ergodic probabilities in the theory of the evolution of conventions. We give an application to the fall of hegemonies in the evolutionary theory of institutions and conflict and illustrate the theory with the fall of the Qing Dynasty and rise of Communism in China.

JEL Classification Numbers: C73 (Evolutionary Games)

Keywords: Evolution, Conventions, Markov Chains, State Power

[☆]First Version: October 4, 2013. We are especially indebted to Juan Block for his many comments and suggestions. We would also like to thank Drew Fudenberg, Kevin Hasker, Matt Jackson, Peyton Young, and five anonymous referees. We are grateful to NSF Grant SES-08-51315 and to the MIUR PRIN 20103S5RN3 for financial support.

[*]Corresponding author David K. Levine, Villa San Paolo, Via della Piazzuola 43, I-50133 Firenze - Italy

Email addresses: `david@dklevine.com` (David K. Levine), `salvatore.modica@unipa.it` (Salvatore Modica)

[1]Department of Economics, European University Institute and Washington University in St. Louis

[2]University of Palermo, Italy

1. Introduction

The modern theory of the evolution of conventions deals with a Markov process in which there are strong forces such as learning towards equilibrium and weaker evolutionary forces such as "mutations" that disturb an equilibrium and lead from one equilibrium to another. There is a large economics literature studying models of this type: just to take some recent contributions, we have Kreindler and Young (2012), Kreindler and Young (2013) and Ellison, Fudenberg and Imhof (2014) showing how convergence in these models may be very fast and Sabourian and Juang (2012) giving a folk theorem for equilibrium selection. To prove theorems, the limit when weak forces are small is analyzed. In the limit equilibria appear as recurrent communicating classes of the Markov process: sets of states all of which are accessible to each other, but grouped into classes which are isolated from each other. Prior to the limit - the situation of interest - the Markov process is ergodic, and puts positive weight on all states. However some states are more equal than others, and in the unique limit of the stationary distributions weight is placed only on the recurrent communicating classes of the limit - and moreover, only some of these classes have positive weight - the stochastically stable classes. In particular, while the limiting Markov process can have many classes, the limit of the Markov processes may place weight on only one or a few classes. The literature, especially Foster and Young (1990), Kandori, Mailath and Rob (1993), Young (1993), Ellison (2000), Cui and Zhai (2010), and Hasker (2014) develops a set of techniques for determining which of these classes get weight in the limit, and gives a useful picture of what the stochastic process looks like when the weak forces are small but not zero. Roughly, the classes that have positive weight in the limit are seen most of the time, but the system will occasionally move away from a class and back again, or transit from one class to another.

The focus of much of the literature has been on determining which classes get weight in the limit - and this is important to understand. But the transitions - the movement from one class to another - are also interesting and important for economics. For example: in a model of evolution such as that of Levine and Modica (2012) where different economic and political institutions compete with each other, the recurrent communicating classes correspond to hegemonies - a single society that controls all economic resources - and the stochastically stable classes are the most powerful hegemonies. There is considerable historical evidence for the existence of hegemonies: China, the Roman Empire and so forth. In the theory - as in reality - hegemonies inevitably fall, and eventually reappear. How the transitions take place, what kind of phases mark the crucial steps in the transitions, is of some interest. Do the eventual winners of the conflict appear on the scene and battle back and forth with the hegemony for a while until they take over and establish their own hegemony (short answer - no) or does something else happen, and if so what? Our results make it possible to answer this question in some detail, describing the rise and fall of hegemony and the warring states period that takes place in between. It may also serve as a guide to policy, showing how different institutions can impact transitions.

The mathematical methods used in analyzing stochastic stability contain clues for what the transitions might look like. In particular, stochastically stable classes can be characterized by trees of recurrent communicating classes where the distance between classes is measured by "resistance" and the stochastically stable classes appear as the root of trees with least total resistance. Because of the role played by least resistant paths in this analysis, a natural conjecture is that least resistant paths are in some sense more likely than higher resistant paths. Here we establish in exactly what sense this is true.

The starting point is to observe that a basic feature of resistance is that if we compare the probability of two specific paths when evolutionary forces are very weak, the lower resistance path is far more likely than the higher resistance path. However, if we look at all paths from one recurrent communicating class to another - the "quasi-direct routes" that include those that may dawdle within the class they start from before moving on - then as a group least resistant paths are far *less* likely than higher resistant paths. The reason for this is simple: it is likely to take a very long time to reach another recurrent communicating class - in the meantime there are likely to be many failed attempts to get there, and these attempts will typically involve some resistance. On the other hand, consider the actual transition from one class to another, that is the paths which leave the starting recurrent communicating class for the final time and which do not pass through a third class - we call these "direct routes[3]." We show that the transition to the other class is likely to happen relatively quickly, in the sense that the set of least resistance direct routes are far more likely than other paths.

We establish the theory in two parts: we first develop a set of bounds for direct routes and then for quasi-direct routes. As a by-product, understanding these transitions also gives us clues about the ergodic probabilities. By examining which recurrent communicating classes are reached "next" from a given starting point we construct a straightforward recursive algorithm that gives precise bounds on the ratio between the ergodic probabilities of all states that are "reasonably close" to recurrent communicating classes.

To illustrate the theory we apply it to a simplified version of the Levine and Modica (2012) evolutionary model of conflict and the emergence of hegemonies - some details of which are motivated by the transition theory of Acemoglu and Robinson (2001) - and provide in particular an account of the fall of the last Qing dynasty in China and the ensuing rise of Communism.

2. Main Results Through an Illustrative Example

We are interested in economic models that can be represented as Markov processes where some transitions are much less likely than others. To illustrate this we start with an example of a "standard" evolutionary model. This simple and familiar example is designed to illustrate the gaps in current knowledge and how the results of this paper fill those gaps. A historical application that

[3] Note that "direct" is sometimes used to mean "in a single transition." Here a direct route can pass through many transitions, but it must not pause in an ergodic class along the way. From Ellison (2000) we know that such paths are not necessarily the quickest way of getting to the target, an issue we carefully account for.

motivates why these gaps are interesting is examined at the end of the paper (Section 7). Consider the 2x2 symmetric coordination game with actions G, B and payoff matrix

	G	B
G	$2,2$	$0,0$
B	$0,0$	$1,1$

This game has two pure Nash equilibria at GG and BB and a mixed equilibrium with probability 1/3 of G.

To put this in an evolutionary context, we assume that there are five players. Each player receives a payoff equal to the average he gets in all matches against his four opponents.[4] In each period a single player is chosen with equal probability to reconsider her move; the other four play as in the previous period. The state of the system is the number of players playing G. So the state space Z has $N = 6$ states. We first define the behavior rule representing "rational" learning[5]: the player who gets to move chooses a best response to the actions chosen by the opposing players. In addition there are independent trembles: with probability $1 - \epsilon$ the behavior rule is followed, while with probability ϵ the player's choice is uniform and random over all possible actions. The presumption is that the chance of "arational" play ϵ is small compared to the probability $1 - \epsilon$ of "rational" play. This arational play is often called a "mutation" in the literature.

This dynamic can be represented as a Markov process on the state space Z defined above with six states representing the number of players playing G. Denoting source states by rows and target states by columns as is standard in the theory of Markov chains, the transition matrix can be computed as[6]

$$P_\epsilon = \begin{pmatrix} 1-\frac{\epsilon}{2} & \frac{\epsilon}{2} & 0 & 0 & 0 & 0 \\ \left(\frac{1}{5}-\frac{\epsilon}{10}\right) & \left(\frac{4}{5}-\frac{3\epsilon}{10}\right) & \frac{4\epsilon}{10} & 0 & 0 & 0 \\ 0 & \left(\frac{2}{5}-\frac{2\epsilon}{10}\right) & \frac{\epsilon}{2} & \left(\frac{3}{5}-\frac{3\epsilon}{10}\right) & 0 & 0 \\ 0 & 0 & \frac{3\epsilon}{10} & \left(\frac{3}{5}-\frac{\epsilon}{10}\right) & \left(\frac{2}{5}-\frac{2\epsilon}{10}\right) & 0 \\ 0 & 0 & 0 & \frac{4\epsilon}{10} & \left(\frac{4}{5}-\frac{3\epsilon}{10}\right) & \left(\frac{1}{5}-\frac{\epsilon}{10}\right) \\ 0 & 0 & 0 & 0 & \frac{\epsilon}{2} & 1-\frac{\epsilon}{2} \end{pmatrix}$$

A critical concept in analyzing this system for ϵ small but not 0 is the notion of resistance.

[4] As Ellison (1993) points out this global interaction model converges much more slowly than if each player is matched only with a neighbor. Here the model is intended for illustrative purposes.

[5] In some models such as Kandori, Mailath and Rob (1993) this rational component of the dynamic is deterministic so can be referred to as the "deterministic dynamic." Here, as in, Binmore and Samuelson (1997) and Blume (2003) the revision rule is stochastic because the player who moves is determined randomly.

[6] Take for example the first diagonal entry $P_\epsilon(0,0) = 1 - \epsilon/2$: if the current state is 0 then whoever is picked next period faces 4 opponents playing B so the best response is B. With probability $1 - \epsilon$ the player is rational and plays B; with probability ϵ the player is arational and plays B with probability 1/2. So with probability $1-\epsilon+\epsilon/2 = 1-\epsilon/2$ the move is B and next state will be 0. Or take the second row: the only G player is drawn with probability 1/5, while with probability 4/5 it is one of the B players; best response is still B in any case so play probability is again $1 - \epsilon/2$ for B and $\epsilon/2$ for G; hence the state remains at one G player if either the G player is chosen and plays G - probability $(1/5) \cdot (\epsilon/2)$ or a B player is chosen and plays B - probability $(4/5) \cdot (1 - \epsilon/2)$; the sum of these two terms is the $4/5 - 3\epsilon/10$ appearing at the second diagonal entry $P_\epsilon(1,1)$ in the matrix. And so on.

Although we give a formal definition below, in this example the resistance is the minimum number of transitions with probability of order ϵ needed to get from one state to another. For example, the resistance of going from $\{1\}$ to $\{0\}$ is 0 since the transition probability is $P_\epsilon(1,0) = \frac{1}{5} - \frac{\epsilon}{10}$, while the resistance of going from $\{0\}$ to $\{1\}$ is 1 since the transition probability is $P_\epsilon(0,1) = \epsilon/2$. Transitions with probability zero independent of ϵ have infinite resistance. The resistance matrix in this case is therefore the following:

$$r = \begin{pmatrix} 0 & 1 & \infty & \infty & \infty & \infty \\ 0 & 0 & 1 & \infty & \infty & \infty \\ \infty & 0 & 1 & 0 & \infty & \infty \\ \infty & \infty & 1 & 0 & 0 & \infty \\ \infty & \infty & \infty & 1 & 0 & 0 \\ \infty & \infty & \infty & \infty & 1 & 0 \end{pmatrix}$$

We can analyze this model giving a heuristic outline of the methods we will develop in the paper.

1. Recurrent communicating classes for $\epsilon = 0$.[7]

   The two recurrent communicating classes consist of the singleton sets $\{0\}, \{5\}$ - these sets are each absorbing and here they correspond to the pure Nash equilibria of the game. We denote the set of recurrent communicating classes by $\Omega = \{\Omega_0, \Omega_5\} = \{\{0\}, \{5\}\}$.

2. Relation between the classes for $\epsilon > 0$.

   Starting in one recurrent communicating class, which comes next, how long will it take and relatively how much time is spent at each recurrent communicating class?

   Here there is only one other recurrent communicating class, so the "next one" is always the "other one." In the general case Corollary 3 tells us that the "next one" is one that can be reached with least resistance. This least resistance is known as the radius. In this example it takes 2 mutations to get from $\{0\}$ to $\{5\}$ and 3 to get back, so the radius of $\{0\}$ is 2 and the radius of $\{5\}$ is 3.

   How long it will take to get from one class to another is known from Ellison (2000) - the main existing result concerning transitions - and is covered here in Theorem 4: it is ϵ^{-1} raised to the power of the radius: for $\{0\}$ the waiting time to $\{5\}$ is of order ϵ^{-2} and for $\{5\}$ the waiting time to $\{0\}$ is of order ϵ^{-3}. Relatively ϵ^{-1} times as long is spent at $\{5\}$ as at $\{0\}$. When there are many recurrent communicating classes computing the relative amount of time at each is complicated: we give a constructive algorithm for finding it in Section 6.3.

[7] A *closed or recurrent communicating class* of a Markov process is a set with the property that there is a positive probability of reaching any point in the set from any other and the probability of leaving the set is zero. The literature also sometimes refers to these as limit sets.

3. Basins.

The basin of a recurrent communicating class is the set of states for which the probability of eventually reaching that class when $\epsilon = 0$ is one. The basin of $\{0\}$ consists of the points $\{0\}, \{1\}$. The basin of $\{5\}$ is $\{3\}, \{4\}, \{5\}$. The state $\{2\}$ is in the "outer range" of both $\{0\}$ and $\{5\}$ but the basin of neither: it has positive probability of reaching either of the two recurrent communicating classes.

4. Relation between basins and classes.

The basin consists of states that are "close" to the corresponding recurrent communicating class.[8] By Theorem 4 during the time before reaching a new class most of the time will be spent in the current recurrent communicating class, but there will be a large number of periods during which the system will move to these close points and back. From Theorem 8 the relative amount of time spent at these states compared to the time spent in the recurrent communicating class is of the order of the difference between the resistance of reaching the point and the radius. For example, starting at $\{5\}$ the radius is 3 and the resistance of getting to $\{3\}$ from $\{5\}$ is 2 so that the system will spend roughly ϵ^{-1} times as much time at $\{5\}$ as at $\{3\}$. From the result above, that also means that the system spends roughly the same amount of time at $\{3\}$ as at $\{0\}$ - but the nature of the time is quite different. The system will remain at $\{0\}$ for long contiguous periods of time with only occasional departures, while the system will remain at $\{3\}$ for very short periods of time, but go there very frequently.

5. Transitions

How do we get from one recurrent communicating class to another? This is covered in Theorem 3. It says that with very high probability the path will have least "peak resistance": such a path may leave the recurrent communicating class and return any number of times, but during each departure from the recurrent communicating class the resistance encountered can be no more than the radius. When the recurrent communicating class is left for the final time the path followed to the new recurrent communicating class must have least resistance and the transition is "very quick." In the example, going from $\{5\}$ to $\{0\}$ the path can leave $\{5\}$ and return many times but the resistance encountered during these departures cannot be greater than 3. So for example, $(5,4,3,2,3,4,5)$ can occur since it has resistance 3, but not $(5,4,3,4,3,4,3,4,5)$ because this has resistance 4. The final transition must have resistance 3, which in this case means it must be monotone: $\{4\}, \{3\}, \{2\}, \{1\}, \{0\}$ must occur in that order. However, any of these states can recur except $\{2\}$ since remaining adds no resistance.

[8] Strictly speaking this is true not of the basin, but only of the inner basin from Definition 5. However in this example the inner basin and the basin coincide.

So, for example, $(5,4,4,3,2,1,1,0)$ is a possible transition, but $(5,4,3,2,2,1,0)$ is not. Despite the fact that the transition paths can have loops of zero resistance in them, the expected length of the path is bounded independent of ϵ. This is shown in Theorem 1.

3. The Model

In the general case we are given a finite state space Z with N elements and a family P_ϵ of Markov chains[9] on Z indexed by $0 \leq \epsilon < 1$. This family satisfies two regularity conditions:

1. $\lim_{\epsilon \to 0} P_\epsilon = P_0$
2. for all $x, z \in Z$ there is a resistance function $0 \leq r(x,z) \leq \infty$ and constants $0 < C < 1 < D < \infty$ such that $C\epsilon^{r(x,z)} \leq P_\epsilon(z|x) \leq D\epsilon^{r(x,z)}$

Notice that zero resistance is equivalent to positive probability with respect to P_0 - a fact we will use all the time - and infinite resistance is zero probability in all P_ϵ's. If $f(\epsilon)$ and $g(\epsilon)$ are non-negative functions a useful notation concerning resistances is to define $f(\epsilon) \sim g(\epsilon)$ if $\liminf_{\epsilon \to 0} f(\epsilon)/g(\epsilon) > 0$ and $\limsup_{\epsilon \to 0} f(\epsilon)/g(\epsilon) < \infty$ with the obvious convention that $0 \sim 0$. With this notation we can then write $P_\epsilon(z|x) \sim \epsilon^{r(x,z)}$. For readability we will state results in the text using this order notation; we restate (and prove) the results with exact bounds in the Appendices.

As in the example of the previous section, we let Ω be the set of the recurrent communicating classes of P_0. We write Ω_x for the recurrent communicating class containing x where $\Omega_x = \emptyset$ if x is not part of a recurrent communicating class. A path a is a finite sequence $(z_0, z_1, \ldots, z_t)$ of at least two not necessarily distinct states in Z and we write $t(a) = t$: this is the number of transitions (z_{s-1}, z_s). The resistance of the path is $r(a) \equiv r(z_0, z_1) + r(z_1, z_2) + \ldots + r(z_{t-1}, z_t)$.

We summarize some well known properties of P_0 and Ω. Non-empty recurrent communicating classes $\Omega_x \neq \emptyset$ are characterized by the property that from any point $y \in \Omega_x$ there is a positive probability path to any other point $z \in \Omega_x$ and that every positive probability path starting at y must lie entirely within Ω_x. Since positive probability in P_0 is the same as zero resistance, we may equally say that from any point $y \in \Omega_x$ there is a zero resistance path to any other point $z \in \Omega_x$ and that every zero resistance path starting at $y \in \Omega_x$ must lie entirely within Ω_x. An additional useful notion is this:

Definition 1. A set W is *comprehensive* if for any point $z \in Z$ there is a zero-resistance path $a = (z, \ldots, w)$ to some point $w \in W$. In particular the set Ω is comprehensive.

We can give the following characterization of a comprehensive set:

Proposition 1. *A set W is comprehensive if and only if it contains at least one point from every non-empty recurrent communicating class, that is, for all $\Omega_x \in \Omega$ there exists $w \in W$ with $w \in \Omega_x$.*

[9] In the literature it is often assumed that for $\epsilon > 0$ the chain is ergodic and in our analysis of the limit ergodic distribution in the later part of the paper we make that assumption. However for the analysis of transitions the assumption is not needed, and it can be useful to apply the analysis to interim dynamics at states that will never be reached again. Moreover, the assumption that the state space is finite is not needed for the analysis of transitions and the bounds given in the Appendix hold for countable state spaces as well as finite state spaces - in this case when $\epsilon > 0$ there may be no long run limit at all.

Proof. Sufficiency: for any point $z \in Z$ there is a zero resistance path to some point y in some recurrent communicating class Ω_y, and from there a zero resistance continuation to the point in $\Omega_y \cap W$ which is assumed to exist. Necessity: if there is a set $\Omega_y \neq \emptyset$ with $\Omega_y \cap W = \emptyset$ then the zero resistance path to W assumption fails: any zero resistance path originating in Ω_y must remain entirely within Ω_y and hence does not reach W. □

Notice that there are a great many comprehensive sets and our analysis is conditional on a particular choice of a comprehensive set. Different comprehensive sets may serve different useful purposes.

4. Direct Routes

In P_0 a path that hits a point in a recurrent communicating class is then trapped in that class, so cannot reach a target outside of that class. When $\epsilon > 0$ this need not be the case. However, if a point in an recurrent communicating class is hit then it is very likely that the path will then linger in that recurrent communicating class passing through every point in the class many times - and in particular through states in any comprehensive set. Hence there is a sense in which paths that do not hit a comprehensive set must be "quick" - they cannot linger in a recurrent communicating class. We will call such paths "direct routes". Now *any* path that leaves a class Ω_x contains a direct route - the route to its first point in $\Omega \setminus \Omega_x$. So we start by studying direct routes. They are a bit like the hare in the story of the tortoise and the hare. Direct routes get to the destination quickly - they must if they are not to fall into the forbidden comprehensive set. Because of this, as Ellison (2000) points out, they are not very reliable: routes that linger in a recurrent communicating class may be far more likely than direct routes to reach their destination. Such quasi-direct routes we will study in Section 5 below.

Formally, define a *forbidden set* $W \subseteq Z$ for a path $a = (z_0, z_1, \ldots, z_t)$ to be a set that the path does not touch except possibly at the beginning and end, that is $z_1, \ldots, z_{t-1} \notin W$.

Definition 2 (Direct Routes)**.** Given an initial point $x \in Z$ and sets $B \subseteq W$, we call a non-trivial path a from x to B with forbidden set W a *direct route* if W is comprehensive and the path has finite resistance $r(a) < \infty$ (equivalently, positive probability for $\epsilon > 0$).

For each x, B and comprehensive W there is a set of direct routes from x to B with forbidden set W, which we denote by A_{xBW}.[10]

We are interested in the following questions: how likely is the set of direct routes A_{xBW}, which paths in A_{xBW} are most likely, what are these paths like and how long are they?

Results on Direct Routes

The intuition behind the results we present next is simple. Direct routes must hit the target without falling into a comprehensive set. This is hard, hence these routes have to be quick - and the quickest way is to make least resistance steps. This will be made precise in the following.

[10]The assumption that $B \subseteq W$ is without loss of generality. We can always define a forbidden set $W' = W \cup B$ without changing the set of direct routes. Note that for given x, B, W the set of direct routes may be empty: it may be impossible to get to B without first hitting $W \setminus B$.

First, to avoid triviality, we assume that $A_{xBW} \neq \emptyset$. An important observation is that there are typically many direct routes. Specifically, if there is a path $(z_0, z_1, \ldots, z_t) \in A_{xBW}$ that contains a *loop*, that is, $z_\tau = z_{\tau'} \notin W$ for $\tau \neq \tau'$ then A_{xBW} is countably infinite since the loop can be repeated an arbitrary number of times. Notice also that if this loop has zero resistance then the length of paths in A_{xBW} is without bound. Never-the-less we shall see the expected length of such paths is quite short. For nonempty $A \subseteq A_{xBW}$, the first important fact proven in Appendix 1 is that $r(A) = \min_{a \in A} r(a)$ is well-defined (and finite) - it is the least resistance of any path in the set A.

The main result on direct paths characterizes their probability and length. The proof along with detailed bounds is in Appendix 1.

Theorem 1. *If $A \subseteq A_{xBW}$ is non-empty then $P_\epsilon(A|x) \sim \epsilon^{r(A)}$ and $E_\epsilon[t(a)|x, A] \sim 1$.*

In particular, positive resistance direct routes are not very likely to occur as ϵ gets small, yet they are unlikely to be terribly long in the sense that the expected length is bounded independent of ϵ. Despite the fact that they are unlikely, these paths are important because they are needed to get from one recurrent communicating class to another. Intuitively the reason these paths are short is that at each point along a direct route there is a zero resistance path that leads to the forbidden set W. The more time is spent along the route, the greater the risk of falling into the forbidden set and fail to reach its destination. By contrast, we will see subsequently that the expected time spent in a recurrent communicating class goes to infinity as $\epsilon \to 0$.

The order of $P_\epsilon(A|x)$ established in Theorem 1 directly implies the other facts characterizing direct routes.

Corollary 1. *Let $A = \{a|r(a) = r(A_{xBW})\}$ denote the least resistance paths in $A_{xBW} \neq \emptyset$. Then $\lim_{\epsilon \to 0} \frac{P_\epsilon(A|x)}{P_\epsilon(A_{xBW} \backslash A|x)} = \infty$.*

In other words, least resistance direct paths are far more likely than other direct paths. It is also the case that all least resistance direct paths have a probability similar to each other.

Corollary 2. *Let $A = \{a|r(a) = r(A_{xBW})\}$ and $a \in A$. Then $\frac{P_\epsilon(a|x)}{P_\epsilon(A|x)} \sim 1$.*

This completes the discussion of the basic results on direct routes.

5. Transitions Between Recurrent Communicating Classes

We next study how the system is most likely to leave a recurrent communicating class, and to which other class it is most likely to transit; then also how long it takes, and where these paths spend their time along the way.

5.1. Quasi-Direct Routes

Ellison (2000) observes that being able to pass through every point in a recurrent communicating class may have a profound impact on the nature of the paths. The results of this section makes this precise by characterizing quasi-direct routes that spend most of the time moving about without resistance within an initial class Ω_x.

We start again with an initial point $x \in Z$ in the recurrent communicating classΩ_x, a forbidden set $W \subseteq Z$ and a target set $B \subseteq W$. Notice that the definition in Section 3 implies that direct routes from x to B with forbidden set W are not allowed to pass through all states in Ω_x, since the forbidden set W was assumed to be comprehensive. We now wish to relax that restriction, and consider routes which are allowed to linger freely inside Ω_x.[11] So we exclude Ω_x from the forbidden set, that is we assume $W \cap \Omega_x = \emptyset$. Thus W cannot be comprehensive. However, we assume that W contains at least one point from every recurrent communicating class except for Ω_x. We then call W *quasi-comprehensive*.

Definition 3 (Quasi-direct Routes). A non-trivial path a from $x \in Z$ to $B \subseteq W$ is a *quasi-direct route* if W is quasi-comprehensive and the path has finite resistance.

We denote the set of such paths by Q_{xBW}. As in the direct case we assume the set Q_{xBW} is non-empty. Again we are interested in the structure of the paths in Q_{xBW}, in particular: which paths in Q_{xBW} are most likely, what do these paths look like, and how long are they?

Consider a path a in Q_{xBW}. The path originates at x and eventually hits B for the first time without first hitting W. The path may return to the start point x a number of times before finally departing and reaching the destination B. Consequently it can be decomposed into a series of loops starting from x and returning to x, followed by an final departure or *exit path* to B. The loops start at x and return to x without hitting x or W in between, hence they are routes from x to x with forbidden set $W \cup \{x\}$. Since W is quasi-comprehensive $W \cup \{x\}$ is comprehensive and we abbreviate these direct routes as $A_{xxW} \equiv A_{xx(W\cup\{x\})}$ - this should not lead to confusion since W is quasi-comprehensive and so A_{xxW} has no other meaning. Following the loops the exit path is a route from x to B that does not hit either x or W - that is the forbidden set is again $W \cup \{x\}$ and so these are again direct routes which we abbreviate as $A_{xBW} \equiv A_{xB(W\cup\{x\})}$. For $a \in Q_{xBW}$ we write $n(a)$ for the number of loops in a (it may be that $n(a) = 0$). Then if $n(a) > 0$ we can uniquely decompose $a \in Q_{xBW}$ as $a_1, a_2, \cdots, a_{n(a)}, a^+$ where the $a_i \in A_{xxW}$ are the loops and $a^+ \in A_{xBW}$ is the exit path; while $a = a^+$ if $n(a) = 0$.[12]

Next we want a measure of the resistance of a quasi-direct path $a \in Q_{xBW}$. As we shall see, for such path it is not the total resistance $r(a)$ that matters. What matters is the *peak resistance* $\rho(a) = \max\{r(a_i)|_{i=1}^{n(a)}, r(a^+)\}$, the greatest resistance of any of the loops or the exit path.

Definition 4. The *least peak resistance* of a set $Q \subseteq Q_{xBW}$ is $\rho(Q) = \min_{a\in Q} \rho(a)$.

The next result shows that quasi direct routes with least peak resistance consist of a least resistance exit path preceded by loops of weakly lower resistance:

[11] Notice that for the case of singleton Ω_x this means the path may remain at x for some time, or leave and return a number of times before hitting the target.

[12] Note that both A_{xxW} and A_{xBW} can contain loops from a point $y \neq x$ in Ω_x back to y provided we do not touch x in between. One can imagine alternative decompositions without this property, but this decomposition is just a tool for understanding how we get from x to B and technically our decomposition works well.

Theorem 2. If $a \in Q_{xBW}$ has least peak resistance so that $\rho(a) = \rho(Q_{xBW})$, then it has peak resistance equal to least exit resistance: $\rho(a) = r(a^+) = r(A_{xBW})$.

Proof. Suppose $\rho(a) = \rho(Q_{xBW})$. From the definition $\rho(a) \geq r(A_{xBW})$ so the lemma can fail only if there is a path $\tilde{a} \in A_{xBW}$ for which $r(\tilde{a}) < \rho(a)$. But $\tilde{a} \in Q_{xBW}$ so this contradicts a having least peak resistance. □

5.2. Leaving a Recurrent Communicating Class

The following result (proved in Appendix 2) plays a role in the theory of quasi-direct routes similar to that played by least resistance in the theory of direct routes in Corollary 1:

Theorem 3. *Let* $A = \{a|\rho(a) = \rho(Q_{xBW})\}$ *denote the least peak resistance paths in* $Q_{xBW} \neq \emptyset$. *Then* $\lim_{\epsilon\to 0} \frac{P_\epsilon(A|x)}{P_\epsilon(Q_{xBW}\backslash A|x)} = \infty$.

Theorem 3 not only tells us the most likely routes from Ω_x to $\Omega\backslash\Omega_x$, by implication it also tells us where we are likely to leave Ω_x from and where we are likely to end up. First, since all states in Ω_x can be reached from x with no resistance, the path must leave Ω_x through a point $z \in \Omega_x$ from which the path to B is of least resistance among the direct routes from Ω_x to B with forbidden set $W \cup \Omega_x$ - since leaving Ω_x through any other point would incur higher resistance. We call such z an *express exit.*

Next we consider where we end up. In case $B = W = \Omega\backslash\Omega_x$, which recurrent communicating class in $\Omega\backslash\Omega_x$ are we likely to move to? Let Ω_{-x}^{LP} be the collection of Ω_y in $\Omega\backslash\Omega_x$ for which there is a quasi-direct route from x to y of least peak resistance and let Ω_{-x}^{GP} be the remainder of $\Omega\backslash\Omega_x$. Let $P_\epsilon(\Omega_{-x}^j|x)$ denote the probability that starting at x the first arrival at $\Omega\backslash\Omega_x$ is in Ω_{-x}^j for $j = LP, GP$. Then we have the following immediate corollary of Theorem 3.

Corollary 3. $\lim_{\epsilon\to 0} \frac{P_\epsilon(\Omega_{-x}^{LP}|x)}{P_\epsilon(\Omega_{-x}^{GP}|x)} = \infty$.

To better understand what these results say about the actual dynamics of the system, recall that Ellison (2000) defines the *basin* of Ω_x as the set of states in Z for which there is a zero resistance path to Ω_x and no zero-resistance paths to $\Omega\backslash\Omega_x$. Said otherwise, it is the set of states for which there is probability one in P_0 of reaching Ω_x. He also defines the *radius* as the least resistance of paths from Ω_x out of the basin. Focus on the case $B = W = \Omega\backslash\Omega_x$. Theorem 2, together with the fact that there are zero resistance paths from x to any other point $z \in \Omega_x$ and from outside the basin to $\Omega\backslash\Omega_x$, shows that the least peak resistance $\rho(Q_{xBW})$ is the same as the radius. Theorem 3 shows that what matters for leaving the basin are the least peak resistance paths in the sense that paths from Ω_x to $\Omega\backslash\Omega_x$ which have peak resistance higher than the radius are very unlikely. That includes both paths with a higher exit resistance than the radius and paths which have loops with higher resistance than the radius. Corollary 3 shows that the recurrent communicating classes Ω_y that will be reached from Ω_x are very likely to be those for which the peak resistance is the same as the radius - which is to say that when we leave the basin on a direct route with resistance equal to the radius there is then a zero resistance path to Ω_y.

5.3. Expected Length and Visits of Quasi-Direct Routes

We know from Theorem 1 that transition paths in the direct route case are short. For paths that are allowed to remain in Ω_x we have the opposite result: these paths are quite long. At this point it is convenient to focus on the case where B is reached with probability one, that is $P_\epsilon(Q_{xBW}|x) = 1$. We assure this by assuming that $B = W$.[13] Our goal is to show that Q_{xBW} has paths of expected length $\epsilon^{-r(A_{xBW})}$; that the fraction of time spent in Ω_x goes to one; and that the absolute time spent outside of Ω_x goes to infinity. We will also show which kind of loops are likely to recur many times. The proofs of the results of this section may be found in Appendix 2.

The first result concerns the amount of time it takes to leave Ω_x and how much of that time is spent in Ω_x:

Theorem 4. *Suppose that $B = W$. For $a = (a_1, a_2, \ldots, a_{n(a)}, a^+) \in Q_{xBW}$ we let $a^- = (a_1, a_2, \ldots, a_{n(a)})$ denote its loops and define $t^-(a)$ to be the amount of time along a^- spent outside of Ω_x.*[14] *Then*

$$E_\epsilon[t(a)|x, Q_{xBW}] \sim E_\epsilon[t(a^-)|x, Q_{xBW}] \sim \epsilon^{-r(A_{xBW})}$$

and

$$\lim_{\epsilon \to 0} E_\epsilon \left[\frac{t^-(a)}{t(a^-)} \middle| x, Q_{xBW} \right] = 0.$$

This says that quasi-direct paths including or excluding the exit path are long and spend most of their time in $\Omega(x)$.

The second result characterizes more exactly what happens while the system spends time outside of Ω_x during a quasi-direct route:

Theorem 5. *Suppose that $B = W$. For $A \subseteq A_{xxW}$ let $M(a, A)$ be the number of loops of a that lie in A. For $A \subseteq A_{xxW}$*

$$E_\epsilon[M(a, A)|x, Q_{xBW}] \sim \epsilon^{r(A) - r(A_{xBW})}$$

if in addition $r(A) < r(A_{xBW})$ and $k \geq 0$

$$\lim_{\epsilon \to 0} P_\epsilon[M(a, A) > k|x, Q_{xBW}] = 1.$$

Also let $A_{xxW}[t]$ be the set of loops which spend at least t consecutive periods outside of Ω_x. If there is a path $a^0 \in A_{xxW}$ that contains a zero resistance loop not touching Ω_x with $0 < r(a^0) < r(A_{xBW})$ then for any $k > 0$

$$\lim_{\epsilon \to 0} P_\epsilon(M(a, A_{xxW}[kt(a^+)])|x, Q_{xBW}) > k) = 1$$

The first two statements say that if there is some $a^0 \in A_{xxW}$ with $0 < r(a^0) < r(A_{xBW})$ then this loop will occur many times even though the fraction of the time spent outside Ω_x in such loops

[13] Recall that the paths in Q_{xBW} by definition have positive probability of reaching B without touching W along the way; when $B = W$ there is no other way of reaching B so this probability becomes one.

[14] This first result is an extension of Ellison (2000)'s result that the waiting time for leaving Ω_x is of order ϵ^{-r} where r is the radius. Ellison mentions both an upper and lower bound in the text, but we have been unable to locate his proof of the lower bound. The extension here is that W can be a general quasi-comprehensive set, for example, it might include states in the basin of Ω_x.

must be small. Or we can say it this way: as $\epsilon \to 0$ loops in A_{xxW} that have resistance strictly less than the radius occur an arbitrarily large number of times before we leave the basin of x and those which have a resistance strictly greater than the radius have a vanishing small probability of occurring before we exit the basin. The third result says that it will often be the case that a will spend more than k times as long outside of $\Omega(x)$ as it takes to get to the final destination following the exit path.

On a more technical note, the first result describes the expected number of occurrences while the second describes the realized number of occurrences of loops in A. The difference between the two is this: the amount of time before leaving the basin is random. It could be that with very high probability the number of occurrences is small (less than k say) and in those rare cases where the length of time before leaving the basin is very large the number of occurrences is very large. In this case the expected number of occurrences may grow as ϵ gets smaller while the probability of seeing more than k occurrences remain unchanged or even falls. The second part of the Theorem 5 shows that this cannot happen.

6. The Big Picture

Reconsider the dynamics of P_ϵ. Starting at any point x, by Theorem 1 we move quickly to one of the recurrent communicating sets Ω_y. Once there, by Theorem 4 it is a long time before we reach a different Ω_z and most of that time is spent in Ω_y. One question we now address is what the dynamics look like during the long period when we are in Ω_y. When we do finally leave Ω_y, by Theorem 2 we move quickly to the next Ω_z and it is most likely the recurrent communicating set that has least exit resistance from Ω_y. The second question we will address is, over the longer run how much time do we spend in the different recurrent communicating sets in Ω? To this end, we assume in this section that P_ϵ is ergodic for $\epsilon > 0$ and denote by μ_ϵ the unique ergodic distribution of the process.

The big picture which we will break down over the next subsections can be visualized by thinking of astronomy. At the bottom level are planets, which correspond to recurrent communicating classes. Movements within recurrent communicating classes can be though of as moving around on the planetary surface - that is relatively quick. Recurrent communicating classes are surrounded by states in their "inner basin" that are tightly bound to them - like moons around a planet. There are also a few states that are either far from recurrent communicating classes, or bound to several of them. For these only we cannot give precise bounds on the ergodic probabilities - we may think of them as comets. Recurrent communicating classes in turn are grouped into "circuits" - think of those as solar systems: movement within a solar system being much more rapid than movement between solar systems. These circuits - solar systems - are grouped into higher order circuits - galaxies, and the "galaxies" in turn to even higher level circuits and so forth. We give tight bounds for the relative ergodic probabilities for elements within a circuit: planets within a solar system, solar systems within a galaxy. In the end all these circuits are contained in a single universe, and this will give us the relative ergodic probabilities of all recurrent communicating classes.

We should emphasize that while the method of circuits can be turned loose on arbitrary models to determine the structure and relative probability of recurrent communicating classes, it can also be useful in building models. Some structures are easier to analyze than others. For example, in the hegemony model of Levine and Modica (2013) all recurrent communicating classes are in the same circuit, so it is trivial to analyze their relative resistances: it is given by the differences of least resistances to leaving - or more simply, in that context: hegemonies can be ranked by their state power in their worst state, and hegemonies with higher state power are relatively much more likely in the ergodic distribution than those with lesser state power. At the next level, if we can establish assumptions that bind recurrent communicating classes into groups where all the groups lie in a single circuit then within each group relative resistances are determined by exit resistances; and the relative resistances between groups is determined by the exit resistances from group to group. More broadly our intuition may tell us what the structure of circuits "ought" to look like so that it would be natural to focus on assumptions that lead to that type of structure.

6.1. Inside Recurrent Communicating Classes: On the Surface of the Planet

When $\epsilon = 0$ we cannot move between recurrent communicating classes but we have well defined and ergodic dynamics within each class. Moreover, these dynamics are fast in the sense that they are independent of ϵ. Hence if we are interested in the approximate probability of events within a class we should consider paths of bounded length. For such events, the probabilities in P_0 are much the same as for P_ϵ for ϵ small. Specifically,

Theorem 6. *Let A_1 and A_2 be any collections of paths of bounded length starting at $x \in \Omega_x$ and for which $P_0(A_2|x) > 0$. Then*

$$\lim_{\epsilon \to 0} \frac{P_\epsilon(A_1|x)}{P_\epsilon(A_2|x)} = \frac{P_0(A_1|x)}{P_0(A_2|x)}.$$

Proof. Since the probabilities are defined by finite sums of finite products of the transition probabilities $P_\epsilon(z|y)$ and the length of the sums and products are bounded independent of ϵ the result follows immediately from the assumption that $\lim_{\epsilon \to 0} P_\epsilon(z|y) = P_0(z|y)$. □

In particular since paths that lie entirely within Ω_x have probability one in P_0 given x, the probability of sets of paths of finite length within Ω_x is roughly the same in P_ϵ as in P_0 when ϵ is small.

The other important characteristic of Ω_x is the amount of time spent at different states. Notice that if we restrict the state space to Ω_x then P_0 is an ergodic Markov process on that space, so has a unique and strictly positive ergodic distribution $\overline{\mu}_0(y)$, where $\sum_{y \in \Omega_x} \overline{\mu}_0(y) = 1$. Notice in particular that if $y \in \Omega_x$ the ratio $\overline{\mu}_0(x)/\overline{\mu}_0(y)$ is well-defined and finite. We can relate this to the ratio of stationary probabilities $\mu_\epsilon(x)/\mu_\epsilon(y)$ for the process when $\epsilon > 0$. In Appendix 3 we show:

Theorem 7. *If $y \in \Omega_x$ then*

$$\lim_{\epsilon \to 0} \frac{\mu_\epsilon(x)}{\mu_\epsilon(y)} = \frac{\overline{\mu}_0(x)}{\overline{\mu}_0(y)}.$$

6.2. Basins and Ranges: Moons and Comets

Recall that the basin of Ω_x consists of the states for which there is probability one in P_0 of reaching Ω_x. For the purpose of relating states to recurrent communicating classes it is useful to define three variations on the basin. First observe that from any point outside the basin there is positive P_0-probability of reaching $\Omega \backslash \Omega_x$, that is from outside the basin there is a zero-resistance path to $\Omega \backslash \Omega_x$. Therefore the radius - least resistance of leaving the basin - is also the least resistance to get to $\Omega \backslash \Omega_x$.

Definition 5. The *outer range* of Ω_x is the set of states for which there is a zero resistance path to Ω_x. The *inner range* of Ω_x is the set of states y in the outer range which can also be reached from x with resistance not larger than the radius of Ω_x.[15] If we further require that the resistance of being reached is strictly less than the radius we have the *inner basin.*

The outer range is the largest set of states affiliated with Ω_x in the sense that any states outside the outer range will never get to Ω_x when $\epsilon = 0$. The outer range contains the basin, but unlike the basin the outer range does not require reaching Ω_x with probability one when $\epsilon = 0$, and a point can be in the outer range of different recurrent communicating classes.[16] By contrast the basins of different recurrent communicating classes must be disjoint.

The inner basin is a subset of the inner range by definition. It is also a subset of the basin, since if a point in the inner basin were not in the basin we could go from x to $\Omega \setminus \Omega_x$ with resistance strictly less than the radius which is impossible by definition. The states in the inner basin are the states most tightly affiliated with Ω_x: this is shown by Theorem 4 which says that these states will be hit many times before moving on to the next recurrent communicating class.

There is no firm relationship between the basin and the inner range. The basin may contain states further from Ω_x than the radius, so states not in the inner range. The inner range contains states at a distance equal to the radius, and some of these states must have zero resistance paths to other recurrent communicating classes so that they they cannot be part of the basin. Never-the-less states in the inner range are still "close" to Ω_x: they are part of least peak resistance paths to other recurrent communicating classes and Theorem 4 says the expected number of times they will be hit before moving on is positive. By contrast the states that are in the basin but not the inner range are "far" from Ω_x and this can be seen precisely in Theorem 3 which shows that these states are unlikely to be reached from Ω_x prior to reaching another recurrent communicating class.

States in the inner basin are like "moons" tightly bound to the recurrent communicating class, while states in the outer range but not the inner basin are more like "comets" that are not tightly bound to any recurrent communicating class.

Let $r^D(x, y)$ be the least resistance of any direct path from $x \in \Omega_x$ to y that does not pass through any recurrent communicating class other than Ω_x, that is, that is we define $r^D(x, y) \equiv$

---

[15] If this is true then there is also a direct route with forbidden set $W = (\Omega \backslash \Omega_x) \cup \{x\} \cup \{y\}$ with this property. Basically within the basin it does not matter whether least resistance is measured along direct routes or all routes since it is not possible to pass through $\Omega \backslash \Omega_x$ while remaining in the basin.

[16] In the introductory example the state $\{2\}$ is in the outer range of both Ω_0 and Ω_5.

$r(A_{xyW})$ with $W = (\Omega\backslash\Omega_x) \cup \{x\} \cup \{y\}$. Take $W = \Omega\backslash\Omega_x$ and define the radius[17] of $\Omega_x$to be $r^0(\Omega_x) = r(Q_{xWW})$. It is shown in Appendix 3 that:

Theorem 8. *If $x \in \Omega_x$ and $r^D(y,x) = 0$ then*

$$\frac{\mu_\epsilon(y)}{\mu_\epsilon(x)} \sim \epsilon^r$$

where $\min\{r^D(x,y), r^0(\Omega_x)\} \leq r \leq r^D(x,y)$.

For the "comets" - states that are not in the inner range of any recurrent communicating class, that is, states that are in the outer range of one or more recurrent communicating class, but are "hard" to reach - Theorem 8 gives bounds for the ergodic probabilities - they cannot be too likely. For the "moons" - states in the inner range - the bound is tight, and says that their relative probability of occurring is inversely proportional to the least resistance of getting to them from Ω_x.

Corollary 4. *If y is in the inner range of $x \in \Omega_x$, then*

$$\frac{\mu_\epsilon(y)}{\mu_\epsilon(x)} \sim \epsilon^{r^D(x,y)}.$$

6.3. Recurrent Communicating Classes: Solar Systems, Galaxies and beyond

Consider again the overview of the dynamics of P_ϵ: starting at any point x we move quickly to one of the recurrent communicating sets Ω_y, and once there it is a long time before we reach a different Ω_z and most of that time is spent in Ω_y. When we do finally leave Ω_y we move quickly - and directly, in our sense - to the next Ω_z and it is most likely the recurrent communicating set that has least exit resistance from Ω_y. Proceeding in this way we get a sequence of recurrent communicating sets Ω_i connected by least exit resistances.[18] Since the set Ω of recurrent communicating classes in P_0 is finite, eventually this sequence must have a cycle.

More general than the notion of a cycle, we introduce the notion of a circuit. The defining property of a circuit is that between any two of its states there is a path within the circuit with least resistance transitions. We start by defining circuits in Ω. For $\Omega_x, \Omega_y \in \Omega$ define transition resistance $r^0(\Omega_x, \Omega_y) = \min\{r^D(x,y) \mid y \in \Omega_y\}$, that is the least resistance of any direct path from x to Ω_y not touching $\Omega \setminus \Omega_x$; least resistance out of Ω_x is then defined as $r^0(\Omega_x) = \min_{\Omega_y \in \Omega\setminus\Omega_x} r^0(\Omega_x, \Omega_y)$ - Ellison (2000)'s radius of Ω_x.

Definition 6 (Circuits)**.** A set $\Omega_x^1 \subseteq \Omega$ is a *circuit* if for any pair $\Omega_1, \Omega_y \in \Omega_x^1$ there is a path $(\Omega_1, \Omega_2, \ldots, \Omega_n)$ in Ω_x^1 with $\Omega_n = \Omega_y$ and $r^0(\Omega_{\tau-1}, \Omega_\tau) = r^0(\Omega_{\tau-1})$ for $\tau = 2, 3, \ldots n$.

Our basic observation is that once we reach a circuit, we remain within the circuit for a long time before going to another circuit. We first compare states within a circuit. Since the probability of leaving Ω_x is of order $\epsilon^{r^0(\Omega_x)}$, the expected length of any visit to Ω_x is $1/\epsilon^{r^0(\Omega_x)}$. We might then

[17] As noted above this is different from Ellison (2000)'s definition but equivalent.
[18] We could equally well say radius.

expect that the amount of time we spend at Ω_x is roughly $\epsilon^{r^0(\Omega_y)-r^0(\Omega_x)}$ as long as the amount of time we spend at Ω_y. In Web Appendix 4 we show that this is indeed true.

Theorem 9. *If the recurrent communicating classes Ω_x and Ω_y are in the same circuit then*

$$\frac{\mu_\epsilon(x)}{\mu_\epsilon(y)} \sim \epsilon^{r^0(\Omega_y)-r^0(\Omega_x)}.$$

We next ask how long do we actually spend in a circuit over Ω? We stay in Ω_x roughly $\epsilon^{-r^0(\Omega_x)}$ periods. On the other hand the probability of going to a fixed $\Omega_y \neq \Omega_x$ out of the circuit is of order $\epsilon^{r^0(\Omega_x,\Omega_y)}$. Hence the probability of going to Ω_y during a visit to Ω_x is of order $(1/\epsilon^{r^0(\Omega_x)})\epsilon^{r^0(\Omega_x,\Omega_y)}$. In order for this to occur with very high probability the number of visits to Ω_x must be roughly k_x where $k_x(1/\epsilon^{r^0(\Omega_x)})\epsilon^{r^0(\Omega_x,\Omega_y)} = 1$. That is $k_x = 1/\epsilon^{r^0(\Omega_x,\Omega_y)-r^0(\Omega_x)}$. Following Ellison (2000) we define the *modified resistance from Ω_x to Ω_y* as $R^0(\Omega_x,\Omega_y) = r^0(\Omega_x,\Omega_y) - r^0(\Omega_x)$. Then the number of visits is least for the element Ω_z in the circuit which has minimum $R^0(\Omega_z,\Omega_y)$ over $\Omega_y \notin \Omega_x^1$. This is the most likely (actually least modified resistant) exit from the circuit. Also, it will exit to a circuit which is easiest to reach. This in turn suggests that we can form circuits of circuits using modified resistances as the measure of resistance in going from one circuit to another. The system moves between circuits of circuits in a longer time horizon. Moreover, as we have seen the crossings between circuits are direct routes, hence we will define resistance in terms of such paths. We spell out these ideas next.

The procedure we will describe is essentially the same as that employed by Cui and Zhai (2010) to compute the stochastically stable state.[19] Here we employ the procedure to determine the relative probabilities of recurrent communicating classes.

We build circuits recursively starting from $\Omega^0 \equiv \Omega$. Assuming Ω has $N_\Omega \geq 2$ elements, we observe in Appendix 4 that there is at least one circuit that is non-trivial in the sense of having at least two elements, and that every singleton element is trivially a circuit. Hence we can form a non-trivial partition of Ω^0 into circuits, and we denote by Ω^1 the collection of elements of this partition - so an element of Ω^1 is a circuit on $\Omega = \Omega^0$. Given two distinct elements $\Omega_x^1, \Omega_y^1 \in \Omega^1$, define transition resistance $r^1(\Omega_x^1,\Omega_y^1) = \min\{R^0(\Omega_x,\Omega_y) \mid \Omega_x \in \Omega_x^1, \Omega_y \in \Omega_y^1\}$, least outgoing resistance $r^1(\Omega_x^1) = \min\{r^1(\Omega_x^1,\Omega_y^1) \mid \Omega_y^1 \in \Omega^1 \backslash \Omega_x^1\}$ and modified resistance $R^1(\Omega_x^1,\Omega_y^1) = r^1(\Omega_x^1,\Omega_y^1) - r^1(\Omega_x^1)$. A *set* in Ω^1 is a circuit if - as before - between any two of its states there is a path where each transition has least outgoing resistance (according to r^1 of course). Continuing, as long as Ω^{k-1} has more than one element we partition it into circuits, call Ω^k the resulting collection of elements - and for $\Omega_x^k, \Omega_y^k \in \Omega^k$ define $r^k(\Omega_x^k,\Omega_y^k) = \min\{R^{k-1}(\Omega_x^{k-1},\Omega_y^{k-1}) \mid \Omega_x^{k-1} \in \Omega_x^k, \Omega_y^{k-1} \in \Omega_y^k\}$, $r^k(\Omega_x^1) = \min\{r^k(\Omega_x^k,\Omega_y^k) \mid \Omega_y^k \in \Omega^k \backslash \Omega_x^k\}$ and modified resistance $R^k(\Omega_x^k,\Omega_y^k) = r^k(\Omega_x^k,\Omega_y^k) - r^k(\Omega_x^k)$. Note that since each partition is non-trivial, this construction has at most N_Ω layers before the

[19]There are a number of known algorithms for computing the stochastically stable state related to that of Cui and Zhai (2010). Hasker (2014) has a good review of the literature. Cui and Zhai (2010) in fact use the terminology of cycle, but as that seems misleading, we prefer the term circuit since movement within a circuit will not necessarily be a cycle.

partition has a single element and the construction stops: $k \leq N_\Omega$.

The crucial function at each layer k turns out to be the following *modified radius* of $x \in \Omega_x$ of order k, defined by

$$\overline{R}^k(x) = \sum_{\kappa=0}^{k} r^\kappa(\Omega_x^\kappa)$$

where $\Omega_x^0 = \Omega_x$ and for each $\kappa > 0$ the element $\Omega_x^\kappa \ni \Omega_x^{\kappa-1}$. It is a measure of the difficulty of traveling far from x and has an intuition similar to that of Ellison (2000)'s modified coradius. We show in Appendix 4 that the following holds:

Theorem 10. *Let k be such that $\Omega_x^k = \Omega_y^k$ - that is x and y are in the same circuit in Ω^k. Then*

$$\frac{\mu_\epsilon(x)}{\mu_\epsilon(y)} \sim \epsilon^{\overline{R}^{k-1}(y) - \overline{R}^{k-1}(x)}.$$

It is useful to define the difference between the modified radii $\overline{R}^{k-1}(y) - \overline{R}^{k-1}(x)$ as the *relative ergodic resistance* of x over y. The theorem says that the relative probabilities within a circuit are proportional to ϵ to the power of the relative ergodic resistance. Note here the implication: if x and y are in the same circuit in Ω^k then of course they have to be in the same circuit in any Ω^κ for $\kappa > k$. Hence in this case $\overline{R}^{\kappa-1}(y) - \overline{R}^{\kappa-1}(x) = \overline{R}^{k-1}(y) - \overline{R}^{k-1}(x)$.

Traditionally interest has focused on the states x that have ergodic probabilities that are bounded away from zero - the *stochastically stable* states. We can see from Theorem 10 that in any circuit these states must have non-positive relative ergodic resistance over any other state and strictly negative over any state that is not also stochastically stable. Since this must be true in any circuit, it must be true in the top level circuit that contains all recurrent communicating classes - that is to say the Ω^k where k is the highest level of the filtration where the partition is a singleton. We thus have the following

Corollary 5. *The stochastically stable states are exactly those with the highest values of $\overline{R}^{k-1}(x)$, where k is the level at which the partition into circuits is a singleton.*

6.4. An Example

To illustrate the application of Theorem 10, let us give a complete analysis of the case where Ω has three elements.[20] Note that there are 9 trees on 3 states, so the analysis by means of trees is already difficult. For simplicity let us make the generic assumption that no two resistances or sums or differences of resistances are equal.

There are two cases: either there is a single circuit, or there is one circuit consisting of two states, and a isolated point. The case of a single circuit is trivial - in this case in Theorem 10 $k = 1$ and since $\overline{R}^0(x) = r^0(\Omega_x)$ the relative ergodic resistances are given simply by the differences in least outgoing resistances between the three states, and the stochastically stable state is the point with least outgoing resistance.

[20] See also Hasker (2014).

Take then, the case of Ω with one two-point circuit and an isolated point, and denote by Ω_x, Ω_y the two states on the circuit and with Ω_z the remaining point. Then $k = 2$ in the theorem. Assume without loss of generality that $r^0(\Omega_x) > r^0(\Omega_y)$ so that within the circuit Ω_x is relatively more likely. Notice that $r^0(\Omega_x, \Omega_y) < r^0(\Omega_x, \Omega_z), r^0(\Omega_y, \Omega_x) < r^0(\Omega_y, \Omega_z)$ since Ω_x, Ω_y are on the same circuit - this also implies $r^0(\Omega_x) = r^0(\Omega_x, \Omega_y), r^0(\Omega_y) = r^0(\Omega_y, \Omega_x)$. Turning to the recursion, we need to work out the least modified resistances. Let $\Omega^1_x = \{\Omega_x, \Omega_y\}$ be the circuit and $\Omega^1_z = \Omega_z$ the isolated point. Then $r^1(\Omega^1_z, \Omega^1_x) = \min\{r^0(\Omega_z, \Omega_x) - r^0(\Omega_z), r^0(\Omega_z, \Omega_y) - r^0(\Omega_z)\} = 0$ while $r^1(\Omega^1_x, \Omega^1_z) = \min\{r^0(\Omega_x, \Omega_z) - r^0(\Omega_x), r^0(\Omega_y, \Omega_z) - r^0(\Omega_y)\}$. Hence $\overline{R}^1(z) = r^0(\Omega_z)$, which is just the radius of Ω_z, while

$$\begin{aligned}\overline{R}^1(x) &= r^0(\Omega_x) + \min\{r^0(\Omega_x, \Omega_z) - r^0(\Omega_x), r^0(\Omega_y, \Omega_z) - r^0(\Omega_y)\} \\ &= \min\{r^0(\Omega_x, \Omega_z), r^0(\Omega_x) + r^0(\Omega_y, \Omega_z) - r^0(\Omega_y)\}\end{aligned}$$

which is to say exactly what Ellison (2000) defines as the modified co-radius of Ω_z.[21] The relative ergodic resistance of x over y is therefore $r^0(\Omega_z) - \overline{R}^1(x)$, while the relative ergodic resistance of y can be recovered from the relative ergodic resistance of y over x which is just $r^0(\Omega_x) - r^0(\Omega_y)$. With respect to stochastic stability, we see that Ω_z is stochastically stable if and only if its radius $r^0(\Omega_z)$ is greater than its co-radius $\overline{R}^1(x)$ which is Ellison (2000)'s sufficient condition, and otherwise Ω_x is not stochastically stable. In short, the entire ergodic picture comes down to computing three numbers: the radius and co-radius of Ω_z and the difference between the radii of Ω_x and Ω_y.

7. The Fall of Hegemonies

We now discuss how our results can be used to interpret historical facts concerning sequences of long-run social events of small probability. This is the natural field of application of our theory, which concerns transitions along paths whose steps are each quite unlikely to occur. We will focus in particular on the fall of the Qing dynasty in twentieth century China. We use a variation of the model of Levine and Modica (2013) and Levine and Modica (2012), identifying the fall of a hegemonic society with its progressive loss of land to a different society.[22] We emphasize that this application is limited in scope - the goal is to illustrate how results on transitions lead to interesting predictions. The sensitivity of the predictions to assumptions and the validity of those predictions across a broad range of data are of independent interest but beyond the scope of this paper and discussed in Levine and Modica (2012).

There are a finite number J of societies. In each period t each society j has one of a finite

[21] In fact the modified co-radius is defined as the larger of $\overline{R}^1(x)$ and $\overline{R}^1(y) = \min\{r^0(\Omega_y, \Omega_z), r^0(\Omega_y, \Omega_x) + r^0(\Omega_x, \Omega_z) - r^0(\Omega_x, \Omega_y)\}$. However, $r^0(\Omega_x, \Omega_z) > r^0(\Omega_y, \Omega_x) + r^0(\Omega_x, \Omega_z) - r^0(\Omega_x, \Omega_z)$ and $r^0(\Omega_x, \Omega_y) + r^0(\Omega_y, \Omega_z) - r^0(\Omega_y, \Omega_x) > r^0(\Omega_y, \Omega_z)$ imply $\overline{R}^1(x) \geq \overline{R}^1(y)$.

[22] Levine and Modica (2013) and Levine and Modica (2012) use a model in which the conflict is distinct from learning. The model here simplifies that model by having learning take place on a single unit of land at a time when a society is "unstable." This greatly simplifies presentation of the model without any important consequences for the results. The model here otherwise weakens the assumptions in those papers.

number of internal states $\xi_{jt} \in \Xi_j$. These states evolve according to a fixed Markov process $\Pi_j(\xi_{jt}|\xi_{j,t-1}) > 0$ independent of ϵ so that all transitions are possible. External forces such as disease, climate, other real shocks to productivity, or the interference of outsiders who are are protected themselves by geographical barriers, or superior technology can lead to changes in the internal state; the state may also represent changes in the internal structure of institutions. A good example of the Π_j process within a given society can be found in Acemoglu and Robinson (2001): there external shocks in the form of recessions drive changes in institutions whereby the voting franchise is extended or contracted. The ability of a society to resist and influence other societies is indexed by "state power" $\gamma_j(\xi_{jt})$. Societies may or may not satisfy incentive constraints: we represent this by a stability index $b_j \in \{0,1\}$ with 1 indicating stability, where societies violating incentive constraints are thought to be unstable. We assume that the strongest unstable society with the most favorable value of ξ_j is stronger than the strongest stable society in its least favorable value of ξ_j. That is $\max_{j|b_j=0,\xi_{jt}\in\Xi_j}\gamma_j(\xi_{jt}) > \max_{j|b_j=1}\min_{\xi_{jt}\in\Xi_j}\gamma_j(\xi_{jt})$. This reflects the idea that unstable societies face weaker incentive constraints.

Societies compete over a single resource called land. Each society j holds an integral number of units of land L_j where there are L units of land in total. A state z is a list of land holding and real shocks of the different societies, $z = (L_1, \xi_1, L_2, \xi_2, \ldots, L_J, \xi_J)$. Land changes ownership between societies due to conflict. We assume that at most one unit of land changes hands each period. The probability that society j loses a unit of land is given by a conflict resolution function with resistance $r_j(z) < \infty$ to j losing one unit of land - that is the resistance of a transition from z to a state where j has one less unit of land. If j looses land the probability the land goes to society k has *land gain resistance* $\lambda_{jk}(z)$ where $\lambda_{jj} = \infty$ but if $k \neq j$ then $\lambda_{jk} < \infty$.

Conceptually in this model there are two distinct types of societies: *active societies* that have positive land holdings and *inactive societies* that do not. Inactive societies represent templates for societies that might exist but do not exist currently: an inactive society may become active because when an active society loses land the land may be lost to an inactive society - that is the loss of land may represent experimentation with new institutions.

We make several specific assumptions about the conflict resolution and land gain resistance. We assume that r and λ depend only on the land holdings, state power and stability of the different societies. Since they are just templates for societies, we assume that the state power of inactive societies does not matter. We assume that conflict resolution resistance is monotone, so that the probability of j losing land decreases with its own state power and land, and increases with that of other societies. In particular, we assume that unstable societies always have zero resistance to losing a unit of land - if incentive constraints are not satisfied, individuals experiment with different actions, and societies experiment with different institutions (albeit just on a single unit of land at a time); and for stable societies we assume the resistance is strictly monotone when non-zero, and that the weakest society with positive land holding has zero resistance to losing a unit of land. Finally we assume for given land holding that resistance is greater when facing more than one opponent with positive land holding than when all enemy land is in the hands of the strongest land holding opponent - also that this is strict if resistance is positive. With respect to land gain resistance we assume that if $k \neq j$ and $L_k > 0$ then $\lambda_{jk}(z) = 0$, that is, active societies all have zero resistance to gaining land.

7.1. The Rise, Fall and Warring States

We apply the same outline of analysis to this model as we did to the simple example of Section 2: we look first for the recurrent communicating classes, then for the relationship between them, then the basins, then the relationship between the basins and recurrent communicating classes and finally examine the transitions.

1. Recurrent communicating classes for $\epsilon = 0$.

A *hegemony* is a single society that controls all the land. The assumption that the weakest society with positive land holding has zero resistance to losing a unit of land plays a key role in determining the recurrent communicating classes: it implies that there is a zero resistance path from every non-hegemonic state to a hegemonic state: monotonicity implies that losing land cannot increase resistance, so the weakest society keeps losing land until hegemony is established. Second, there is a zero resistance path from any unstable hegemony to a stable hegemony: by assumption the unstable hegemony has zero resistance to losing land and by symmetry the first unit of land lost has zero resistance to being "taken" by a stable society - which once it becomes active continues to have zero resistance to taking over the next unit of land and so forth.

There are three cases, depending on what the stable hegemonies are like. It may be that the stable hegemony has so little state power that it also has zero resistance to losing land. In this case there is a zero resistance path from a hegemony to any state in which there are no more than two active societies (and in particular from any hegemony to any other) and from there to other states without an hegemonic society so that there is only one recurrent communicating class and it is "large" in the sense that it includes within it unstable as well as stable societies, and societies with all levels of state power. Second, it may be that there is a single stable hegemony that has the greatest state power and that only this hegemony is strong enough to resist losing land. In this case that single hegemony Ω_x is the only recurrent communicating class. If $y \in \Omega_x$ then $j(x)$ can denote the single stable society $j(x)$ that controls all the land, while the different states $y \in \Omega_x$ correspond to different shocks $\xi_{j(x)}$. Finally, it may be that two or more stable societies are strong enough to have resistance when hegemonic. These different societies may have the same or different state power. Regardless, these societies each constitute a recurrent communicating class. Our interest here is in the third case - about how hegemonies fall - that is, how we move from a hegemony Ω_x to $B = \Omega \backslash \Omega_x, W = B$.

2. Relation between the classes for $\epsilon > 0$.

Again we ask: starting in one recurrent communicating class, which comes next, how long will it take and relatively how much time is spent at each recurrent communicating class?

The first step is the standard one of understanding the radius and least resistance paths out of the basin which by Corollary 3 we know are the most likely ways of leaving the basin. Because of monotonicity a single invader with the greatest state power in the most favorable state always has the least resistance to gaining a unit of land from the hegemon. Hence a path in which such an invader repeatedly takes land from the hegemon is a least resistance direct route out of the basin. There is a threshold level of land such that once the hegemon loses this amount of land it loses resistance - because this "optimal" invader is necessarily at least as strong as the hegemon and because the weaker of the two societies always has zero resistance to losing land, this threshold is no more than 50% of the land.

Next we observe that we have assumed that unstable societies can generate more state power than any stable society. This means that the "optimal" invader must be unstable. We refer to the unstable society that generates the greatest state power as "zealots." These are the "optimal" invaders.

Notice that once the basin has been exited, zealots can continue to gain land with zero resistance until there is a hegemony of zealots. At this point, since zealots are unstable and have zero resistance to losing land themselves their hegemony is transient, and any stable society can enter and take all the land away from the zealots until they themselves form a hegemony. Hence: from any $\Omega_x \in \Omega$ there is a least resistance path to any other $\Omega_y \in \Omega$. That means that all the recurrent communicating classes are in the same circuit, and so by Theorem 9 their relative probabilities simply depend on the differences in their radii. Since the radius is determined entirely by state power and is strictly increasing in state power, this means that hegemonies with greater state power are far more likely in the long run than societies with less state power. Notice how the (very reasonable) assumption that greater power is generated by ignoring incentive constraints than by satisfying them leads to this very simple long-run dynamic.

By the earlier work of Ellison (2000) and Theorem 4 we know that the waiting time for a hegemony to fail is determined by the radius and hence by the state power. Hegemonies with greater state power are more durable.

3. Basins.

The states in the basin of a hegemony Ω_x are simply those states in which the society has positive resistance to losing land for all internal states $\xi_{j(x)}$: since this means that $j(x)$ is not the weakest society, it follows that there is zero resistance to some other society losing land, and since $j(x)$ has zero resistance to gaining that land and continues to have positive resistance to losing land when it increases its land holding, this means a probability one path to the hegemony when $\epsilon = 0$. The inner basin is just that subset of the basin which can be hit with less resistance than the radius: this is typically a much smaller set than the basin.[23]

The outer range consists of states that are in the basin but have resistance strictly greater than

[23] A point is in the basin but not the inner basin if the hegemony is reached with probability 1 when $\epsilon = 0$ but the resistance to reaching the point is greater than the radius. So, for example, if a sufficiently weak opponent occupies a large fraction of the land then it will be overcome by the hegemony when $\epsilon = 0$ with probability 1, yet starting at the hegemony the resistance to such a weak opponent grabbing so much land can be much greater than the resistance to determined zealots taking enough land to put an end to the hegemony.

the radius, plus a large collection of states which have zero resistance both to returning to the hegemony and some other recurrent communicating class. This is true of any state in which an unstable society is the strongest (in its strongest internal state) - such as those states reached after invasion by zealots.[24] The inner range just adds those states in the outer range that have resistance equal to the radius.

4. Relation between basins and classes.

By Theorem 4, during the time before reaching a new class most of the time will be spent in the current recurrent communicating class, but there will be a large number of periods during which the system will move to states in the inner basin and back. From Theorem 8 the relative amount of time spent at these states compared to the time spent in the recurrent communicating class is and exponential function of the difference between the resistance of reaching the point and the radius: further states are less likely.

5. Transitions

How we get from one recurrent communicating class to another is covered in Theorem 3. This says that when ϵ is small it is nearly certain that this transition will take place along a least peak resistance path, and we want to describe such paths in the model at hand: such a path may leave the recurrent communicating class and return any number of times, but during each departure from the recurrent communicating class the resistance encountered can be no more than the radius.

Notice that from an empirical point of view we probably do not know what the inner basin looks like - but we may have a pretty good idea of a bound on the amount of land that the hegemon needs to be in the inner basin. For example we may think that if the hegemon loses 30% or less of its land it remains "safe" in the sense that it still has resistance. This shows some of the strength of using arbitrary quasi-comprehensive sets W. We can simply take the forbidden set and the target set $B = W$ to be the set of states in which the hegemon controls at least 70% of the land - all the theorems about least peak resistance apply equally well to this W, except that instead of radius, we simply say "resistance to losing 30% of land" which has more empirical content.

Prior to the fall of the hegemony, by Theorem 4 the theory predicts that there should be a small fraction but large number of periods where there are failed rebellions: lands that are lost to other societies but quickly regained. These failed rebellions may or may not involve zealots, and need not take place when ξ_j is at its nadir. However, prior to the actual exit, ξ_j must be such that state power is at a nadir, and this means that resistance to rebellions is lower, so there should be more frequent and larger rebellions prior to the final fall of the society.

The exit path must have least resistance to leaving the basin. We know one such path in which the zealots simply grab one unit of land after another. This implies that the hegemon can lose

[24] As we remarked above: because the unstable society is strongest there is zero resistance to the unstable society taking over. However, once the unstable society has taken over it has zero resistance to losing all its land to any hegemony. In other words from such a state there is a zero resistance path to every recurrent communicating class. Consequently these states are in the outer range, yet they are certainly not in the basin.

land only to the zealots - since any other loss would incur greater resistance - and that it can never regain land, since then the zealots would have to regain it. This is exactly like in the discussion of transitions in the example of Section 2.

Once the inner basin is breached we enter the outer range. We call this the *fall of the hegemony*. Once the hegemony loses resistance there will be several societies competing each with an appreciable chance of success. We refer to this turbulent period before the basin of another hegemony is reached as the period of *warring states*. During this period there will be many societies which may rise and fall, and swap land back and forth - it is a chaotic and turbulent phase. The exit path to another recurrent communicating class will then be concluded with a *rise* of a new hegemony. The rise of the new hegemony is in some respects opposite of the fall. Once a stable society has enough land that it has positive resistance to any opponents (implied if they have positive resistance to an opponent consisting entirely of zealots), least resistance implies it can only gain land and not lose it.

The entire period of transition we know (Theorem 4) to be short relative to the length of the hegemony, and of course that must be true individually for each of the three phases. We would like to say something about the relative length of the three phases. The fall and the rise are both monotone: during the fall the hegemony only loses land; during the rise the new hegemony only gains land. Hence we expect these phases to be relatively fast. By contrast during the warring states land may swap back and forth many times before a victor is established. This leads us to suspect that the warring states period should last considerably longer than the fall and the rise. However, if we simply fix the number of units of land, no such result is possible: the warring states period may involve only a few units of land, while the fall and the rise could involve many units of land. In this case - since just a few units of land are involved - the back and forth during the warring states period does not much matter, since there is a good chance one society quickly gains the small number of units of land needed to establish a new hegemony.

Warring States with Many Units of Land

It is natural to think in terms of relatively small units of land - for example the Alsace-Lorraine region, a relatively small area - swapped back and forth between France and Germany four times in less than a century. To capture "small units" of land, we model the idea that the number of units of land L is large. In this case we might expect that the length of the rise and fall - being monotone - will last an amount of time roughly proportional to L, while the warring states which is more like a random walk would last an amount of time more nearly proportional to L^2.

To make this precise, define the share of land held by society j as $\theta_j = L_j/L$ and suppose that the conflict resolution and land gain resistance functions are continuous functions of these shares. We also suppose that there is a threshold $\overline{\theta} > 1/2$ such that if no society has this fraction of land then the land gain resistance of all active societies must be zero. Notice that this implies that any path from one hegemony to another must enter the warring states phase, since at some point the original hegemon falls below $\overline{\theta}L$ units of land for the first time and at this point the warring state phase is entered, since certainly no other society has that much land.[25]

As we increase L the number of states, which consists of all ways of dividing L units of land among a fixed number J of societies, grows very rapidly. Hence bounds on the expected length of direct routes that depend on the number of states are not going to be particularly useful. In Appendix 1 specific bounds for least resistance paths are given that do not depend on the number of states and a stronger version of this is given in Web Appendix 1. These bounds show that the least resistance paths corresponding to the rise and fall have an expected length proportional to L.

We also want to argue that the warring state phase last much longer than the fall or the rise. To do this we need an additional assumption: we need to know not just resistance during the warring states phase, but something about the actual transition probabilities. We assume that during the warring states phase when no society has more than $\overline{\theta}L$ units of land each active society has the same probability $\beta < 1/J$ of losing or gaining a unit of land. Notice that at some point a new would-be hegemon must have $L/2$ units of land. Hence take an initial condition in which the society with the most land is j and $L_{jt} = L/2$ units of land. Then $L_{j\tau}$ is a random walk with β chance of increasing by one or decreasing by one at least until either $L_{j\tau} \geq \overline{\theta}L$ or $L_{j\tau} \leq (1-\overline{\theta})L$.[26] In Web Appendix 2 we establish that the expected passage time is of order L^2, which can be compared to the expected length of the rise and fall in Web Appendix 2 which are of order L. We put these results together in Web Appendix 3 to establish

Proposition 2. *For any K there exists an $\overline{L}$ such that for all $L \geq \overline{L}$ there exists an $\overline{\epsilon}$ such that for all $\epsilon \leq \overline{\epsilon}$ the expected length of the warring states phase exceeds that of either the fall or rise by K periods.*

Note the order of limits here: for larger L we will generally have to choose smaller ϵ.

7.2. The Fall of the Last Qing Dynasty in China

An interesting exercise is to compare the theoretical predictions of the transition to the fall of an actual hegemony. As a case study for which there is quite a bit of historical information, we take the fall of the final Qing dynasty in China and the subsequent rise of the communist hegemony.[27] The basic fact is that Chinese institutions that lasted from roughly the introduction of the Imperial Examination System in 605 CE until 1911 CE were swept away in less than a year. It is useful to begin the story about 1838, before the First Opium War. At that time the Qing dynasty held a hegemony over China proper: the area bordered by the difficult terrain of Indochina in the Southeast, the Himalayan mountains in the South, the inhospitable deserts in the West, the Pacific Ocean in the East and the wasteland of Mongolia in the North. It also held a number of

[25]Since we are assuming L is large we are ignoring the rounding off needed due to the integer constraint.

[26]Even if $L_{j\tau} \leq (1-\overline{\theta})L$ no other society may have enough land to become hegemon, but certainly no other society can have enough land to become hegemon until this condition is satisfied, so we can use this condition to derive a lower bound on the expected length of the warring states period.

[27]There are of course many accounts of this period, and while they sometimes disagree on exactly who did what to whom when, all agree on the basic facts we describe below. One readable account by a journalist is that of Fenby (2008).

outlying areas not part of China proper - the Korean Peninsula, Indochina and Taiwan. As these are not so easily defended, are not Chinese, and have only been part of the Chinese hegemony at certain times - and moreover, the current government claims only Taiwan among these territories - we do not count them as part of the hegemony.

Several independent sources of instability concurred to the fall of the hegemony. In the early 1800s China fell into a severe economic depression from which it did not recover prior to the fall of the hegemony. Outsiders, most notably the English, French, and Japanese actively intervened in China, sometimes fighting for and other times against the Qing, but in any case certainly piling on pressure. Opium consumption, induced by the English to correct trade imbalances, increased as well.

From 1839 to 1910 there were a series of unsuccessful attempts to overthrow the Qing dynasty including local rebellions and acts of defiance by committed revolutionaries. During this time the outlying territories were lost: Korea became independent, Indochina was lost to the French, and Taiwan to the Japanese further weakening the hegemon. Roughly speaking the state ξ_j became increasingly worse. However, each internal rebellion was successfully repressed, each war brought to an end, and in each case the Qing hegemony over China proper - tax collecting authority, control of institutions local and global - remained intact. There were institutional changes that took place during this period, some forced by the outsiders, and an attempt to placate the revolutionaries, such as the abolition of the imperial examination system in 1905. These can be viewed as shocks ξ_j that further weakened the state. Although it is hard to measure the relative frequency of failed rebellions before and after the economic weakness of the 19th Century, in the earlier periods there seem not to have been such dramatic episodes as the Boxer rebellion and the less known Duggan revolt (which lasted for fifteen years). As Theorem 4 predicts, before the actual fall the state ξ_j is very bad, and there are many and probably increasing failed attempts at rebellion.

The actual fall of the Qing occurred in 1911 and as Theorem 1 suggests, it was very quick. There were again a series of revolts - now however they succeeded. Also as the theory suggests, the length of the successful revolt - less than a year - is considerably shorter than the longest failed rebellions - the Boxer and Duggan rebellion lasting many years. The final successful revolt is coordinated by Sun Yat Sen. The groups carrying out the various revolts can reasonably be described as zealots: they share in common a dedication to overthrowing the Qing, they are willing to suffer severe risk and live under unpleasant circumstances in order to achieve that goal. Such behavior is power maximizing - but is not stable in the sense that no society has ever lasted very long based on the fanatical devotion of its members - nor was it the case in China. Hence the theoretical description of the fall of the hegemony is relatively accurate: zealots quickly capture the land, and do so without a serious setback. In some cases land is seized by other groups, but they quickly join Sun Yat Sen as the theory suggests. By the end of 1911 the Qing Emperor abdicated and Sun Yat Sen became the provisional President of China, which however no longer was hegemonic in any reasonable sense of the word.

Next is the period of warring states, both in theory and in fact. The theory says that there can be many competing societies, land may be lost and gained, zealots may or may not play a role. Again, this is an accurate description of the situation in China between 1911 and 1946. Sun Yat Sen was quickly deposed by a less fanatical and more materialistic warlord Yuan Shikai, but until about 1927, and even after, there are many warlords in various parts of China who rise and fall, many revolutions, some successful and other unsuccessful. There is also the Sino-Tibetan war and the Soviet invasion of Xinjiang during this period. Basically the theory predicts chaos (in the non-technical sense) and that is what we see. Beginning in about 1927 things settle down slightly with two relatively more powerful groups, the Nationalists and the Communists, fighting a civil war - but there remain many warlords who continue to rise and fall, at times forming alliances or professing allegiance to the two more significant groups. These two groups, unlike the earlier revolutionaries appear to have coherent and potentially stable institutions. Then in 1936 the Japanese seize control of most of the country, an occupation that lasts until 1945. Notice that as the theory suggests the length of this warring states period - 35 years - is much longer than either the fall (less than a year) and the rise (about three years).

The final stage of a least resistance transition is the rise of the new hegemon. Again all transitions must have zero resistance, but now we are in the basin of the hegemony so the least resistance path consists of the hegemony gaining territory - without losing any - until hegemony is again established. Notice that since in this model once the basin is left there are zero resistance transitions to any particular hegemony breaching the threshold, the model makes no prediction about which hegemony eventually emerges - in particular there is a non-negligible probability that even a very weak hegemony emerges. In China, the threshold appears to be reached about 1946 when the Communists controlled about a quarter of the country and about a third of the population. They quickly overran the remaining areas held by the Nationalists, who retreated to Taiwan in 1949.

8. Conclusion

This paper is about events and combinations of events that are unlikely and that can be modeled as a finite Markov process, in particular how such a process moves from one relatively stable long-run state to another. Examples are transitions between different equilibria in a game or different political regimes. We show that these systems exhibit long periods of stability punctuated by brief episodes of change, and we give a detailed description of the probabilities and frequencies of these different outcomes. Within the literature on "evolution of conventions" we complement the results of Kandori, Mailath and Rob (1993), Young (1993), Ellison (2000), Cui and Zhai (2010) and Hasker (2014) on long run dynamics in games. When applied to the context of social evolution, our theory has implications both for the societies we are likely to see and for the design of institutions: institutions that will persist for long periods of time must be robust against multiple failures, and it is these multiple failures that lead a society to fall.

It may be useful to look at smaller systems about which we have a great deal of information - also subject to small unlikely shocks, and subject to the same type of Markov analysis - to see what is involved. For example commercial airlines, which despite the vast number of flights and miles flown crash relatively infrequently. As our theory predicts, when they do crash, it is typically due to multiple near simultaneous failures. To take a specific example, on November 24, 2001, en route from Berlin on approach to Zurich Crossair Flight 3597 crashed near Bassersdorf, Switzerland killing 24 of the 33 people on board. According to the flight investigation seven independent unfortunate events occurred on that occasion.[28] These multiple failures seem typical of commercial aviation crashes. Each individual failure is unlikely, but none terribly so. What is highly unlikely is that all occur in combination. In general airplanes are designed with a high degree of redundancy to provide insulation against failure of one or even several components: multiple pilots, multiple navigation systems, multiple engines, multiple independent hydraulic systems and so forth. So it is with human societies. For example penal codes and the legal systems have a high degree of redundancy (appeals procedures) to prevent the punishment of the innocent. Societies that survive for long periods of time must be well cushioned against even multiple failures. For example, the fall of the Roman Empire has been attributed to many factors: religious ferment, the plague, corruption, the forced migration of hostile outsiders, economic recession, and so forth. Despite the effort of historians to establish each as “the” cause of the fall, as is the case with Flight 3597 all of these things happened - and while each is uncommon, none is particularly unlikely, and the Roman Empire had suffered through each of these, often in combination, many times before. What is unique about the fall is that all these things occurred at once. When a system or society is well designed it takes a perfect storm - everything going wrong at once - to bring it down. But - as this paper shows - it is the least unlikely combination of things - the least resistance direct route - that will typically lead - for good or ill - to abrupt and sudden change.

[28] See AAIB (2002): (1) the pilot had a bad record of following procedures during landing and was inadequately trained, but was allowed never-the-less to transport passengers; (2) the flight was behind schedule and consequently the pilot was in a hurry to land; (3) due to noise regulations the plane was diverted to a less safe runway; (4) the runway had inadequate instrumentation and the airport parameters and protocols for landing on the runway were inadequate; (5) the range of hills the plane crashed into was not marked on the chart; (6) the pilot put the plane into an overly steep descent and descended too low without proper visual contact with the ground; (7) the pilot did not monitor the proper instruments during the attempted landing.

Appendix 1: Direct Routes

Our goal is to establish probability and expectations bounds on subsets $A \subseteq A_{xBW} \neq \emptyset$. Define $t(A) = \min\{t(a)|a \in A, r(a) = r(A)\}$ to be the minimum number of transitions of any least resistance path in the set A. This Appendix is devoted to proving the following main result on direct paths which characterizes their probability and length. Theorem 1 in the text follows directly.

Theorem 11. *For $A \subseteq A_{xBW}$ non-empty and $t \geq 0$ there are bounds $G(t), H(t) > 0$ non-decreasing in t such that $C^{t(A)}\epsilon^{r(A)} \leq P_\epsilon(A|x) \leq G(t(A))\epsilon^{r(A)}$ and $E[t(a)|x, A] \leq H(t(A))$.*

The bounds G, H will be specifically computed.

First we establish that $r(A) = \min_{a \in A} r(a)$ exists. Then getting a lower bound on $P_\epsilon(A|x)$ is relatively easy: it is bounded below by the probability of a path $a \in A$ with resistance $r(a)$, which is to say, it is of order $\epsilon^{r(A)}$. The main goal is to establish a similar upper bound. The problem is that A can easily contain infinitely many paths with resistance $r(A)$ as well as paths of greater resistance. However, there are only finitely many paths of any given length, so if there are infinitely many paths most of them must be very long. The idea is that since paths in A must avoid the comprehensive set W they are not likely to be very long since there are zero resistance routes to W. To make this precise we construct a finite set of template paths of relatively low resistance and show that all the paths in A can be constructed by adding loops to the template paths. We then show that the probability of all paths constructed from a given template is bounded by the probability of the template times a constant that does not depend on ϵ. This same method also yields bounds on the expected length of the direct routes.

Loop Cutting and a Lower Bound

In general paths $a \in A$ contain loops, and since an analysis of loops form a key part of the analysis we begin by introducing the notion of loop-cutting. The idea is to construct templates for a from which a can be reconstructed by adding loops. If $a = (z_0, z_1, \ldots, z_t)$ we say that a' is a *loop-cut* of a at $z_\tau = z_{\tau'}$ for $t > \tau' > \tau \geq 0$ if $a' = (z_0, z_1, \ldots, z_\tau, z_{\tau'+1}, \ldots, z_t)$, that is if the loop at z_τ has been cut out.[29] Note the obvious fact that $r(a') \leq r(a)$.

Definition 7. A map m from the set of all paths to itself is a *loop-cutting algorithm* if there is a sequence $a_1, a_2, \ldots, a_M$ with $a_1 = a$, $a_M = m(a)$ and a_{j+1} a loop-cut of a_j for $j = 1, 2, \ldots, M-1$.

Note that $r(m(a)) \leq r(a)$. The path $m(a)$ is a template for a from which a can be reconstructed by adding loops. A loop cutting algorithm is *maximal* if $m(m(a)) = m(a)$.

We can establish the existence of $\min_{a \in A} r(a)$ using the *zero-cut* algorithm defined as follows. For any $a = (z_0, z_1, \ldots z_t)$ if there is no loop of zero resistance stop. Otherwise cut the first and shortest loop of zero resistance and repeat. This is obviously maximal. Note that $m(a)$ is a *no-zero-loop* path in the sense that it contains no zero-resistance loops, and that $r(m(a)) = r(a)$. Our first

[29]Note that there cannot be a loop-cut that begins with the final element z_t. Indeed $z_t \in B$, when the path gets there it stops.

step is to give a bound on the length of no-zero-loop paths. Let Z_A to be the set of non-end points touched by paths in A, that is the set of z_τ such that there is some $(x, z_1, z_2, \ldots, z_\tau, \ldots, z_t) \in A$ with $\tau < t$ - leading and trailing elements not counted. Let N_A be the number of elements in Z_A and N_{AB} the maximum of N_A and the number of elements in the target B. These are both bounded above by N, but may be much smaller: using computations based only on A allows our results to be extended from a finite state space to a countable state space. Let $\underline{r}(A)$ be the smallest finite non-zero resistance of any transition in any path in A.

Lemma 1. *If $a \in A$ is a no-zero-loop path*[30] *then $t(a) \leq N_A r(a)/\underline{r}(A)$.*

Proof. Observe that since non-zero resistance transitions have resistance at least $\underline{r}(A)$ there are at most $r(a)/\underline{r}(A)$ such transitions in a, and the remaining transitions must have zero resistance. Since there are no zero-resistance loops, the number of zero-resistance transitions between each positive resistance transition is at most N_A. □

We can now apply the zero-cut algorithm to prove the basic fact that

Lemma 2. *$r(A) = \min_{a \in A} r(a)$ is well-defined.*

Proof. Fix $a \in A$. Recall that by assumption a has positive probability for $\epsilon > 0$, so that $r(a) < \infty$. Let m be the zero-cut algorithm. Consider that for any $a' \in A$ with $r(a') \leq r(a)$ we have $r(a') = r(m(a'))$. Let $\overline{A}$ be the set of all finite resistance paths $(x, z_1, z_2, \ldots, z_\tau, \ldots, z_t)$ with $z_t \in B$ and $z_\tau \in Z_A$. Since $m(a') \in \overline{A}$ by Lemma 1 $t(m(a')) \leq N_{\overline{A}} r(a)/\underline{r}(\overline{A})$. But there are only finitely many paths of length t that begin at x and take values in Z_A and end in B, so only finitely many possible values of $r(a') \leq r(a)$. Hence $\min_{a \in A} r(a)$ exists. □

Having established that $r(A)$ is well-defined we can easily establish a lower bound on $P_\epsilon(A|x)$.

Lemma 3. *$P_\epsilon(A|x) \geq C^{t(A)} \epsilon^{r(A)}$.*

Proof. Let $a \in A$ satisfy $r(a) = r(A)$ and $t(a) = t(A)$.Then

$$P_\epsilon(A|x) \geq P_\epsilon(a|x) = \prod_{\tau=1}^{t(a)} P_\epsilon(z_\tau|z_{\tau-1}) \geq \prod_{\tau=1}^{t(a)} C\epsilon^{r(z_\tau, z_{\tau-1})} = C^{t(a)} \epsilon^{r(a)}.$$

□

More About Loops

To establish an upper bound we need further results about loops. First we give a useful refinement on Lemma 1 using the *all-cut* algorithm, defined as follows. For any $a = (z_0, z_2, \ldots z_t)$ if there is no loop stop. Otherwise cut the first and shortest loop and repeat. This is obviously maximal.

Lemma 4. *If $a \in A$ is a no-zero-loop path then*

$$t(a) \leq 2N_A \left(1 + \frac{r(a) - r(A)}{\underline{r}(A)}\right).$$

[30] Note that this results for any set of paths A not just direct routes.

Proof. Let m be the all-cut algorithm. Observe that to get from $m(a)$ to a we must add loops containing at least $t(a)-t(m(a))$ elements. Suppose that there are fewer than $(t(a)-t(m(a)))/2N_A$ loops that cannot be subdivided into two non-overlapping sub-loops. Then one of these loops must have length at least $2N_A$. But the first N_A of the loop must contain a loop and so must the second N_A, so that the loop can be divided into two non-overlapping sub-loops. By assumption none of the loops have zero resistance, so each has resistance at least $\underline{r}(A)$, hence

$$r(a) \geq r(m(a)) + (t(a) - t(m(a)))\frac{\underline{r}(A)}{2N_A}$$

and the result follows from $r(m(a)) \geq r(A)$ and $t(m(a)) \leq N_A$. □

Next we consider a loop-cutting algorithm that produces templates with resistance no smaller than $r(A)$. We say that m *preserves* r if $r(a) \geq r$ implies $r(m(a)) \geq r$. One such algorithm is the *r-preserving* algorithm. For any $a = (x, z_1, \ldots, z_t)$ if no loop can be cut without reducing the resistance of a below r stop. Otherwise cut the first and shortest such loop and repeat. Observe that for an r-preserving algorithm the image $m(A)$ consists of no-zero-loop paths and is maximal since removing any loop would necessarily reduce the resistance below r. The key property of this algorithm is that it produces templates with resistance not too much bigger than r and of bounded length (by Lemma 1)- and in particular that means there are finitely many templates.

Definition 8. $\overline{r}(A)$ to be the greatest (finite) resistance of any transition in any path in A, and

$$\bar{t}(A) \equiv 2N_A\left(2 + \frac{r(A) - r(A_{xBW})}{\underline{r}(A)}\right).$$

Lemma 5. *For $a \in A$ and the $r(A)$-preserving loop-cut algorithm $r(m(a)) \leq r(A) + N_A\overline{r}(A)$ and $t(m(a)) \leq \bar{t}(A)$; and $m(A)$ is finite with at most $N_{AB}^{\bar{t}(A)}$ elements.*

Proof. Observe first if m is the $r(A)$-preserving loop-cut algorithm if $r(m(a)) > N_A\overline{r}(A)$ then $m(a)$ must have a loop of resistance greater than zero, and hence must have a loop of resistance no greater than $N_A\overline{r}(A)$. If $r(m(a)) > r(A) + N_A\overline{r}(A)$ removing such a loop leaves resistance greater than or equal to $r(A)$ contradicting the fact that the $r(A)$-preserving algorithm can leave no such loop.

Now let m be the $r(A)$-preserving algorithm and suppose that $r(m(a)) > r(A_{xBW})$. Remove the shortest loop from $m(a)$ to get a' so that $r(m(a)) \geq r(A) > r(a') \geq r(A_{xBW})$. Since a' has no zero resistance loops, by Lemma 4

$$t(a') \leq 2N_A\left(1 + \frac{r(a') - r(A_{xBW})}{\underline{r}(A)}\right).$$

On the other hand since the a' is $m(a)$ with the shortest loop - so less than N_A in length - removed, we must have $t(m(a)) \leq t(a') + N_A$. Hence

$$t(m(a)) \leq 2N_A\left(2 + \frac{r(a') - r(A_{xBW})}{\underline{r}(A)}\right) \leq 2N_A\left(2 + \frac{r(A) - r(A_{xBW})}{\underline{r}(A)}\right) = \bar{t}(A)$$

If $r(m(a)) = r(\mathcal{A})$ then m removes all the loops, so we have the bound $t(m(a)) \leq N_A$ so that certainly $t(m(a)) \leq \bar{t}(a)$ and that bound holds in all cases.

Finally, since $a \in m(A)$ are constructed of elements from $Z_A \cup B$ and have length at most $\bar{t}(A)$, there are at most $N_{AB}^{\bar{t}(A)}$ elements. $\square$

We also can give a bound in terms of $t(a)$

Lemma 6. *We have*

$$\bar{t}(A) \leq 2N_A \left(2 + \frac{t(A)\bar{r}(A)}{\underline{r}(A)}\right)$$

Proof. Clearly $r(A) \geq 0$. Moreover, since there is a path $a \in A$ with length $t(a) = t(A)$, such a path can have resistance at most $t(A)\bar{r}(A)$, so $r(A)$ cannot be greater than this. $\square$

Upper Bounds

Once the loops have been removed to create templates, we must put them back in to compute probabilities. It is convenient at this point to work with loops with first element removed, so that a loop at z_τ is a sequence of the form $\emptyset, (z_\tau)$ or $(\zeta_\tau, \ldots, \zeta_{\tau+k}, z_\tau)$ where $\zeta_\tau \in Z_A$.[31] If we have a transition $(z_\tau, z_{\tau+1})$ and ℓ_τ is a loop at z_τ then the corresponding path is $a = (z_\tau, \ell_\tau, z_{\tau+1})$, with number of transitions $t(a)$ equal to the number of elements of ℓ_τ plus 1.

For any path $a = (x, z_1, z_2, \ldots, z_t)$ and $0 \leq \tau \leq t$ let $a[\tau] = (x, z_1, z_2, \ldots, z_\tau)$ be the corresponding τ-truncation. Then for any $a = (z_0, z_1, z_2, \ldots, z_t) \in A$, $\tau < t$ and path a' consider the path $(a[\tau], a') = (z_0, z_1, z_2, \ldots, z_\tau, a')$ that starts along a and deviates to a' at z_τ. Define $t_F(a, \tau)$ to be the least length of any path (z_τ, a') that has zero resistance and $(a[\tau], a') \notin A$ (so that the deviation does not reach B). Notice that $t_F(a, \tau) \leq N - 1$: since W is comprehensive there is a path (z_τ, a') of zero resistance with no loops that ends in W, and such a path can have at most N elements, hence at most $N - 1$ transitions.

Definition 9. The *failure time* $t_F(A) = \max_{a \in A, \tau < t(a)} t_F(a, \tau)$.

Hence $t_F(A) \leq N - 1$ but may be much smaller: in one of the examples in the text $t_F(A) = 1$ regardless of N.

To establish upper bounds on the probability of A and on the expected length of its paths, we start by establishing the fact that, given that W is comprehensive, long loops are not very likely. Let $\lfloor R \rfloor$ be the largest integer not greater than R. We can now give the following bound on the probability of long loops.

Lemma 7. *Suppose $z \in Z_A$ and that $L_A(z)$ is a set of loops at z. Then for $t \geq 0$ we have $P_\epsilon\big(t((z, \ell, y)) = t + 1,\ \ell \in L_A(z) | z\big) \leq P_\epsilon(y|z)(1 - C^{t_F(A)})^{\lfloor t/t_F(A) \rfloor}$.*

Proof. If $t = 0$ the result is immediate since in fact the left and right sides are equal. Since (z, ℓ, y) always ends with the transition (z, y) we have

$$P_\epsilon\big(t((z, \ell, y)) = t + 1,\ \ell \in L_A(z) | z\big) \leq P_\epsilon(y|z)$$

[31] Note that we must include $\emptyset$ because when we compute probabilities of transitions from z to y along a set of loops, we must also include the probability that we go directly from z to y without a loop, that is, along a null loop.

for all t and in particular for $t < t_F(A)$ which is the desired bound in that case. If $t \geq t_F(A)$ we may define $L_k$to be the set of paths (z, ℓ) with $\ell \in L_A(z)$ truncated at $kt_F(A)$ for $1 \leq k \leq \lfloor t/t_F(A) \rfloor$ and we have

$$P_\epsilon\big(t((z,\ell,y)) = t+1,\ \ell \in L_A(z)|z\big) \leq P_\epsilon(L_{\lfloor t/t_F(A)\rfloor}|z)P_\epsilon(y|z)$$

Moreover $L_{\lfloor t/t_F(A)\rfloor} \subseteq L_k$ for $1 \leq k \leq \lfloor t/t_F(A)\rfloor$, so it suffices to prove recursively that $P_\epsilon(L_k|z) \leq (1 - C^{t_F(A)})^k$ for $1 \leq k \leq \lfloor t/t_F(A)\rfloor$.

First we take $k = 1$. Observe that starting at z there is a zero resistance path of length no longer than $t_F(A)$ that reaches a point y that is contained in no loop. Since the probability of each zero-resistance transition is at least C there is at most a probability $1 - C^{t_F(A)}$ of remaining in the set L_1.

Now we suppose the result is true for k and prove it for $k+1$. Each loop in L_{k+1} has the form (a_1, z', a_2) where $(a_1, z') \in L_k$ and $t(a_2) = t_F(A)$. Let $L^+(a_1, z')$ be the set $\{a_2|(a_1, z', a_2) \in L_{k+1}\}$. Then

$$\begin{aligned}
P_\epsilon(L_{k+1}|x) &= \sum_{(a_1,z')\in L_k} \sum_{a_2 \in L^+(a_1,y)} P_\epsilon((a_1,z')|x)P_\epsilon(a_2|z') \\
&= \sum_{(a_1,z')\in L_k} P_\epsilon((a_1,z')|x) \sum_{a_2 \in L^+(a_1,y)} P_\epsilon(a_2|z') \\
&= \sum_{(a_1,z')\in L_k} P_\epsilon((a_1,z')|x)P_\epsilon(L^+(a_1,z')|z').
\end{aligned}$$

Moreover, since again there is a zero resistance path starting at z' of length no longer than $t_F(A)$ that reaches a point y' that is contained in no loop we have $P_\epsilon(L^+(a_1, z')|z') \leq 1 - C^{t_F(A)}$. Hence

$$\begin{aligned}
P_\epsilon(L_{k+1}|x) &\leq \sum_{(a_1,z')\in L_k} P_\epsilon((a_1,yz')|x)(1 - C^{t_F(A)}) \\
&= P_\epsilon(L_k|x)(1 - C^{t_F(A)}) \\
&\leq (1 - C^{t_F(A)})^{k+1}
\end{aligned}$$

by the inductive hypothesis. □

We are now ready to reverse the loop-cutting procedure by adding loops to templates to construct the paths in A. Opposite to a loop cutting algorithm is the idea of a *loop insertion set*. Let $L_A(z_\tau)$ be a set of loops at z_τ. We suppose we are given an $a = (z_0, z_1, \ldots, z_t) \in A$. A loop set $\overline{m^{-1}}(a)$ is defined by a sequence of sets of loops $L_\tau \subseteq L_A(z_\tau), \tau = 0, 1, \ldots t-1$ and consists of all paths of the form $(z_0, \ell_0, z_1, \ell_1, \ldots, z_{t-1}, \ell_{t-1}, z_t)$ such that $\ell_\tau \in L_\tau$. If m is a loop-cutting algorithm and $a \in m(A)$ if $\overline{m^{-1}}(a) \supseteq A \cap m^{-1}(a)$ we say that $\overline{m^{-1}}(a)$ *covers* m, a.

We now define

$$S_k(\overline{m^{-1}}(a)|x) \equiv \sum_{t_0=0}^{\infty}\sum_{t_1=0}^{\infty}\cdots\sum_{t_{t(a)-1}=0}^{\infty}\left[\left(\sum_{s=0}^{t(a)-1} t_s\right)^k\right]\prod_{\tau=0}^{t(a)-1} P_\epsilon\big(t((z_\tau,\ell,z_{\tau+1})) = t_\tau + 1,\ \ell \in L_\tau|z_\tau\big).$$

The significance of these numbers is given by the following

Lemma 8. *If $\overline{m^{-1}}(a)$ covers m, a then*

$$\sum_{a' \in \overline{m^{-1}}(a)} P_\epsilon(a'|x) \leq S_0(\overline{m^{-1}}(a)|x)$$

$$\sum_{a' \in \overline{m^{-1}}(a)} t(a') P_\epsilon(a'|x) \leq S_1(\overline{m^{-1}}(a)|x)$$

and both holds with equality if every $a' \in \overline{m^{-1}}(a)$ has a unique representation of the form $a' = (z_0, \ell_0, z_1, \ell_1, \ldots z_t)$ where $\ell_\tau \in L_\tau$.

Proof. By definition every $a' \in \overline{m^{-1}}(a)$ has a representation of the form $a' = (z_0, \ell_0, z_1, \ell_1, \ldots z_t)$ where $\ell_\tau \in L_\tau$.[32] Hence for any non-negative function $f(a')$ we have

$$\sum_{a' \in \overline{m^{-1}}(a)} f(a') P_\epsilon(a'|x) \quad \leq \quad \sum_{\ell_0 \in L_0} \sum_{\ell_1 \in L_1} \cdots \sum_{\ell_{t(a)-1} \in L_{t(a)-1}} f((z_0, \ell_0, z_1, \ell_1, \ldots z_t)) \prod_{\tau=0}^{t(a)-1} P_\epsilon(\ell_\tau | z_\tau)$$

with equality if the representation is unique - that is, if each a' has a unique representation then it appears exactly once in the sum. Rearranging the sum by adding over the length of the loops then gives the desired result. □

We can now compute the desired upper bounds. First we have

Lemma 9. $S_0(\overline{m^{-1}}(a)|x) \leq P_\epsilon(a|x)\left[t_F(A)/C^{t_F(A)}\right]^{t(a)}$ and $S_1(\overline{m^{-1}}(a))/S_0(\overline{m^{-1}}(a)) \leq 3t(a)t_F(A)^2/C^{2t_F(A)}$

Proof. From Lemma 7 we have

$$\begin{aligned}
S_0(\overline{m^{-1}}(a)|x) &= \sum_{t_0=0}^{\infty} \sum_{t_1=0}^{\infty} \cdots \sum_{t_{t(a)-1}=0}^{\infty} \prod_{\tau=0}^{t(a)-1} P_\epsilon(t((z_\tau, \ell, z_{\tau+1})) = t_\tau + 1,\ \ell \in L_\tau | z_\tau) \\
&= \prod_{\tau=0}^{t(a)-1} \sum_{t=0}^{\infty} P_\epsilon(t((z_\tau, \ell, z_{\tau+1})) = t + 1,\ \ell \in L_\tau | z_\tau) \\
&\leq \prod_{\tau=0}^{t(a)-1} \sum_{t=0}^{\infty} P_\epsilon(z_{\tau+1}|z_\tau)(1 - C^{t_F(A)})^{\lfloor t/t_F(A) \rfloor} \\
&= \left(\prod_{\tau=0}^{t(a)-1} P_\epsilon(z_{\tau+1}|z_\tau) \right) \prod_{\tau=0}^{t(a)-1} \sum_{t=0}^{\infty} (1 - C^{t_F(A)})^{\lfloor t/t_F(A) \rfloor} \\
&= P_\epsilon(a|x)\left[t_F(A)/C^{t_F(A)}\right]^{t(a)}
\end{aligned}$$

Next to simply notation set

$$P_\tau(t_\tau) \equiv P_\epsilon(t((z_\tau, \ell, z_{\tau+1})) = t_\tau + 1,\ \ell \in L_\tau | z_\tau).$$

[32] A simple example shows that there can be more than one representation. Suppose $(z_0, z_1, z_2) \in m(A)$. If $\ell_0 = (z_1, z_0), \ell_1 = \emptyset$ then $a' = (z_0, \ell_0, z_1, \ell_1, z_2) = (z_0, z_1, z_0, z_1, z_2)$. If $\ell_0 = \emptyset, \ell_1 = (z_0, z_1)$ then $(z_0, \ell_0, z_1, \ell_1, z_2) = (z_0, z_1, z_0, z_1, z_2) = a'$.

Then

$$\begin{aligned}
S_1(\overline{m^{-1}}(a)|x) &= \sum_{t_0=0}^{\infty}\sum_{t_1=0}^{\infty}\cdots\sum_{t_{t(a)-1}=0}^{\infty}\left(\sum_{s=0}^{t(a)-1} t_s\right)\prod_{\tau=0}^{t(a)-1} P_\tau(t_\tau) \\
&= \sum_{s=0}^{t(a)-1}\sum_{t_0=0}^{\infty}\sum_{t_1=0}^{\infty}\cdots\sum_{t_{t(a)-1}=0}^{\infty} t_s \prod_{\tau=0}^{t(a)-1} P_\tau(t_\tau) \\
&= \sum_{s=0}^{t(a)-1}\left(\frac{\sum_{t=0}^{\infty} t_s P_s(t)}{\sum_{t=0}^{\infty} P_s(t)}\right)\prod_{\tau=0}^{t(a)-1}\sum_{t=0}^{\infty} P_\tau(t) \\
&= \sum_{s=0}^{t(a)-1}\left(\frac{\sum_{t=0}^{\infty} t_s P_s(t)}{\sum_{t=0}^{\infty} P_s(t)}\right) S_0(\overline{m^{-1}}(a)|x).
\end{aligned}$$

Moreover,

$$\begin{aligned}
\sum_{t=0}^{\infty} P_s(t)) &\geq P_s(1) = P_\epsilon(t((z_s,\ell,z_{s+1})) = 1,\ \ell \in L_\tau | z_s) \\
&= P_\epsilon(z_{s+1}|z_s)
\end{aligned}$$

and applying again Lemma 7 and using a summation formula proven in Web Appendix 1 Lemma 11 we have

$$\begin{aligned}
S_1(\overline{m^{-1}}(a)|x)/S_0(\overline{m^{-1}}(a)|x) &\leq \sum_{s=0}^{t(a)-1}\left(\frac{\sum_{t_s=0}^{\infty} t_s P_\epsilon(z_{s+1}|z_s)(1-C^{t_F(A)})^{\lfloor t_s/t_F(A)\rfloor}}{P_\epsilon(z_{s+1}|z_s)}\right) \\
&= t(a)\sum_{t_s=0}^{\infty} t_s(1-C^{t_F(A)})^{\lfloor t_s/t_F(A)\rfloor} \\
&\leq t(a)3t_F(A)^2/C^{2t_F(A)}
\end{aligned}$$

which is the desired bound. □

We can now establish the upper bounds stated in Theorem 11.

Theorem 12. *If $A \subseteq A_{xBW}$ is non-empty then* $P_\epsilon(A|x) \leq \left[N_{AB}Dt_F(A)/C^{t_F(A)}\right]^{\bar{t}(A)} \epsilon^{r(A)}$.

Proof. Take m to be the $r(A)$-preserving loop cut algorithm, and for $a \in m(A)$ take $\overline{m^{-1}}(a)$ to be defined by the sequence $L_\tau = L_A(z_\tau)$. Let $T(A) = \max_{a\in m(A)} t(a) \leq \bar{t}(A)$. Since $P_\epsilon(\overline{m^{-1}}(a)|x) \leq S_0(\overline{m^{-1}}(a)|x)$ we may apply Lemma 9 and use the fact that $\overline{m^{-1}}(a)$ covers m,a to conclude $P_\epsilon(A|x) \leq \#m(A)\epsilon^{r(A)}\left[Dt_F(A)/C^{t_F(A)}\right]^{T(A)}$. Moreover, by Lemma 5 we have $\#m(A) \leq N_{AB}^{\bar{t}(A)}$ giving the desired result. □

Theorem 13. *If $A \subseteq A_{xBW}$ is non-empty then*

$$E[t(a)|x,A] \leq \bar{t}(A)\left[\frac{3t_F(A)^2}{C^{2t_F(A)}}\right]\frac{\left[N_{AB}Dt_F(A)/C^{t_F(A)}\right]^{\bar{t}(A)}}{C^{t(A)}}$$

and if A are the least resistance paths and $\tilde{t}(A)$ is the longest least resistance path containing no loops then

$$E[t(a)|x,A] \quad \leq \quad \tilde{t}(A)\left[\frac{3t_F(A)^2}{C^{2t_F(A)}}\right].$$

Proof. In all cases by Lemma 9 we have for any loop cutting algorithm that preserves $r(A)$

$$\begin{aligned} E[t(a)|x,A] &\leq \frac{\sum_{a\in m(A)} S_1(\overline{m^{-1}}(a)|x)}{P_\epsilon(A|x)} \\ &= \sum_{a\in m(A)} \frac{S_1(\overline{m^{-1}}(a)|x)}{S_0(\overline{m^{-1}}(a))}\frac{S_0(\overline{m^{-1}}(a))}{P_\epsilon(A|x)} \\ &\leq \frac{3t_F(A)^2}{C^{2t_F(A)}}\sum_{a\in m(A)}\frac{t(a)S_0(\overline{m^{-1}}(a))}{P_\epsilon(A|x)}. \end{aligned}$$

In general we can take m to be the $r(A)$-preserving loop cut algorithm, and for $a \in m(A)$ take $\overline{m^{-1}}(a)$ to be defined by the sequence $L_\tau = L_A(z_\tau)$. Then we observe that there are at most $N_{AB}^{\bar{t}(A)}$ templates and apply Theorems 3 and 12 to get

$$E[t(a)|x,A] \quad \leq \quad \frac{3t_F(A)^2}{C^{2t_F(A)}} N_{AB}^{\bar{t}(A)} \frac{\bar{t}(A) D^{\bar{t}(A)}\epsilon^{r(A)}\left[t_F(A)/C^{t_F(A)}\right]^{\bar{t}(A)}}{C^{t(A)}\epsilon^{r(A)}}$$

giving the first result.

Now suppose that A are the least resistance paths. We can now take m to be the all-cut algorithm, which, since no least resistance path can have a positive resistance loop in it, is the same as the zero-cut algorithm when applied to A. For $a \in m(A)$ we now take $\overline{m^{-1}}(a)$ to be defined by the sequence L_τ of zero-resistance loops in $L_A(z_\tau)$ which do not contain z_s for $s < \tau$. [33]

First we observe that $\overline{m^{-1}}(a) = m^{-1}(a) \cap A$. Starting at z_τ the all-cut algorithm cuts the longest loop ending at z_τ, hence z_τ cannot subsequently appear in the template. Moreover, the loops added back are exactly the ones that were cut. Second we observe that for a given template $(z_1, z_2, \ldots, z_t) \in m(A)$ and two states $a' = (z_0, \ell'_0, z_1, \ell'_1, \ldots, z_t), a" = (z_0, \ell_0", z_1, \ell_1", \ldots, z_t) \in \overline{m^{-1}}(a)$ then $a' = a"$ only if $\ell'_0, \ell'_1, \ldots, \ell'_{t-1} = \ell_0", \ell_1", \ldots, \ell_{t-1}"$. Hence $S_0(\overline{m^{-1}}(a)) = P_\epsilon(m^{-1}(a) \cap A)$. So

$$\begin{aligned} E[t(a)|x,A] &\leq \frac{3t_F(A)^2}{C^{2t_F(A)}}\sum_{a\in m(A)}\frac{t(a)P_\epsilon(m^{-1}(a)\cap A)}{P_\epsilon(A|x)} \\ &= \frac{3t_F(A)^2}{C^{2t_F(A)}}\sum_{a\in m(A)}\frac{t(a)P_\epsilon(m^{-1}(a)\cap A)}{\sum_{a\in m(A)} P_\epsilon(m^{-1}(a)\cap A)} \\ &\leq \tilde{t}(A)\frac{3t_F(A)^2}{C^{2t_F(A)}}. \end{aligned}$$

giving the second result. □

[33]Notice that the construction in the example of Footnote 32 is ruled out in the present case.

Appendix 2: Quasi-Direct Routes

We fix a start point $x \in \Omega_x$, a target set B and quasi-comprehensive forbidden set $W \supseteq B$, and let Q_{xBW} be the corresponding non-empty set of quasi-direct routes. Recall that the paths $a \in Q_{xBW}$ can be decomposed as $a_1, a_2, \cdots, a_{n(a)}, a^+$ where the $a_i \in A_{xxW}$ are loops from x to x that do not touch x or W in between and $a^+ \in A_{xBW}$ is the exit path, a direct route from x to B that does not touch W nor x in between.

Theorem. *[Theorem 3 in text] Let* $A = \{a \in Q_{xBW} | \rho(a) = \rho(Q_{xBW})\}$ *denote the least peak resistance paths in* Q_{xBW}*. Then* $\lim_{\epsilon \to 0} \frac{P_\epsilon(A|x)}{P_\epsilon(Q_{xBW} \setminus A|x)} = \infty$.

Proof. Fix $\rho = \rho(Q_{xBW})$. By Theorem 2 the set A consists of the set of paths of the form $a_1, a_2, \ldots, a_n, a^+$ where $a_i \in A_{xxW}$ and $a^+ \in A_{xBW}$ and $r(a_i) \leq \rho, r(a^+) = \rho$. Define the set of paths $A_{>\rho}$ as those having the form $a_1, a_2, \ldots, a_n, a_{n+1}$ where $a_i \in A_{xxW}$ and $a_{n+1} \in A_{xxW} \cup A_{xBW}$ and $r(a_i) \leq \rho, r(a_{n+1}) > \rho$. We claim that $P_\epsilon(A_{>\rho}|x) \geq P_\epsilon(Q_{xBW} \setminus A|x)$ so that it will suffice to prove that$\lim_{\epsilon \to 0} \frac{P_\epsilon(A|x)}{P_\epsilon(A_{>\rho}|x)} = \infty$. To prove this claim observe that if $a \in Q_{xBW} \setminus A$ then the first part of the path must necessarily lie in $A_{>\rho}$ so the event $Q_{xBW} \setminus A$ implies the event $A_{>\rho}$.

Now let $A^\rho_{xxW}, A^\rho_{xBW}, A^{>\rho}_{xxW}, A^{>\rho}_{xBW}$ denote the subsets of resistance exactly equal to ρ and strictly bigger than ρ respectively. We compute

$$\begin{aligned}
\frac{P_\epsilon(A|x)}{P_\epsilon(A_{>\rho}|x)} &= \frac{\sum_{n=0}^\infty P_\epsilon^n(A^\rho_{xxW}|x) P_\epsilon(A^\rho_{xBW}|x)}{\sum_{n=0}^\infty P_\epsilon^n(A^\rho_{xxW}|x) P_\epsilon(A^{>\rho}_{xxW} \cup A^{>\rho}_{xBW}|x)} \\
&= \frac{P_\epsilon(A^\rho_{xBW}|x)}{P_\epsilon(A^{>\rho}_{xxW} \cup A^{>\rho}_{xBW}|x)} \geq \frac{P_\epsilon(A^\rho_{xBW}|x)}{P_\epsilon(A^{>\rho}_{xxW}|x) + P_\epsilon(A^{>\rho}_{xBW}|x)}
\end{aligned}$$

and the result now follows directly from Theorem 1 on the probability of direct paths. □

When $B = W$ the decomposition also makes it easy to do computations since the loops a_i are independent identically distributed random variables. Specifically, for $f : A_{xxW} \to \Re$ and $a \in A$ define $F(a) \equiv \sum_{i=1}^{n(a)} f(a_i)$. Then for any function $g(n)$ of the number of loops we have

Lemma 10. *if* $B = W$ *then* $E(Fg|x, Q_{xBW}) = E(f|x, \mathcal{A}^0) E(ng|x, Q_{xBW})$*, and* $P_\epsilon(n|x, Q_{xBW})$ *is geometric with* $E(n|x, Q_{xBW}) = (1/P_\epsilon(A_{xBW}|x)) - 1$

Proof. Since $B = W$ we have $P_\epsilon(A_{xxW}|x) + P_\epsilon(A_{xBW}|x) = 1$, while A_{xxW} and A_{xBW} are disjoint. Then abbreviating $Q = Q_{xBW}$ we get

$$E(Fg|x, Q) = E[\sum_{i=1}^n f(a_i) g|x, Q] = E[\sum_{i=1}^n E[f(a_i) g|x, Q, n]|x, Q] = E[\sum_{i=1}^n g E[f(a_i)|x, Q, n]|x, Q].$$

The event (Q, n) is exactly the event $a_i \in A_{xxW}$ for $i = 1, 2, \ldots, n$ and $a_{n+1} \in A_{xBW}$ and conditional on x these are independent events. Hence $E(f(a_i)|x, Q, n) = E(f|x, A_{xxW})$.We conclude that

$$\begin{aligned}
E(Fg|x, Q) &= E[\sum_{i=1}^n g E(f|x, \mathcal{A}^0)|x, Q] \\
&= E(f|x, \mathcal{A}^0) E[\sum_{i=1}^n g|x, Q] = E(f|x, \mathcal{A}^0) E[ng|x, Q].
\end{aligned}$$

This is the first result. Also since $P_\epsilon(A_{xxW}|x) + P_\epsilon(A_{xBW}|x) = 1$ it follows that n is geometrically distributed with success probability $P_\epsilon(A_{xBW}|x)$ which gives the stated expected value. □

Recall that $M(a, A)$ is the number of loops of a that lie in $A \subseteq A_{xxW}$. That is, if $f : A_{xxW} \to \Re$ is the indicator of A (taking the value one if $a^0 \in A$ and zero otherwise) then $M(a, A)$ is the aggregate F defined by $F(a, A) \equiv \sum_{i=1}^{n(a)} f(a_i)$. Also, recall that $t^-(a)$ is the amount of time along a spent outside of $\Omega(x)$. Let Z_{xW} be the subset of Z that is reachable by finite resistance paths that start at x and touch W at most once and Let N_{xW} be the number of elements in Z_{xW}. This is bounded above by N, and sets A of direct routes that we consider satisfy $Z_A \subseteq Z_{xW}$ so that $N_{AB} \leq N_{xW}$. Define

$$G_1 \equiv \left[N_{xW}^2 D / C^{N_{xW}}\right]^{N_{xW}}, \; H_1 = \frac{6N_{xW}^3 G_1}{C^{3N_{xW}}}$$

Theorem. *[Theorem 4 and first two parts of Theorem 5 in text] If $B = W$ then*

$$\epsilon^{-r(A_{xBW})}/G_1 \quad \leq \quad E_\epsilon[t(a)|x, Q_{xBW}], E_\epsilon[t(a^-)|x, Q_{xBW}] \leq H_1 \left(1 + C^{-2N_{xW}} \epsilon^{-r(A_{xBW})}\right) \epsilon^{-r(A_{xBW})}$$

while

$$\lim_{\epsilon \to 0} E_\epsilon[\frac{t^-(a)}{t(a^-)}|x, Q_{xBW}] = 0.$$

For $A \subseteq A_{xxW}$

$$G_1^{-1} C^{2N_{xW}} \epsilon^{r(A) - r(A_{xBW})} \quad \leq \quad E_\epsilon[M(a, A)|x, Q_{xBW}] \leq G_1 C^{-2N_{xW}} \epsilon^{r(A) - r(A_{xBW})}$$

and if $r(A) < r(A_{xBW})$ then for all $k \geq 0$

$$\lim_{\epsilon \to 0} P_\epsilon[M(a, A) > k|x, Q_{xBW}] = 1.$$

Proof. Observe that for $A \in \{A_{xxW}, A_{xBW}\}$ $t_F(A) \leq N_{xW}$ and $\bar{t}(A) = 4N_A \leq 4N_{xW}$ so that the bound from Theorem 12 is in turn bounded by G_1 and that from Theorem 13 by H_1.

From Lemma 10 $E_\epsilon[t(a^-)|x, Q_{xBW}] \leq E_\epsilon[t|x, A_{xxW}]/P_\epsilon(A_{xBW}|x)$. Moreover, recall that $t(A)$ is the number of transitions in the shortest of the least resistance paths in A; since A_{xxW} contains all zero resistance loops $r(A_{xxW}) = 0$, the shortest of these loops is no longer than N_{xW}; so $t(A_{xxW}) \leq N_{xW}$. Analogously, A_{xBW} contains templates without loops hence $t(A_{xBW}) \leq N_{xW}$. Hence from Theorems 12 and 13 we have $1 \leq E_\epsilon[t|x, A_{xBW}] \leq H_1$, $1 \leq E_\epsilon[t|x, A_{xxW}] \leq H_1$ and $C^{N_{xW}} \epsilon^{r(A_{xBW})} \leq P_\epsilon(A_{xBW}|x) \leq G_1 \epsilon^{r(A_{xBW})}$. This gives the stated bounds on $E_\epsilon[t|x, Q_{xBW}]$.

Now set $\tau^- = t^-(a)$ and $\tau = t(a^-)$. For the next part observe that $E_\epsilon[\tau^-/\tau|x, Q_{xBW}] \leq E_\epsilon[\tau^-/n|x, Q_{xBW}] = E_\epsilon[t^-(a)|x, A_{xxW}]$. Now split A_{xxW} into two disjoint sets $\mathcal{A}_0^0$ of paths of zero resistance and $\mathcal{A}_r^0$ of positive resistance, where r is the least positive resistance in A_{xxW}. Then $E_\epsilon[t^-|x, A_{xxW}] = E_\epsilon[t^-|x, \mathcal{A}_0^0] P_\epsilon[\mathcal{A}_0^0|x, A_{xxW}] + E_\epsilon[t^-|x, \mathcal{A}_r^0] P_\epsilon[\mathcal{A}_r^0|x, A_{xxW}]$. However $E_\epsilon[t^-|x, \mathcal{A}_0^0] = 0$ by definition, while by Theorems 12 and 13

$$E_\epsilon[t^-|x, \mathcal{A}_r^0] P_\epsilon[\mathcal{A}_r^0|x, A_{xxW}] \leq [H_1 G_1 / C^{t(\mathcal{A}_r^0)}] \epsilon^r \to 0$$

which implies the result.

Next, given $A \subseteq A_{xxW}$ and f the indicator of A, Lemma 10 gives

$$\begin{aligned} E_\epsilon[M(a,A)|x,Q_{xBW}] &= E_\epsilon[F|x,Q_{xBW}] = E_\epsilon(f|x,A_{xxW})/P_\epsilon(A_{xBW}|x) \\ &= P_\epsilon(A|x,A_{xxW})/P_\epsilon(A_{xBW}|x). \end{aligned}$$

Applying Theorem 12 (and recalling that $r(A_{xxW})=0$) then gives the stated bounds on $M(a,A)$.

Lastly, since $P_\epsilon(n|x,Q_{xBW})$ is geometric, we have $P_\epsilon(n(a)\geq n|x,Q_{xBW}) = (1-P_\epsilon(A_{xBW}|x))^n$ and

$$\begin{aligned} P_\epsilon[M(a,A)>k|x,Q_{xBW},n(a)=n] &= 1-\sum_{i=0}^{k}\left(\begin{array}{c} n \\ i \end{array}\right) P_\epsilon(A|x)^i(1-P_\epsilon(A|x))^{n-i} \\ &\geq 1-(k+1)n^k(1-P_\epsilon(A|x))^n \end{aligned}$$

By hypothesis we can choose $r(A)<r<r(A_{xBW})$. Take $n=\epsilon^{-r}$. Then by Theorem 12 we have $P_\epsilon(n(a)\geq n|x,Q_{xBW}) \geq \left(1-G_1\epsilon^{r(A_{xBW})}\right)^{\epsilon^{-r}}$ and $P_\epsilon[M(a,A)>k|x,Q_{xBW},n(a)=n] \geq 1-(k+1)\epsilon^{-rk}(1-G_1\epsilon^{r(A)})^{\epsilon^{-r}}$. Taking the log of the first expression and using de l'Hospital's rule gives as $\epsilon \to 0$

$$\lim \frac{\log\left(1-G_1\epsilon^{r(A_{xBW})}\right)}{\epsilon^r} = \lim -\frac{G_1 r(A_{xBW})\epsilon^{r(A_{xBW})-1}}{r\epsilon^{r-1}\left(1-G_1\epsilon^{r(A_{xBW})}\right)} = \lim -\frac{G_1 r(A_{xBW})}{r}\epsilon^{r(A_{xBW})-r} = 0$$

so that $P_\epsilon(n(a)\geq n|x,Q_{xBW}) \to 1$. Next take the log of $\epsilon^{-rk}(1-G_1\epsilon^{r(A)})^{\epsilon^{-r}}$ to find

$$-rk\log\epsilon + \epsilon^{-r}\log(1-G_1\epsilon^{r(A)}) = \left[1-\frac{rk\log\epsilon}{\epsilon^{-r}\log(1-G_1\epsilon^{r(A)})}\right]\epsilon^{-r}\log(1-G_1\epsilon^{r(A)}).$$

Then by de l'Hospital's rule

$$\begin{aligned} \lim \frac{rk\log\epsilon}{\epsilon^{-r}\log(1-G_1\epsilon^{r(A)})} &= \lim \frac{rk}{-\frac{r\log(1-G_1\epsilon^{r(A)})}{\epsilon^r} - \frac{G_1 r(A^0)\epsilon^{r(A)-r}}{\log(1-G_1\epsilon^{r(A)})}} \\ &\leq -\lim \frac{rk\log(1-G_1\epsilon^{r(A)})}{G_1 r(A)\epsilon^{r(A)-r}} = 0 \end{aligned}$$

and

$$\lim \frac{\log(1-G_1\epsilon^{r(A)})}{\epsilon^r} = \lim -\frac{G_1 r(A)\epsilon^{r(A)-r}}{r\left(1-G_1\epsilon^{r(A)}\right)} = -\infty$$

so $-rk\log(\epsilon)+\epsilon^{-r}\log(1-G_1\epsilon^{r(A)}) \to -\infty$ and $\epsilon^{-rk}(1-G_1\epsilon^{r(A)})^{\epsilon^{-r}} \to 0$. Hence $P_\epsilon[M(a,A)>k|x,Q_{xBW},n(a)=n] \to 1$. □

Recalling that $A_{xxW}(t)$ are the loops which spend at least t consecutive periods outside of Ω_x we now prove the corollary.

Corollary 6. *[Final part of Theorem 5 in text] If there is a path $a^0 \in A_{xxW}$ that contains a zero resistance loop not touching $\Omega(x)$ with $0<r(a^0)<r(A_{xBW})$ then for any $k>0$*

$$\lim_{\epsilon\to 0} P_\epsilon(M(a,A_{xxW}(kt(a^+)))>k|x,Q_{xBW}) = 1$$

Proof. Fix a $\delta>0$. By Theorem 11 and Chebychev's inequality we may choose K such that

$P_\epsilon(t(a^+) > K) < 1 - \delta$. Given a^0 as described, we can insert zero resistance loops to get a path $a^K \in A_{xxW}(K)$ with resistance $r(a^K) < r(A_{xBW})$. Hence we may apply Theorem 8 to conclude that for all sufficiently small ϵ

$$P_\epsilon[M(a, A_{xxW}(K)) > K|x, Q_{xBW}] < 1 - \delta;$$

since conditional on x the loops and exit path are independent, we then have

$$P_\epsilon(M(a, A_{xxW}(kt(a^+))) > K|x, Q_{xBW}) < (1 - \delta)^2,$$

which implies the desired result. □

Appendix 3: Ergodic Probabilities and Bounds

Theorem. *[Theorem 7 in text] If $y \in \Omega_x$ then*

$$\lim_{\epsilon \to 0} \frac{\mu_\epsilon(x)}{\mu_\epsilon(y)} = \frac{\overline{\mu}_0(x)}{\overline{\mu}_0(y)}.$$

Proof. Partition the matrix P_ϵ with rows corresponding to source states and columns to target states into P_ϵ^{ij} where $i, j = 1$ corresponds to Ω_x and $i, j = 2$ corresponds to $\Omega \backslash \Omega_x$. In particular P^{11} is square, the size of Ω_x. Correspondingly let e^i be the column vectors of ones with dimension corresponding to $i = 1, 2$. Define the row vector $\overline{\mu}_\epsilon(z) = \mu_\epsilon(z)/\sum_{y \in \Omega(x)} \mu_\epsilon(y)$, and partition this vector conformally. Since $\overline{\mu}_\epsilon$is normalized to one on Ω_x and $\overline{\mu}_0$ is strictly positive, it suffices to prove that as $\epsilon \to 0$ every limit point $\overline{\mu}_\epsilon^1$ is equal to $\overline{\mu}_0^1$ where we include the superscript to emphasize that we are dealing only with the invariant distribution on Ω_x. The invariance condition is $\overline{\mu}_\epsilon^1 = \overline{\mu}_\epsilon^1 P_\epsilon^{11} + \overline{\mu}_\epsilon^2 P_\epsilon^{21}$. Multiplying this on the right by e^1 we get $1 = \overline{\mu}_\epsilon^1 P_\epsilon^{11} e^1 + \overline{\mu}_\epsilon^2 P_\epsilon^{21} e^1$ while the fact that P_ϵ is a Markov kernel means that $P_\epsilon^{11} e^1 + P_\epsilon^{12} e^2 = e^1$ or $P_\epsilon^{11} e^1 = e^1 - P_\epsilon^{12} e^2$. Substituting we see that $1 = \overline{\mu}_\epsilon^1 (e^1 - P_\epsilon^{12} e^2) + \overline{\mu}_\epsilon^2 P_\epsilon^{21} e^1 = 1 - \overline{\mu}_\epsilon^1 P_\epsilon^{12} e^2 + \overline{\mu}_\epsilon^2 P_\epsilon^{21} e^1$ so that $\overline{\mu}_\epsilon^2 P_\epsilon^{21} e^1 = \overline{\mu}_\epsilon^1 P_\epsilon^{12} e^2$, which says roughly that the steady state flow into Ω_x must equal the steady state flow out. Now $P_\epsilon^{12} \to 0$ as $\epsilon \to 0$ since these are the probabilities of leaving the recurrent communicating set Ω_x; hence $\overline{\mu}_\epsilon^2 P_\epsilon^{21} e^1 \to 0$. But $\overline{\mu}_\epsilon^2 P_\epsilon^{21}$ is a non-negative vector, so $\overline{\mu}_\epsilon^2 P_\epsilon^{21} e^1 \to 0$ is possible only if $\overline{\mu}_\epsilon^2 P_\epsilon^{21} \to 0$. Then in the invariance condition $\overline{\mu}_\epsilon^1 = \overline{\mu}_\epsilon^1 P_\epsilon^{11} + \overline{\mu}_\epsilon^2 P_\epsilon^{21}$ as $P_\epsilon^{11} \to P_0^{11}$ if $\overline{\mu}_{00}^1$ is a limit point of $\overline{\mu}_\epsilon^1$ it must satisfy the limiting condition that $\overline{\mu}_{00}^1 = \overline{\mu}_{00}^1 P_0^{11}$. However, as Ω_x is recurrent communicating this equation has only one solution $\overline{\mu}_0^1$, so we conclude that in fact $\overline{\mu}_\epsilon^1 \to \overline{\mu}_0^1$. □

Recall that $r(\Omega_x) = \min_{\Omega_y \in \Omega} r(\Omega_x, \Omega_y)$. Also let $\overline{r}, \underline{r}$ be the largest resistance of any transition and smallest positive resistance, and set $\overline{G} \equiv G_1((N-1)(1 + (\overline{r}/\underline{r})))$.

Theorem. *[Theorem 8 in text] Allowing that Ω_x may be empty, if $A = A_{xyW}$ are the direct routes from x to y with forbidden set $W = \{x\} \cup \{y\} \cup (\Omega \backslash \Omega_x)$ then $\mu_\epsilon(y) \geq \mu_\epsilon(x) C^N \epsilon^{r(A)}$. If $x \in \Omega_x$ and there is a zero resistance path from y to x then also $\mu_\epsilon(y) \leq \mu_\epsilon(x) C^{-N} \overline{G}^2 \epsilon^{\min\{r(A), r(\Omega_x)\}}$.*

Proof. We use the standard fact about Markov ergodic probabilities as used for example by Ellison (2000): if we let $N_\epsilon(y, x|x)$ be the expected number of times y occurs before x starting at x then $\mu_\epsilon(y) = \mu_\epsilon(x) N_\epsilon(y, x, |x)$.

The lower bound is immediate: since with probability $P_\epsilon(A)$ we have y hit once without returning to x we have from Theorem 11 $\mu_\epsilon(y) = \mu_\epsilon(x) N_\epsilon(y, x|x) \geq \mu_\epsilon(x) P_\epsilon(A) \geq \mu_\epsilon(x) C \epsilon^{r(A)}$.

Next we suppose that y has zero resistance for getting to $x \in \Omega_x$. We use the reverse condition $\mu_\epsilon(x) = \mu_\epsilon(y) N_\epsilon(x, y, |y)$, so we must find a lower bound on $N_\epsilon(x, y|y)$. Let $A_1 = A_{yx(x \cup y \cup (\Omega \backslash \Omega_x))}$.

Observe that $N_\epsilon(x,y|y) \geq P_\epsilon(A_1)N_\epsilon(x,y|x)$. Since there is a zero resistance path from y to x we have from Theorem 11 the bound $P_\epsilon(A_1) \geq C^N$, so $N_\epsilon(x,y|y) \geq C^N N_\epsilon(x,y|x)$.

Now define sets $B = \{y\} \cup (\Omega \backslash \Omega_x)$ and $A_2 = A_{xB(x \cup B)}$. Then $N_\epsilon(x,y|x) \geq N_\epsilon(x,B|x))$. Since starting at x the events B and $\sim B = A_2$ are mutually exclusive independent events, $N_\epsilon(x,B|x)) = 1/P_\epsilon(\sim B) = 1/P_\epsilon(A_2)$. From Theorem 11 $P_\epsilon(A_2) \leq \epsilon^{r(A_2)}$, and we get $N_\epsilon(x,y|y) \geq C^N \epsilon^{-r(A_2)}/\overline{G}$, or $\mu_\epsilon(y) \leq \mu_\epsilon(x) C^{-N} \overline{G}^2 \epsilon^{r(A_2)}$.

Finally the event A_2 is contained in the event $A_{xy(x \cup y \cup (\Omega \backslash \Omega_x)} \cup A_{x(\Omega \backslash \Omega_x)(x \cup (\Omega \backslash \Omega_x))}$. Hence

$$r(A_2) = \min\{r\left(A_{xy(x \cup y \cup (\Omega \backslash \Omega_x)}\right), r\left(A_{x(\Omega \backslash \Omega_x)(x \cup (\Omega \backslash \Omega_x))}\right)\}.$$

However $r\left(A_{x(\Omega \backslash \Omega_x)(x \cup (\Omega \backslash \Omega_x))}\right) = r(\Omega_x)$ and $r\left(A_{xy(x \cup y \cup (\Omega \backslash \Omega_x)}\right) = r(A)$ which gives the desired upper bound. □

References

Acemoglu, D. and J. A. Robinson (2001): "A Theory of Political Transitions," *The American Economic Review* **91**: 938-963.

Aircraft Accident Investigation Bureau (2002), "Final Report concerning the accident to the aircraft AVRO 146-RJ100, HB-IXM, operated by Crossair under flight number CRX 3597, on 24 November 2001 near Bassersdorf/ZH," Report 1793.

Binmore, K., and L. Samuelson (1997): "Muddling Through: Noisy Equilibrium Selection," *Journal of Economic Theory*, 74: 235-265.

Blume, L. E. (2003): "How noise matters," *Games and Economic Behavior*, 44: 251.

Cui, Z. W. and J. A. Zhai (2010): "Escape dynamics and equilibria selection by iterative cycle decomposition" *Journal of Mathematical Economics* 46: 1015—1029.

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

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

Ellison, G., D. Fudenberg and L. A. Imhof (2014): "Fast convergence in evolutionary models: A Lyapunov approach", mimeo

Fenby, J. (2008): *Modern China: The Fall and Rise of a Great Power, 1850 to the Present,* Penguin Press.

Foster, D. P. and H. P. Young (1990): "Stochastic Evolutionary Game Dynamics," *Theoretical Population Biology* **38**: 219-232.

Freidlin, M. I. and A. D. Wentzell (2012): *Random Perturbations of Dynamical Systems*, 3rd Edition, Springer.

Hasker, K.(2014): "The Emergent Seed: A Representation Theorem for Models of Stochastic Evolution and two formulas for Waiting Time," mimeo Bilkent University.

Kreindler, G. E. and H. P. Young (2012): "Rapid Innovation Diffusion with Local Interaction," Oxford Economics Discussion Papers Series.

Kreindler, G. E. and H. P. Young (2013) "Fast Convergence in Evolutionary Equilibrium Selection," *Games and Economic Behavior*, 80: 39-67.

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

Levine, D. K. and S. Modica (2013): "Anti-Malthus: Conflict and the Evolution of Societies," *Research in Economics*, 67: 289-306

Levine, D. K. and S. Modica (2012): “Conflict, Evolution, Hegemony and the Power of the State,” working paper, www.dklevine.com.

Sabourian, H. and W.-T. Juang (2012) “A Folk Theorem on Equilibrium Selection,” mimeo, Cambridge.

Young, H. P. (1993): “The Evolution of Conventions,” *Econometrica* 61: 57-84.

Web Appendix 1: Further Bounds on the Length of Direct Paths

We first prove two lemmas.

Lemma 11. *For $0 < C < 1$ and $t_F \geq 0$ we have $\sum_{t_s=0}^{\infty} t_s(1-C^{t_F})^{\lfloor t_s/t_F \rfloor} \leq 3t_F^2/C^{2t_F}$*

Proof. We have

$$\begin{aligned}
\sum_{t_s=0}^{\infty} t_s(1-C^{t_F})^{\lfloor t_s/t_F \rfloor} &= \sum_{k=0}^{\infty}\sum_{h=1}^{t_F}(kt_F+h)(1-C^{t_F})^k \\
&= \sum_{h=1}^{t_F}\left(t_F\sum_{k=0}^{\infty}k(1-C^{t_F})^k + h\sum_{k=0}^{\infty}(1-C^{t_F})^k\right) \\
&= \sum_{h=1}^{t_F(A)}\left(t_F(1-C^{t_F})/C^{2t_F} + h/C^{t_F}\right) \\
&= \left(t_F^2(1-C^{t_F})/C^{2t_F} + (t_F^2+t_F)/(2C^{t_F})\right) \\
&\leq 3t_F^2/C^{2t_F}
\end{aligned}$$

giving the desired result. □

Lemma 12. . *If $A \subseteq A_{xBW}$ is not empty and W is comprehensive then for $t \geq 0$ we have*

$$P_\epsilon\big(t(a)=t+1,\ a\in A|x\big) \leq (1-C^{t_F(A)})^{\lfloor (t+1)/t_F(A)\rfloor}$$

and if B is a singleton then

$$P_\epsilon\big(t(a)=t+1,\ a\in A|x\big) \leq \max_{(x,z_1,z_2,\ldots,z_{t-1},z_t)\in A} P_\epsilon(z_t(a)|z_{t-1}(a))(1-C^{t_F(A)})^{\lfloor t/t_F(A)\rfloor}$$

Proof. The first inequality was proven in the course of proving Lemma 7. The second makes use of the fact that in the course of proving Lemma 7 we used only the fact that all the loops ended at the same target and that all had the same transition probability at the end. If we replace the unique final transition probability with the maximum over all final transition probabilities the same argument goes through. □

In Theorem 13 a better bound is given for least resistance paths that exploits the fact that they have a special structure. The idea is that long least resistance paths are not likely to be very long because to be long they must contain long loops, and long loops are not very likely. For least resistance paths these loops must have zero resistance, however in a large state space we could have zero resistance pieces of least resistance paths that are "unnecessarily" long but do not in fact loops. Our goal is to show that these too are unlikely. To do so, we introduce the idea of a *waypoint* of a path $a = (z_0, z_1, \ldots, z_t)$. Let $(z_{\tau-1}, z_\tau)$ be the first transition in the path that has positive resistance. The first waypoint is defined as z_τ. Similarly, the second waypoint is defined to be the end of the second transition in the path that has positive resistance and so forth. We say that two paths a, a' are equivalent, written $a \sim a'$ if they have the same waypoints. The idea is now to give conditions for least resistance paths under which the amount of time between waypoints is bounded independent of the size of the state space, and consequently get a bound on the expected length of least resistance paths of order equal to the number of waypoints. Let $Y(A)$ be the set of sequences of waypoints derived from paths in A, and for any given sequence of waypoints $y \in Y(A)$ let $A_{\tau-1}(y)$ be the set of least resistance paths from $z_{\tau-1}$ to z_τ.

Theorem 14. *If W is comprehensive and $A \subseteq A_{xBW}$ not empty is the set of all least resistance paths then*

$$E_\epsilon(t(a)|x,A) \leq \max_{y=(z_0,z_1,\ldots,z_{t-1})\in Y(A)} t\left[\max_{0\leq s\leq t-1} 3Dt_F(A_s(y))^2/C^{2t_F(A_s(y))+t(A_S(y))}\right]$$

Proof. Pick $y = (z_0, z_1, \ldots, z_{t-1}) \in Y(A)$, that is a sequence of waypoints, and let A_y be the paths with those waypoints. Notice these sets form a partition of A. If a_τ is a sequence of states (indexed starting with 1), let $z_s(\tau)$ be the sth element of the sequence and $s(\tau)$ the length of the sequence. Since that paths in question are least resistance paths, they are exactly paths of the form $(a_0, a_1, \ldots, a_{t-1})$ where

* $z_1(0) = x$
* either $z_{s(t-1)}(t-1) \in B$ or $z_{t-1} \in B, a_{t-1} = \emptyset$
* any transitions in a_τ have zero resistance
* transitions $z_{s(\tau-1)}(\tau-1), z_1(\tau)$ have positive resistance $r_{\tau-1}$ that depends only on τ
* $(a_{\tau-1}, z_1(\tau))$ is a least resistance path from $z_1(\tau-1)$ to $z_1(\tau)$ (with forbidden set W).

Put differently, setting $A_{\tau-1} = A_{\tau-1}(y)$ (the set of least resistance paths from $z_{\tau-1}$ to z_τ) then a path is a least resistance path if and only if $a_{\tau-1} \in A_{\tau-1}$and $a_\tau \in A_\tau$ implies that any transitions in a_τ have zero resistance, and the transition $z_{s(\tau-1)}(\tau-1), z_1(\tau)$ has positive resistance equal to $r_{\tau-1}$ independent of which path in $A_{\tau-1}$ is chosen. Let $P_\tau(t) \equiv P_\epsilon\big(t((a, z_{\tau+1})) = t+1,\ a \in A_\tau|x\big)$. Then (using the same algebra as Lemma 9) we have

$$\begin{aligned} E(t(a)|x,A_y) &= \frac{\sum_{t_0=0}^{\infty}\sum_{t_1=0}^{\infty}\cdots\sum_{t_{t-1}=0}^{\infty}\left(\sum_{s=0}^{t-1} t_s\right)\prod_{\tau=0}^{t-1} P_\tau(t_\tau)}{\prod_{\tau=0}^{t(a)-1}\sum_{t=0}^{\infty} P_\tau(t)} \\ &= \sum_{s=0}^{t-1}\frac{\sum_{t=0}^{\infty} t_s P_s(t)}{\sum_{t=0}^{\infty} P_s(t)}. \end{aligned}$$

As in Lemma 7 by using Lemma 12 and Lemma 11 we find

$$\begin{aligned} \sum_{t=0}^{\infty} t_s P_s(t) &\leq \sum_{t_s=0}^{\infty} t_s D\epsilon^{r_s}(1 - C^{t_F(A_s)})^{\lfloor t_s/t_F(A_s)\rfloor} \\ &\leq D\epsilon^{r_s} 3t_F(A_s)^2/C^{2t_F(A_s)}. \end{aligned}$$

On the other hand, $\sum_{t=0}^{\infty} P_s(t) \geq C^{t(A_s)}\epsilon^{r_s}$, which gives the desired bound. □

As we move away from a recurrent communicating class along a least resistance path, initially we are in the basin of the class and we encounter resistance. This gives a natural monotonicity to this part of the path: each time we encounter resistance we cannot go back and do it again because to do so would add unnecessary resistance. The bounds in Theorem 14 exploits this monotonicity and so is useful in bounding the time it takes to get out of the basin. However, once we leave the basin there will be zero resistance paths to other recurrent communicating classes, and so there will be no more waypoints and the bound is not useful. Indeed, as Web Appendix 2 shows, the length of time in this region may not scale. However, in applications such as the model of hegemony, once we get close enough to the recurrent communicating class that will be the end of the least resistance path, there may be a form of monotonicity: in the example there is a point at which the eventual hegemon can only gain land (along a least resistance path) and not lose it. If in place of the natural monotonicity of Theorem 14 we assume monotonicity then we can get a bound for this final segment of the least resistance path.

To formalize this, we first give a bound on the probability of zero resistance paths in the basin. Suppose that for comprehensive W the set $A \subseteq A_{xBW}$ of least resistance paths is not null. Define $r_{xBW} \equiv \min\{r(A_{x(W\backslash B)W}), r(A_{xBW}\backslash A)\}$ and $t_{xBW} \equiv \max\{t(A_{x(W\backslash B)W}), t(A_{xBW}\backslash A)\}$. Notice that $r_{xBW} > 0$ means that $r(A) = 0$ since there must be some zero resistance path from x to W, and that x is in the basin of B since all 0 resistance direct routes from x to B are in A.

Theorem 15. *If $r_{xBW} > 0$ then$P_\epsilon(A|x) \geq 1 - 2G(t_{xBW})\epsilon^{r_{xBW}}$.*

Proof. Since W is comprehensive, with probability 1 every path originating at x hits W with probability 1. Hence $P_\epsilon(A_{x(W\backslash B)W}|x) + P_\epsilon(A_{xBW}\backslash A|x) + P_\epsilon(A|x) = 1$. However, by Theorem 11 we have $P_\epsilon(A_{x(W\backslash B)W}|x), P_\epsilon(A_{xBW}\backslash A|x) \leq G(t_{xBW})\epsilon^{r_{xBW}}$ giving the desired result. □

Now consider a sequence of targets $B_1, B_2, \ldots, B_t$ where $B_t = B$. Also set $B_0 = \{x\}$. For any a starting at x we may consider $t_1(a)$ the first time B_1 is hit before hitting W, possibly infinite, and if $B_1$is hit before W we may consider $t_2(a)$ the additional amount of time from first hitting B_1 until $B_2$is hit before hitting W, again infinite if either target is not hit before reaching W, and so forth. We say that the sequence is a *Liapunoff* sequence for A if for every a we have $t_\tau(a) < \infty$. In this case the sequence of states $(z_1, z_2, \ldots, z_t)$ that are hit are similar to waypoints. For $y \in B_\tau$ let $A_\tau(y) \equiv \mathcal{A}(y, B_{\tau+1}, W)$. Let $t_{FF}(A) \equiv \max_{0 \leq \tau < t} t_F(A_\tau)$. Then

Theorem 16. *If $B_1, B_2, \ldots, B_t$ is a Liapunoff sequence for least resistance paths A then*

$$E_\epsilon(t(a)|x, A) \leq t\,\frac{1}{P_\epsilon(A|x)}\frac{3t_{FF}(A)^2}{C^{2t_{FF}(A)}}$$

Proof. Define $\underline{t}_\tau(a)$ to be $t_\tau(a)$ if it is finite, zero otherwise, and observe that for $a \in A$ we have $t_\tau(a) = \underline{t}_\tau(a)$. Hence we may write

$$\begin{aligned} E_\epsilon(t(a)|x, A) &= \sum_{\tau=0}^{t-1} E_\epsilon(\underline{t}_\tau(a)|x, A) \\ &= \sum_{\tau=0}^{t-1} \frac{E_\epsilon(\underline{t}_\tau(a)|x, A)P_\epsilon(A|x)}{P_\epsilon(A|x)} \\ &\leq \frac{1}{P_\epsilon(A|x)} \sum_{\tau=0}^{t-1} E_\epsilon(\underline{t}_\tau(a)|x). \end{aligned}$$

Moreover $E_\epsilon(\underline{t}_\tau(a)|x) \leq \max_{y \in B_\tau} E_\epsilon(\underline{t}_\tau(a)|y)$ as either $\underline{t}_\tau(a)$ is zero or a hits some $y \in B_\tau$ before hitting $B_{\tau+1}$ by definition. The desired bound now follows from Lemma 12 and the summation formula Lemma 11. □

Web Appendix 2: Expected Passage Time Bounds

Let V_t a standard Weiner process with 0 drift and instantaneous variance 1 that starts at 0. Now let T be the first time that V_t leaves the region $[-A, +A]$. As usualΦ is the standard normal. First we prove

Lemma 13. $ET \geq \frac{1}{2[\Phi^{-1}(1/8)]^2} A^2$

Proof. Let τ^+ be the first passage time for $A > 0$. We first establish a standard result: $Pr(V_t > A) = Pr(V_t > A \,\&\, \tau^+ < t) = (1/2)Pr(\tau^+ < t)$. The first equality follows from the fact that if $V_t > A$ then certainly $\tau^+ < t$. The second follows from the reflection principle: starting at $V_{\tau^+} = A$ there is an equal probability of $1/2$ that $V_t > A$ and $V_t < A$ hence if $\tau^+ < t$ the probability that $V_t > A$ also is half the probability that $\tau^+ < t$.

Our goal is to establish a lower bound on the expectation of T. Let τ^- be the first passage time of $-A$. First we observe that

$$Pr(\tau^+ < t) = Pr(\tau^+ < t \,\&\, \tau^- > t) + Pr(\tau^+ < t \,\&\, \tau^+ < \tau^- < t) + Pr(\tau^+ < t \,\&\, \tau^- < \tau^+).$$

Using the reflection principle we have

$$Pr(\tau^+ < t \,\&\, \tau^- < \tau^+) = Pr(\tau^- < t \,\&\, \tau^+ < \tau^-) = Pr(\tau^+ < t \,\&\, \tau^+ < \tau^- < t)$$

so that

$$\begin{aligned} Pr(\tau^+ < t) &= Pr(\tau^+ < t \,\&\, \tau^- > t) + 2Pr(\tau^+ < t \,\&\, \tau^+ < \tau^- < t) \\ &\geq Pr(\tau^+ < t \,\&\, \tau^- > t) + Pr(\tau^+ < t \,\&\, \tau^+ < \tau^- < t) \end{aligned}$$

Moreover

$$\begin{aligned} Pr(T < t) &= 2 \quad Pr(\tau^+ < t \,\&\, \tau^- > t) + 2Pr(\tau^+ < t \,\&\, \tau^+ < \tau^- < t) \\ &\leq 2 \quad Pr(\tau^+ < t) = 4Pr(V_t > A) = 4\Phi(-A/\sqrt{t}) \end{aligned}$$

Finally $ET \geq t(1 - Pr(T < t)) \geq t(1 - 4\Phi(-A/\sqrt{t}))$ for all t and in particular for $t = A^2/\left[\Phi^{-1}(1/8)\right]^2$ which gives $ET \geq \frac{1}{2[\Phi^{-1}(1/8)]^2} A^2$. □

Now we consider a random walk with probability β of moving up or down by one and passage time K to $\pm\overline{\theta}L$.

Theorem 17. *The expected hitting time is bounded below by*

$$E\kappa \geq \frac{(\overline{\theta}/(2\beta))^2}{6\left[\Phi^{-1}(1/8)\right]^2} L^2$$

Proof. Let L_k be the random walk and consider the sums $S_L(t) = \sum_{k=1}^{t/L^2} (L_k - L_{k-1})/(2\beta L)$ as $L \to \infty$ converges weakly to a Weiner process with instantaneous variance 1. The random walk passes $\pm\overline{\theta}L$ when $S_L(t)$ passes $\pm\overline{\theta}/(2\beta)$. Consider the $\overline{T}$ truncated hitting time $\tilde{T}$, we have

$$E_S T \geq E_S \tilde{T} \geq E_W T - |E_W T - E_W \tilde{T}| - |E_W \tilde{T} - E_S \tilde{T}|.$$

where the final inequality is just the triangle inequality. However $\lim_{L\to\infty} E_S\tilde{T} = E_W\tilde{T}, \lim_{\overline{T}\to\infty} E_W\tilde{T} = E_W T$. So for all sufficiently large $L, \overline{T}$ we can make $|E_W T - E_W\tilde{T}|, |E_W\tilde{T} - E_S\tilde{T}|$ both less than or equal to 1/3rd the bound in Lemma 13 giving the bound

$$E_S T \geq (1 - \frac{1}{3} - \frac{1}{3})\frac{(\overline{\theta}/(2\beta))^2}{2\,[\Phi^{-1}(1/8)]}.$$

Finally observe that the number of periods corresponding to T is $L^2 T$. □

Web Appendix 3: Length of the Fall, Rise and Warring States

Here we prove

Proposition 3. *[Proposition 2 in the text] For any K there exists an $\overline{L}$ such that for all $L \geq \overline{L}$ there exists a $\overline{\epsilon}$ such that for all $\epsilon \leq \overline{\epsilon}$ the expected length of the warring states period exceeds that of either the fall or rise by K periods.*

Proof. First the fall. From Web Appendix 1 we see that the waypoints are where the hegemon loses a unit of land to opponents that consist entirely of a single society of zealots. Hence there are no more than $\overline{\theta}L$ waypoints. The time to failure is 1 since the hegemon can gain a unit of land with zero resistance and game over, and the least length of a least resistance path from the state after a waypoint to the next waypoint is 2: one transition to replace the society that initially gained the land with the zealots, and one transition for the zealots to take a unit of land from the hegemon. Hence from theorem 14 we have the bound

$$E_\epsilon(t(a)|x, A) \leq \overline{\theta}LD3/C^6.$$

Turning to the rise, fix x such that a would be hegemon j has enough land $\theta_0 L$ to resist an opponent consisting entirely of zealots. Let r_z be that resistance. By Theorem 15 we have the bound $P_\epsilon(A|x) \geq 1 - 2G(t_{xBW})\epsilon^{r_z}$. Moreover the sets B_τ such that the hegemon has $\theta_0 L + \tau$ units of land form a Liapunoff sequence. Notice that for this sequence $t_{FF}(A) = 1$ since there is always zero resistance to the hegemon gaining a single unit of land, and along a least resistance path starting at x he can never lose any land. Hence by Theorem 16 we also have the bound

$$\begin{aligned} E_\epsilon(t(a)|x, A) &\leq (1-\theta_0)L\frac{1}{P_\epsilon(A|x)}\frac{3}{C^2} \\ &\leq (1-\theta_0)L\frac{1}{1-2G(t_{xBW})\epsilon^{r_z}}\frac{3}{C^2} \end{aligned}$$

during the rise.

Recall that at some point during the warring states period there is a society with $L_{j\tau}$ units of land that follows a random walk with β chance of increasing by one or decreasing by one at least until either $L_{j\tau} \geq \overline{\theta}L$ or $L_{j\tau} \leq (1-\overline{\theta})L$. From Theorem 17 we have the expected passage time bound

$$E_\epsilon\kappa \geq \frac{(\overline{\theta}/(2\beta))^2}{6\,[\Phi^{-1}(1/8)]^2}L^2.$$

Hence for L sufficiently large the expected amount of time in the warring states is $3K$ larger than an upper bound $\overline{\theta}LD3/C^6$ on the expected amount of time during a least resistance path during the fall and larger than $(1-\theta_0)L3/C^2$ which is not quite an upper bound on the expected amount of time during the rise. This is not quite the end of the story, since it is the expected amount of time of all paths during the rise or the fall that matters, and because we must account for dividing by the probability of the rise. However, the expected length of all non-least-resistance paths is bounded above by Theorem 13 as is $G(t_{xBW})$ and while that bound increases quite rapidly with L it is also weighted according to Theorem 12 by a probability that goes to zero with ϵ. Hence once we fix L we can choose a small enough ϵ that the expected length of all paths (during the rise or fall) is at most K larger than that of the length of least resistance paths - that is of total length at most $2K$. Hence the expected amount of time in the warring states period is at least K larger than during the rise or fall. □

Web Appendix 4: Ergodic Probabilities and Circuits

We are given a finite set of nodes Ω^k and for $\psi, \phi \in \Omega^k$ a resistance function $r^k(\psi, \phi)$. For any $\psi \in \Omega^k$ we define the *least resistance* $r^k(\psi) = \min_{\phi \in \Omega^k \setminus \psi} r^k(\psi, \phi)$. We are interested in trees T on Ω^k. For any such tree and any ψ let $T(\psi)$ denote the unique predecessor of ψ on the tree (which is null for the unique root). Note that we follow the standard game theory terminology that the predecessor is closer to the root - in contrast to Young who follows the logic of the Markov process in imagining that the node closer to the root is the successor node. The *resistance of the tree T* is defined to be $r^k(T) = \sum_{\psi \in \Omega^k} r^k(\psi, T(\psi))$ where $r^k(\psi, \emptyset) \equiv 0$.

Our goal is to characterize least resistance trees by showing how they are constructed out of groups of nodes that we call circuits. As in the text, $\Omega_x^{k+1} \subseteq \Omega^k$ is a *circuit* if for each pair $\psi_1, \psi_y \in \Omega_x^{k+1}$ there is a path $\psi_1, \psi_2, \ldots, \psi_n \in \Omega_x^{k+1}$ with $\psi_n = \psi_y$ such that for $\tau = 2, 3, \ldots n$ we have $r^k(\psi_{\tau-1}, \psi_\tau) = r^k(\psi_{\tau-1})$, that is, there is a path from ψ_1 to ψ_y within the circuit such that each connection has least resistance.

Definition 10 (Consolidation)**.** A circuit Ω_x^{k+1} is *consolidated within the tree T* if there is a $\phi \in \Omega_x^{k+1}$ that precedes all other $\psi \in \Omega_x^{k+1}$, and for these other $\psi \neq \phi$ we have $T(\psi) \in \Omega_x^{k+1}$ and $r^k(\psi, T(\psi)) = r^k(\psi)$.

In other words, in the consolidated tree the circuit Ω_x^{k+1} forms a subtree with root ϕ, and each connection within the circuit has least resistance. We refer to ϕ as the *top of the circuit.*

Intuitively if we think of the circuit as a circle of least resistance connections then we will break that circle after ϕ to make a subtree and use ϕ to connect this subtree to the the rest of the tree. Breaking the connection saves at least $r^k(\phi)$, while making the new connection costs $r^k(\phi, T(\phi))$, hence we define the *modified resistance from ϕ to ψ* as $R^k(\phi, \psi) = r^k(\phi, \psi) - r^k(\phi)$.

In the next lemma we consolidate a circuit within a tree by breaking it after the node that minimizes modified resistance. By so doing, the resistance of the tree cannot increase.

Lemma 14. *Suppose that T on Ω^k has root ψ and that Ω_x^{k+1} is a circuit on Ω^k. Then there is a tree T' with root ψ such that $r^k(T') \leq r^k(T)$ and Ω_x^{k+1} is consolidated in T' with the additional properties that (1) if $\phi' \notin \Omega_x^{k+1}$ then $T'(\phi') = T(\phi')$ and (2) if ϕ is the top of Ω_x^{k+1} in T' then $R^k(\phi, T'(\phi)) = \min\{R^k(\phi', T'(\phi)) \mid \phi' \in \Omega_x^{k+1}\}$.*

Proof. Let T have root ψ and let $\phi^* \in \Omega_x^{k+1}$ be such that the unique path from ϕ^* to the root ψ contains no element of Ω_x^{k+1}. If $\phi^* = \psi$ take $\phi = \phi^*$. Otherwise choose as top a $\phi \in \Omega_x^{k+1}$ such that $r^k(\phi, T(\phi^*)) - r^k(\phi) = \min\{r^k(\phi', T(\phi^*)) - r^k(\phi') \mid \phi' \in \Omega_x^{k+1}\}$. We now use tree surgery to create a sequence of new trees ending in the desired tree T'. As we proceed we never cut a connection originating in any set other than Ω_x^{k+1} so that property (1) will be satisfied.

At each step Ω_x^{k+1} will be divided into two sets $\Phi_\phi, \Phi_{\sim\phi} = \Omega_x^{k+1} \backslash \Phi_\phi$. The first set Φ_ϕ will contain at least ϕ and consist of those elements of Ω_x^{k+1} that are already consolidated with ϕ at the top, and such that no element of $\Phi_{\sim\phi}$ appears between ϕ and the root. We will proceed constructing new trees by moving one element from $\Phi_{\sim\phi}$ to Φ_ϕ at a time making sure that all properties are preserved.

We start the process. If $\phi = \psi$ or $\phi = \phi^*$ we do nothing. Otherwise cut ϕ from the tree and paste it to $T(\phi^*)$. Observe that this increased the resistance of the tree by at most $r^k(\phi, T(\phi^*)) - r^k(\phi)$. Let Φ_ϕ be the maximal set consolidated with ϕ at the top: this set now contains at least ϕ.

We now continue the process until $\Phi_{\sim\phi}$ is empty. Pick an element $\phi' \in \Phi_{\sim\phi}$. Because Ω_x^{k+1} is a circuit there is a least resistance path in Ω_x^{k+1} from ϕ' to ϕ. Let ϕ_τ be the last element in $\Phi_{\sim\phi}$ that is reached on this path. Then cut ϕ_τ from the tree and paste it to $\phi_{\tau+1}$. Notice that this cannot increase the resistance of the tree since the connection from ϕ_τ to $\phi_{\tau+1}$ has least resistance. Moreover, if $\phi \neq \phi^*$ then at some step $\phi_\tau = \phi^*$ and at this step the resistance of the tree is decreased by exactly $r^k(\phi^*, T(\phi^*)) - r^k(\phi^*)$. Once again let Φ_ϕ be the maximal set consolidated with ϕ at the top: this set now contains at least one more element ϕ_τ.

When we are finished we end up with the new tree T'. Now observe that either $\phi = \phi^*$ or the resistance over the original tree was increased only in the first step, by at most $r^k(\phi, T(\phi^*)) - r^k(\phi)$, and it was decreased by $r^k(\phi*, T(\phi^*)) - r^k(\phi^*)$ when we pasted ϕ^*. By the choice of ϕ we have $r^k(\phi, T(\phi^*)) - r^k(\phi) \leq r^k(\phi^*, T(\phi^*)) - r^k(\phi^*)$, and in all other cases the resistance did not increase. Therefore $r^k(T') \leq r^k(T)$. Since by construction $T'(\phi) = T(\phi^*)$ we have $R^k(\phi, T'(\phi)) = \min\{R^k(\phi', T'(\phi)) \mid \phi' \in \Omega_x^{k+1}\}$. □

We now focus on least resistance trees. Let $\mathscr{T}(\psi)$ be the set of trees with root ψ, $r_\psi^k = \min_{T \in \mathscr{T}(\psi)} r^k(T)$ be the least resistance of any tree with root ψ and $\mathscr{T}_\psi^k = \arg\min_{T \in \mathscr{T}(\psi)} r^k(T)$ be the set of least resistance trees with root ψ. First we prove a simple relation between least resistance of trees and of their roots:

Lemma 15. *If ψ, ϕ are in the same circuit on Ω^k then $r_\psi^k - r_\phi^k = r^k(\phi) - r^k(\psi)$.*

Proof. Suppose $\psi, \phi \in \Omega_x^{k+1}$ where Ω_x^{k+1} is a circuit. Then we can choose a path $\phi_1, \ldots, \phi_\nu, \ldots, \phi_n \in \Omega_x^{k+1}$ with $\phi_1 = \psi, \phi_\nu = \phi, \phi_n = \psi$ such that for $\tau = 2, 3, \ldots n$ we have $r^k(\phi_{\tau-1}, \phi_\tau) = r^k(\phi_{\tau-1})$. Choose $T_1 \in \mathscr{T}_{\phi_1}$, and supposing that $T_{\tau-1}$ has root $\phi_{\tau-1}$ define T_τ as the tree in which we cut ϕ_τ from $T_{\tau-1}$, make it the root of T_τ and paste the root of $T_{\tau-1}$ to ϕ_τ. This tree has root ϕ_τ and resistance $r^k(T_\tau) \leq r^k(T_{\tau-1}) + r^k(\phi_{\tau-1}, \phi_\tau) - r^k(\phi_\tau) = r^k(T_{\tau-1}) + r^k(\phi_{\tau-1}) - r^k(\phi_\tau)$. Hence $r^k(T_\tau) \leq r^k(T_1) + r^k(\phi_1) - r^k(\phi_\tau)$. Since $\phi_n = \phi_1$, we conclude that $r^k(T_n) \leq r^k(T_1)$ and since T_1 had least resistance, it must be that $r^k(T_n) = r^k(T_1)$. Hence all the inequalities must hold with equality, that is, $r^k(T_\tau) = r^k(T_1) + r^k(\phi_1) - r^k(\phi_\tau)$. Choosing $\tau = \nu$ we then have $r^k(T_\tau) = r_\psi^k + r^k(\psi) - r^k(\phi)$, whence $r_\phi^k \leq r_\psi^k + r^k(\psi) - r^k(\phi)$; but by interchanging ϕ and ψ and rearranging we get $r_\phi^k \geq r_\psi^k + r^k(\psi) - r^k(\phi)$; this gives the conclusion. □

We now assume that for $\epsilon > 0$ P_ϵ is ergodic so that there is a unique ergodic probability distribution μ_ϵ on the state space Z. Let $\mathscr{T}_S(x)$ denote all trees over a set S with root x and set

$$\mathscr{M}_\epsilon(x) = \sum_{T \in \mathscr{T}_Z(x)} \prod_{z \in Z} P_\epsilon(T(z)|z).$$

Following Young (1993) and Friedlin and Wentzell (2012)we observe that

$$\mu_\epsilon(x) = \frac{\mathscr{M}_\epsilon(x)}{\sum_{z\in Z}\mathscr{M}_\epsilon(z)}.$$

Let the resistance $r(x, y)$ on Z be the ordinary resistance. Let r_x be the least resistance of trees on Z with root x. Observing from Cayley's formula that N^{N-2} is the number of trees with the same root over N nodes it follows that

Theorem 18. *The ratio of ergodic probabilities satisfies the bounds*

$$\frac{C^N}{N^{N-2}D^N}\epsilon^{r_x - r_y} \leq \frac{\mu_\epsilon(x)}{\mu_\epsilon(y)} \leq \frac{N^{N-2}D^N}{C^N}\epsilon^{r_x - r_y}.$$

Proof. We may rearrange the Friedlin and Wentzell (2012) result to get

$$\mu_\varepsilon(x)\sum_{z\in Z}\mathscr{M}_\varepsilon(z) = \mathscr{M}_\varepsilon(x)$$

so that

$$\frac{\mu_\epsilon(x)}{\mu_\epsilon(y)} = \frac{\mathscr{M}_\epsilon(x)}{\mathscr{M}_\epsilon(y)}$$

Recall the bounds $C\epsilon^{r(x,z)} \leq P_\epsilon(z|x) \leq D\epsilon^{r(x,z)}$ on transition probabilities. Hence we have

$$C^N\varepsilon^{r_x} \leq \sum_{T\in\mathscr{T}_Z(x)} C^N \prod_{x\in Z}\varepsilon^{r(x,z)} \leq \mathscr{M}_\varepsilon(x) \leq \sum_{T\in\mathscr{T}_Z(x)} D^N \prod_{x\in Z}\varepsilon^{r(x,z)} \leq D^N\varepsilon^{r_x}N^{N-2}.$$

Dividing by $\mathscr{M}_\epsilon(y)$ and using the corresponding bounds then gives the result. □

These bounds are in terms of resistances of least resistance trees. The next goal is to translate them in terms of appropriate resistances of least resistance paths.

Applying Lemma 15 give as immediate corollary the following result, where recall that $r^0(\Omega_x)$ is defined in Section 6.3 in terms of direct routes:

Theorem. *[Theorem 9 in text] If the recurrent communicating classes Ω_x and Ω_y are in the same circuit on $\Omega^0 \equiv \Omega$ then*

$$\frac{C^N}{N^{N-2}D^N}\epsilon^{r^0(\Omega_y)-r^0(\Omega_x)} \leq \frac{\mu_\epsilon(x)}{\mu_\epsilon(y)} \leq \frac{N^{N-2}D^N}{C^N}\epsilon^{r^0(\Omega_y)-r^0(\Omega_x)}.$$

This goes one step in the desired direction but applies only to elements of a given circuit. In general, we can find the least resistance of trees in Z by finding the least resistance of trees in Ω. Recall that $r^0_{\Omega_x}$ is the least resistance of trees on Ω with root Ω_x, and r_x is the least resistance of trees on Z with root x. We next show that they are equal:

Lemma 16. *If $x \in \Omega_x \in \Omega$ then $r_x = r^0_{\Omega_x}$.*

Proof. Young (1993) proves this lemma (Lemma 2 in his Appendix) for the case where the resistance, call it $r^*(\Omega_x, \Omega_y)$, is the least resistance of any path from Ω_x to Ω_y - that is, he allows the path to pass through recurrent communicating classes Ω_z which are neither Ω_x nor Ω_z (Ellison

(2000) does the same in his definition of the modified co-radius). Our resistance is in general larger than Young's since we do not allow paths to pass through these other recurrent communicating classes. However, his proof requires only minor modification to yield the stronger result. Young first shows that the least resistance $r^*_{\Omega_x}$ of any tree on Ω with root Ω_x is greater than or equal to r_x. Since $r^0_{\Omega_x} \geq r^*_{\Omega_x}$ we have the immediate implication that $r^0_{\Omega_x} \geq r_x$.

The second part of Young's proof shows that $r^*_{\Omega_x} \leq r_x$. Following Young we show how to transform a least resistance tree $T \in \mathscr{T}_x$ on Z into a tree $T' \in \mathscr{T}(\Omega_x)$ over Ω such that $r^0(T') \leq r^0(T)$. The easiest way to do this would be by simply taking one point from each irreducible class and using the resistance between those states to get a tree over Ω. However, this does not work because there can be double-counting if paths in T join between irreducible classes. Young shows how to avoid double-counting by reorganizing the tree. We can use his construction if we can avoid having or creating paths between irreducible classes that contain elements of a third irreducible class. This is the case if we start by choosing the "right" least resistance tree and the "right" point from each irreducible class before we apply Young's method.

Observe that each $\phi \in \Omega$ is a circuit, so by consolidating where needed as from Lemma 14 we can assume that each $\phi \in \Omega$ is already consolidated in T. The first step of Young's proof is to choose one point $y' \in \phi$ for each $\phi \in \Omega$ - these are what Young calls *special vertices*. We do this by choosing for each $\phi \in \Omega$, the top of ϕ in the tree. Observe that because the tree is consolidated the path from any special vertex to the next special vertex y in the direction of the root cannot contain elements of any irreducible class other than Ω_y.

Now apply Young's construction to eliminate junctions (a *junction* in a tree T is any vertex y with at least two incoming T-edges). Observe that when Young cuts a subtree T^* from a vertex y that is not in a recurrent communicating class this preserves the consolidated structure, because those $\phi' \in \Omega$ that lie further from the root than y are necessarily entirely contained in T^*. Consequently we never need to cut junctions at y that are in recurrent communicating classes, for T is consolidated and therefore the path from y to the top of the circuit has zero resistance and no double-counting is involved.

Finally, when Young pastes cuts T^* from the junction y back into the tree T, he implicitly introduces new paths $a = (y, z_1, \ldots, z_{t-1}, z)$ from y to a special vertex z with $r(a) = 0$. However, these implicit paths cannot contain elements of any recurrent communicating class Ω_y other than Ω_z. If they did the path could not have zero resistance since there is no path from $\Omega_y \neq \Omega_z$ to Ω_z that has zero resistance. Hence at the end of Young's procedure we find that the paths along which resistance is computed - those from one special vertex to the next special vertex in the direction of the root - do not contain a vertex from a third recurrent communicating class. By this procedure we then obtain a tree in $\mathscr{T}(\Omega_x)$ with resistance not larger than T, whence $r^0_{\Omega_x} \leq r_x$. $\square$

Our next goal is to recursively compute r^k and by doing so find bounds on $\mu_\epsilon(x)/\mu_\epsilon(y)$ - without the restriction that Ω_x and Ω_y be in the same circuit.

We take $\Omega^0 = \Omega$, so an element $\psi^1 \in \Omega^1$ will be a circuit of recurrent communicating classes and for $\psi, \phi \in \Omega^0$ the resistance $r^0(\psi, \phi)$ is just the least resistance along a direct route. We recursively define on Ω^{k-1} the modified resistance function $R^{k-1}(\psi^{k-1}, \phi^{k-1}) = r^{k-1}(\psi^{k-1}, \phi^{k-1}) - r^{k-1}(\psi^{k-1})$, and we define a resistance function on Ω^k by the least modified resistance: $r^k(\psi^k, \phi^k) = \min_{\psi^{k-1} \in \psi^k, \phi^{k-1} \in \phi^k} R^{k-1}(\psi^{k-1}, \phi^{k-1})$. Then the following formula holds, where notice that the term $\sum_{\phi^{k-1} \in \Omega^{k-1}} r^{k-1}(\phi^{k-1})$ is a constant independent of the tree in question.

Lemma 17. *If* $\psi^{k-1} \in \psi^k$ *then* $r^{k-1}_{\psi^{k-1}} = r^k_{\psi^k} - r^{k-1}(\psi^{k-1}) + \sum_{\phi^{k-1} \in \Omega^{k-1}} r^{k-1}(\phi^{k-1})$.

Proof. Suppose we have a tree T^{k-1} on Ω^{k-1} that is consolidated with respect to all the circuits in Ω^k, and let ψ^{k-1} be its root. The fact that T^{k-1} is consolidated means that the top of each circuit has a predecessor which belongs to a different circuit. For $\psi^k \in \Omega^k$ denote by $\Gamma(T^{k-1}, \psi^k) \in \Omega^{k-1}$ the top of circuit ψ^k in T^{k-1}. Then if $T^{k-1}(\Gamma(T^{k-1}, \psi^k)) = \phi^{k-1} \in \phi^k \neq \psi^k$ (where if ϕ^{k-1} is null we set $\phi^k = \emptyset$ as well), we may define $T^k(\psi^k) = \phi^k$. In this way we define a tree on Ω^k. We have $r^{k-1}(T^{k-1}) = \sum_{\phi^{k-1}\in\Omega^{k-1}} r^{k-1}(\phi^{k-1}, T^{k-1}(\phi^{k-1}))$. However, since the tree is consolidated, for any ϕ^{k-1} not at the top of the corresponding circuit ϕ^k we have $r^{k-1}(\phi^{k-1}, T^{k-1}(\phi^{k-1})) = r^{k-1}(\phi^{k-1})$, hence we may write

$$r^{k-1}(T^{k-1}) = \sum_{\phi^{k-1}\in\Psi^{k-1}} r^{k-1}(\phi^{k-1}) - r^{k-1}(\psi^{k-1}) + \sum_{\phi^k\in\Psi^k} \rho^{k-1}(\Gamma(T^{k-1},\phi^k), T^{k-1}(\Gamma(T^{k-1},\phi^k))).$$

Now start with a least resistance tree $T^{k-1} \in \mathscr{T}_{\psi^{k-1}}$. By Lemma 14 we may consolidate this tree T^{k-1} with respect to all the circuits in Ω^k to get another least resistance tree $\tilde{T}^{k-1} \in \mathscr{T}_{\psi^{k-1}}$. By the previous computation and the definition of r^k we see that

$$\begin{aligned} r^{k-1}_{\psi^{k-1}} = r^{k-1}(\tilde{T}^{k-1}) &= \sum_{\phi^{k-1}\in\Psi^{k-1}} r^{k-1}(\phi^{k-1}) - r^{k-1}(\psi^{k-1}) + \sum_{\phi^k\in\Psi^k} \rho^{k-1}(\Gamma(T^{k-1},\phi^k), T^{k-1}(\Gamma(T^{k-1},\phi^k))) \\ &\geq \sum_{\phi^{k-1}\in\Psi^{k-1}} r^{k-1}(\phi^{k-1}) - r^{k-1}(\psi^{k-1}) + \sum_{\phi^k\in\Psi^k} r^k(\phi^k, T^k(\phi^k)) \\ &\geq \sum_{\phi^{k-1}\in\Psi^{k-1}} r^{k-1}(\phi^{k-1}) - r^{k-1}(\psi^{k-1}) + r^k_{\psi^k}. \end{aligned}$$

Next start with a least resistance tree $T^k \in \mathscr{T}_{\Omega^k_x}$, where $\psi^{k-1} \in \psi^k$, and construct a tree on Ω^{k-1} as follows. For the root $\phi^k = \psi^k$ define $\phi^{k-1} = \psi^{k-1}$. For given non-root ϕ^k and $T^k(\phi^k)$ there are points $\phi^{k-1} \in \phi^k$ and $\tilde{\phi}^{k-1} \in T^k(\phi^k)$ such that $r^k(\phi^k, T^k(\phi^k)) = r(\phi^{k-1}, \tilde{\phi}^{k-1}) - r(\phi^{k-1})$. For each ϕ^k consolidate the tree over ϕ^k with root ϕ^{k-1} to get a tree $T[\phi^k, \phi^{k-1}]$. Now define a tree on Ω^{k-1} by putting together these subtrees as follows: if $\hat{\phi}^{k-1}$ is in $T[\phi^k, \phi^{k-1}]$ but is not the root, set $T^{k-1}(\hat{\phi}^{k-1}) = T[\phi^k, \phi^{k-1}](\hat{\phi}^{k-1})$. For the root ϕ^{k-1} set $T^{k-1}(\hat{\phi}^{k-1}) = \tilde{\phi}^{k-1}$. This is clearly a tree with root ψ^{k-1}, and we see that the resistance is

$$\begin{aligned} r^{k-1}_{\psi^{k-1}} \leq r^{k-1}(T^{k-1}) &= \sum_{\phi^{k-1}\in\Psi^{k-1}} r^{k-1}(\phi^{k-1}) - r^{k-1}(\psi^{k-1}) + \sum_{\phi^k\in\Omega^k} r^k(\phi^k, T^k(\phi^k)) \\ &= \sum_{\phi^{k-1}\in\Omega^{k-1}} r^{k-1}(\phi^{k-1}) - r^{k-1}(\psi^{k-1}) + r^k_{\psi^k}. \end{aligned}$$

Putting together the two inequalities gives the desired result. □

Lemma 18. *If Ω^k has at least two elements it has at least one non-trivial circuit.*

Proof. Starting at an arbitrary point $\psi^k \in \Omega^k$ choose a path of least resistance. Since Ω^k is finite, this must eventually have a loop, and that loop is necessarily a circuit. □

We can now recursively define a class of reverse filtrations with resistances over the set $\Omega^0 = \Omega$ of recurrent communicating classes for P_0; assume Ω has N_Ω elements, with $N_\Omega \geq 2$. Starting with Ω^{k-1} we observe that there is at least one non-trivial circuit, and that every singleton element is trivially a circuit. Hence we can form a non-trivial partition of Ω^{k-1} into circuits, and denote this partition Ω^k. All the resistances are defined as before. Note that since each partition is non-trivial, this construction has at most $k \leq N_\Omega$ layers before the partition has a single element and the construction stops.

The *modified radius* of $x \in \Omega_x$ of order k is defined by

$$\overline{R}^k(x) = \sum_{\kappa=0}^{k} r^\kappa(\Omega_x^\kappa)$$

where $\Omega_x^0 = \Omega_x$ and for each $\kappa > 0$ the element $\Omega_x^\kappa \ni \Omega_x^{\kappa-1}$. Then

Theorem. *[Theorem 10 in the text] Let k be such that $\Omega_x^k = \Omega_y^k$; then $r_x - r_y = \overline{R}^{k-1}(y) - \overline{R}^{k-1}(x)$ and consequently*

$$\frac{C^N}{N^{N-2}D^N}\epsilon^{\overline{R}^{k-1}(y)-\overline{R}^{k-1}(x)} \leq \frac{\mu_\epsilon(x)}{\mu_\epsilon(y)} \leq \frac{N^{N-2}D^N}{C^N}\epsilon^{\overline{R}^{k-1}(y)-\overline{R}^{k-1}(x)}.$$

Proof. From Lemma 16 we know that $r_x - r_y = r^0_{\psi^0(x)} - r^0_{\psi^0(y)}$. Applying Lemma 17 iteratively, we see that if $\psi^{k-1} \in \psi^k$ then

$$r^0_{\psi^0} = r^k_{\Omega^k_x} + \sum_{\kappa=0}^{k-1}\left[\sum_{\phi^\kappa \in \Omega^\kappa} r^\kappa(\phi^\kappa)\right] - \sum_{\kappa=0}^{k-1} r^\kappa(\psi^\kappa)$$

from which

$$r^0_{\psi^0(x)} - r^0_{\psi^0(y)} = -\sum_{\kappa=0}^{k-1} r^\kappa(\psi^\kappa(x)) + \sum_{\kappa=0}^{k-1} r^\kappa(\psi^\kappa(y)) = \overline{R}^{k-1}(y) - \overline{R}^{k-1}(x).$$

□