

A characterization of stochastic stability and waiting time.*

Kevin Hasker[†]

June 26, 2018

Abstract

In stochastic evolution models, we show that there is an intermediate structure—the *emergent seed*—that simplifies analyses.

Conditional on knowing this graph and the cost function, stochastic potential can be found with path optimization. This makes finding two measures of waiting time—the *coheight*, (the precise waiting time,) and the *censored coradius*, (a natural generalization of the modified coradius)—is immediate. We illustrate the technique in several applications, one of which is novel—the speed of evolution on the three dimensional lattice. Among other results, we derive the first case where the true waiting time (coheight) is strictly lower than the modified (censored) coradius.

Key words: Coradius, Edmond’s Algorithm, Emergent Seed, Matching Games, Minimal Cost Spanning Trees, Radius, Stochastic Evolution.

JEL codes: C63 C73 C78 C79

1 Introduction

That people involved in large repeated matching games do not calculate the law of motion is self-evident. The law of motion is the current state of social behavior and how it will change in the future. Although this information is needed to optimize, finding it is difficult. For example, the rules of dating used to be clear but have been in constant flux in the past fifty years. One doubts that young people construct the large sample across space and time that is necessary for formulating an optimal strategy—and even this would be insufficient. To predict what others will do, they also need to understand others’ information collection and processing techniques, which are usually not observable.

We cannot structurally model this type of interaction; however, we can make some deductions. First, it is sensible and necessary to assume *inertia*. Second, we should observe at a variety of simple *decision rules*. The simplest is imitation (Robson and Vega-Redondo, 1996). A more complex model would have agents take a sample across space and best respond or best respond with mutations (hereafter, BRM; Kandori, Mailath, and Rob, 1993). More complex still would be using this sample to calculate expected payoffs and take actions with probabilities proportional to payoffs, such as the logit model (Blume, 1993).

This final model has a wide basis in the psychological and experimental literature (see Alos-Ferrer and Netzer, 2010). Additionally, it can be fully rational if we follow Harsanyi (1973) by accepting that preferences

*This paper was previously circulated under various titles: "The Emergent Seed: Simplifying the Analysis of Stochastic Evolution," "The Emergent Seed: A Representation Theorem for Models of Stochastic Evolution," and "The Emergent Seed: A Representation Theorem and Two Formulas for Waiting Time for Models of Stochastic Evolution."

[†]Bilkent University, Ankara, Turkey, hasker@bilkent.edu.tr

[‡]I would like to thank Emin Karagozoglu, Maxwell Pak, Sudipta Sarangi, William Sandholm, Francesco Squintani, Tom Watson, Alison Watts, and Peyton Young for their comments. I would like to express special appreciation to David Levine. Any remaining errors are the responsibility of the author.

are heterogeneous (see Myatt and Wallace, 2003). Finally, it means that agents may take actions that (seem) like *errors*. Because the first two models are too parsimonious for this implication, we always assume that people make errors.

We now have models of stochastic evolution: inertia, a decision rule, and rare errors. We are interested in the steady-state or long-run implications of these models. Our analysis relies only on the system being strongly ergodic and with rare errors, and our primitive is *resistance*—the key determinant of the unlikelihood of a direct transition as errors become rare (Young, 1993a).¹

Finding the maximum likelihood (or *stochastically stable*) state(s) in this distribution is simplified by three insights. First, when errors become unlikely, the distribution over a transition is dominated by the one that determines its *cost*—the least resistance method. Freidlin and Wentzell (2012) show that the steady-state likelihood can be analyzed as a static tree minimization problem. Finally, Young (1993a) shows that only *limit sets*—sets with positive short-run likelihood—need to be analyzed. We find a minimal cost-spanning tree on a directed graph for each limit set and the solution is the *stochastic potential*.

This paper recommends an intermediate step—finding the emergent seed. The resulting representation provides two measures of waiting time or the speed of evolution. The *coheight* is an exact measure, and is the height (exit time) for all other limit sets (Beggs, 2005). The *censored coradius* is a generalization of the modified coradius (Ellison, 2000).

A *seed* is a graph over the states such that every state has an exiting transition and some states are transitioned to from all other states. The largest set of states that is transitioned to from all other states is called the *hub* of that seed.² One then constructs trees by including a transition from this hub to the state in question. It is an *optimal seed* if the resulting tree always has the stochastic potential of the state in question. These structures exist and, indeed, in general there will be a class of them; however, we also wish this optimal seed to arise from local analysis or be emergent.

To do this, we will iterate the concept of the limit set. In abstract, one finds limit sets by constructing a graph of the zero resistance transitions. The probability of these transitions increases as errors become rarer. Of this graph, the states, cycles, and circuits³ that have no exiting transitions are the limit sets. A technical but important point is that a resistance of zero is only a normalization; mathematically, we could normalize the resistance by a state-dependent constant and the analysis would be unchanged except for replacing "zero" with "minimal."

If only one limit set exists, then the analysis terminates; however, there are frequently more limit sets. At this point, one usually simply calculates stochastic potentials. We recommend iterating the limit set methodology. Bortz, Kalos, and Lebowitz (1975) point out that these limit sets must transition to each other and that the most likely (least resistance) one will become infinitely likely as errors become rare. This suggests that we should renormalize the resistance for states in the limit sets. An *exiting transition* is one that is not likely to return to a limit set in the short run after it occurs once. Our new resistance will reflect the most likely exiting transition having zero resistance. The most likely exiting transition(s) determine the *radius* (Ellison, 2000). Thus, we define a *first difference resistance* as the cost of an exiting transition minus the radius of the limit set. This is the same as the modified cost (Ellison, 2000). This new resistance will

¹The simplest model is to count the number of errors needed for the transition (BRM). In a more complex model, resistance might be affected by the loss from using a strategy (Logit).

²We would like to thank Tom Watson for his assistance with this terminology. Prior versions referred to it as the core.

³These are sets of intersecting cycles. This terminology is from Levine and Modica (2014b).

have new (and fewer) limit sets, which we call *first iteration limit sets*. If there is only one, then we stop; otherwise, we iterate this procedure until there is only one. The graph we generate to find this limit set is the *emergent seed* and the *hub* is the unique limit set in our final iteration.

The benefit of this approach is that we replace the global restriction of transitioning to a state with the local restriction of exiting a limit set. For example, using this approach, we will find the speed of evolution on a three-dimensional lattice. Ellison (1993) found the speed of evolution on a one-dimensional lattice, and Ellison (2000) extends this to two dimensions, but three or higher is still an open question. This literature considers a classic coordination game using BRM. Because matching must be done on the lattice, there can be geometric areas of agents using the risk-dominant strategy and others using the other strategy. For example, in two dimensions, a square of agents using the risk-dominant strategy is stable. We solve the problem for three dimensions by using only local analysis. All we need to know is to where a limit set is most likely to transition. It is fairly obvious that it will either go "up" (to a limit set in which more agents play the risk-dominant action) or "down" (in which fewer do). In three dimensions, this results in a tipping point, above which one transitions to the first iteration limit set containing the state in which everyone plays the risk-dominant action, and below which one transitions to a first iteration limit set containing the state in which no one does. Because one knows what will happen at the next iteration when there are only two first iteration limit sets, further analysis is not necessary. Indeed, the speed of convergence is given by the modified coradius. For a detailed analysis, the reader should turn to Section 6.3. Note that at no point did we actually use any new technology to solve this problem. Rather, the emergent seed gave our analysis a new focus. Hasker (2018) is working on the general problem but had not finalized the waiting time estimates as of this writing.

This result highlights an aspect of the emergent seed. Although it *may* be enlightening, it is useful. This point has been proven by the literature, in which the vast majority of applications implicitly use this technique, such as the first application: the Nash Demand game (Young, 1993b). The article finds the most likely transition from each limit set (the radii) and shows that a graph of these transitions has one cycle (the hub) and, finally, that something in this cycle has the highest radius (*hub dominance*). Binmore, Samuelson, and Young (2003) propose this as a test (the *naive minimization test*) but also have no recommendations if it fails. We recommend continuing to analyze the emergent seed. In the Contract game (Young, 1998), we find a closed form objective (Section 6.1).

The other common solution methods are either sufficient (radius/(modified) coradius) or guess and verify (root switching). *Root switching* is our terminology for the standard method. In this argument, one hypothesizes that a given state is stochastically stable and then either verifies or contradicts this by switching the root of its minimal cost tree. Young (1993a) uses this technique, which is very powerful. For example, Binmore, Samuelson, and Young (2003) use it to prove the radius/coradius theorem (Ellison, 2000). In the Online Appendix (Section C), we use it to prove the radius/modified coradius theorem. We expect that this technique could construct the emergent seed.

The radius/(modified) coradius technique (Ellison, 2000) is a sufficient methodology. The radius being higher than the (modified) coradius is a sufficient condition for stochastic stability. This is at its best when little is known about the model. For example, Bergin and Bernhardt (2009) use it to prove a general result. In most other cases, we have found that either the analysis identifies the emergent seed or the distance between sufficiency and necessity is significant. One example in the latter class is the Gift Giving game; see

Section 6.2.

We know of three other solution techniques in the literature. Beggs (2005) proposes an iterative height algorithm. Height is the expected exit time for a set of states, and the algorithm suggests that we discard states with a low height in each step. Rozen (2008) transforms the primal problem into a dual problem. Cui and Zhai (2010) propose a cyclic decomposition methodology that finds the most likely cycles in the process of evolution and continues iteration until all limit sets are linked into one grand cycle. Unfortunately, these methodologies have yet to be used in applications.

The emergent seed is not always the best methodology. A better one is to directly characterize the limiting distribution, which can be done in the logit model when the game has a potential (Monderer and Shapley, 1986), and sometimes in other models. Examples are Fudenberg and Imhof (2006), Sandholm (2007), and Kandori, Serrano, and Volij (2008). However, the emergent seed may be useful to find the speed of evolution; see Section 5. Furthermore, robustness arguments often will not benefit from the emergent seed. Kandori and Rob (1998, half dominance), Ellison (2000, half dominance), Peski (2010), and Sandholm (2010) find local characteristics of an equilibrium that imply that it will be stochastically stable. The emergent seed is a global characteristic. As well, the results of the speed of evolution in Montanari and Saberi (2010), Young (2011), and Kreindler and Young (2012) are robust to the graph and, thus, the emergent seed. Robust stochastic stability might benefit (Alos-Ferrer and Netzer, 2012), which generalizes the radius/coradius test, and the total radius/censored coradius might be used in these arguments.

The reader familiar with the literature on the minimal cost spanning tree will recognize that our algorithm is a modification of Edmonds' Algorithm (Edmonds, 1967—first published by Chiu and Liu, 1965). This is not surprising because it is a unique algorithm in mathematics. Rozen (2008) is the first paper in the stochastic evolution literature to explore the link. Earlier papers—Noeldke and Samuelson (1993), Samuelson (1994), and Kandori and Rob (1995)—use one or two iterations of the algorithm, and Troger (2002) uses it in an application. We are not the first to use it to analyze the global problem—this was done by Humblet (1983). What is novel is our modification. At each step, we drop states that are outside of cycles, resulting in fewer iterations and, thus, greater analytic efficiency.

Given that Edmonds' algorithm may be the unique optimum, that it has been rediscovered, as noted by Rozen (2008), is not surprising. It is well known that Bock (1971) rediscovered it in computer science. Our methodology benefited from a rediscovery of Monte Carlo simulations—Bortz, Kalos, and Lebowitz (1975). Freidlin and Wentzell (2012) rediscovered it in the theory of large deviations. Cui and Zhai (2010) rediscovered it in stochastic evolution. The last two are similar to Humblet (1983) because they do not specify a root.

In the Edmonds' algorithm, the initial step is to have each state point at the state(s) in which it has the least resistance for transitioning to (which might be itself). One then finds the limit sets or the (possibly degenerate) cycles in this graph and treats them as states. For each of these (sets of) states, one then finds the least resistance transition outside of that set and iterates. The modified algorithm in Cui and Zhai (2010) is interesting because it first drops states that are not in limit sets (the cycles from the initial step), and then follows the Edmonds' algorithm. The difference between the limit set algorithm and Edmonds' algorithm is that all states not in a (possibly degenerate) cycle are listed as being in the *outer basin of attraction* of one or more of these limit sets. Each of these states has a least resistance sequence of transitions that leads to a limit set(s). We note that this modification introduces an asymmetry in their approach. First, they use

the limit set algorithm, and then they switch to a different one. Instead, we iterate the limit set algorithm. The first iteration of both techniques will result in the same graph. In the next stage, we call the cycles in this graph first iteration limit sets, and all other states are in the first iteration outer basin(s) of attraction of one (or more) of them. We then continue to only analyze the new limit sets, indicating that we have fewer objects to analyze at each iteration, and finishing our algorithm requires fewer iterations. The cost is that we cannot solve the problem using only this methodology. In the end, we must calculate the cost of a path from the hub to a limit set to find its stochastic potential. In contrast, Cui and Zhai (2010) derive a completely novel characterization theorem, but their methodology requires more iterations. For more details on this topic, please turn to Section 7.

We know of two alternative algorithms. Beggs (2005) shows that one can iteratively discard sets of states with a low height (expected exit time). Trygubenko and Wales (2006) find an algorithm that does not require iteration in the field of Monte Carlo simulations. That paper analyzes waiting time for a given root; however, similar to Bortz, Kalos, and Lebowitz (1975), it may be generalizable.

Kandori and Rob (1995) introduce *optimized cost*, which we simply call cost. The other key tools in our analysis—the basin of attraction, radius, and modified cost—were all introduced in Ellison (2000). The emergent seed is constructed by iterating these concepts. Rozen (2008) is the closest paper to ours. It uses Edmonds’ algorithm to transform the primary problem into a dual problem and mentions that the algorithm only needs to be implemented once, and derives a restrictive version of local hub dominance and an alternative coradius measure.

We next turn to describing the general model in Section 2. We then describe how to find the emergent seed in Section 3; Section 4 presents the characterization; and Section 5 presents two measures of waiting time. In Section 6, we turn to a survey of the applied literature, including one new application and a re-analysis of two others. Next, we turn to a discussion of other methodologies, including an analysis in which our methodology fails in Section 7. Finally, we conclude in Section 8. Proofs of propositions and theorems are in Section 9, and all other proofs and some supplementary materials are in the Online Appendix.

2 The Model

Our notation follows Ellison (2000) with changes where necessary. The fundamental of our model is a finite set of states of the world, which we denote as Z . These states will often be social, or the strategies of all agents. For example, if we have uniform random matching, then it can be a distribution over the strategies. We endow the states of the world with a Markov transition matrix, P_β , which must be (strongly) ergodic.⁴ This restriction is satisfied because agents make errors. The goal of our analysis is to identify the steady state distribution, μ_β :

$$\mu_\beta = \mu_\beta P_\beta . \quad (1)$$

This will also be the long-run distribution but is of interest because it is self contained.⁵

For arbitrary P_β , μ_β might be very dispersed and uninformative; however, in the problems in which we are interested, as $1/\beta$ becomes small, μ_β becomes concentrated. To see this, we decompose $P_\beta(x, y)$ into two

⁴Ellison (2000) denoted the states by Z and denoted all matrices by a Latin letter. We reserve X and Y for subsets of Z , and x , y , and z for states. Any other Latin letter is a matrix or, equivalently, a function with domain $Z \times Z$ and sometimes one other parameter (β). Functions may be denoted with calligraphic font for clarity or because their domain is not Z or $Z \times Z$.

⁵Note that μ_β is written as a row matrix, thus $P_\beta(x, y)$ is the probability of transitioning from x to y .

functions. The first is a *resistance function* (Young, 1993a) $r : Z \times Z \rightarrow \mathbb{R} \cup \infty$. The second is a *weighting function* $W : Z \times Z \times (0, \infty) \rightarrow [\underline{w}, \bar{w}]$ where $0 < \underline{w} < \bar{w} < \infty$. Then, for $\beta > 0$, we write:

$$P_\beta(x, y) = \frac{W(x, y, \beta) e^{-\beta r(x, y)}}{\sum_{z \in Z} W(x, z, \beta) e^{-\beta r(x, z)}}. \quad (2)$$

Note that $r(x, y) > r(x, \tilde{y})$, then the relative likelihood of transiting to y rather than \tilde{y} is $P_\beta(x, y) / P_\beta(x, \tilde{y}) = e^{-\beta[r(x, y) - r(x, \tilde{y})]} W(x, y, \beta) / W(x, \tilde{y}, \beta)$. Thus, as β gets large, $P_\beta(x, y) / P_\beta(x, \tilde{y}) \rightarrow 0$. This results in μ_β being more concentrated. The weighting function is usually not important. We normalize the resistance such that it is non-negative and, for all x , there is a y such that $r(x, y) = 0$.⁶ We illustrate resistances using the classic coordination game with $\sigma \in [0, 1]$:

	A	B	
A	1 - σ , 1 - σ	0, 0	
B	0, 0	σ , σ	

(3)

Assume one population and uniform random matching. Let $z = n_z/n$ where n_z is the number of agents playing A and n is the population size. Note that A is the best response if $z \geq \sigma$.

Then, resistance in the best response with mutations (hereafter BRM) is:

$$r(x, y) = \begin{cases} 0 & \text{if } y \geq x \geq \sigma \text{ or } y \leq x \leq \sigma \\ n|x - y| & \text{else} \end{cases}. \quad (4)$$

If $y \geq x \geq \sigma$ (or $y \leq x \leq \sigma$), then agents are switching to A (respectively, B) and it is a best response. Thus, the resistance is zero. Otherwise, the resistance is the number of agents who must error.

In the logit model, only one agent at a time can change strategy and the state affects the resistance. The resistance is:

$$r(x, y) = \begin{cases} \infty & \text{if } n|x - y| > 1 \\ 0 & \text{else if } y \geq x \geq \sigma \text{ or } y \leq x \leq \sigma \\ |x - \sigma| & \text{else} \end{cases}. \quad (5)$$

Remember that $x = n_x/n$ and $y = n_y/n$; thus, if $n|x - y| > 1$, then more than one agent must have changed strategy. Thus, resistance is infinite and $P_\beta(x, y) = 0$. Resistance is zero in the same cases as BRM. Otherwise, if x is close to σ , then both strategies have nearly the same expected utility, and the resistance is small.

The definition of a limit set is:

Definition 1 (Limit Set) A limit set is a minimal set $\theta \subseteq Z$ such that $\forall s \in \mathbb{N} \lim_{\beta \rightarrow \infty} \Pr(z_{t+s} \notin \theta | z_t \in \theta) \rightarrow 0$.

We denote the family of these limit sets as Θ . Note that because the system is ergodic for all β , $\lim_{s \rightarrow \infty} \Pr(z_{t+s} \notin \theta | z_t \in \theta) > \varepsilon > 0$. The terminology references the noiseless system. See the Online Appendix (Section A) for a discussion of this and other terminology topics.

⁶For examples in which the resistance is naturally strictly positive or sometimes negative, consider transportation problems.

Holding a good in a warehouse has a positive cost. Thus, in some analyses, every action will be costly since one pays for either shipment or storage.

If resistance is the amount of energy used to transition, then a state on a hill would have a negative resistance to nearby lower states. The transition would create energy.

Because the emergent seed is simply found by iterating the concept of a limit set, we present two characterizations of limit sets. The first one is based on graphs, and we use this technique in practice when we iterate the analysis. In the literature, a graph is either a list of ordered pairs or a matrix. We choose the matrix representation because matrix operations are well defined. Thus, a graph is a matrix G that has dimensions $\#(Z) \times \#(Z)$, and if we transition from x to y , then $G(x, y) = 1$, and $G(x, y) = 0$ otherwise.⁷ We remind the reader that Z is the set of states, and X and Y are subsets. All other Latin letters in this analysis are matrices. For example, r is a function with domain $Z \times Z$ and, thus, also a matrix with dimension $\#(Z) \times \#(Z)$. We have already introduced the Markov transition matrix as P_β . The resistance of a graph is:

$$\begin{aligned} r(G) &= \sum_{z \in Z} \sum_{\hat{z} \in Z} r(\hat{z}, z) G(\hat{z}, z) \\ &= \text{vec}(r)' \text{vec}(G) \end{aligned} \quad (6)$$

For a $m \times n$ matrix G , $\text{vec}(G)$ is the $mn \times 1$ matrix achieved by stacking the columns of G on top of each other, and G' is the transpose of G . To complete the transformation, we describe $X \subseteq Z$ as a row matrix, where $x \in X$ means that the x 'th entry is one, and zero if $x \notin X$. Then, we can work neatly with Boolean algebra, the notation is for $\{x, y\} \in \{1, 2, 3, \dots, \#(Z)\}^2$ $[G \cup \tilde{G}](x, y) = \max\{G(x, y), \tilde{G}(x, y)\}$, $G \subseteq \tilde{G}$ means $G(x, y) \leq \tilde{G}(x, y)$, $[G \setminus \tilde{G}](x, y) + \tilde{G}(x, y) \leq 1$, and $[G \cap \tilde{G}](x, y) = \min\{G(x, y), \tilde{G}(x, y)\}$.

A limit set can be characterized as a set of zero resistance cycles. A graph that has zero resistance has an underbar; thus, \underline{G} has $r(\underline{G}) = 0$. A cycle is a path that begins and ends at the same state. The standard notation is Q . It is a *path* from x to y if there is a sequence $(z_s)_{s=1}^S$ with $z_1 = x$, $z_S = y$ and $\prod_{s=1}^{S-1} Q(z_s, z_{s+1}) = 1$, and for $z \in Z$, $\hat{z} \in Z \setminus (z_s)_{s=1}^S$ $Q(z, \hat{z}) = 0$. Let the set of these paths be $\{Q(x, y)\}$. Then, a *cycle* for x is a path that begins and ends at x : $Q \in \{Q(x, x)\}$.⁸

The most elegant characterization uses the optimized resistance or *cost* function. Because we focus on the steady state, we are interested in $\Pr(z_{t+s} = y | z_t = x)$ for $s \leq \#(Z)$ instead of $\Pr(z_{t+1} = y | z_t = x)$. For fixed β , this probability is a distribution over the paths from x to y . However, as $\beta \rightarrow \infty$, this distribution will be dominated by the most likely (least resistance) path(s):

$$c(x, y) = \min_{Q \in \{Q(x, y)\}} \text{vec}(r)' \text{vec}(Q) . \quad (7)$$

Note that $c(x, y) < \infty$ by strong ergodicity and that for subsets, $X \subseteq Z$ and $Y \subseteq Z$, the transition will be dominated by $c(X, Y) = \min_{x \in X, y \in Y} c(x, y)$. It is fairly immediate that:

Lemma 1 *A limit set $\theta \subseteq Z$ can be characterized as either:*

1. *A (degenerate) cycle or set of intersecting cycles:*

- (a) $\forall x \in \theta$, $\theta = \bigcup_{Q \in \{\underline{Q}(x, x)\}} \left[\bigcup_{z \in Z} Q(z, \cdot) \right]$ and
- (b) $\forall z \in Z \setminus \theta$, $\{\underline{Q}(x, z)\} = \emptyset$.⁹

⁷For a set X , $\#(X)$ is the number of elements in the set.

⁸These definitions are minimal. The finiteness of the path is guaranteed by our graph being a finite matrix. They allow for states to appear multiple times in the sequence. An implication is that all circuits are cycles.

⁹ $Q(z, \cdot)$ is the row vector associated with the element z .

2. Using the cost function for all $x \in \theta$:

- (a) for all $y \in \theta$ $c(x, y) = 0$
- (b) for all $z \in Z \setminus \theta$ $c(x, z) > 0$.

Because limit sets are important, we will show the reader how to find them both in a familiar application—the Nash Demand Game (Young, 1993b)—and for an arbitrary resistance.

Example 1 The Limit Sets in the Nash Demand game: We will use BRM with two population uniform matchings as our underlying dynamics. There will be n agents in each role that will be uniformly matched into pairs. With probability $\rho \in (0, 1)$, agents use the strategy they used during the last period. With probability $1 - \rho$, they will choose a new strategy. If they choose a new strategy, it will be a best response to the current distribution of strategies with probability $1 - e^{-\beta}$, and it will be a strategy chosen at random with probability $e^{-\beta}$.

The bargaining problem is a pair of concave and strictly increasing utility functions, $u_i(x)$, $A_1 = A_2 = [0, 1]$, and a pair (a_1, a_2) is feasible if $a_1 + a_2 \leq 1$. For a feasible pair, i gets $u_i(a_i)$; otherwise, i gets zero. We have normalized the disagreement point to zero and assume that there is an open set of feasible (a_1, a_2) such that $\min_{i \in \{1, 2\}} u_i(a_i) > 0$. We need a finite number of strategies; thus, for $\delta > 0$ such that $1/\delta$ is an integer, let $A_i(\delta) = \{0, \delta, 2\delta, \dots, 1\}$. In the Nash Demand game, the strategy sets (S_i) are simply A_i , or each role submits demand s_i . If (s_1, s_2) is feasible, then i gets $u_i(s_i)$; otherwise, they receive zero.

It is easy to prove that every strict Nash equilibrium is a limit set. Because both parties are getting a strictly positive payoff, the unique best response at a strict Nash equilibrium is that equilibrium. The more difficult step is to show that there are no others. There are two cases to consider. The first case is that the best one party can get is zero in the current state. In this case, we can have all agents in that population change strategy this period with positive probability. With positive probability, they will all choose the same new strategy, which can be their part of one strict Nash equilibrium. In the next period with positive probability, everyone in the other role will choose a new strategy and the unique best response, and we are done. The second case is that, in the current state, there are at most two optimal strategies. Again, with positive probability, we choose everyone in that population to change strategy, and they all choose the same best response. At this point, there are two sub-cases. The first sub-case is that what they are now doing is part of a strict Nash equilibrium and, similar to before, in the next period everyone in the other population will choose the unique best response. The second sub-case is that the best that the other role can do is zero, a case that we already analyzed. Thus, we showed that there is always a zero cost path from every other state to a strict Nash equilibrium and these can be the only limit sets.

The problem with this example and, indeed, most examples in the literature, is that the set of strict (pure strategy) Nash equilibria are the limit sets. An arbitrary example allows us to illustrate other possibilities. In this arbitrary case, it is best to characterize limit sets as cycles in the graph of zero resistance transitions. This graph is:

This resistance is sparse. In most analyses, points will be similar to x_a or x_b with multiple zero resistance paths. This resistance allows us to discuss many important types of limit sets. The first type, θ_a , is similar to a strict pure strategy equilibrium in BRM or Logit. From all nearby states, one moves toward it. The

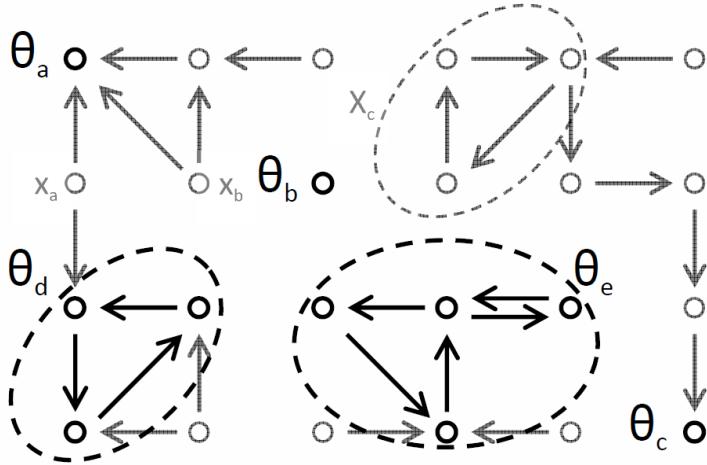


Figure 1: $x \rightarrow y$ means $r(x, y) = 0$, states are circles, limit sets are bolded circles (if states) and bolded dashed ovals (if sets). Assume $r(x, x) = 0$.

second type, θ_b , is an unstable mixed strategy equilibrium in BRM if agents only change strategy when it strictly improves their payoff. It is also the default in a transportation problem. The third type, θ_c , probably only occurs in a transportation problem. It is a "river mouth"—many states are connected to it but few of them are nearby. The fourth type, θ_d , is a simple cycle. The fifth type, θ_e , is a proper *circuit* (Levine and Modica, 2015) or a set of intersecting cycles. Note that the set X_c is a limit set except for the second part of the characterization. There is a zero resistance path from X_c to θ_c .

A key step in our analysis of the Nash Demand game was to show that a zero cost path exists from every other state to a limit set. Formally, we showed that they were all in the outer basin of attraction of a limit set.

Definition 2 *The outer basin of attraction of $X \subseteq Z$ is $\bar{\mathcal{D}}(X) = \{z \in Z | c(z, X) = 0\}$.*

This is not the more common concept of the *basin of attraction* (Ellison, 2000). In that definition, states must reach θ with a probability that converges to one, or $\mathcal{D}(X) = \{x \in \bar{\mathcal{D}}(X) : \forall z \in Z \setminus \bar{\mathcal{D}}(X) \ c(x, z) > 0\}$. In Figure 1, x_a is in no basin of attraction but is in the outer basin of θ_a and θ_d , whereas x_b is in the basin of θ_a . In the Nash Demand game, we showed that every state is in the outer basin of attraction of a limit set, and the basins of attraction are much smaller. These are states from which the unique best response for both populations is the strict Nash equilibrium. Many states are in no basin of attraction, such as all states in which the best that one party can get is zero. This illustrates that the basins of attraction ($\cup_{\theta \in \Theta} \mathcal{D}(\theta)$) may be a strict subset of Z , whereas with the outer basin, $\bar{\mathcal{D}}(\theta) \cap \bar{\mathcal{D}}(\tilde{\theta}) \neq \emptyset$ is possible for $\theta \neq \tilde{\theta}$.

When characterizing the most likely transition(s) out of X , the cost is the same using either concept; however, the path(s) will only be complete with the outer basin. This is determined by the *radius* (Ellison,

2000):¹⁰

$$\mathcal{R}(X) = \begin{cases} \min_{z \in Z \setminus \bar{\mathcal{D}}(X)} c(X, z) & \bar{\mathcal{D}}(X) \subset Z \\ \infty & \bar{\mathcal{D}}(X) = Z \end{cases}, \quad (8)$$

Because the radii of limit sets are very important for our analysis, we characterize them in the Nash Demand game.

Example 2 The Radii in the Nash Demand game: First, we mention that when one analyzes BRM with uniform matching, a cost is generally characterized as $c(x, y) = \sum_{s=1}^S \lceil np_s \rceil$ where $\lceil x \rceil$ is the least greater integer than x and $p_s \in [0, 1]$. Because the size of the population is not important, it is convenient to normalize this by n and write $c(x, y) / n = \sum_{s=1}^S p_s$.

To determine the radii, we need to find the "best invaders," in other words, the agents who most quickly make one part of the equilibrium strategy not a best response. In the Nash Demand game, an invading population can demand either more or less. Let a limit set θ be $\theta = (\theta_1, \theta_2) = (\theta_1, 1 - \theta_1)$. First consider demanding more and let $p_1^+(\theta, k)$ be the mass of agents demanding $1 - \theta_1 + k\delta$ to make $\theta_1 - k\delta$ as good of a response as θ_1 for role one, then $p_1^+(\theta, k)$ is:

$$(1 - p_1^+(\theta, k)) u_1(\theta_1) = u_1(\theta_1 - k\delta) \quad (9)$$

because if the players in population one reduce their demand they get it from everyone. Therefore:

$$p_1^+(\theta, k) = (u_1(\theta_1) - u_1(\theta_1 - k\delta)) / u_1(\theta_1), \quad (10)$$

and it is clear the best candidate in this class has $k = 1$, they should demand only a little more. In contrast, if they demand less, the asymmetry is reversed. If a player keeps her current demand, she will get it from everyone. Thus if invaders demand $s_2, 1 - s_2$ is a best response if:

$$\begin{aligned} u_1(\theta_1) &= p_1^-(\theta, s_2) u_1(1 - s_2) \\ p_1^-(\theta, s_2) &= u_1(\theta_1) / u_1(1 - s_2) \end{aligned} \quad (11)$$

and thus the optimal demand is $s_2 = 0$. Let $p_i^+(\theta) = p_i^+(\theta, 1)$, $p_i^-(\theta) = p_i^-(\theta, 0)$ for $i \in \{1, 2\}$. Then for all θ , $\mathcal{R}(\theta) / n = \min [p_1^+(\theta), p_2^+(\theta), p_1^-(\theta), p_2^-(\theta)]$. Note that as $\delta \rightarrow 0$ $\max [p_1^+(\theta), p_2^+(\theta)] \rightarrow 0$ while $\min [p_1^-(\theta), p_2^-(\theta)]$ is large and constant. This allows us to conclude that for small enough δ , $\mathcal{R}(\theta) / n = \min [p_1^+(\theta), p_2^+(\theta)]$.

Lemma 2 If δ is small enough and $p_1^+(\theta) < p_2^+(\theta)$ or

$$(u_1(\theta_1) - u_1(\theta_1 - \delta)) u_2(1 - \theta_1) - (u_2(1 - \theta_1) - u_2(1 - \theta_1 - \delta)) u_1(\theta_1) < 0 \quad (12)$$

then $\mathcal{R}(\theta) / n = (u_1(\theta_1) - u_1(\theta_1 - \delta)) / u_1(\theta_1)$ or the cost of transitioning from $(\theta_1, 1 - \theta_1)$ to $(\theta_1 - \delta, 1 - \theta_1 + \delta)$.

Condition 12 is the derivative of the Nash Bargaining objective $(u_1(s_1) u_2(1 - s_1))$ in difference form.

¹⁰The radius and the basin of attraction both predate Ellison (2000); however, this is the most familiar paper to introduce them in economics.

The standard representation theorem requires that we estimate the probability of getting to x from every $y \in Z \setminus x$. We estimate this using *trees with root $X \subseteq Z$* . In a tree, once we get to X , we do not move, and from every $y \in Z \setminus X$, there is a path to X . Mathematically, $\sum_{x \in X} \sum_{y \in Z} T(x, y) = 0, \forall y \in Z \setminus X \exists (z_s)_{s=1}^S$ with $z_1 = y, z_S \in X$ and $\prod_{s=1}^{S-1} T(z_s, z_{s+1}) = 1$. Let the set of these graphs be $\{T(X)\}$. Then, the *stochastic potential* of $X \subseteq Z$ is:

$$sp(X) = \min_{T \in \{T(X)\}} \text{vec}(r)' \text{vec}(T) = \min_{T \in \{T(X)\}} r(T); \quad (13)$$

and x is *stochastically stable* if $x \in \arg \min_{z \in Z} sp(z)$. This implies that $\lim_{\beta \rightarrow \infty} \mu_\beta(x) > 0$.

A comment on notation: in the sequel, we will derive a sequence of resistances, $\Delta^m r$ for $m \in \{1, 2, 3, \dots, M\}$, and we denote the limit sets with regard to these resistances as θ^m and the set as Θ^m . A subset of states may be X, Y , or θ^m (where $\theta^0 = \theta$). A set of sets of states will be a set of limit sets, denoted as Θ^m ($\Theta^0 = \Theta$).

3 The Emergent Seed

It is easiest to understand the emergent seed using a dynamic argument. After one has found the limit sets, are there more simple insights we can generate about the dynamics? Remember that the cost function is derived by arguing that all other paths have essentially zero probability when $1/\beta$ is small. Likewise, if one transition has a higher cost than another, it has essentially zero probability. Thus, a limit set will transition to a limit set that determines its radius with probability one—compared with all other transitions. This suggests that we look at a graph of the radii. Let us consider this in our abstract example.

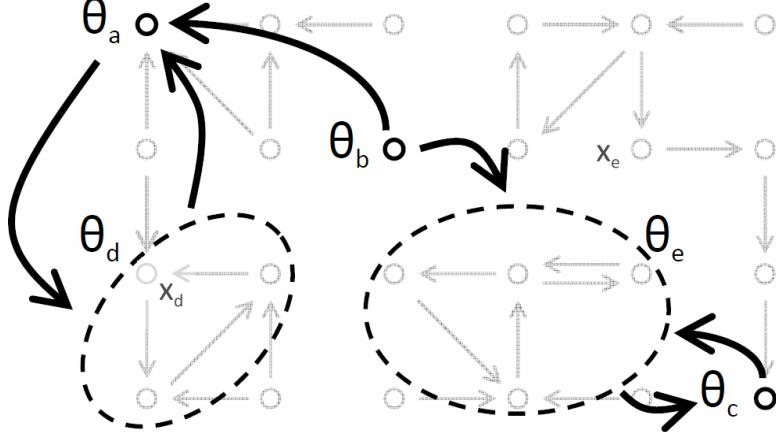


Figure 2: $\theta \rightsquigarrow \tilde{\theta}$ if $c(\theta, \tilde{\theta}) = \mathcal{R}(\theta)$

The radii are arbitrary and the similarity between Figure 2 and Figure 1 is obvious. Now, we have cycles or circuits in radii and it seems clear that these should be limit sets regarding some resistance function. Which resistance function? Recognizing that the convention $\forall x \in Z \exists y \in Z r(x, y) = 0$ is nothing more than a normalization, perhaps we should normalize the cost of exiting a limit set to zero, i.e., renormalize the resistance for exiting transitions.

Definition 3 An exiting transition is one that starts at a state in a limit set and leaves its outer basin of attraction.

We recognize that this method might not tell us everything that we can learn about dynamics. Given this new resistance, we have two first iteration limit sets: $\theta_a^1 \supseteq \{\theta_a, \theta_d\}$ and $\theta_b^1 \supseteq \{\theta_c, \theta_e\}$. However, this is no problem. We have a resistance and limit sets, and we simply need to iterate this analysis. We continue until there is only one M 'th iteration limit set. The formal algorithm is:

Algorithm 1 The emergent seed is found by:

0. For each state normalize the least resistance to zero.

For $m \geq 1$:

- m. If there is more than one $m-1$ iteration limit set, let the m iteration resistance be the $m-1$ iteration cost, except that for states in limit sets and exiting transitions normalize the least cost to zero.

It should be clear that all standard steps in the analysis must be completed at each iteration. We must find the m 'th iteration limit sets, their basins of attraction, and the cost function. Our next task is to specify the m iteration resistance.

Explicitly E^0 is the graph of all zero cost paths. Then:

Definition 4 The first iteration of the emergent seed is E^1 and includes all paths from θ to all $\tilde{\theta} \in \arg \mathcal{R}(\theta) = \arg \min_{\tilde{\theta} \in \Theta \setminus \theta} c(\theta, \tilde{\theta})$.

We now renormalize the resistance of limit sets by the least cost exiting transition. By definition the cost of this path must be $\mathcal{R}(\theta)$, thus our new resistance is:

$$\Delta r(x, y) = \begin{cases} c(x, y) - \mathcal{R}(x) & \text{if } x \in \theta, y \in Z \setminus \bar{\mathcal{D}}(\theta) \\ c(x, y) & \text{else} \end{cases} . \quad (14)$$

We refer to this as the *first difference resistance*. It is the same as that found in Ellison (2000) to understand the dynamics of evolution. We define the *first difference cost* ($\Delta c(\cdot, \cdot)$) like we did $c(\cdot, \cdot)$, the *first iteration limit sets* are denoted θ^1 with the set Θ^1 , the *first difference outer basins of attraction* are denoted $\Delta \bar{\mathcal{D}}(\theta^1)$, and the *first difference radius* is $\Delta \mathcal{R}(\theta^1) = \min \Delta c(\theta^1, \tilde{\theta}^1)$. The iterative step begins with:

Definition 5 The m 'th iteration of the emergent seed is E^m and includes all paths from θ^{m-1} to all $\tilde{\theta}^{m-1} \in \arg \Delta^{m-1} \mathcal{R}(\theta^{m-1})$.

The m 'th difference resistance is:

$$\Delta^m r(x, y) = \begin{cases} \Delta^{m-1} c(x, y) - \Delta^{m-1} \mathcal{R}(\theta^{m-1}) & \text{if } x \in \theta^{m-1}, y \notin \Delta^{m-1} \bar{\mathcal{D}}(\theta^{m-1}) \\ \Delta^{m-1} c(x, y) & \text{else} \end{cases} . \quad (15)$$

We stop at the first M such that $\Delta^M c(\cdot, \cdot)$ has only one limit set, and $E^* = \cup_{m=0}^M E^m$. Note that this construction requires only path optimitzation. While the size of the neighborhoods is increasing we are only interested in a least cost path. The impact of one path on others is irrelevant.

A critical point is that if $\Delta^0 c(\cdot, \cdot) = c(\cdot, \cdot)$ and $\Theta^0 = \Theta$ then for all $m \geq 0$, $\Delta^m c(x, y) \leq \Delta^{m+1} c(x, y)$, this implies that $\#(\Theta^m) \geq \#(\Theta^{m+1})$, indeed since we exit every m 'th difference basin of attraction $\#(\Theta^m) \geq \frac{1}{2} \#(\Theta^{m+1})$ and finding the emergent seed will take at most $\ln \#(\Theta) / \ln 2$ iterations. It is a rare paper in the literature where the emergent seed has more than one iteration.

Example 3 *The emergent seed in the Nash Demand game:* When finding the radii in this game, we specified the limit set that was transitioned to. Specifically, if $p_1^+(\theta) < p_2^+(\theta)$, then we transitioned to $\tilde{\theta} = (\theta_1 - \delta, 1 - \theta_1 + \delta)$. If $p_1^+(\tilde{\theta}) < p_2^+(\tilde{\theta})$, then we transition to $(\theta_1 - 2\delta, 1 - \theta_1 + 2\delta)$. Thus, we transition in a linear fashion to the point at which $(\theta_1, 1 - \theta_1)$ can transition to $(\theta_1 - \delta, 1 - \theta_1 + \delta)$ and $(\theta_1 - \delta, 1 - \theta_1 + \delta)$ can transition back. This is the unique first iteration limit set and it took one iteration to find the emergent seed.

The emergent seed is useful because it contains an *optimal seed*.

Definition 6 A seed is a graph, S , over Z such that:

1. Every $x \in Z$ has a transition to some $y \in Z \setminus \bar{\mathcal{D}}(x)$.
2. Some $x \in Z$ are transitioned to from all states.

Let the maximal set of states that meet the second criterion be the hub of S .

An *optimal* seed is then simply $\hat{S} \in \arg \min_{S \in \{S\}} c(S)$. One can look at the optimal tree with a root at x as a two step process. First every $y \in Z \setminus x$ exits $\bar{\mathcal{D}}(z)$, and then it enters $\bar{\mathcal{D}}(x)$. An optimal seed guarantees that the first criterion is met as best as possible, thus our remaining task is connecting the hub to x (a path optimization problem). Thus we can proceed once we have a measure of the cost of the emergent seed.

Proposition 1 $\exists S^* \subseteq E^*$ such that $c(S^*) = \sum_{m=1}^M \sum_{\theta^m \in \Theta^m} \Delta^{m-1} \mathcal{R}(\theta^{m-1})$, and if \hat{S} is an optimal seed then $\hat{S} \subseteq E^*$ and $c(\hat{S}) = c(S^*) - \max_{\theta^M} \Delta^{M-1} \mathcal{R}(\theta^{M-1})$, where the maximization is over θ^{M-1} that are in the hub of E^* . Furthermore E^* is found using only path optimization.

Thus only difference between an S^* and an \hat{S} is that the latter does not include a exiting transition from some θ^{M-1} . It is more convenient to work with S^* because for a given x we might need the transition that has been dropped from \hat{S} . Notice two things about S^* . First finding one requires solving a tree minimization problem and second there may be many of them. Fortunately we do not need S^* in analysis, only its cost.

4 A Characterization

We now use this initial structure to find the stochastic potential. Given the x of which we wish to find the stochastic potential, we then choose an $S^* \subseteq E^*$. This choice is arbitrary thus we define the *seed cost* of the emergent seed as $c_s(E^*) = c(S^*) = \sum_{m=1}^M \sum_{\theta^m \in \Theta^m} \Delta^{m-1} \mathcal{R}(\theta^{m-1})$.

Given this initial structure we now have a pair of shortest path problems instead of a tree minimization problem. First any state that is in $\Delta^m \bar{\mathcal{D}}(\theta)$ will transition to θ and then stop, thus we can drop the transitions from θ to the core, these will decrease the stochastic potential by $\sum_{m=0}^{M-1} \Delta^m \mathcal{R}(\theta)$. Please recall

that if $\theta \subseteq \theta^m$ then $\Delta^m \mathcal{R}(\theta) = \Delta^m \mathcal{R}(\theta^m)$, thus these terms will only be non-zero if $\theta \subseteq \theta^m$. Second we need to make sure that all other states also go to θ , which we can simply do by constructing a path from the hub to θ , this will have the cost $\Delta^M c(\theta^M, \theta)$.

Before writing our representation theorem, notice that there is a convenient change of basis. Since $c_s(E^*)$ is a constant it will not affect analysis, so we define the *likelihood potential* of a state as $lp(x) = c_s(E^*) - sp(x)$. For clarity we restate our definitions.

Theorem 1 *Given $c_s(E^*) = \sum_{m=1}^M \sum_{\theta^m \in \Theta^m} \Delta^{m-1} \mathcal{R}(\theta^{m-1})$, and $lp(x) = c_s(E^*) - sp(x)$ then the likelihood potential of $\theta \in \Theta$ is:*

$$lp(\theta) = \sum_{m=0}^{M-1} \Delta^m \mathcal{R}(\theta) - \Delta^M c(\theta^M, \theta). \quad (16)$$

Intuitively, the likelihood potential is similar to the radius/(modified) coradius theorem (Ellison, 2000). If the radius is larger than the (modified) coradius, then a limit set must be stochastically stable. We find that having a high *total radius* ($\sum_{m=0}^{M-1} \Delta^m \mathcal{R}(\theta)$) and a low *hub attraction rate* ($\Delta^M c(\theta^M, \theta)$) increases the likelihood potential. The (total) radius is a measure of how long it takes to leave a set, and the (modified) coradius or hub attraction rate are both measures of how long it takes to get there—although neither is precise.

A common solution method in the literature is to implicitly find the hub and then note that something in the hub has a high radius. We call this *hub dominance* and note that it is easily sufficient.

Corollary 1 (Hub Dominance) *If θ exists such that $\theta \subseteq \theta^M$ and $\sum_{m=0}^{M-1} \Delta^m \mathcal{R}(\theta) \geq \sum_{m=0}^{M-1} \Delta^m \mathcal{R}(\tilde{\theta})$ for all $\tilde{\theta} \in \Theta$, then θ is stochastically stable.*

Note that the Nash Demand game is solved by hub dominance. However, we need a further restriction to make it simple to characterize the likelihood potentials.

Example 4 Stochastic Stability and Likelihood Potentials in the Nash Demand game: Note that the function $\mathcal{R}(\theta)/n = \min[p_1^+(\theta), p_2^+(\theta)]$ is tent-shaped: $p_2^+(\theta)$ is strictly increasing and $p_1^+(\theta)$ is strictly decreasing. The maximum is characterized as the point(s) at which θ $p_2^+(\theta) \leq p_1^+(\theta)$ and at $\hat{\theta} = (\theta_1 + \delta, 1 - \theta_1 - \delta)$ $p_2^+(\hat{\theta}) \geq p_1^+(\hat{\theta})$. This also characterizes the hub, and hub dominance tells us that something in the hub is stochastically stable. It will be the Nash Bargaining solution(s) on the finite grid.

Finding likelihood potentials is more difficult because we must know whether taking two small steps is better than taking one large step.

Lemma 3 *If for all $\theta \in \Theta$*

$$\max[p_1^+(\theta) p_1^+(\theta_1 + \delta, 1 - \theta_1 - \delta), p_2^+(\theta) p_2^+(\theta_1 - \delta, 1 - \theta_1 + \delta)] < \min[p_1^+(\theta), p_2^+(\theta)] \quad (17)$$

then for θ such that $\theta_1 > \max_{\tilde{\theta} \subseteq \theta^1} \tilde{\theta}_1$ the likelihood potential is:

$$lp(\theta)/n = p_1^+(\theta) - \sum_{k=0}^{K(\theta)-1} \left(p_2^+(\hat{\theta}) - p_1^+(\hat{\theta}) \middle| \hat{\theta} = \left(\max_{\tilde{\theta} \subseteq \theta^1} \tilde{\theta}_1 + k\delta, 1 - \max_{\tilde{\theta} \subseteq \theta^1} \tilde{\theta}_1 - k\delta \right) \right) \quad (18)$$

where $K(\theta) = \frac{\theta_1 - \max_{\tilde{\theta} \subseteq \theta^1} \tilde{\theta}_1}{\delta}$ is the number of steps between this limit set and the hub.

Note that as $\delta \rightarrow 0$, $\max [p_1^+(\theta), p_2^+(\theta)] \rightarrow 0$ thus, Condition 17 will be satisfied for small δ . This example illustrates a point. Knowing the emergent seed does not immediately characterize the family of minimal cost trees. A simple emergent seed may be paired with a heterogeneous set of minimal cost trees.

5 Two Measures of Waiting Time

Waiting time is linked to the likelihood (or stochastic) potential; thus, the characterization allows us to find the speed of evolution. Indeed, we give two different measures. The coheight is precise but sometimes difficult to use and understand. The censored coradius is a generalization of the modified coradius (Ellison, 2000), often easier to use, and sometimes a method to establish stochastic stability. Our objective is now the log waiting time of θ , or:

$$\ln \tau(\theta) = \lim_{\beta \rightarrow \infty} \frac{\ln E_\beta(\min s | z_{t+s} \in \theta, z_t \in \Theta \setminus \theta)}{\beta}. \quad (19)$$

Beggs (2005) derives a general formula for this. The *height* of a set is the expected waiting time to exit that set. The log waiting time is the *coheight* or the expected waiting time to exit $\Theta \setminus \theta$. We write this as:

$$\ln \tau(\theta) = Ch(\theta) = H(\Theta \setminus \theta) = \max_{\tilde{\theta} \in \Theta} lp(\{\theta, \tilde{\theta}\}) - lp(\theta), \quad (20)$$

where $lp(X)$ is the likelihood potential of $X \subseteq Z$.

It would seem that finding $lp(\{\theta, \tilde{\theta}\})$ will be difficult. How do we know whether to have limit sets transition to θ or $\tilde{\theta}$? The emergent seed provides a simple answer to this question. Since it specifies the least cost existing transitions we should continue use the transitions in the emergent seed as much as possible. This leaves only the choice of whether to have the hub transition to θ or $\tilde{\theta}$. A small detail is that sometimes we might also have double counting, if for some $m \{\theta, \tilde{\theta}\} \subseteq \theta^m$ then we can not drop both the transitions from θ and $\tilde{\theta}$. Thus define:

$$m(\theta, \tilde{\theta}) = \begin{cases} \min \{m | \exists \theta^m \in \Theta^m, \{\theta, \tilde{\theta}\} \subseteq \theta^m\} & \text{if one exists} \\ M & \text{else} \end{cases}, \quad (21)$$

using this term our characterization is:

$$lp(\{\theta, \tilde{\theta}\}) = \sum_{m=0}^{M-1} \Delta^m \mathcal{R}(\theta) + \sum_{m=0}^{m(\theta, \tilde{\theta})-1} \Delta^m \mathcal{R}(\tilde{\theta}) - \min \{\Delta^M c(\theta^M, \theta), \Delta^M c(\theta^M, \tilde{\theta})\}.$$

Using this and the likelihood potential (Equation 16), we can then define the coheight as:

Proposition 2 For given $\tilde{\theta} \in \Theta \setminus \theta$ let:

$$Ch(\tilde{\theta}, \theta) = \sum_{m=0}^{m(\theta, \tilde{\theta})-1} \Delta^m \mathcal{R}(\tilde{\theta}) + \max \{0, \Delta^M c(\theta^M, \theta) - \Delta^M c(\theta^M, \tilde{\theta})\}, \quad (22)$$

then the coheight of θ is $Ch(\theta) = \max_{\tilde{\theta} \in \Theta \setminus \theta} Ch(\tilde{\theta}, \theta)$.

We note that this is the same measure as found in Beggs (2005) and rediscovered in Cui and Zhai (2010). We use this characterization to find the coheight in all of our applications, including two in which it is strictly lower than the modified (or censored) coradius.

Notice that if $\Delta^M c(\theta^M, \theta) \leq \Delta^M c(\theta^M, \tilde{\theta})$ then θ has no impact on the amount of time it takes to transition to θ . This apparent puzzle can be explained using the *critical droplet* from Physics and helps understand evolutionary time. Mathematically the critical droplet is the $X(\theta, \tilde{\theta}) \subseteq Z$ "closest" to $\tilde{\theta}$ such that from $X(\theta, \tilde{\theta})$ one is infinitely more likely to go to θ than $\tilde{\theta}$. Let $m(\theta, \tilde{\theta}) = M$, then in our function the first term summarizes the amount of time it will take for $\tilde{\theta}$ to reach the hub, and what $\Delta^M c(\theta^M, \theta) - \Delta^M c(\theta^M, \tilde{\theta}) \leq 0$ means is that from the hub one will be infinitely more likely to go to θ than $\tilde{\theta}$. Thus in a horse race between θ and $\tilde{\theta}$ it will take zero evolutionary time to go from the hub to θ . It may take a great deal of calendar time and require many unlikely events, but once evolution reaches the hub one is essentially there.

If one does not consider this reasoning, one arrives at a measure such as the censored coradius. The censored coradius is proposed as a generalization of the modified coradius (Ellison, 2000). This is the most common measure of waiting time used in economics. Define $\theta^m(\tilde{\theta}) = \{\theta^m \in \Theta^m | \theta \subseteq \Delta^m \bar{D}(\theta^m)\}$ and for simplicity assume $\theta^m(\tilde{\theta}) \in \Theta^m$ for all m , then the *censored coradius* is:

$$\overline{CR}(\theta) = \max_{\tilde{\theta} \in \Theta \setminus \theta} \sum_{m=0}^{M-1} \Delta^m \mathcal{R}(\theta^m(\tilde{\theta})) + \Delta^M c(\theta^M, \theta) .^{11} \quad (23)$$

If $E^1 = E^*$ it is the modified coradius and the difference is that the second term requires no analysis. One simply finds the maximizer of the first term. A benefit of the censoring is that we have found it in all of our applications. Ellison (2000) has doubts about whether the modified coradius will be simple to apply.

Because we have an exact measure of waiting time, all we need to show is that the censored coradius is a bound for the coheight (log waiting time).

Lemma 4 *For all $\theta \in \Theta$, $\overline{CR}(\theta) \geq Ch(\theta) = \ln \tau(\theta)$, a sufficient condition for them to be equal is if both are determined by a $\tilde{\theta}$ which is in the hub and $m(\theta, \tilde{\theta}) = M$.*

In Section C, we derive this measure from a waiting time argument using Bortz, Kalos, and Lebowitz (1975). One use of the censored coradius is that a total radius that is higher than the censored coradius is still sufficient for stochastic stability. To illustrate how one uses these techniques, we return again to the Nash Demand game.

Example 5 Dynamics in the Nash Demand game: *Because we solved the game using hub dominance after one iteration, the waiting time to get to the stochastically stable state is simply the second highest radius. Denote the Nash Bargaining solution as $\theta_* = (\gamma_{NBS}, 1 - \gamma_{NBS})$, then:*

$$Ch(\theta_*) = \overline{CR}(\theta_*) = \max [\mathcal{R}(\gamma_{NBS} - \delta, 1 - \gamma_{NBS} + \delta), \mathcal{R}(\gamma_{NBS} + \delta, 1 - \gamma_{NBS} - \delta)] . \quad (24)$$

As $\delta \rightarrow 0$, $Ch(\theta_) \rightarrow 1$ or evolution will be very fast. However, at some point, $R(\theta_*) = 1$ and all limit sets will be stochastically stable. An analysis at the point right before this limit is interesting. From any limit set, the most likely event is to move toward the hub. Furthermore, the likelihood of transitioning toward the*

hub increases the further away one is from it. Cui and Zhai (2010) implies that the most likely medium-run prediction is that one is in the hub near the stochastically stable state, and that the likelihood that one is in the state $\gamma_{NBS} + k\delta$ or $\gamma_{NBS} - k\delta$ will be strictly decreasing in k . Thus, the θ series will appear very similar to a price series. It will be similar to a random walk with a bias toward the long-run price $((\gamma_{NBS}, 1 - \gamma_{NBS}))$.

6 Examples of Emergent Seeds

The emergent seed will simplify our (re) analysis of several applications. Throughout the paper, we have reanalyzed the Nash Demand game and now (re)analyze four further problems. The Nash Demand game and our further examples span the three common types of emergent seeds, and the final two are more exotic. It is quite common for the emergent seed to be one of three simple classes: *lines*, *stars*, or *circuits*. Fully 74% of the applications we survey are one of these three classes: 38% linear, 26% stars, and 17% circuits. (See the Online Appendix, Section B and note that three applications appear twice). When considering these examples, we hope that readers believe that they already knew about these structures and simply lacked terminology.

In this section, the reader will probably hope for a "smoking gun." However, the emergent seed is not that messy. Its core directive is to find radii and to construct a graph of these radii. This directive is enough that we can refine the results in the Gift Giving game (Johnson, Levine, and Pesendorfer, 2000) and derive a novel one—the speed of evolution on Three Dimensional Lattices. The closest to a result that cannot be derived without the emergent seed are two close form characterizations—the Contract game (Young, 1998) and the Contribution game (Myatt and Wallace, 2008b). The Nash Demand game is our representative of a *linear* emergent seed. In linear emergent seeds, there is an implicit linear order over the limit sets and from each limit set one transitions either one up or down. The Contract game is our representative of a *star* emergent seed. In a Star, the radius of every limit set is determined by a limit set in the hub. The Gift Giving game is our representative of a *circuit* emergent seed. In this structure, every limit set is in the hub. Please note that Example 4 showed that knowing the emergent seed is not equivalent to knowing the family of minimal cost trees. Although it will be the basis for every tree, a given tree might change a significant portion of this structure.

6.1 The Contract Game, a Star Emergent Seed

The Contract game is based on the bargaining problem discussed in Example 1, and we will use the same evolutionary dynamics. The difference is that a contract is complete and lists the payoffs to all parties. Thus, now $S_i = A_1 \times A_2$ and if $s_1 = s_2$ then there is agreement; otherwise, both parties get zero. The set of strict pure strategy equilibria (and limit sets) is any $s_1 = s_2 = s$ as long as $\min [u_1(s_{11}), u_2(s_{22})] > 0$ and $s_{11} + s_{22} \leq 1$. We write the limit set s as θ .

Now, invaders will offer a party more or less, but no opportunity exists to compromise. If invaders offer \tilde{s}_1 , accepting \tilde{s}_1 is the best response if:

$$\begin{aligned} (1 - p_1(\tilde{s}_1)) u_1(\theta_1) &= p_1(\tilde{s}_1) u_1(\tilde{s}_1) \\ p_1(\tilde{s}_1) &= u_1(\theta_1) / (u_1(\theta_1) + u_1(\tilde{s}_1)) \end{aligned} \tag{25}$$

thus, the best invaders offer $\tilde{s}_1 = 1$ and:

$$\mathcal{R}(\theta)/n = \min \left[\frac{u_1(\theta_1)}{u_1(\theta_1) + u_1(1)}, \frac{u_2(\theta_2)}{u_2(\theta_2) + u_2(1)} \right]. \quad (26)$$

At this point, we must consider two separate cases. The simpler one is when either $(0, 1)$ or $(1, 0)$ are not strict equilibria, and we address this second. We will now analyze the case in which *disagreement is irrelevant*, or $u_1(0) > 0$ and $u_2(0) > 0$.

6.1.1 If Disagreement is Irrelevant

Young (1998) focuses on this case, and we suggest that the irrelevance of disagreement is reasonable; however, the relative value of the worst agreement ($u_i(0)/u_i(1)$) should be small. In a dynamic matching model, this would be the cost of finding a new partner. When disagreement is irrelevant, all contracts are strictly pure strategy equilibria, specifically $\{(0, 1), (1, 0)\}$, and these are the hub. The direct cost of going from the hub to θ is:

$$p_i(\theta) = u_i(0)/(u_i(0) + u_i(\theta_i)), \quad (27)$$

and the likelihood potential is:

$$\begin{aligned} lp(\theta)/n &= \mathcal{R}(\theta)/n - \Delta c(\theta^M, \theta)/n = \min \left[\frac{u_1(\theta_1)}{u_1(\theta_1) + u_1(1)}, \frac{u_2(\theta_2)}{u_2(\theta_2) + u_2(1)} \right] \\ &\quad - \min \left[\frac{u_1(0)}{u_1(0) + u_1(\theta_1)} - \frac{u_1(0)}{u_1(0) + u_1(1)}, \frac{u_2(0)}{u_2(0) + u_2(\theta_2)} - \frac{u_2(0)}{u_2(0) + u_2(1)} \right]. \end{aligned} \quad (28)$$

The emergent seed delivers a closed form objective. We cannot provide a full characterization for smooth utility functions because the objective function is piecewise Leontief—placing no restrictions on u'_1/u'_2 . We can see that it is Pareto efficient and independent of δ , and one can show that the function is locally either strictly concave or monotonic—in the latter case, the function is increasing as we go towards the Kalai-Smorodinsky solution: $u_1(\gamma_{KS})/u_1(1) = u_2(1 - \gamma_{KS})/u_2(1)$.

Below, we have a case in which the stochastically stable limit set is far from the Kalai-Smorodinsky solution. However, we can provide two conditions under which it will, at least, be near the Kalai-Smorodinsky solution—in the sense that it will be the allocation either just above or below the Kalai-Smorodinsky solution on the grid. If utility is symmetric:

Lemma 5 *If $u_1(x) = u_2(x)$ and utility is differentiable, then the stochastically stable limit set is near the Kalai-Smorodinsky solution.*

Alternatively, if the value of agreement ($u_i(0)/u_i(1)$) is small. This argument uses the normalization:

$$u_i(x)/u_i(1) = (1 - \beta_i)v_i(x) + \beta_i \quad (29)$$

where $v_i(0) = 0$, $v_i(1) = 1$, and $\beta_i = u_i(0)/u_i(1)$.

Lemma 6 *For all $\{v_1(\cdot), v_2(\cdot)\}$ if $\min[\beta_1, \beta_2] \rightarrow 0$, then the stochastically stable limit set is near the Kalai-Smorodinsky solution.*

For a counter example, allow:

$$u_1(s) = \left(1 - \frac{1}{2}\right)s_1 + \frac{1}{2}, \quad u_2(s) = \left(1 - \frac{3}{10}\right)s_2 + \frac{3}{10} \quad (30)$$

The Kalai-Smorodinsky solution is $\gamma_{KS} = \frac{5}{12} \sim .42$ but the stochastically stable limit set has $\theta_{1*} \sim .32$. Because we have a closed form objective, we can easily search for counter examples to stability being nearly Kalai-Smorodinsky. We searched for a more reasonable one for the case $v_1(x) = v_2(x) = x^\alpha$, $\beta_1 \geq \beta_2$. In any counter example, the smallest value of β_1 is .38, and we needed $\beta_2 \geq .24$.¹² Thus, it seems that usually stochastic stability and Kalai-Smorodinsky coincide, but if not we can identify it.

The Censored Coradius and the Coheight Because the stochastically stable limit set is usually not in the hub, the coheight is strictly lower than the censored coradius. The censored coradius is:

$$\overline{CR}(\theta_*) = \max_{\gamma \in A_1(\delta) \setminus \theta_{1*}} \mathcal{R}(\gamma, 1 - \gamma) + \Delta c(\theta^M, \theta_*) . \quad (31)$$

The coheight is lower because we have the choice between $\tilde{\theta}$ and θ_* going to the hub.

$$Ch(\theta_*) = \max_{\tilde{\gamma} \in A_1(\delta) \setminus \theta_{1*}} [\mathcal{R}(\tilde{\gamma}, 1 - \tilde{\gamma}) + \max [\Delta c(\theta^M, \theta_*) - \Delta c(\theta^M, (\tilde{\gamma}, 1 - \tilde{\gamma})), 0]] . \quad (32)$$

In this model, evolution is counter intuitive. If one does not stay at the current contract, then one goes to an extreme contract. For a while, society flips back and forth between extreme contracts and then settles on a more reasonable contract. Society will not generally be near θ_* , it simply will show up more often and stay around longer.¹³

6.1.2 If Disagreement is Relevant

The model is much simpler to solve if $u_1(0) \leq 0$ or $u_2(0) \leq 0$, assume that $u_1(0) \leq 0 < u_2(0)$. This has no impact on the best invaders; however, when we transition to $(0, 1)$, this is not a strict pure strategy equilibrium. Thus, role one agents' (weak) best response is to choose any strategy, and we can transition to any θ . We transition to $(0, 1)$ in one or two steps; thus, every limit set is in the hub. Stochastic stability then means the maximal radius, or:

$$\max_{\gamma \in A_1(\delta)} \min \left[\frac{u_1(\gamma)}{u_1(1)}, \frac{u_2(1 - \gamma)}{u_2(1)} \right] , \quad (33)$$

which is the Kalai-Smorodinsky objective function. Note that now the emergent seed is a circuit.

6.2 The Gift Giving Game, a Circuit Emergent Seed

In the Gift Giving game, the critical question is when agents will cooperate (give the gift). The article uses the radius/coradius test and finds sufficient conditions. It precisely characterizes when the selfish strategies (not giving a gift) will be stochastically stable, and the subsequent analysis shows that, otherwise, cooperation

¹²We conducted a grid search in 10^{-2} increments over $.99 \geq \beta_1 \geq \beta_2 \geq .01$ for $v_i(x) = x^\alpha$ $\alpha \in [.01, 1]$. We allowed the increment for $\omega_1(d\omega_1)$ to be as small as 10^{-4} . We must have $|\omega_{1*} - \gamma_{KS}| > d\omega_1$ to have a verified counter example. The smallest β_1 was .38, $\beta_2 \in [.24, .30]$, $\alpha \in [.98, 1]$.

¹³"Near" is in the sense of the Euclidean metric.

will be. This article is one of the few in the literature to analyze extensive form games and to innovate the assumption that all information sets have to be reached with positive probability.

Agents live for two periods. In period t , they are *young* and in period $t + 1$ they are *old*. When they are young, they have a choice between giving a gift (1) or not (0). When they are old, they either receive or do not receive a gift. Giving a gift costs 1 and receiving a gift gives a benefit of α , where $\alpha > 1$. If there is no link between giving a gift in period t and receiving one in $t + 1$, an agent will never give the gift. This linkage is established using a *social status*, agents are either green (g) or red (r). The social status of old agents will be determined by their action when they were young.

Thus, a strategy has two elements: an action conditional on social status $a : \{r, g\} \rightarrow \{0, 1\}$ and a transition rule $\tau : \{r, g\} \times \{0, 1\} \rightarrow \{r, g\}$. Although there are 64 strategies, many are equivalent. First, either red or green could be good. Evolution cannot determine the language; thus, we usually assume that green is good ($a(g) \geq a(r)$). If $a(g) = a(r)$, the transition rule does not matter. If $a(g) = a(r) = 0$, these are the *selfish* strategies. If $a(g) = a(r) = 1$, these are the *generous* strategies. A *cooperative* strategy ($a(g) > a(r)$) can only be an equilibrium if $\tau(g, 1) = g$ and $\tau(g, 0) = r$. There are only four cooperative strategies to consider:¹⁴

$\tau(g, 1)$	$\tau(g, 0)$	$\tau(r, 1)$	$\tau(r, 0)$	Name
g	r	r	g	team
g	r	g	g	weak team
g	r	r	r	insider
g	r	g	r	tit for tat

We need to allow for agents to use different strategies. Johnson, Levine, and Pesendorfer (2001) assumes that each agent has a *flag* of social statuses: $f \in \{r, g\}^{16}$ —one for each transition rule. An agent using strategy s then uses the appropriate social status.

We will insert noise into the flag process: with probability $\eta > 0$, a player's f will be replaced with another one at random. Naturally, η is small and must satisfy $\eta < \bar{\eta}$, where $(1 - \bar{\eta})\alpha = 1$. Let Φ_t be the distribution of flags in t , then in t agents know Φ_{t-1} . The noise guarantees that for all f and Φ_{t-1} , $\Pr(\Phi_t(f) | \Phi_{t-1}) > 0$. When there is noise, it should be clear that tit-for-tat is not an equilibrium. The best response is a generous strategy, and a selfish strategy is the best response to a generous one. The other cooperative strategies are equilibria for small η .

Evolutionary dynamics will be determined by one population BRM. There will be n agents that will be uniformly matched into pairs— n is even. With probability $\rho \in (0, 1)$, agents use the strategy they used last period. If they choose a new strategy, it will be a best response to the current distribution of strategies and flags with probability $1 - e^{-\beta}$. Otherwise, it will be a strategy chosen at random.

In the Online Appendix (Section C), we provide a detailed analysis of the value functions; here, we provide an overview. Let $v(s, p)$ be the value function of someone using strategy s when with probability p someone is using strategy s' . If $v(s, p|f_s)$ for $f_s \in \{r, g\}$ is the value function conditional on a player's social status being f_s , then obviously:

$$v(s, p) \geq \min \{v(s, p|g), v(s, p|r)\}$$

and the strategy is not in equilibrium if either $v(s, p, g)$ or $v(s, p, r)$ is too low relative to the invader. Thus,

¹⁴The strategy names are from a working paper version of Johnson, Levine, and Pesendorfer (2001).

we either need $a(g) = 0$ ($v(s, p, g)$ is low) or $a(r) = 1$ ($v(s, p, r)$ is low). Comparing $v(s, p, g)$ when s is cooperative ($s \in \{\text{team, weak team, insider}\}$) and s' is selfish is enough to isolate the probability of going to and from the selfish strategies. To get $a(r) = 1$, the invader must be a cooperative strategy for which red is good. Remember that this cooperative strategy could be tit-for-tat, which is in the basin of attraction of the selfish strategies. The analysis then shows:

Lemma 7 {selfish, team, weak team, insider} are all strict equilibria for small enough η . The radii are:

$$\begin{aligned}\mathcal{R}(\text{selfish})/n &= \frac{1}{(1-\eta)\alpha}, \quad \mathcal{R}(\text{team})/n = \min\left[1 - \frac{1}{(1-\eta)\alpha}, \frac{1}{2}\left(1 + \frac{1}{(1-\eta)\alpha}\right)\right] \\ \mathcal{R}(\text{insider})/n &= \mathcal{R}(\text{weak team})/n = \min\left[1 - \frac{1}{(1-\eta)\alpha}, \frac{1}{(1-\eta)\alpha}\right].\end{aligned}\quad (34)$$

from the selfish one transitions to any cooperative ({team, weak team, insider}), and from a cooperative one can always transition to a selfish one.

Because the selfish can always transition to any cooperative equilibrium and any cooperative can always transition to the selfish (possibly via tit-for-tat), all limit sets are in the hub and we only need to find which has the maximum radius.

Lemma 8 If $(1-\eta)\alpha < 2$, then selfish strategies are stochastically stable; if $(1-\eta)\alpha = 2$, then all equilibrium strategies are; and if $(1-\eta)\alpha > 2$, then team strategies are.

Like always, when the emergent seed has one iteration and hub dominance holds, the censored coradius and the coheight are the same—the second highest radius.

6.3 Exotic Emergent Seeds

If the emergent seed does not fall into one of our simple classes, we refer to it as *exotic*. Naturally, in this class, the two problems that cannot be solved using the emergent seed fall, but many can be solved. We illustrate two that can be. The first, the three-dimensional lattice under BRM, is a novel contribution. We are able to precisely pin down the speed of evolution in the three dimensional lattice. In the second, the contribuition game, our contribution is to remove a simplifying assumption and give an analytic characterization of stochastic stability.

6.3.1 The Three-Dimensional Lattice with BRM.

Economists were not the first to turn to the lattice as a model of local interaction. Ising (1925) first analyzed the one-dimensional lattice in physics to explain the fast and uneven manner in which ice forms. In a similar vein, Ellison (1993) turned to the lattice when he wanted to show that evolution might be fast. Similar to ice crystals, local interactions could lead to fast propagation of stochastically stable limit sets. Both studies turned to more general models of local interaction after progress stalled. Ellison (2000) extended the analysis in economics to two dimensions. Arous and Cerf (1996) extended it in physics to three dimensions—and relied on a potential function. One needs to find a *critical path*—a path from the risk dominated equilibrium to the risk dominant one. Section 5 shows how the emergent seed can be of assistance. Using this, Hasker

(2018) found the emergent seed for all finite number of dimensions in both the BRM and the Logit models. Here, we explain the three dimensional lattice under BRM.

One population BRM is that agents choose a new strategy with probability $1 - \rho \in (0, 1)$ and choose a best response to the current distribution with probability $1 - e^{-\beta}$. Otherwise, they choose a strategy at random. In this analysis, the innovation is that agents only interact with their neighbors in a lattice. Thus, the population of agents, I , has n^3 members for $n \geq 6$. Each $i \in I$ will be endowed with a three-dimensional location, $(\chi_1(i), \chi_2(i), \chi_3(i))$, where for $d \in \{1, 2, 3\}$ $\chi_d(i) \in \{0, 1, 2, \dots, n-1\}$.¹⁵ Each agent will interact only with their *neighbors*. We say that j is a neighbor of i (denoted $j \sim i$) if there is a $d \in \{1, 2, 3\}$ such that $(\chi_d(i) \pm 1) \bmod (n-1) = \chi_d(j)$ and for $\tilde{d} \in \{1, 2, 3\} \setminus d$ $\chi_{\tilde{d}}(i) = \chi_{\tilde{d}}(j)$. In essence, we are taking a cube of agents and wrapping it at the edges to avoid a boundary effect. Each period an agent plays a classic coordination game with all of its neighbors:

	A	B	
A	1 - σ , 1 - σ	0, 0	
B	0, 0	σ , σ	

(35)

We now require that $\sigma \in (\frac{1}{3}, \frac{1}{2})$. The upper bound makes (A, A) risk dominant, and the lower bound makes the problem non-degenerate. Given our normalization, (A, A) is also Pareto efficient.

A state is a subset of agents: $x \subseteq I$, if $i \in x$, then i is using the strategy A . Let $\#(i, x) = \#\{j \in x | j \sim i\} \in \{0, 1, 2, \dots, 6\}$. If $BR(i, x)$ is the best response of i given the state x , then $\sigma \in (\frac{1}{3}, \frac{1}{2})$ means:

$$BR(i, x) = \begin{cases} A & \text{if } \#(i, x) \geq 3 \\ B & \text{if } \#(i, x) \leq 2 \end{cases}. \quad (36)$$

In this problem, all of the limit sets will be strictly pure strategy Nash equilibria, and Peski (2010) proved that the state in which everyone plays A ($\theta_A = I$) is stochastically stable. Our key question is the speed of evolution from $\theta_B = \emptyset$ —where everyone plays B . We will go *up* from θ if we transition to the set: $\Theta_+(\theta) = \{\tilde{\theta} \in \Theta | \theta \subset \tilde{\theta}\}$, and *down* if we go to the set $\Theta_-(\theta) = \{\tilde{\theta} \in \Theta | \theta \supset \tilde{\theta}\}$. In the emergent seed, this will always occur.

Lemma 9 *For $\theta \in \Theta \setminus \{\theta_A, \theta_B\}$ $\mathcal{R}(\theta) = \min \{c(\Theta_+(\theta), \theta), c(\Theta_-(\theta), \theta)\}$.*

Our analysis shall rest on two particular types of *boxes* (or *orthotopes*). We will analyze one-, two-, and three-dimensional boxes, and will require that there are at least two agents in each dimension. Thus, we denote a box: $box(d, l_1, l_2, l_3)$ where $d \in \{1, 2, 3\}$ is the dimension and $l_{\tilde{d}} \geq 2$ ($\tilde{d} \in \{1, \dots, d\}$) are the lengths. If we do not mention a length, then it is two. Thus, $box(3)$ is a set of eight agents arranged in a cube. The limit sets where the least (excluding θ_B) and the most (excluding θ_A) agents play A can be characterized with boxes as long as $n > 4$.

Lemma 10 *If $\#(\theta) = \min \{\#(\theta) | \theta \in \Theta \setminus \{\theta_A, \theta_B\}\}$, then θ is a $box(3)$; likewise, if $\#(\theta) = \max \{\#(\theta) | \theta \in \Theta \setminus \{\theta_A, \theta_B\}\}$, then θ is $I \setminus box(3, n)$.*

From now on, we will want to use $box(3)$ and $box(3, n)$ as the bases of our analysis. To make the cost measurements precise, we make some simplifying restrictions on the limit sets we analyze.

¹⁵There is a one-to-one mapping between locations and agents.

Definition 7 We say that a limit set is:

1. **small** if $\theta \subseteq \text{box}(3, n-2, n-2, n-2)$
2. **convex** if for all $i \in \theta$ and $j \in \theta$, either for all relevant $\lambda \in (0, 1)$, $\lambda(\chi_1(i), \chi_2(i), \chi_3(i)) + (1-\lambda)(\chi_1(j), \chi_2(j), \chi_3(j))$ is or is not in θ .¹⁶
3. **orbicular** if there is a sequence of boxes $(x_s)_{s=1}^S$, where $x_1 \in \text{box}(3)$ and for $s > 1$ $x_s \in \text{box}(2)$ such that $\theta = \cup_{s=1}^S x_s$ and for all $\hat{S} \leq S$, for $i \in \cup_{s=1}^{\hat{S}} x_s$ $BR(i, \cup_{s=1}^{\hat{S}} x_s) = A$.

These assumptions are without loss of generality because all relevant limit sets in the critical path must satisfy them. Convexity rules out holes in limit sets, such as the "bagel," which is constructed by taking eight $\text{box}(3)$ and arranging them in a circle. This limit set is orbicular. Orbicular rules out the "pair of dice." Take two $\text{box}(3)$ that have only one agent in common. This limit set is convex but not orbicular. Both imply that there are not separate areas of agents playing A , such as two $\text{box}(3)$ that have no common neighbors.

With these assumptions, we can be precise about the cost of going up and down. To go up, we have to append a $\text{box}(d)$ for $d \in \{1, 2, 3\}$ such that all agents playing A in the new state are in a $\text{box}(3)$. To be precise, for $\theta \in \Theta \setminus \theta_A$, we need to append a box of dimension:

$$d(\theta) = \min \{d | \text{box}(d) \not\subseteq \theta, \forall i \in \theta \cup \text{box}(d), BR(i, \theta \cup \text{box}(d)) = A\}, \quad (37)$$

which we refer to as the *dimension* of θ . Likewise, for $\theta \in \Theta \setminus \theta_B$ going down, we need to remove:

$$l(\theta) = \min \{\#(\theta \cap \text{box}(3, n)) | \theta \cap \text{box}(3, n) \neq \emptyset, \forall i \notin \theta \setminus \text{box}(3, n), BR(i, \theta \setminus \text{box}(3, n)) = B\}, \quad (38)$$

which we call the *length* of θ . Note that $d(\theta) = 3$ only if $\theta = \theta_B$, otherwise $d(\theta) \in \{1, 2\}$. Likewise, $l(\theta) \geq 2$. It is fairly immediate that:

Lemma 11 Assume that θ is small, convex, and orbicular. Then, $c(\Theta_+(\theta), \theta) = 2^{d(\theta)-1}$ and $c(\Theta_-(\theta), \theta) = \lfloor l(\theta)/2 \rfloor$.

Our key result is then:

Proposition 3 In the emergent seed, θ_A and θ_B are in the two first iteration limit sets. Furthermore:

$$\ln \tau(\theta_A) = \mathcal{R}(\theta_B) + \Delta \mathcal{R}(\theta_B) = 2^{3-1} + 3 * 2 = 10.$$

This result is because evolution will proceed through a sequence of three-dimensional boxes. One goes up by adding a $\text{box}(2)$ to a (largest) side. Once one has done this, the new limit set has a dimension of one, and one can fill in that side. Going in the other direction, we remove a (smallest) edge from the box. The tipping point is $\text{box}(3, 4, 4, 4)$. Above this point, removing an edge is at least as costly as adding a $\text{box}(2)$.

¹⁶A λ is relevant if $\lambda(\chi_1(i), \chi_2(i), \chi_3(i)) + (1-\lambda)(\chi_1(j), \chi_2(j), \chi_3(j))$ is in the lattice.

6.3.2 Contribution Game

In the contribution game, there are $n = \#(I)$ agents and a public good that requires l people to contribute, where $1 < l \leq n$. Agents are willing to contribute if and only if necessary. Instead of specifying a model of evolution, an agent can be described by two parameters, $(b_i, d_i) \in (0, 1)^2$. With probability b_i^β , an agent will contribute when there is no benefit. With probability d_i^β , they will stop contributing when it means the public good will not be provided. Without loss of generality, we will assume that $b_s > b_{s+1}$ for $s \in \{1, 2, 3, \dots, n-1\}$ and that (b_i, d_i) are generic.¹⁷ Myatt and Wallace (2008b) provide a characterization when $d_s < d_{s+1}$, and otherwise uses robustness. We extend the characterization to all $(d_i)_{i=1}^n$. Let a state, z , be the agents who are contributing. The strict equilibria (and limit sets) are $\theta_\emptyset = \emptyset$ —no one contributes, and $\hat{\Theta}$ where if $\theta \in \hat{\Theta}$, then $\#(\theta) = l$; therefore, $\hat{\Theta} = \{\theta_\emptyset, \hat{\Theta}\}$. If $i \in \theta \in \hat{\Theta}$, then $r(\theta, \theta \setminus i) = \ln \frac{1}{d_i}$ and if $i \notin z \in \hat{\Theta}$ or $\#(z) < l-1$ $r(z, z \cup i) = \ln \frac{1}{b_i}$.

To exit θ_\emptyset , we need to get $l-1$ agents to contribute when there is no benefit, thus $\mathcal{R}(\theta_\emptyset) = \min_{z: \#(z)=l-1} \sum_{i \in z} \ln \frac{1}{b_i} = \sum_{i=1}^{l-1} \ln \frac{1}{b_i}$. If we let $z_+ = \{1, 2, 3, \dots, l-1\}$, then we can go from θ_\emptyset to any $\hat{\theta}_k$ where $n \geq k \geq l$ and $\hat{\theta}_k = z_+ \cup k$.

For $\theta \in \hat{\Theta}$, the radius is $\mathcal{R}(\theta) = \min \left[\min_{i \in \theta} \ln \frac{1}{d_i}, \min_{j \in I \setminus \theta} \ln \frac{1}{b_j} \right]$.

The problem is that there are many forms of indifference. For example, if you add $j(\theta) = \min \{j | j \in I \setminus \theta\}$, then you can drop anyone, easily creating a cycle. In cases like this, the best course is to find limit set(s) that can be in any unconstrained least cost path—where every transition is determined by the radius. These limit set(s) will be in the hub.

Lemma 12 *From every $\theta \in \Theta \setminus \hat{\theta}_l$, there is an unconstrained least-cost path from θ to $\hat{\theta}_l$.*

Thus, $\hat{\theta}_l$ must be in the hub. Its radius is $\mathcal{R}(\hat{\theta}_l) = \min \left[\ln \frac{1}{d_+}, \ln \frac{1}{d_l}, \ln \frac{1}{b_{l+1}} \right]$, where $\ln \frac{1}{d_+} = \min_{i \in z_+} \ln \frac{1}{d_i}$. Using $\hat{\theta}_l$ and θ_\emptyset , we can now rule out any $\theta \in \Theta \setminus \{\hat{\theta}_k\}_{k=l}^n$.

Lemma 13 *For all $\theta \in \Theta \setminus \{\hat{\theta}_l, \theta_\emptyset\}$ $\ln \frac{1}{b_{l+1}} > \mathcal{R}(\theta)$, and if $\theta \neq \hat{\theta}_k$ then $\mathcal{R}(\theta_\emptyset) > \mathcal{R}(\theta)$. Thus, only $\{\hat{\theta}_k\}_{k=l}^n$ or θ_\emptyset can be stochastically stable.*

The radii of the $\hat{\theta}_k$ for $k > l$ are $\mathcal{R}(\hat{\theta}_k) = \min \left[\ln \frac{1}{d_+}, \ln \frac{1}{d_k}, \ln \frac{1}{b_l} \right]$. Stochastic stability is easy to characterize when not contributing (θ_\emptyset) is in the hub. Similar to the article, we focus on when θ_\emptyset is stochastically stable. If θ_\emptyset is in the hub, it requires:

$$\mathcal{R}(\theta_\emptyset) = \sum_{i=1}^{l-1} \ln \frac{1}{b_i} \geq \max_{k \geq l} \mathcal{R}(\hat{\theta}_k) . \quad (39)$$

The more difficult problem is when it is not. It is not if and only if $\mathcal{R}(\hat{\theta}_l) = \ln \frac{1}{b_{l+1}}$ and $\mathcal{R}(\hat{\theta}_{l+1}) = \ln \frac{1}{b_l}$ and the hub is $\{\hat{\theta}_l, \hat{\theta}_{l+1}\}$. If θ_\emptyset is not in the hub then we must find θ_\emptyset 's core attraction rate and compare θ_\emptyset 's likelihood potential to that of $\hat{\theta}_l$. This path can begin at $\hat{\theta}_l$, the key question is how many intermediate steps we might want. Notice that if $i \in z_+$ $\ln \frac{1}{b_i} < \ln \frac{1}{b_{l+1}}$, since $\mathcal{R}(\hat{\theta}_l) = \ln \frac{1}{b_{l+1}}$ we now that $\ln \frac{1}{b_{l+1}} \leq \ln \frac{1}{d_+}$. Combining these facts tells us that for all $i \in z_+$, $\ln \frac{1}{b_i} < \ln \frac{1}{d_i}$. This implies there is no benefit to removing

¹⁷A property is *generic* if it is true for a dense open set. Here, we rule out ties between the $(b_i, d_i)_{i=1}^n$.

someone from z_+ because this will only increase the probability of adding someone else in the following step. Thus there can be at most one step, and the core attraction rate is:

$$\Delta c(\theta^1, \theta_\emptyset) = \min \left[\min \left[\ln \frac{1}{d_+}, \ln \frac{1}{d_l} \right], \min_{k>l} \left(\min \left[\ln \frac{1}{d_+}, \ln \frac{1}{d_k} \right] - \mathcal{R}(\hat{\theta}_k) + \ln \frac{1}{b_k} \right) \right] - \ln \frac{1}{b_{l+1}}. \quad (40)$$

The first term handles the direct jump and the second allows for an intermediate step. When θ_\emptyset is not in the hub, it is stochastically stable if and only if:

$$\begin{aligned} \mathcal{R}(\theta_\emptyset) - \min \left[\min \left[\ln \frac{1}{d_+}, \ln \frac{1}{d_l} \right], \min_{k>l} \left(\min \left[\ln \frac{1}{d_+}, \ln \frac{1}{d_k} \right] - \mathcal{R}(\hat{\theta}_k) + \ln \frac{1}{b_k} \right) \right] - \ln \frac{1}{b_{l+1}} &\geq \ln \frac{1}{b_{l+1}} \\ \mathcal{R}(\theta_\emptyset) &\geq \min \left[\min \left[\ln \frac{1}{d_+}, \ln \frac{1}{d_l} \right], \min_{k>l} \left(\min \left[\ln \frac{1}{d_+}, \ln \frac{1}{d_k} \right] - \mathcal{R}(\hat{\theta}_k) + \ln \frac{1}{b_k} \right) \right] \end{aligned}$$

interestingly enough one can show that the right hand side is $\overline{CR}(\theta_\emptyset)$. To reach the formula in the article, assume that $\ln \frac{1}{d_+} > \max_{k \geq l} \ln \frac{1}{d_k}$, in which case the two cases merge. The coheight of θ_\emptyset is always $\max_{k \geq l} \mathcal{R}(\hat{\theta}_k)$, which is easier to find and strictly lower than the censored coradius when θ_\emptyset is not in the hub.

7 Limitations and Comparisons

The emergent seed is not always the best methodology. In our survey, we found that it could have solved at least 95%, but not all, of the applications in the literature. For example, Bergin and Bernhardt (2009) show that the cooperative outcome is stochastically stable in arbitrary symmetric games with long-memory imitation. Essentially, that paper uses the radius/coradius test and the emergent seed might be arbitrarily complex.

An example in which only root switching seems to work is Ben-Shoham, Serrano, and Volij (2004). This paper analyzes housing allocations and rank-based errors. Similar to proper equilibrium (Myerson, 1978), if one mistake is worse than another, then it gets an order of magnitude lower probability. For example, if i trades her second best house for her third best, it has a probability of $e^{-(3-2)\beta}$. If she trades for her fifth best, it has a probability of $e^{-(5-2)\beta}$. It is also the first class of problems analyzed in economics in which the emergent seed can have more than one iteration.¹⁸ The article establishes a *reversible paths* property. The optimal path from θ to $\tilde{\theta}$ is also the optimal path from $\tilde{\theta}$ to θ , and the difference in the costs is the difference in the levels of envy. Root switching then establishes that stochastic stability and minimal envy are equivalent.

In terms of computational complexity, we doubt that we have improved on Edmonds' algorithm, but we doubt this about any paper in the field. Gabow, Galil, Spencer, and Tarjan (1986) find the optimal computational method for a given root and use Edmonds' algorithm. A trade off exists between the simplicity of each iteration and the number of iterations. Our algorithm decreases the number but increases the complexity. In contrast, it is undeniable that finding the emergent seed will take fewer computations using our method than that proposed in Cui and Zhai (2010). As we previously stated, the first iteration of

¹⁸Assume that two agents have the same preferences and that in θ one has his or her favorite house and the other has his or her second favorite. Then, there is a $\tilde{\theta}$ in which the other has the favorite house. These allocations have a resistance of one and will be part of a cycle. One can easily construct others, making the number of first iteration limit sets as large as desired.

both methods will find the same set of cycles. In the next iteration, the Cui and Zhai (2010) algorithm often will have a first iteration limit set pointing to something in its first iteration outer basin of attraction. This method finds the first iteration radii and, thus, cannot be more efficient than ours. However, once the emergent seed is found, a clear comparison cannot be made. Cui and Zhai (2010) continue connecting the hub to other limit sets, and we recommend solving the shortest path problems. We are interested to see an application for which the Cui and Zhai (2010) algorithm is superior. In terms of computational complexity, there is a technical problem in both Beggs (2005) and Cui and Zhai (2010). Both ask one to repeatedly solve minimal cost tree problems. Finding a minimal cost tree takes many computations and our method finds stochastic potentials without this.

Let us compare our algorithm to Cui and Zhai (2010). In the Nash Demand game, the first iteration of both methods will find the hub. Our methodology stops but theirs continues. Note that all of the other limit sets will continue to point at the limit set that they previously pointed at, and in each iteration the circuit will (generically) pick up one of them until they are all in one grand cycle. Although each iteration is simple and one can use the characterization in Cui and Zhai (2010), this reaches the upper bound on the number of iterations that their algorithm can require: $\#(\Theta) - 1$. In the Contract game, our method again stops at one iteration and Cui and Zhai (2010) must continue. In the second iteration, their method will pick up either the limit sets $(x, 1 - \delta)$ or $(1 - \delta, x)$, where $x \in \{0, \delta\}$. If it picks up $(x, 1 - \delta)$, then in the next iteration it will pick up either $(1 - \delta, x)$ or $(y, 1 - 2\delta)$, where $y \in \{0, \delta, 2\delta\}$. We hope that the process is clear and note that one will need at most $1/\delta - 1$ steps.¹⁹ Note that if disagreement is relevant, both methods require only one iteration. Any problem with a circuit emergent seed will require only one iteration with either method.

We feel that the contribution of Cui and Zhai (2010) is a characterization of stochastic stability based on cycles. The paper never claims to reduce computational complexity. Likewise, in Beggs (2005), although there is an algorithm, the goal is to increase our understanding of waiting time. Our goal is to explain and use an underlying architecture in stochastic evolution.

Unfortunately, the algorithm in Beggs (2005) is so dissimilar to Edmonds' algorithm that a direct comparison is difficult. To see that it is not the same algorithm, recognize that it proposes iteratively dropping sets of state with a low height. In general, height is not trivial to compute, but for a limit set, it is the radius. Thus, in the Contract game with irrelevant disagreement, one of the first limit sets dropped is an extreme contract, which is in the first iteration limit set of Edmonds' algorithm (our hub). We would be fascinated to see an application of this methodology. Trygubenko and Wales (2006) present an algorithm that improves on Bortz, Kalos, and Lebowitz (1975) because it does not require generating the iterated resistance. Because the method in Bortz, Kalos, and Lebowitz (1975) is similar to ours, this might be better as well.

8 Conclusion

We hope that the emergent seed has helped the reader understand stochastic evolution. It is an intuitive method, essentially iterating the concept of the limit set, and gives us both a characterization of stochastic potential and waiting time. We also make a second claim on the basis of the fact that a vast majority of the applications implicitly used this method. This methodology is self evident. We have been implicitly using it and this paper's contribution is to explain this use.

¹⁹In this problem, $\#(\Theta) = \frac{1}{2\delta} (\frac{1}{\delta} - 1)$.

Even if analysts decide not to use our methodology, we hope they have benefitted from this paper. This paper is the first to lay out the general methodologies currently in use. These are deriving the limiting distribution, root switching, and the radius/(modified) coradius test. All of these methodologies have cases in which they are better than the emergent seed.

We hope that this paper has brought some clarity to the study of stochastic evolution. Although this is a promising field, the methodology used and the rationale for results are often confusing and opaque. The most popular current methodology is root switching. This will always be a guess-and-verify methodology but the literature is a tribute to its success. The emergent seed is another methodology, one that provides a characterization and formulas for waiting time. However, its true value will be measured by future applications.

9 Appendix—Proofs

Proof of Proposition 1. First we will explain how to construct S^* . We begin at the M 'th level with S^M . To construct this for each θ^{M-1} choose an exit path in E^M , these choices can be arbitrary. After this point we are contructing trees with roots in θ^M . Fortunately we do not actually need S^* , only to know that it exists. For a given θ^{M-1} , there will be some $\theta^{M-2} \subseteq \theta^{M-1}$ that are needed to construct the path between the entering and exiting θ^{M-2} of θ^{M-1} in S^M using paths from E^{M-1} (note $\theta^{M-2} \subseteq \theta^{M-1}$). For $\theta^{M-1} \subseteq \Delta^M \bar{\mathcal{D}}(\theta^M)$ that are not in these paths and one must select a sequence of paths from E^{M-1} so that all of these limit sets are connected to the path just found. We then iterate this process until we have found S^* . Note that in the last step we consider states instead of limit sets.

In order to show that for optimal seeds, \hat{S} , that $\hat{S} \subseteq E^*$ we shall argue by construction that this must be true. First of all it is clear that in \hat{S} from each $x \notin \Theta$ we must include a path from x to some θ such that $x \in \bar{\mathcal{D}}(\theta)$. The question is which to include. We notice that at later dates we might have reason to use one or another of these paths and thus we should include them all, this gives us E^0 . If $\#(\Theta) = 1$ then we are done, otherwise we must have a transition from θ that exits $\bar{\mathcal{D}}(\theta)$, and obviously we should choose the one that has the least cost or $\tilde{\theta} \in \arg \mathcal{R}(\theta)$. Again we should delay the decision of which to include and we derive E^1 . There may, of course, be new cycles in these limit sets, which we denote θ^1 's and the set Θ_1 . If $\#(\Theta_1) = 1$ we are done, if not the cost function we should use for the next step is $\Delta c(\cdot, \cdot)$ because we want to choose paths that will increase the cost of our \hat{S} the least, and this is precisely what the difference cost function isolates. We now proceed by iteration and see that $\hat{S} \subseteq E^*$. Given that $\hat{S} \subseteq E^*$ it is clear for θ^M what we wish to do is construct a tree with a root at the one with a highest radius, and thuse whatever $c(S^*)$ is, $c(\hat{S}) = c(S^*) - \max_{\theta^M \in \Theta^m} \Delta^M \mathcal{R}(\theta^M)$.

Now we show that $c(S^*) = \sum_{m=0}^M \sum_{\theta^m \in \Theta^m} \Delta^m \mathcal{R}(\theta^m)$. First the cost of any cost minimizing seed on E^0 is zero, likewise if we also include E^1 it will be $\sum_{\theta \in \Theta} \mathcal{R}(\theta)$. If we further include E^2 then the cost should be $\sum_{m=0}^1 \sum_{\theta^m \in \Theta^m} \Delta^m \mathcal{R}(\theta^m) = \sum_{\theta^1 \in \Theta^1} \Delta \mathcal{R}(\theta^1) + \sum_{\theta \in \Theta} \mathcal{R}(\theta)$. To see that this is correct assume that from θ we transition to $\tilde{\theta}$, then the total cost associated with θ is $[c(\theta, \tilde{\theta}) - \mathcal{R}(\theta)] + [\mathcal{R}(\theta)]$ where the second term is from E^2 and the first from E^1 . The addition is simply $c(\theta, \tilde{\theta})$ and this is true. For the inductive step assume some $\theta^{\tilde{m}}$ is not used in a new path until iteration $m > \tilde{m}$. In this case the current cost associated with $\theta^{\tilde{m}}$ is $[\Delta^{\tilde{m}} \mathcal{R}(\theta^{\tilde{m}})]$ and the cost of transitioning to some $\tilde{\theta}^{\tilde{m}}$ is $[\Delta^{\tilde{m}} c(\theta^{\tilde{m}}, \tilde{\theta}^{\tilde{m}}) - \Delta^{\tilde{m}} \mathcal{R}(\theta^{\tilde{m}})]$ and again the accounting is correct.

Finally, note that by construction E^* is only found using path optimization. ■

Proof of Theorem 1. We begin by noticing that since we need an exiting path from every state except θ , we should use the exiting paths in E^* as much as possible, thus we should choose an $S^* \subseteq E^*$. Clearly we need to include a path from θ^M to θ , the least additional cost way to do this is given by $\Delta^M c(\theta^m, \theta)$, finally for those states which are in $\Delta^m \bar{\mathcal{D}}(\theta)$ we do not need a path from these states to the hub (θ^M) this

is addressed by subtracting $\sum_{m=0}^{M-1} \Delta^m \mathcal{R}(\theta)$, noticing that if $\theta \not\subseteq \theta^m$ then $\Delta^m \mathcal{R}(\theta) = 0$. ■

Compliance with Ethical Standards: The author declares that he has no conflict of interest.

References

- [1] Alos-Ferrer, C and N Netzer (2010) "The logit-response dynamics" *Games and Economic Behavior* 68:413-427.
- [2] Alos-Ferrer, C and N Netzer (2012) "Robust stochastic stability" *University of Zurich Working Paper Series*
- [3] Arous, GB; and Cerf, R. (1996) "Metastability of the Three Dimensional Ising Model on a Torus at Very Low Temperatures." *Electronic Journal of Probability*, 1:1-55.
- [4] Beggs, AW (2005) "Waiting times and equilibrium selection" *Economic Theory*, 25:599-628
- [5] Bergin, J; and D. Bernhardt (2009) "Cooperation through Imitation." *Games and Economic Behavior*, 67:376-388.
- [6] Ben-Shoham, A; Serrano, R; and O. Volij (2004) "The evolution of exchange" *Journal of Economic Theory* 114:310-328.
- [7] Binmore, KG; Samuelson, L; and HP Young. (2003) "Equilibrium Selection in Bargaining Models." *Games and Economic Behavior*, 45:296-328.
- [8] Blume, LE (1993) "The Statistical Mechanics of Strategic Interaction." *Games and Economic Behavior* 5:387-424
- [9] Bock F (1971) "An algorithm to construct a minimum directed spanning tree in a directed network" in *Developments in Operations Research*, Gordon and Breach, New York, pp. 29–44.
- [10] Bortz, AB; Kalos, MH; and JL Lebowitz (1975) "A New Algorithm for Monte Carlo Simulation of Ising Spin Systems" *Journal of Computational Physics*, 17:10-18.
- [11] Chiu, YJ; and TH Liu (1965) "On the Shortest Arborescence of Directed Graphs." *Scientia Sinica*, 14:1390-1400.
- [12] Cui, ZW; and JA Zhai (2010) "Escape dynamics and equilibria selection by iterative cycle decomposition" *Journal of Mathematical Economics* 46:1015–1029.
- [13] Edmonds, J (1967). "Optimal Branchings." *Journal of Research of the National Bureau of Standards*, 71B: 233-240.
- [14] Ellison, G (1993) "Learning, Local Interaction and Coordination." *Econometrica*, 61(5),1047-1071.
- [15] Ellison, G (2000) "Basins of Attraction, Long-Run Stochastic Stability, and the Speed of Step-by-Step Evolution." *Review of Economic Studies*, 67:1-23.
- [16] Freidlin, M; and A Wentzell (2012) *Random Perturbations of Dynamical Systems* Springer Verlag, Berlin Heidelberg *third edition*.
- [17] Fudenberg, D; and LA Imhof (2006) "Imitation processes with small mutations" *Journal of Economic Theory* 51:251-262.
- [18] Gabow, HN; Galil, Z; Spencer, T; and RE Tarjan, (1986) "Efficient algorithms for finding minimum spanning trees in undirected and directed graphs," *Combinatorica* 6:109-122
- [19] Harsanyi, JC (1973) "Games with randomly disturbed payoffs: a new rationale for mixed-strategy equilibrium points" *International Journal of Game Theory* 2:1-23.

- [20] Hasker, K (2018) "The Speed of Evolution on Lattices" *work in progress*
- [21] Humblet, P (1983) "A distributed algorithm for minimum weighted directed spanning trees", *IEEE Transactions on Communications*, 6:756-762.
- [22] Ising, E. (1925), "Beitrag zur Theorie des Ferromagnetismus", *Z. Phys.* 31: 253–258.
- [23] Johnson, P; Levine, D; and W Pesendorfer. (2001): "Evolution and Information in a Gift-Giving Game." *Journal of Economic Theory*, 100:1-21.
- [24] Kandori, M; Mailath, G; and R Rob. (1993): "Learning Mutation, and Long Run Equilibria in Games." *Econometrica*, 61:9-56.
- [25] Kandori, M; and R Rob. (1995): "Evolution of Equilibria in the Long Run: A General Theory and Applications." *Journal of Economic Theory*, 65:383-414.
- [26] Kandori, M; and R Rob. (1998): "Bandwagon Effects and Long Run Technology Choice." *Games and Economic Behavior*, 22:30-60.
- [27] Kandori, M; Serrano, R; and O Volij (2008) "Decentralized trade, random utility and the evolution of social welfare" *Journal of Economic Theory* 140:328-338.
- [28] Kreindler, GE; and HP Young (2012) "Rapid Innovation Diffusion with Local Interaction," *Oxford Economics Discussion Papers Series*.
- [29] Levine, DK; and S Modica (2015) "Interim and Long-Run Dynamics in the Evolution of Conventions," forthcoming in *Theoretical Economics*.
- [30] Monderer, D; and LS Shapley (1996) "Potential Games," *Games and Economic Behavior* 14:124–143.
- [31] Montanari, A; and A Saberi (2010) "The spread of innovations in social networks," *Proceedings of the National Academy of the Sciences*, 107:20196–20201.
- [32] Myatt, DP and C Wallace (2003) "A multinomial probit model of stochastic evolution," *Journal of Economic Theory*, 113:286-301.
- [33] Myatt, DP; and C Wallace (2008b) "When does one bad apple spoil the barrel? An evolutionary analysis of collective action" *Review of Economic Studies* 75:499-527.
- [34] Myerson, R (1978) "Refinements of the Nash equilibrium concept," *International Journal of Game Theory*, 15:133-154.
- [35] Noeldke, G; and Samuelson, L (1993) "An Evolutionary Analysis of Backward and Forward Induction." *Games and Economic Behavior* 5:425-454.
- [36] Peski, M (2010) "Generalized risk-dominance and asymmetric dynamics" *Journal of Economic Theory* 145:216-248.
- [37] Robson, AJ; and F Vega-Redondo (1996) "Efficient equilibrium selection in evolutionary games with random matching" *Journal of Economic Theory* 70:65-92.
- [38] Rozen, K (2008) "A ‘dual’-improved shortcut to the long run" *Yale University Economics Department Working Paper* No. 41
- [39] Samuelson, L (1994) "Stochastic stability in games with alternative best replies." *Journal of Economic Theory* 65:383-414.
- [40] Sandholm, WH (2007) "Simple formulas for stationary distributions and stochastically stable states" *Games and Economic Behavior*, 59:154-162.
- [41] Sandholm, WH (2010) "Orders of limits for stationary distributions, stochastic dominance, and stochastic stability" *Theoretical Economics*, 5:1–26.

- [42] Troger, T (2002) "Why Sunk Costs Matter for Bargaining Outcomes: An Evolutionary Approach" *Journal of Economic Theory* 102:375-402
- [43] Trygubenko, SA; and DJ Wales (2006) "Graph transformation method for calculating waiting times in Markov chains" *Journal of Chemical Physics*, 124:234110.
- [44] Young, HP (1993a): "The Evolution of Conventions." *Econometrica*, 61:29-56.
- [45] Young, HP (1993b) "An Evolutionary Model of Bargaining." *Journal of Economic Theory*, 59:145-168.
- [46] Young, HP (1998) "Conventional Contracts." *Review of Economic Studies*, 65:776-792.
- [47] Young, HP (2011) "The dynamics of social innovation," *Proceedings of the National Academy of the Sciences*, 108:21285–21291.

Online Appendix for "A characterization of stochastic stability and waiting time" By Kevin Hasker June 26, 2018

A A Discussion of Etymology and Notation

As is typical in any literature, a wide variety of terminology is used. This literature has the additional challenge that it has roots in graph theory, large deviations, and Markov chains. For example, although the term limit set is not precise, it is used by Freidlin and Wentzell (2012) based on the noiseless system. In the continuum, when noise is explicit, there may be other states in the stable sets. Freidlin and Wentzell (2012) call these *compacta* and clarify that each contains a limit set. The theory of Markov chains refers to these sets as *recurrence classes* (Young, 1993a) or *communication classes* (Robles, 1988). Young (1993a) uses the term *absorbing state* when the set is a state.

Likewise, Young (1993a) refers to the primitive in the analysis as resistance; Kandori, Mailath, and Rob (1993) refer to it as a cost; and Beggs (2005) as a *rate*. Alos-Ferrer and Netzer (2010) call it *waste* in the logit model. Kandori and Rob (1995) then define a path-minimizing function as *optimized cost*. Freidlin and Wentzell (2012) use our cost but also no terminology. We have chosen resistance as our primitive and cost for path minimizing because a cost is usually optimized. Hwang and Newton (2014) use our terminology except that they say *overall cost*. In graph theory, the primitive is often cost and, in line with this tradition, Robles (1998) and Pak (2008) call the primitive cost and the optimized function resistance.

Our use of the term iteration resistance instead of modified cost (Ellison, 2000) is fortunate because modified cost is the terminology used in graph theory when applying the standard implementation of Edmonds' algorithm. Levine and Modica (2015) extend this concept to the k 'th order modified radius. Because our m 'th iteration radius has different terminology, there can be no confusion. We thank Rozen (2008) for the $\mathcal{R}(\cdot)$ notation for the radius, which allows no confusion between the radius and resistance.

When deriving the resistance, a standard methodology (Ellison, 2000) is to derive it as the limit of a sequence of matrices. We are confident that our methodology is equivalent. Another normalization is to write $P_\beta(x, y) = \tilde{W}(x, y, \beta) e^{-\beta r(x, y)}$ where $\tilde{W}(x, y, \beta) = W(x, y, \beta) / \sum_{z \in Z} (W(x, z, \beta) e^{-\beta r(x, z)})$. We do not because the $\tilde{W}(x, y, \beta)$ functions must satisfy constraints. Alternatively, we could let $\varepsilon = e^{-\beta}$ and analyze $\varepsilon \rightarrow 0$ instead of $1/\beta \rightarrow 0$. We prefer β because $\varepsilon = 0$ is mathematically possible but outside of the analysis.

Our methodology is not appropriate if the infinite past matters, such as in Robles (1998) and Pak (2008). Of course, our model is fine if only the finite past matters, such as in adaptive play (Young, 1993a).

B Categorization of Applications by Emergent Seed

At first pass, a survey of applications is perplexing. Authors seem to need to innovate novel methodologies with almost every paper. This is a natural result of the stochastic stability paradigm; however, a lot of this apparent complexity is only on the surface. At a deeper level, a great deal of similarity exists and almost all of these papers have an emergent seed that is from one of the three simplest classes—linear, star, or circuit.

In this survey, we first recognize that if one can analytically derive the limiting distribution, then further analysis is unnecessary. Finding the emergent seed can only assist in finding the speed of evolution. We then dispense with problems that were found after the analysis to be unsolvable. Additionally, we will not explicitly analyze any problem in which one or two limit sets exist because these are circuits by definition. This leaves us with 42 analyses, including a novel one in this paper. Of these 42 analyses, we find that 16 have linear emergent seeds, 11 have a star, and seven have a circuit, leaving 11 that do not fit into any of these three simple classes. Note that two applications (Young, 1998; Myatt and Wallace, 2008a) sometimes have a star and otherwise a circuit, and that Jackson and Watts (2002a) can be classified as either linear or a circuit.

Many of these papers implicitly use the emergent seed. This means that the original article identifies the emergent seed, and this information is sufficient to identify stochastic stability. Under this criteria, 35 (83%) implicitly used the emergent seed and, in fact, 40 (95%) could have used the emergent seed. However other methodologies could also be used. For example, 36 (86%) could have used the naive minimization test (Binmore, Samuelson, and Young, 2003).

B.1 Analyses with Linear Emergent Seeds

When linear emergent seeds occur, the analyst is usually quick to recognize and use this fact. For example, Young (1993b) clearly explains that the emergent seed is linear and that hub dominance holds. As well, in the analysis of population games with a summary statistic, Robles (1997) first recognizes that the emergent seed is linear and then uses this to characterize the solution.

Many of these papers are variations of the classic bargaining game, or show how the game under analysis is essentially this game. Saez-Marti and Weibull (1999) consider the impact of smart agents in bargaining. Andreozzi (2012), Dawid and Macleod (2008), Ellingsen and Robles (2002), and Troger (2002) all look at variations on the bargaining problem when there is a priori investment. Robles (2008) considers bargaining when there is an outside option of waiting. Other linear emergent seeds are found by showing that the model is equivalent to bargaining. For example, Agastya (2004) shows that a double auction is equivalent to bargaining; and Agastya (1999) and Newton (2012) both show that coalitional bargaining is equivalent.

Others seem unrelated, but still end up with a linear emergent seed. For Example, Jackson and Watts (2002a) look at network formation when people are playing a coordination game. Agents choose either the strategy or the network. With moderate link formation costs, if at least two are playing each strategy, it is a limit set. A transition only takes one error to the other strategy, and this agent then will choose to join the other network. All that play one or the other strategy require two strategy errors to exit; thus, both are stochastically stable. Note that because all limit sets are in the hub, it is also a circuit; however, we find the linear nature more striking. Anwar (2002) considers the coordination game played in which agents can choose one of two (capacity constrained) islands as well. Here, there are three limit sets, both islands playing the Pareto efficient equilibrium, both playing risk dominant, and each playing a different one. From the intermediate state, the radius is determined by going to one of the extremes and, thus, the emergent seed is linear. A similar structure is found in the evolution of markets in Gerber and Bettzüge (2007). An interesting pair of papers—Naidua, Hwang, and Bowles (2010) and Hwang and Newton (2014)—show that, in the contract game, if one only considers errors that might have a potential benefit, then the emergent seed transforms from a star to linear.

B.2 Analyses with Star Emergent Seeds

The Contract game (Young, 1998) is somewhat unique in the class of star emergent seeds. Only that paper and the volunteer game (Myatt and Wallace, 2008a) lack the *global attractor property*. When a game has the global attractor property, there is a limit set with a maximal radius that also determines the radii of all others. Regarding imitation, Alos-Ferrer and Netzer (2012) point out that this is the same as having a GESS (Schaeffer, 1988). When there is a global attractor, many methodologies will work, such as both the naive minimization test and the radius/coradius test.

The classic example of a star with a global attractor is the pure coordination game (Kandori and Rob, 1995). If all off-diagonal payoffs are constant, it is clear that the Pareto efficient equilibrium will both have the highest radius and determine the radii of all other limit sets.

In imitation, the equivalence of a GESS and stochastic stability has resulted in numerous papers. For example, the seminal analysis of the Cournot game (Vega-Redondo, 1997) is a result of this fact. Tanaka (1999) extends this to when there are low or high costs. Alos-Ferrer, Ania, and Schenk-Hoppe (2000) extend this to a Bertrand competition, although if the number of firms is odd, there will be a set of stochastically stable limit sets. Hohenkamp and Wambach (2010) consider product differentiation.

Interestingly, many papers that analyze the evolution of networks are stars. For Example, Goyal and Vega-Redondo (2004) modify Jackson and Watts (2002a), which changes the emergent seed from linear to a star. As well, Feri (2007) analyzes a model in which direct and indirect links are equally valuable and, again, the emergent seed is a star with a global attractor. Jackson and Watts (2002b) show that there is a star in the coauthor problem.

Dutta and Prasad (2002) analyze risk sharing with moral hazard. This has a global attractor but uses half dominance. This is the only example in which the emergent seed is not actually identified. In the volunteer game, Myatt and Wallace (2008a) identify the emergent seed and show that one agent will always be in the hub; however, that agent may not be stochastically stable.

B.3 Analyses with Circuit Emergent Seeds

If one excludes games with two limit sets (as we have), then circuits are the rarest of the simple emergent seeds. Of course, when they are found they make the analysis transparent. Finding either a circuit or a global attractor makes almost every known method simple to use. However, our more precise results in the Gift Giving game are from our use of the emergent seed.

Interestingly, a common situation exists when they occur. If there is a (class of) *null space invader(s)*, then a circuit automatically follows. A null space invader is always the best invader. Furthermore, any strategy is a best response when society is in that state. For example, in the analysis of state competition (Levine and Modica, 2013), this is the "fanatic band"—a band that is always the strongest when overturning the current regime but that cannot form a stable government itself. After the fanatic band destroys the current government, any viable state can replace it. Hwang and Newton (2014) show that in a Contract game, if disagreement is relevant (for both parties), then the extreme contracts are null space invaders and the emergent seed is a circuit. This is also true for Young (1998) under these conditions. In the volunteer game (Myatt and Wallace, 2008a), if the most likely way to transition from the (most likely) volunteer is to the state in which no one volunteers, then likewise the emergent seed is a circuit. Again, if there is no volunteer,

then everyone is willing to become the volunteer.

In the Gift Giving game (Johnson, Levine, and Pesendorfer; 2000), no null space invader exists; however, there is a limit set to which the best invader is any other limit set. If the selfish strategies are being used, then any cooperative strategy can invade. When switching the language of the strategy, because the invader can be tit-for-tat, we can always transition from any other equilibrium strategy to the selfish strategies. The only other example without a globally optimal invader is the three-action coordination game (Kandori and Rob, 1998). In that paper, the analysis focuses on the conditions under which the emergent seed would be a circuit.

B.4 Analyses with Exotic Emergent Seeds.

We would like to emphasize that 9 of the 11 exotic emergent seeds (82%) could have used the emergent seed to solve the problem. As an interesting example, consider Alos-Ferrer (2004), who applies imitation with short memory to the Cournot game. Any quantity between everyone producing the Cournot quantity (q_c) and the Walrasian quantity (q_w) has maximal radius, and the radius/coradius test establishes that only quantities in this range may be stochastically stable. However, are they all stochastically stable? To answer this question, the paper shows that all of these quantities are in the hub. Going from q_c to q_w is quite simple by having one error be producing more output and a second error be producing less, such that the total quantity (and, hence, the price) is the same. Then, everyone will increase their imitation of error production. The opposite is more difficult, and the paper shows that one can seesaw above q_w and then to below q_w , repeating this process until everyone produces q_c . Thus, without intention, the analyses exactly show that all of these limit sets are in a circuit, and then the result follows by hub dominance.

Indeed, if we wished, we could develop another category that would encompass five of the eleven exotic emergent seeds. We would call this class *short step stars*. To be in this class, there must be a unique limit set in the hub with maximal radius and there must be one iteration to the emergent seed. In this case, one will approach this maximal radius limit set in a series of steps. Ellison (2000) is an example of a short step star. To see this, one simply notes that the limit set for which everyone plays the risk-dominant strategy (θ_A) has maximal radius (for a large enough population), then notes that the radius of every other limit set can be achieved by going to a limit set that has more agents playing A ("up"). It is then quite clear that the modified coradius is two or the radius of the limit set for which everyone plays the risk-dominated strategy (θ_B).

Other analyses with this nature are Fisher and Vega-Redondo (2003)—who analyze imitation in a model of comparative advantage; Alos-Ferrer and Kirchsteiger (2010)—imitation in trading institutions; Kim and Wong (2011)—imitation in an exchange economy; and Zhang, Cui, and Zu (2014)—who analyze free trade networks under a best response dynamic. Even Myatt and Wallace (2008b) is a short step star for a range of parameters. The reason we are not happy with this category is that all circuits are also short step stars. Furthermore, the uniqueness requirement is troubling. However, if we dropped that requirement, any problem with one iteration emergent seeds and hub dominance would be short step stars. This category is intended to show that these analyses are similar to ones with the global attractor property. Because there are multiple stochastically stable limit sets, we cannot include the analysis in Ben-Shoham, Serrano, and Volij (2004), who consider simple BRM and housing allocations. That analysis implicitly finds the emergent seed.

Of the two examples that cannot be solved using the emergent seed technique, we could argue that Bergin and Bernhardt (2009) truly find a robustness result, but we prefer not to. This result applies to imitation with long memories in any symmetric game. The stochastically stable limit set will be the joint utility maximizing strategy, and with long enough memory making agents forget that this great payoff will be too costly.

The key element of Ben-Shoham, Serrano, and Volij (2004) (rank-based errors) is reversible paths. One should recognize that proving the reversible paths property and characterizing the difference between costs root switching will be superior. One hypothesizes that θ is stochastically stable and then switch the root with some $\tilde{\theta}$ that has a negative difference in costs. This allows precise identification of stochastic stability without knowing anything about the emergent seed.

C Proofs and Auxiliary Results

In this section we include minor proofs and establish two results that are referred to in the main text. Lemma 14 extends the argument in Binmore, Samuelson, and Young (2003) from the radius/coradius test to the radius/modified coradius test. Proposition 4 is of value only in establishing that there is a dynamic argument behind the censored coradius. Proofs will be presented in the order of the corresponding statement.

Using the cost function the modified coradius is $CR^*(X) = \max_{\tilde{\theta} \in \Theta} \{ \Delta c(X, \tilde{\theta}) + \mathcal{R}(\tilde{\theta}) \}$. Define the limit set transitioned to after X in a graph G as $\theta(X, G)$.

Lemma 14 *Using root switching one can prove that if $\mathcal{R}(X) > CR^*(X)$ then no $z \in Z \setminus X$ may be stochastically stable.*

Proof. Let $\tilde{\theta} \in \Theta$ be arbitrary, and T^* be a tree with root $\tilde{\theta}$ that achieve its stochastic potential. We wish to switch the root of this tree with X and prove that $sp(X) < sp(\tilde{\theta})$. This means there must be a $Q \in \{Q(\tilde{\theta}, X)\}$ such that the change in cost is strictly negative, or:

$$\sum_{\theta' \cap Q \neq 0} c(\theta', \theta(\theta', Q)) - \left[\sum_{\theta' \subseteq \mathcal{D}(X)} c(\theta', \theta(\theta', T^*)) + \sum_{\theta' \cap Q \neq 0} c(\theta', \theta(\theta', T^*)) \right] < 0$$

Now clearly the term in brackets is higher than:

$$\mathcal{R}(X) + \sum_{\theta' \cap Q \neq 0} \mathcal{R}(\theta') - \mathcal{R}(\tilde{\theta}) ,$$

where the special treatment for $\tilde{\theta}$ is because in T^* it did not have a transition. Thus we have:

$$\sum_{\theta' \cap Q \neq 0} [c(\theta', \theta(\theta', Q)) - \mathcal{R}(\theta')] + \mathcal{R}(\tilde{\theta}) < \mathcal{R}(X)$$

and now we can minimize the left hand side over paths, giving us:

$$\Delta c(X, \tilde{\theta}) + \mathcal{R}(\tilde{\theta}) < \mathcal{R}(X) ,$$

if this is true for all $\tilde{\theta} \in \Theta$ then the statement is proven. ■

Proof of Lemma 1. We first prove the equivalence of the two characterizations. If $\theta = \bigcup_{Q \in \{\underline{Q}(x,x)\}} \left[\bigcup_{z \in Z} Q(z, \cdot) \right]$

then every $y \in \theta$ is an element of some $Q \in \{\underline{Q}(x,x)\}$, thus $c(x,y) = 0$. If for all $y \in \theta$ $c(x,y) = 0$ we know there is a zero resistance path from x to y , since x was arbitrary there is also one from y to x , and thus y is in some $Q \in \{\underline{Q}(x,x)\}$. Likewise if for all $z \in Z \setminus \theta$, $\{\underline{Q}(x,z)\} = \emptyset$ this means $c(x,z) > 0$ and if for all $z \in Z \setminus \theta$ $c(x,z) > 0$ this means $\{\underline{Q}(x,z)\} = \emptyset$. A transition from x to y will occur with a probability on the order of $\exp[-\beta c(x,y)]$ thus if $c(x,y) = 0$ and we are currently in the state x the probability of transitioning to y is increasing in β . Thus if $x \in \theta$ then we must have $y \in \theta$, thus θ is minimal. Furthermore if $c(x,z) > 0$ then this transition occurs with vanishing probability as β gets large. This establishes that for fixed s $\Pr(z_{t+s} \notin \theta | z_t \in \theta) \rightarrow 0$ as $\beta \rightarrow \infty$. ■

Proof of Lemma 2. We need $(u_1(\theta_1) - u_1(\theta_1 - \delta)) / u_1(\theta_1) < (u_2(1 - \theta_1) - u_2(1 - \theta_1 - \delta)) / u_2(1 - \theta_1)$, and the Condition 12 is a simple rearrangement of this condition. ■

Proof of Lemma 3. The key question is when is it better to take two steps of size δ instead of one large one of size 2δ . Consider two steps going up, then this will be true if:

$$\begin{aligned} & \left[\frac{u_2(1 - \theta_1) - u_2(1 - \theta_1 - \delta)}{u_2(1 - \theta_1)} - \frac{u_1(\theta_1) - u_1(\theta_1 - \delta)}{u_1(\theta_1)} \right] + \left[\frac{u_2(1 - \theta_1 - \delta) - u_2(1 - \theta_1 - 2\delta)}{u_2(1 - \theta_1 - \delta)} - \frac{u_1(\theta_1 + \delta) - u_1(\theta_1)}{u_1(\theta_1 + \delta)} \right] \\ & \leq \left[\frac{u_2(1 - \theta_1) - u_2(1 - \theta_1 - 2\delta)}{u_2(1 - \theta_1)} - \frac{u_1(\theta_1) - u_1(\theta_1 - \delta)}{u_1(\theta_1)} \right]. \end{aligned} \quad (41)$$

Notice that if we take two short steps we subtract two radii, if we only take one long one then we only subtract one. After some algebra this becomes the condition:

$$\frac{u_2(1 - \theta_1 - \delta) - u_2(1 - \theta_1 - 2\delta)}{u_2(1 - \theta_1 - \delta)} \frac{u_2(1 - \theta_1) - u_2(1 - \theta_1 - \delta)}{u_2(1 - \theta_1)} \leq \frac{u_1(\theta_1 + \delta) - u_1(\theta_1)}{u_1(\theta_1 + \delta)}. \quad (42)$$

Let $\theta = (\theta_1 + \delta, 1 - \theta_1 - \delta)$ then this is equivalent to $p_2^+(\theta) p_2^+(\theta_1 - \delta, 1 - \theta_1 + \delta) \leq p_1^+(\theta)$. We note that $p_1^+(\theta) < p_2^+(\theta)$ but since as $\delta \rightarrow 0$ $\max\{p_1^+(\theta), p_2^+(\theta)\} \rightarrow 0$ this condition will be satisfied for small δ . The condition in the lemma is a sufficient generalization. ■

Proposition 4 For $x \in Z$ $\ln \tau(x) \leq \overline{CR}(x)$.

Proof. We begin with the Markov transition matrix, P_β , where $P_\beta(x,y)$ is defined in Equation 2. Since we are interested in the case where $\beta \rightarrow \infty$ and $t \rightarrow \infty$ we can replace this with a transition matrix based on c , denoted $P_{\beta,c}$, and then we can define a new matrix on Θ by letting $P_{\beta,c}(\theta, \tilde{\theta}) = \sum_{x \in \bar{\mathcal{D}}(\theta)} \sum_{y \in \bar{\mathcal{D}}(\tilde{\theta})} P_{\beta,c}(x, y)$. Now Bortz, Kalos and Lebowitz (1975) shows we can construct a new transition matrix where $\tilde{P}_{\beta,c}(\theta, \tilde{\theta}) = P_{\beta,c}(\theta, \tilde{\theta}) / (1 - P_{\beta,c}(\theta, \theta))$, and from Ellison (2000) we know that $(1 - P_{\beta,c}(\theta, \theta))$ is on the order of $e^{-\beta \mathcal{R}(\theta)}$. Thus since we are interested in large β , without loss of generality we can let $\hat{P}_{\beta,c}(\theta, \tilde{\theta}) = e^{-\beta(c(\theta, \tilde{\theta}) - \mathcal{R}(\theta))}$. At this point we can iterate the process with $\hat{P}_{\beta,c}$ as our initial matrix and $\Delta r(\theta, \tilde{\theta}) = c(\theta, \tilde{\theta}) - \mathcal{R}(\theta)$ as our resistance function.

For a given initial $\tilde{\theta}$ this tells us that we can characterize the path from $\tilde{\theta}$ to θ as two sub-sequences. One goes to the hub, $(\tilde{\theta}^m)_{m=0}^M$ where $\tilde{\theta}^0 = \tilde{\theta}$ and $\theta^M = \theta^M$, and one that comes from the hub $(\theta^m)_{m=M}^0$, where $\theta^0 = \theta$ and $\theta^M = \theta^M$. The former must have the cost $\sum_{m=0}^{M-1} \Delta^m \mathcal{R}(\theta^m(\tilde{\theta}))$ by likelihood maximization.

The latter has the cost $\Delta^M c(\theta^M, \theta)$. Thus an estimate for the likelihood for the transition from $\tilde{\theta}$ to θ is the censored coradius conditional on beginning at $\tilde{\theta}$, and as $\beta \rightarrow \infty$ the maximum such waiting time is the critical one. ■

Proof of Lemma 4. One immediately notes that $\sum_{m=0}^{M-1} \Delta^m \mathcal{R}(\theta^m(\tilde{\theta})) \geq \sum_{m=0}^{m(\theta, \tilde{\theta})-1} \Delta^m \mathcal{R}(\tilde{\theta})$ and since $\Delta^M c(\theta^M, \tilde{\theta}) \geq 0$ the inequality follows. If $\tilde{\theta}$ is in the hub and $m(\theta, \tilde{\theta}) = M$, then $\sum_{m=0}^{M-1} \Delta^m \mathcal{R}(\theta^m(\tilde{\theta})) = \sum_{m=0}^{m(\theta, \tilde{\theta})-1} \Delta^m \mathcal{R}(\tilde{\theta})$ and $\Delta^M c(\theta^M, \tilde{\theta}) = 0$, proving the sufficient condition. ■

Proof of Lemma ??. Note that both $\min(p_1^-(\theta), p_2^-(\theta))$ and $\min(p_1^+(\theta), p_2^+(\theta))$ are first increasing and then decreasing, and that if θ is a limit set then for small enough δ $\min(p_1^-(\theta), p_2^-(\theta)) > \min(p_1^+(\theta), p_2^+(\theta))$. If $\min(p_1^-(\theta), p_2^-(\theta)) > \min(p_1^+(\theta), p_2^+(\theta))$ then we go from θ to $(\theta_1 + \delta, 1 - \theta_1 - \delta)$ if $\theta_1 < \gamma_{NBS}$ and to $(\theta_1 - \delta, 1 - \theta_1 + \delta)$ if $\theta_1 > \gamma_{NBS}$. This means that the limit set transitioned to has a higher radius.

Of course it is possible that for a given δ and θ that $\min(p_1^-(\theta), p_2^-(\theta)) \leq \min(p_1^+(\theta), p_2^+(\theta))$. If the state transitioned to is not a limit set then we can transition to any other limit set with positive probability. If the state is then as mentioned above for small enough δ $\min(p_1^-(\theta), p_2^-(\theta)) > \min(p_1^+(\theta), p_2^+(\theta))$. Thus even if some limit sets go to an extreme solution either from that extreme we take small steps back towards γ_{NBS} or we can go to the hub in one step.

Finally notice that $\min(p_1^+(\theta), p_2^+(\theta)) \geq \mathcal{R}(\theta)/n$ by definition. Thus δ must be small enough to satisfy two conditions. First we must have $\min(p_1^-(\theta), p_2^-(\theta)) > \min(p_1^+(\theta), p_2^+(\theta))$ for limit sets in the hub—or $\underline{\gamma} < \gamma_{NBS} - \delta < \gamma_{NBS} + \delta < \bar{\gamma}$. Second if the extreme contracts ($\theta \in \{(0, 1), (1, 0)\}$) are strict equilibria then we must also be certain that $\min(p_1^-(\hat{\theta}), p_2^-(\hat{\theta})) > \min(p_1^+(\hat{\theta}), p_2^+(\hat{\theta}))$. ■

Proof of Lemma 5. We point out that the likelihood potential is not a function of δ , thus we proceed by analyzing allowing $\theta_1 \in [0, 1]$. Normalizing $u_i(1) = 1$ the likelihood potential is:

$$lp(\theta)/n \in \begin{cases} \frac{u_1(\theta_1)}{u_1(\theta_1)+1} - \frac{u_2(0)}{u_2(0)+u_2(1-\theta_1)} + \frac{u_2(0)}{u_2(0)+1} & (1) \\ \frac{u_1(\theta_1)}{u_1(\theta_1)+1} - \frac{u_1(0)}{u_1(0)+u_1(\theta_1)} + \frac{u_1(0)}{u_1(0)+1} & (2) \\ \frac{u_2(1-\theta_1)}{u_2(1-\theta_1)+1} - \frac{u_2(0)}{u_2(0)+u_2(1-\theta_1)} + \frac{u_2(0)}{u_2(0)+1} & (3) \\ \frac{u_2(1-\theta_1)}{u_2(1-\theta_1)+1} - \frac{u_1(0)}{u_1(0)+u_1(\theta_1)} + \frac{u_1(0)}{u_1(0)+1} & (4) \end{cases} . \quad (43)$$

We then notice that if at θ either case (2) or (3) holds then θ can not be stochastically stable because the function is respectively strictly increasing or strictly decreasing in θ . In cases (1) and (4) one can show the objective function is strictly concave. Let us find the first derivatives of $lp(\theta)/n$ in these cases:

$$\frac{\partial lp(\theta)/n}{\partial \theta_1} \in \begin{cases} u'_1(\theta_1) \frac{1}{(u_1(\theta_1)+1)^2} - u'_2(1-\theta_1) \frac{u_2(0)}{(u_2(0)+u_2(1-\theta_1))^2} & (1) \\ -u'_2(1-\theta_1) \frac{1}{(u_2(1-\theta_1)+1)^2} + u'_1(\theta_1) \frac{u_1(0)}{(u_1(0)+u_1(\theta_1))^2} & (4) \end{cases} \quad (44)$$

Solving these in general is impossible because we have no restriction on $u'_1(\theta_1)/u'_2(1-\theta_1)$, but if we apply symmetry then at the Kalai-Smorodinsky solution ($\gamma_{KS} = \frac{1}{2}$) we have:

$$\frac{\partial lp(\gamma_{KS})/n}{\partial \theta_1} \in \begin{cases} u'(\frac{1}{2}) \left(\frac{1}{(u(\frac{1}{2})+1)^2} - \frac{u(0)}{(u(0)+u(\frac{1}{2}))^2} \right) & (1) \\ u'(\frac{1}{2}) \left(-\frac{1}{(u(\frac{1}{2})+1)^2} + \frac{u(0)}{(u(0)+u(\frac{1}{2}))^2} \right) & (4) \end{cases}$$

Now since $u(0) < u(x) < 1$, $\frac{1}{(u(x)+1)^2} > \frac{(0)}{(u(0)+u(x))^2}$ thus $\frac{\partial \ln(\gamma_{KS})/n}{\partial \theta_1} > 0$ for all $\theta_1 \leq \gamma_{KS}$ and $\frac{\partial \ln(\gamma_{KS})/n}{\partial \theta_1} < 0$ for all $\theta_1 \geq \gamma_{KS}$, thus the stochastically stable limit set is one of the closest to γ_{KS} in $A_1(\delta)$. ■

Proof of Lemma 6. Using the normalization notice that

$$\Delta c(\theta^M, \theta) = \min \left[\frac{\beta_2}{\beta_2+1} \frac{1 - [(1-\beta_2)v_2(1-\theta_1)+\beta_2]}{[(1-\beta_2)v_2(1-\theta_1)+\beta_2+\beta_2]}, \frac{\beta_1}{\beta_1+1} \frac{1 - [(1-\beta_1)v_1(1-\theta_1)+\beta_1]}{[(1-\beta_1)v_1(1-\theta_1)+\beta_1+\beta_1]} \right] \quad (45)$$

and

$$\mathcal{R}(\theta) = \min \left[\frac{\frac{(1-\beta_1)v_1(\theta_1)+\beta_1}{(1-\beta_1)v_1(\theta_1)+\beta_1+1},}{\frac{(1-\beta_2)v_2(\theta_2)+\beta_2}{(1-\beta_2)v_2(\theta_2)+\beta_2+1}} \right]. \quad (46)$$

This means that as $\min[\beta_1, \beta_2] \rightarrow 0$ $\Delta c(\theta^M, \theta) \rightarrow 0$ but $\mathcal{R}(\theta) \rightarrow \min \left[\frac{v_1(\theta_1)}{v_1(\theta_1)+1}, \frac{v_2(\theta_2)}{v_2(\theta_2)+1} \right] > 0$, thus for small enough $\min[\beta_1, \beta_2]$ the solution will be near the Kalai-Smorodinsky solution, which maximizes the radius. ■

Proof of Lemma 7. First we establish that the only strict equilibria of the game with noise are {selfish, team, weak team, insider}. If $\Pr(f_s = r) > 0$ then the strict best response to the tit-for-tat transition rule is the generous strategy ($a(g) = a(r) = 1$). And the unique best response to generous is selfish. This is also the unique best response when $\tau(g, 1) = r$ or $\tau(g, 1) = \tau(g, 0)$ because you either will always or never be rewarded for giving the gift.

To establish that the other strategies are strict equilibria with small η we first have to consider the relevant states. The states in this model are the social status a player will have with one transition rule and the status they would have with a different one. Let f_s be the element of f associated with strategy s , and $f_{s'}$ be the same for s' . Then define $\Pr(f_s, f_{s'})$ as the probability $f_s \in \{r, g\}$ and $f_{s'} \in \{r', g'\}$. The alternative strategy will always be a strict equilibrium, so when we write (s, s') both strategies are the team, weak team, or insider strategy and $s \neq s'$. When the alternative strategy is selfish we will write "selfish." The payoffs are:

$$v(s', s') = \Pr(g, g') \left(\left(1 - \frac{\eta}{2}\right) \alpha - 1 \right) + \Pr(r, g') \left(\left(1 - \frac{\eta}{2}\right) \alpha - 1 \right) + \Pr(g, r') \Pr(1|s', r', 0) \alpha + \Pr(r, r') \Pr(1|s', r', 0) \alpha \quad (47)$$

$$v(s, s') = \Pr(g, g') \left(\left(1 - \frac{\eta}{2}\right) \alpha - 1 \right) + \Pr(r, g') \frac{\eta}{2} \alpha + \Pr(g, r') (\Pr(1|s', r', 1) \alpha - 1) + \Pr(r, r') \Pr(1|s', r', 0) \alpha \quad (48)$$

$$v(\text{selfish}, s') = \Pr(g, g') \frac{\eta}{2} \alpha + \Pr(r, g') \frac{\eta}{2} \alpha + \Pr(g, r') \Pr(1|s', r', 0) \alpha + \Pr(r, r') \Pr(1|s', r', 0) \alpha \quad (49)$$

Where $\Pr(1|s', f_{s'}, a) \in \left\{ \frac{\eta}{2}, 1 - \frac{\eta}{2} \right\}$ is the probability of receiving the gift given the strategy s' , the color of the flag, $f_{s'} \in \{r', g'\}$, and the action of the agent $a \in \{0, 1\}$. Clearly $v(\text{selfish}, s') < v(s', s')$ if $\left(1 - \frac{\eta}{2}\right) \alpha - 1 > \frac{\eta}{2} \alpha$, or η is small enough. For $v(s, s') < v(s', s')$ we must have $\Pr(r, g') > \Pr(g, r')$, and the ratio is large enough. Since everyone is following the strategy s' this means that $\Pr(r, g')$ is on the order of $1 - \eta$ and $\Pr(g, r')$ is on the order of η , thus as long as η is small enough we are fine.

Now we turn to the task of finding the optimal invaders, s_I and let $s' = s_I$. Let p be the probability of the invader in this strategy, and let $v(s, p)$ be the expected utility of using strategy s . First we notice that if $\Pr(f_s)$ is the probability that given s $f_s \in \{r, g\}$ occurs it is obvious that:

$$v(s, p) = \Pr(g) v(s, p|g) + \Pr(r) v(s, p|r) \geq \min[v(s, p|g), v(s, p|r)]. \quad (50)$$

And if the right hand side is low enough (compared to another strategy) then one of the actions for the strategy s is no longer optimal. Thus we should minimize either $v(s, s_I|g)$ or $v(s, s_I|r)$, and we need either $a(g) = 0$, or $a(r) = 1$ to be optimal.

If we need $a(g) = 0$ to be optimal then the new strategy will be selfish. Thus the critical probability is:

$$\begin{aligned} v(s, p|g) &= -1 + (1-p) \left(1 - \frac{\eta}{2}\right) \alpha \leq (1-p) \frac{\eta}{2} \alpha = v(\text{selfish}, p|g) . \\ p &= 1 - 1/\alpha(1-\eta) . \end{aligned} \quad (51)$$

Now assume that we need $a(r) = 1$ to be optimal, or we minimize $v(s, p|r)$. The selfish strategies do not give an incentive for $a(r) > 0$, thus we need either $-s$, $-w$, $-i$, or $-t$ for tat. Where $-s$ is the strategy that treats red as good—the language has changed. Next notice that if (r, g') occurs then both strategies will call for the same action and this can not affect the choice of strategy. Thus what we care about is $v(s, p|r, r')$.

$$\begin{aligned} v(s, p|r, r') &= (1-p) \Pr(1|s, r, 0) \alpha + p \Pr(1|s', r', 0) \alpha \\ v(s', p|r, r') &= -1 + (1-p) \Pr(1|s, r, 1) \alpha + p \Pr(1|s', r', 1) \alpha \end{aligned} \quad (52)$$

Thus we are looking for the critical s' such that for the minimal p :

$$\begin{aligned} v(s, p|r, r') &\leq v(s', p|r, r') \\ (1-p) \Pr(1|s, r, 0) \alpha + p \Pr(1|s', r', 0) \alpha &\leq -1 + (1-p) \Pr(1|s, r, 1) \alpha + p \Pr(1|s', r', 1) \alpha \\ 1 + (1-p) \alpha (\Pr(1|s, r, 0) - \Pr(1|s, r, 1)) &\leq p \alpha (\Pr(1|s', r', 1) - \Pr(1|s', r', 0)) \end{aligned} \quad (53)$$

and we see the choice of s' does not matter, for all of them $\Pr(1|s', r', 1) - \Pr(1|s', r', 0) = 1 - \eta$. For both the insider and the weak team strategy $\Pr(1|s, r, 0) - \Pr(1|s, r, 1) = 0$. Thus for these two equilibria the critical p is:

$$p = 1/\alpha(1-\eta) . \quad (54)$$

For the team strategy $\Pr(1|s, r, 0) - \Pr(1|s, r, 1) = 1 - \eta$ so:

$$p = 1/2 + 1/[2\alpha(1-\eta)] . \quad (55)$$

Now we turn to the selfish limit sets. Like before it doesn't matter if the alternative social status is red. All of our equilibrium strategies are selfish in this state. But then we notice that all of the other equilibrium strategies react in the same way when the social status is green. Thus:

$$\begin{aligned} v(\text{selfish}, p|g') &= p \alpha \frac{\eta}{2} \leq -1 + p \left(1 - \frac{\eta}{2}\right) \alpha = v(s', p|g') \\ p &= 1/\alpha(1-\eta) . \end{aligned} \quad (56)$$

We have derived the radii of all the limit sets. ■

Proof of Lemma 8. Notice that if the optimal transition is to $-s$ then it can, without loss of generality, be $-t$ for tat, which is in the basin of attraction of the selfish strategies. Thus we can transition to the selfish strategies in any transition. From the selfish strategy we can transition to any cooperative strategy

with equal likelihood, thus all limit sets are in the hub and the stochastically stable limit set is the one with maximal radius, and the result follows from simple analysis of the radii. ■

Proof of Lemma 9. If this is not true then $\mathcal{R}(\theta) = c(\theta, \tilde{\theta})$ for $\tilde{\theta} \in \Theta \setminus (\Theta_+(\theta) \cup \Theta_-(\theta) \cup \theta)$. Let $x_+ = \tilde{\theta} \setminus \theta$, $x_- = \theta \setminus \tilde{\theta}$. Now for all $i \in I$, $\#(i, \theta \cup x_+) \geq \#(i, \tilde{\theta})$ and likewise $\#(i, \theta \setminus x_-) \leq \#(i, \tilde{\theta})$, but this means $\theta \cup x_+$ is contained in a limit set in $\Theta_+(\theta)$ and respectively for $\theta \setminus x_-$ in $\Theta_-(\theta)$. One must have a weakly lower cost and the claim is established. ■

Proof of Lemma 10. First consider a $box(3)$, for $i \in box(3)$ $\#(i, box(3)) \geq 3$ and since every $j \notin box(3)$ has at most one neighbor in a $box(3)$, $\#(j, box(3)) \leq 2$ thus this is a strict Nash equilibrium and a limit set. Clearly to satisfy the first characteristic the box can not be any smaller. On the other hand consider an θ which is $I \setminus box(3, n)$. We notice for any $j \notin I \setminus box(3, n)$ can only have three neighbors in a $box(3)$, and by induction this implies that j must be in a $box(3, n)$ of agents all playing B . Since all $i \in I \setminus box(3, n)$ can have at most one neighbor in a $box(3, n)$ we can be sure that $\#(i, I \setminus box(3, n)) \geq 3$. Finally we have to check other geometric objects. The only restriction we have placed is that there must be at least two agents in every direction, thus we have to check planes. A two dimensional plane has four neighbors for every agent in it, but everyone not in the plane has at most one neighbor in the plane, thus in both cases a plane would work. However if $n^2 > 8$ there will be more agents in the plane, and $n^2 > 4n$ there will be fewer agents not in the plane. This clearly requires $n > 4$, which we have already assumed. ■

Proof of Lemma 11. By construction if we append the correct $box(d(\theta))$ then the best response of everyone in that box will be A . Since θ is orbicular none of the agents in $box(d(\theta))$ are already playing A , and since θ is small all of the agents playing B need at least $d(\theta)$ of their neighbors to switch before they will change their best response. This implies that they all have $3 - d(\theta)$ neighbors playing A , thus if half of them error the other half will switch by best response, deriving our first formula. Likewise for small θ the set of agents we need to remove will be in a line, and since θ is convex there is no shorter segment we can remove to reach a smaller limit set. If we have every other agent in that line error to B then the rest will switch by best response, thus we derive our second formula.

Now for θ_A we actually have to change the strategy of a $box(3, n)$, for any pair of adjacent lines this will require n errors, and for the four lines that make up a $box(3, n)$ it will require $2n$. ■

Proof of Proposition 3. By comparing $c(\Theta_+(\theta), \theta)$ and $c(\Theta_-(\theta), \theta)$ we realize that we can go up from a given limit set any time if $d(\theta) = 1$ and if $d(\theta) = 2$ any time $l(\theta) \geq 4$. We may go down from a limit set any time $l(\theta) \leq 3$. Now assume that $d(\theta) = 2$ and $l(\theta) \geq 4$. We will go up from this limit set by appending a $box(2)$, since the length of this box is two the result will be a limit set $\tilde{\theta}$ with $d(\tilde{\theta}) = 1$, and we may continue to go up. Thus from any $box(3, 4, 4, 4)$ the first iteration of the emergent seed will connect us to θ_A at zero first iteration resistance. Now consider a $\theta \in box(3, 4, 4, 3)$, when we go to $\tilde{\theta} \subset \theta$ we will not increase the length, thus in the first iteration of the emergent seed we can go from this limit set to θ_B . Now for θ which is a $box(3)$, $c(\Theta_-(\theta), \theta) = 1 < 2 = c(\Theta_+(\theta), \theta)$ thus θ_B is in a first iteration limit set. Likewise for θ which is $I \setminus box(3, n)$ $c(\Theta_+(\theta), \theta) = 1 < n = c(\Theta_-(\theta), \theta)$ thus θ_A is in a first iteration limit set. Since there are only two first iteration limit sets our analysis of the emergent seed is done.

Finally we establish the log-waiting time. Notice that since θ_A is in the hub for any $\tilde{\theta}$, the censored

coradius conditional on starting at $\tilde{\theta}$ — $\overline{CR}(\tilde{\theta}, \theta_A)$ —is:

$$\overline{CR}(\tilde{\theta}, \theta_A) = \begin{cases} \mathcal{R}(\tilde{\theta}) + \Delta\mathcal{R}(\theta^1(\tilde{\theta})) & \text{if } \tilde{\theta} \in \Delta\mathcal{D}(\theta^1(\theta_B)) \\ \mathcal{R}(\tilde{\theta}) & \text{else} \end{cases}. \quad (57)$$

We then notice that for $\theta \in \Theta \setminus \theta_A$ $\mathcal{R}(\tilde{\theta}) \leq c(\Theta_+(\tilde{\theta}), \tilde{\theta}) < \mathcal{R}(\theta_B)$. Thus the censored coradius is $\overline{CR}(\theta_A) = \mathcal{R}(\theta_B) + \Delta\mathcal{R}(\theta^1(\theta_B)) = Ch(\theta_B, \theta_A) = Ch(\theta_A)$. The second equality is due to the fact that $lp_+(\theta_A|\theta_B) = lp(\theta_A)$ and $lp_+(\theta_B|\theta_A) = lp(\theta_B)$, the third because for any $\tilde{\theta}$ and θ : $Ch(\tilde{\theta}, \theta) \leq \overline{CR}(\tilde{\theta}, \theta)$.

Thus we need to find the censored coradius of θ_B . First we must exit θ_B , this costs $\mathcal{R}(\theta_B) = 2^{3-1}$. Next we have to calculate the first iteration cost of getting to a *box* (3, 4, 4, 4). Notice that each step in this path has the cost of $\max(c(\Theta_+(\theta), \theta) - c(\Theta_-(\theta), \theta), 0)$. Since $c(\Theta_+(\theta), \theta) \in \{1, 2\}$ the only relevant cases are when $c(\Theta_+(\theta), \theta) = 2$ or $d(\theta) = 2$, when this is true $c(\Theta_-(\theta), \theta) = 1$ and the iteration cost is one. One can easily see that there are six such steps, thus the total cost of these steps is six, and the total cost is ten.

■

Lemma 15 of Lemma 12. *In a least cost path, every transition must attain the radius. We note that if the path contains θ_\emptyset , then the next transition can be to $\hat{\theta}_l$. Thus, if $\min_{i \in \theta} \ln \frac{1}{d_i} < \min_{j \in I \setminus \hat{\theta}_l} \ln \frac{1}{b_j}$ at any point in a path, then we can go to $\hat{\theta}_l$ with one more transition. Thus, assume that we do not do this. Then, at every step, we must add $j^*(\theta) = \min\{j | j \in I \setminus \theta\}$. To make the path non-cyclic, we should remove $\hat{i}(\theta) = \max_{i \in \theta} i$. If we do not transition to θ_\emptyset at some point in this path, then we must arrive at $\hat{\theta}_l$ in a finite number of steps.* ■

of Lemma 13. For all $\theta \in \Theta \setminus \{\hat{\theta}_l, \theta_\emptyset\}$, $\ln \frac{1}{b_{l+1}} > \mathcal{R}(\theta)$ is immediate because we must add $j^*(\theta) = \min_{j \in I \setminus \theta} j$, thus $\ln \frac{1}{b_{l+1}} > \min_{j \in I \setminus \theta} \ln \frac{1}{b_j} \geq \mathcal{R}(\theta)$. The latter statement is because for such a θ , $\mathcal{R}(\theta) \leq \min_{j \in I \setminus \theta} \ln \frac{1}{b_j} = \min_{j \in z_+ \setminus \theta} \ln \frac{1}{b_j} < \sum_{j \in z_+} \ln \frac{1}{b_i} = \mathcal{R}(\theta_\emptyset)$.

The conclusion is reached by considering the two cases. First, if $\mathcal{R}(\hat{\theta}_l) = \ln \frac{1}{b_{l+1}}$, then $\hat{\theta}_l$ has strictly higher radius than any $\theta \in \Theta \setminus \theta_\emptyset$, and none of them can be stochastically stable by hub dominance. Second, if $\mathcal{R}(\hat{\theta}_l) < \ln \frac{1}{b_{l+1}}$, then we can transition from $\hat{\theta}_l$ to θ_\emptyset . θ_\emptyset is in the hub and only a $\hat{\theta}_k$ can be stochastically stable by hub dominance. ■

References

- [1] Agastya, M (1999) "Perturbed adaptive dynamics in coalition form games" *Journal of Economic Theory* 89:207-233.
- [2] Agastya, M (2004) "Stochastic stability in a double auction" *Games and Economic Behavior* 48:203–222.
- [3] Alos-Ferrer, C (2004) "Cournot versus Walras in dynamic oligopolies with memory," *International Journal of Industrial Organization*, 22:193-217
- [4] Alos-Ferrer, C; Ania, AB; and KR Schenk-Hoppe (2000) "An Evolutionary Model of Bertrand Oligopoly" *Games and Economic Behavior* 33, 1-19.
- [5] Alos-Ferrer, C; and G Kirchsteiger (2010) "General equilibrium and the emergence of (non)market clearing trading institutions," *Economic Theory*, 44:339-360.

- [6] Alos-Ferrer, C and N Netzer (2010) "The logit-response dynamics" *Games and Economic Behavior* 68:413-427.
- [7] Andreozzi, L (2012) "Property rights and investments: An evolutionary approach" *Games and Economic Behavior* 74:1-11.
- [8] Anwar, AW (2002) "On the co-existence of conventions" *Journal of Economic Theory*, 107:145-155.
- [9] Beggs, AW (2005) "Waiting times and equilibrium selection" *Economic Theory*, 25:599-628.
- [10] Ben-Shoham, A; Serrano, R; and O. Volij (2004) "The evolution of exchange" *Journal of Economic Theory* 114:310-328.
- [11] Bortz, AB; Kalos, MH; and JL Lebowitz (1975) "A New Algorithm for Monte Carlo Simulation of Ising Spin Systems" *Journal of Computational Physics*, 17:10-18.
- [12] Dawid, H; and WB Macleod (2008) "Hold-up and the evolution of investment and bargaining norms" *Games and Economic Behavior* 62:26-52.
- [13] Dutta, J; and K Prasad (2002) "Stable risk-sharing" *Journal of Mathematical Economics* 38:411-439.
- [14] Ellingsen, T; and J Robles (2002) "Does Evolution Solve the Hold-Up Problem?" *Games and Economic Behavior* 39:28-53.
- [15] Ellison, G (2000) "Basins of Attraction, Long-Run Stochastic Stability, and the Speed of Step-by-Step Evolution." *Review of Economic Studies*, 67:1-23.
- [16] Feri, F (2007) "Stochastic stability in networks with decay" *Journal of Economic Theory* 135:442-457.
- [17] Fisher, EO and F Vega-Redondo (2003) "The Evolution of Comparative Advantage" *working paper*
- [18] Freidlin, M; and A Wentzell (2012) *Random Perturbations of Dynamical Systems* Springer Verlag, Berlin Heidelberg *third edition*.
- [19] Gerber, A; and MO Bettzüge (2007) "Evolutionary choice of markets" *Economic Theory* 30:453-472.
- [20] Goyal, S; and F Vega-Redondo (2004) "Network Formation and Social Coordination," *Games and Economic Behavior*, 50:178-207.
- [21] Hehenkamp, B; and A Wambach (2010) "Survival at the center-The stability of minimum differentiation" *Journal of Economic Behavior and Organization*, 76:853-858
- [22] Hwang, SH; and J Newton (2014) "A Classification of Bargaining Solutions by Evolutionary Origin" *U. of Sidney Economics Working Paper 2014-02*.
- [23] Jackson, MO; and A Watts (2002a) "On the formation of interaction networks in social coordination games," *Games and Economic Behavior*, 41:265-291
- [24] Jackson, MO; and A Watts (2002b) "The Evolution of Social and Economic Networks." *Journal of Economic Theory*, 106, 265-295.
- [25] Kandori, M; Mailath, G; and R Rob. (1993): "Learning Mutation, and Long Run Equilibria in Games." *Econometrica*, 61:9-56.
- [26] Kandori, M; and R Rob. (1995): "Evolution of Equilibria in the Long Run: A General Theory and Applications." *Journal of Economic Theory*, 65:383-414.
- [27] Kim, C; and KC Wong (2011) "Evolution of Walrasian equilibrium in an exchange economy" *Journal of Evolutionary Economics* 21:619-647.
- [28] Levine, DK; and S Modica (2013) "Conflict, Evolution, Hegemony, and the Power of the State," *working paper*.

- [29] Levine, DK; and S Modica (2015) "Interim and Long-Run Dynamics in the Evolution of Conventions," forthcoming in *Theoretical Economics*.
- [30] Myatt, DP; and C Wallace (2008a) "An evolutionary analysis of the volunteer's dilemma" *Games and Economic Behavior* 62:67-76.
- [31] Myatt, DP; and C Wallace (2008b) "When does one bad apple spoil the barrel? An evolutionary analysis of collective action" *Review of Economic Studies* 75:499-527.
- [32] Naidua, S; Hwang, S; and S Bowles (2010) "Evolutionary bargaining with intentional idiosyncratic play." *Economics Letters* 109:31-33.
- [33] Newton, J (2012) "Recontracting and stochastic stability in cooperative games" *Journal of Economic Theory* 147:364-381.
- [34] Pak, M (2008) "Stochastic stability and time-dependent mutations" *Games and Economic Behavior* 64:650-665.
- [35] Robles, J (1997) "Evolution and long run equilibria in coordination games with summary statistic payoff technologies" *Journal of Economic Theory* 75:180-193.
- [36] Robles, J (1998) "Evolution with changing mutation rates" *Journal of Economic Theory* 79:180-193.
- [37] Robles, J (2008) "Evolution, bargaining, and time preferences" *Economic Theory* 35:19-36.
- [38] Rozen, K (2008) "A 'Dual'-Improved Shortcut to the Long Run" *Yale University Economics Department Working Paper* No. 41
- [39] Saez-Marti, M; and JW Weibull (1999) "Clever Agents in Young's Evolutionary Bargaining Model" *Journal of Economic Theory*, 86:268-279
- [40] Schaeffer, M. (1988): "Evolutionarily Stable strategies for a Finite Population and a Variable Contest Size," *Journal of Theoretical Biology*, 132:469-478.
- [41] Tanaka, Y (1999) "Stochastically stable states in an oligopoly with differentiated goods: equivalence of price and quantity strategies" *Journal of Mathematical Economics*, 34:235-253.
- [42] Troger, T (2002) "Why Sunk Costs Matter for Bargaining Outcomes: An Evolutionary Approach" *Journal of Economic Theory* 102:375-402
- [43] Vega-Redondo, F (1997) "The evolution of Walrasian behavior" *Econometrica*, 65:375-384.
- [44] Young, HP (1993a): "The Evolution of Conventions." *Econometrica*, 61:29-56.
- [45] Young, HP (1993b) "An Evolutionary Model of Bargaining." *Journal of Economic Theory*, 59:145-168.
- [46] Young, HP (1998) "Conventional Contracts." *Review of Economic Studies*, 65:776-792.
- [47] Zhang, J; Cui, Z; and L Zu (2014) "The evolution of free trade networks." *Journal of Economic Dynamics & Control* 38:72-86.