

The Emergent Seed: A Representation Theorem for Models of Stochastic Evolution and two formulas for Waiting Time.

Kevin Hasker*
Economics Department
Bilkent University
06800 Bilkent
Ankara, Turkey
E-mail: hasker@bilkent.edu.tr
Phone: +90 (312) 290-1875
Fax: +90 (312) 266-5140

First version: July 11, 2002
Current version: April 23, 2014

Abstract

There is a fundamental underlying structure in models of stochastic evolution called the *emergent seed*. Relative to this structure the stochastic potential of a limit set can be found without optimization, and the stochastically stable limit set can be found using local analysis. This representation can be used to find two measures of waiting time—the *coheight* and the *censored coradius*.

We show the usefulness of the emergent seed by reanalyzing several applications in the literature and by showing how most applications in the literature could have been solved using this technique.

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

It is self evident that people involved in large repeated matching games do not calculate the law of motion of society. From a statistical point of view this would require constructing a large sample across both time and space. Space to estimate current behavior and time to detect trends. People must also know that others are doing this, and thus they need to sample other methods. How should they do this? Consider the classic battle of the sexes. When was the last time you (or your sons or daughters) ever tried to estimate even the current distribution of behavior?

What might a structural model look like? We must accept that we, like the people in our game, need an unavailable empirical sample of how agents estimate the law of motion. However we can make some deductions. First of all we should assume *inertia*. This is necessary if there is a natural law of motion and a standard equilibrium assumption. Second of all we should look at a variety of simple *decision rules*. The

*Acknowledgments: The author would like to thank Kivanc Akoz, Gil Eris, Philip Johnson, Emin Karagozoglul, David Levine, Sudipta Sarangi, Francesco Squintani, and Peyton Young for their comments. It goes without saying that any remaining errors are the responsibility of the author.

simplest—and probably the most common in the battle of the sexes—is to find someone who is successful and do what they do, or imitation (Robson and Vega-Redondo, 1996). A more complex model would have agents take a sample across space and best respond to this sample. This is the model in Kandori, Mailath, and Rob (1993) and we follow Sandholm (2010) by calling it BRM (best response with mutations). The next level of complexity would be to hypothesize agents construct this sample and then take an action with a probability that is proportional to its expected payoff—for example the logit model (Blume, 1993).

Now this final model is interesting for several reasons. First, as Alos-Ferrer and Netzer (2010) discusses, it has a wide basis in the psychological and experimental literature. It is also commonly used in particle physics. It can be fully rational if we follow Harsanyi (1973) by accepting that preferences are heterogeneous. Finally it means that all strategies may be used with non-zero probability. The first two models are too parsimonious to have such a rich implication, so we will explicitly assume that rarely people make *errors*—use a strategy that is not optimal for their decision rule.¹

We now have models of stochastic evolution, which is formally a triplet: inertia, a decision rule, and rare errors. It is relatively easy to find the short run implications of these models. Society should be near limit sets, which usually include the strict pure strategy Nash equilibria. However the long run implications—the prediction we need if we have no information—are harder to make. It is this problem that this paper addresses. The results rely only on the system being strongly ergodic, each model will imply a different resistance function (Young, 1993a) and this is our primitive.

We show that this problem is greatly simplified by finding an intermediate structure, the *emergent seed*. Once one has this structure in hand all one has to do is find the most likely transition from the heart of this structure—the *core*—to the limit set in question. We show that there is no choice in this transition—it is predetermined by the emergent seed.² This results in a measure of the long run viability of a limit set—its stochastic potential (Young, 1993a)—that requires no optimization. This representation has several implications. First we can use it to find the likely long run states of the system—stochastically stable states (Young, 1993a)—while we are constructing the emergent seed. Second we can use this representation to find two different estimates of the speed of evolution. One, the *coheight*, is an exact measure as the probability of errors goes to zero. This is the height of all other limit sets from Freidlin and Wentzell (1998). Second, the *censored coradius*, is simpler to apply and is a generalization of the modified coradius (Ellison, 2000). Again, the benefit of knowing the emergent seed is illustrated by this second measure. For most applications this is the modified coradius. Ellison (2000) states that this measure may not be simple to find, but with the emergent seed in hand we can find it easily.

While these results are interesting these are not the primary reason analysts might be interested in this methodology. Finding the probability of transitions is difficult, the simplest one to find is the probability of going anywhere else, the radius (Ellison, 2000). The emergent seed is constructed by finding the radii of limit sets with regards to a sequence of resistance functions.

The likelihood of each direct transition is inversely related to a resistance function. Young (1993a) shows that the likelihood of a limit set as the probabilities of errors goes to zero can be found by finding a minimal cost spanning tree on a directed graph (a limit set’s stochastic potential). A limit set is stochastically stable if

¹These errors are sometimes called experiments or mutations. We prefer error because an experiment should not be a dominated strategy and a mutation is a specific argument for why agents error.

²This requires that the matrix of resistances is *generic*. We make frequent use of this term, and we mean that the results do not change over an open set. In this case this implies the emergent seed is unique and constant over an open set.

it has a minimal minimal cost spanning tree. Thus what we need to do is find minimal minimal cost spanning trees. There is only one known algorithm for finding the minimal cost spanning tree for a given root (limit set in our analysis), Edmonds' algorithm (Edmonds, 1967—first published by Chiu and Liu, 1965). Humblet (1983) shows this can be used to find all minimal cost spanning trees simultaneously. However this does not solve our problem. We need a measure of the cost of the graph relative to a limit set. We show that the Humblet graph can be broken into two complimentary graphs. The *emergent seed* tells us how to get from every limit set to the core, the *core attraction star* tells us how to get from the core to a given limit set. The stochastic potential is then the cost of the emergent seed, plus the cost of the correct transitions from the core attraction star, minus a sequence of radii measures for the given limit set.

The emergent seed is a minimal resistance graph over the limit sets such that we avoid cycles as much as possible—the core will be our only one. So how does one find the emergent seed? The key insight is that we do not care about where we are going. We are not focused on a particular limit set, thus we just want to go somewhere. But this makes the first approximation obvious. We should go from each limit set to one that determines its radius. Obviously we can not be sure that this will succeed because we might not have a core. Instead we might have a set of local cores—from all nearby states one transitions to one of these states. Thus we must transform this graph, and again since we don't care about where we are going all we need to do is transform our first attempt (the first level) in a way that increases the cost the least. If we go from a given limit set to one that does not determine the radius this will increase the cost of the graph by the cost of that transition, but it will decrease it by the radius of the limit set. We have seen this function before. Ellison (2000) calls it the modified cost; we refer to it as the *first difference resistance*.

The second insight of this paper is to recognize that now we should consider what we have just done. We started out with a resistance function and derived a graph over the limit sets. Now we have a new resistance function. The local cores we found with the original one are just limit sets with regards to our new resistance function, the old limit sets that transition to one are in the (first difference) basin of attraction (Ellison, 2000). And what is the least cost graph? The one that is determined by first difference radii of the new limit sets. One continues to iterate these steps until there is one local core, which will be the core.

Now we need to find the stochastic potential of a limit set conditional on the cost of this graph. We need to solve a shortest path problem from the core to the given limit set. At this point we realize that if we take a long step—transitioning between local cores—we should do this early in the path. This in turn implies that each of these steps is within the same local core—and there is no choice on how to proceed.

This method provides us with, for the first time, an intuition about what makes a limit set stochastically stable. If this process takes T iterations then the t 'th difference radius of a limit set will be determined by the t 'th level limit set it is in the basin of attraction of. Its *total radius* will be the sum of the t 'th level radii for $t < T$. Then a limit set is more likely to be stochastically stable if:

1. It is easy to get to (from the core).
2. It is hard to leave (in terms of its total radius).

It is common in the literature that being easy to get to also implies being hard to leave, which we refer to as *core dominance*, but this is not necessary. Kandori and Rob (1995) analyze symmetric pure coordination games, there being easy to get to (Pareto dominant) also implies being hard to leave. Young (1998a) is a

type of asymmetric pure coordination game, and there being easy to get to (giving a lot to one party) means being easy to leave (the other party gets little).

This intuition is already in the literature. Ellison (2000) includes a radius/modified coradius theorem which says that if the radius is higher than the modified coradius then a limit set must be stochastically stable. The modified coradius is one measure of "it is easy to get to" thus that result clearly foreshadows ours. Indeed we do not claim that this is a novel approach to solving these problems. Rather we claim that this approach has or could have been used in almost all papers in the literature and that formalizing this intuition makes it easier to apply.

In applications we have only found a few papers where the emergent seed is not transparent. Indeed we have only found a handful where the naive minimization test proposed in Binmore, Samuelson, and Young (2003) would not have solved the problem. This test proposes constructing the first level of the emergent seed, checking to see that it is the emergent seed, and then finally checking to see if core dominance holds.

To show this we will reanalyze four applications in the literature, making contributions to each one. The first we will analyze—the Nash Demand or bargaining game (Young, 1993)—has created a paradigm of bargaining in the literature. As the article explains the emergent seed in this problem is *linear*. There is an implicit ordering over the limit sets and each either transitions one up or one down. In this problem the radius is increasing as one goes towards the middle, and core dominance holds. Our contribution is to analyze the medium run dynamics, in which discussion we will be assisted by Cui and Zhai (2010). The second problem we analyze is the contract game (Young, 1993b). Here the emergent seed is a *star*, from every limit set one transitions directly to a limit set in the core. In this case our contribution is to characterize stochastic stability when the number of contracts is small. The final two papers fall into our grab-bag of emergent seeds that do not have a simple characterization, which we call *other*. In the gift giving game (Johnson, Levine, and Pesendorfer, 2001) there is always a family of emergent seeds, and sometimes they do not fall into a simple class. However core dominance always holds making it simple to characterize stochastic stability. The article only determines stochastic stability for a range of parameters. In the contribution game (Myatt and Wallace, 2008b) the paper uses an unnecessary assumption to characterize stochastic stability. After each of these exposes we will briefly discuss other games in the literature with similar emergent seeds.

We have found that in most papers in the literature the emergent seed is either linear, a star, other, or a *cycle*. We do not analyze any game in this last class. In a cycle all limit sets are in the core. This is obviously necessary if there are only two limit sets, and only Kandori and Rob (1998) has found an example where there are more. In three action coordination games if both the total and marginal bandwagon properties hold then all three limit sets are in a cycle.

With some embarrassment we must admit that we rediscovered Edmonds' algorithm, but we take comfort in the fact that we are not the first. It is well known that Bock (1971) rediscovered it. Our research has revealed that Bortz, Kalos, and Lebowitz (1975) rediscovered it in the theory of Monte Carlo simulations, and that Cui and Zhai (2010) have done it in the field of stochastic evolution. Rozen (2008) credits an early version of this paper with being the first to apply it to stochastic evolution, but Rozen (2008) was the first to recognize it. That paper develops a dual methodology for analyzing stochastic potential, and proves that Edmonds' algorithm characterizes stochastic potential.

It is quite likely that we have a novel implementation, and this is the reason for our analytic clarity. We credit Bortz, Kalos, and Lebowitz (1975) with the basic insight. As Humblet (1983) discusses in the standard

methodology all states that are not included in cycles are analyzed at the next iteration. In contrast we connect all of these states to one of our cycles, formally they are in a basin of attraction. Thus we have fewer objects in each iteration. The difference can sometimes be dramatic. In both the bargaining and the contract game the standard methodology would require the number of limit sets minus one iterations. In contrast our methodology requires one.

We doubt our method is computationally superior. First of all the optimal program for a given root (Gabow, Galil, Spencer, and Tarjan, 1986) uses the standard methodology, as does Humblet (1983). We also know that our method will have a lot of backtracking. It is well known that there will be sequences of states that are not limit sets but will be used in analysis. Once one has found the limit sets one has to find sequences of states that describe the optimal transitions between them. We must do this at each level, thus some limit sets will originally be going towards one set of states and then end up going to a different one.

Cui and Zhai (2010) uses the standard implementation and develops a link between the cycles in that method and the medium run. Given an initial starting point the most likely medium run predictions are the cycles that methodology generates. Beggs (2005) discusses a different algorithm for finding the stochastically stable states. We are not certain if this is an implementation of Edmonds' algorithm, and it might be superior to the emergent seed. The article shows one can iteratively discard sets of states with a low height, and after having done this one will be left with the minimal minimal cost spanning trees. Height requires the solution of minimal cost spanning tree problems, but in some applications this might be straightforward. We have not seen this method in application.

Peski (2010) uses a similar method to ours for a specific purpose. The article defines a generalized risk dominant strategy as a strategy that will always be selected under weak conditions on the resistance function. The proof first restricts the analysis to a two action game and then iteratively removes the least cost transition. Notice that this article is finding a characteristic of a limit set that implies the emergent seed does not matter. The emergent seed will depend on the game, but a generalized risk dominant strategy relies only a local characterization. Many robustness arguments will not benefit from the emergent seed. For example the concept of stochastic dominance in Sandholm (2010). As well the results on the speed of evolution in Montanari and Saberi (2010), Young (2011), and Kreindler and Young (2012) are robust to the graph under analysis and we are not certain what the emergent seed is in these cases. Hasker (2014) analyzes the problem when the graph is a (square) lattice of a known dimension, but that paper requires knowing the dimension of the graph. On the other hand robust stochastic stability in Alos-Ferrer and Netzer (2012) might benefit from the emergent seed. That paper generalizes the radius/coradius test to a worst case for the radius and a best case for the coradius. It would seem that one could use the modified or censored coradius instead of the coradius in these arguments.

The two general alternatives to the emergent seed methodology are directly characterizing the limiting distribution and *root switching*. Directly characterizing the limiting distribution is superior because it gives you a closed form function for analysis and holds with high levels of errors. For example in the logit model when the game has a potential (Moderer and Shapley, 1986) the state which maximizes the potential will be stochastically stable (Blume, 1997, and Young, 1998b). This can also be done if agents use binary strategies (Sandholm, 2007). Fudenberg and Imhof (2006) show that in a model of imitation if only one error is needed for a transition this can be done. Kandori, Serrano, and Volij (2008) develops a model of decentralized trading with random utility, where the errors have the type 1 extreme values (Gumbel) distribution. In

this case the limiting distribution is one of a class of welfare functions. We should mention that Hasker (2014) illustrates that finding the emergent seed in this type of problem can be useful in finding the speed of evolution.

Root switching is our term for the standard methodology in the literature if one can not characterize the limiting distribution. In this argument one hypothesizes that a given state is stochastically stable, and then checks to see whether switching the root of its minimal cost tree with another state decreases the cost. We would like to credit Young (1993a) with being the first to use this technique, at least (as we shall) one can use this technique to explain why only limit sets can be stochastically stable. This is also the technique that both Kandori and Rob (1998) and Ellison (2000) use to show that in BRM if a strategy is half dominant (a best response when the given strategy has probability one half) then it is stochastically stable. The radius/coradius theorem in Ellison (2000) can also be proven using this technique—see Binmore, Samuelson, and Young (2003). The radius/modified coradius theorem can also be proven using this technique, one simply restricts consideration to the appropriate path. Let us call ω_* a limit set which has a radius that is higher than its modified coradius, and consider the stochastic potential of some other limit set $\tilde{\omega}$. Now to switch the root to ω_* we must add a path from $\tilde{\omega}$ to ω_* . We can assume that the first step is the one that determines the radius of $\tilde{\omega}$, and then we recognize that the change in cost for each of the following steps must be lower than the modified cost. Furthermore in the tree with root ω_* we can subtract the cost of the radius of ω_* , and thus ω_* is stochastically stable.

Let us briefly review the standard models. The first two models suggested are both *ordinal* best response models—like equilibrium analysis they are independent of affine transformations of the utility function. These are BRM (Kandori, Mailath, and Rob, 1993) and adaptive play (Young, 1993a). In BRM inertia is simply a probability $\rho \in (0, 1)$ that an agent does not update their strategy, when they optimize they best respond with the probability $1 - e^{-\beta}$, otherwise they are equally likely to take any action. In adaptive play one agent from each population interacts at a time, and they best respond to a random sample of size k from the last m periods of play with probability $1 - e^{-\beta}$ and otherwise are equally likely to take any action. The two other best response models are *cardinal*. To explain the terminology notice that they can both be based on a model of random utility (Myatt and Wallace, 2003), and thus any affine transformation of the utilities would also affinely transform the error distribution. Blume (1993) was the first to analyze the logit model—where the relative log probability of choices depends on the difference in expected utilities. Myatt and Wallace (2003) assumes the errors are normally distributed, resulting in the probit model. We note that this is the only best response model where a half-dominant strategy is not always selected. The last popular class of models are imitation models, first studied by Robson and Vega-Redondo (1996).³ In this model agents choose strategies based on comparisons with other strategies—relative payoffs are all that matter.

The rest of the paper is organized as follows. We present the model in section 2. Section 3 then describes how to find the emergent seed, section 4 presents the representation theorem, and section 5 presents two measures of waiting time. We then turn to analyzing several applications in section 6. The next to last section (section 7) gives examples where the emergent seed is either worse than other methods or will not be enough, and then section 8 concludes.

³Kandori, Mailath, and Rob (1993) discuss their model as one of optimization or imitation, but Robson and Vega-Redondo (1996) show that this equivalence holds only in the limit.

2 The Model

The fundamental of our model is a finite set of states of the world, which we will denote Z . In most applications these states will be social, or the strategies of all agents. For example if society is playing a strategic form game $G = (I, S, u)$, where I is the set of roles, S is the set of strategies, and n is the number of agents in role $i \in I$ then $Z = \times_{i \in I} (S_i)^n$. Often times there is a simpler representation for Z , if we assume uniform random matching then it can also be represented as $Z = \times_{i \in I} \times_{s_i \in S_i} \frac{n_{s_i}}{n}$ where for all s_i , $0 \leq \frac{n_{s_i}}{n} \leq 1$ and $\sum_{s_i \in S_i} \frac{n_{s_i}}{n} = 1$.

We endow the states of the world with a Markov transition matrix, M_β —which must be (strongly) ergodic. In other words for $x \in Z$, $y \in Z$ we allow for $M_\beta(x, y) = 0$ but there must be a finite $S < \infty$ such that $(M_\beta)^S(x, y) > 0$. The goal of our analysis is then to identify the limiting distribution over the states, $\mu(\beta)$, which can either be defined as $\forall z \in Z \mu(\beta) = \lim_{S \rightarrow \infty} (M_\beta)^S z$ or:

$$\mu(\beta) = M_\beta \mu(\beta) . \quad (1)$$

The representation explains our interest in this object. Given that we have no information what is our best guess? This would be determined by $\mu(\beta)$.

For arbitrary M_β $\mu(\beta)$ might be very dispersed and uninformative, but in the problems we are interested in there are two functions that allow us to concentrate this distribution. The first is a *resistance function* (Young, 1993b) such that for all $x \in Z$ and $y \in Z$ $r(x, y) \geq 0$ and for all $y \in Z$ there is an $x \in Z$ such that $r(x, y) = 0$.⁴ The second is a *weighting function* $W : Z \times Z \times \mathbb{R}_+ \rightarrow (0, 1)$ such that there is an $\varepsilon > 0$, $\lim_{\beta \rightarrow \infty} W(x, y, \beta) > \varepsilon$ for all $x \in Z$ and $y \in Z$. Then we can write:

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

where $\beta > 0$. Recognize that the standard in this literature is that the source (the current state) is written second in all functions and the target (the state we are transitioning to) is written first. Now as $\beta \rightarrow \infty$ the behavior of this system will be completely determined by $r(x, y)$. This will result in the long run distribution at β being more concentrated, and thus we analyze $\lim_{\beta \rightarrow \infty} \mu(\beta)$.

Most applications usually first define a limiting transition matrix— M_∞ —that is based only on the decision rule. Some abstract papers (Ellison, 2000; Cui and Zhai, 2010) follow this tradition but we will not because it is unnecessary. Beggs (2005) presents a more general model in which the cost function is derived from a sequence of transition matrices. We do not because we want to assume that all information necessary to determine the likelihood of a transition is included in the state. Notice that a state may include information about the recent past, the more general formulation is required when the infinite past may have an impact. Our model is identical to the model in Pak (2008) except for notational differences.

Note that $r(x, y) = \infty$ is possible because this would imply that a given transition is not possible. It is also worth mentioning that we could start with any function $f : Z \times Z \rightarrow \mathbb{R}$. We would then define $r(x, y) = f(x, y) - \min_{z \in Z} f(z, y)$ and proceed. At this point it is worthwhile to explain what our genericity

⁴In the literature the resistance function has been called the cost function (Kandori, Mailath, and Rob, 1993), or the waste function (Alos-Ferrer and Netzer, 2010). These papers (usually) later on defines a path minimizing function. We call this latter function the cost function because a cost function is usually minimized. We note with apology that Pak (2008) uses the opposite terminology.

assumption on the matrix of resistances—denoted $[r]$ —gains us. Because of the operation we just explained for generic $[\hat{r}]$ we consider $[r]$ where $[r]_{x,y} = \hat{r}(x,y) - \min_{z \in Z} \hat{r}(z,y)$. Genericity implies there is a strict order over the resistance of transitions from a given y .

Let us give two examples to clarify our methodology. Assume G is a two action symmetric game, $S_i = \{A, B\}$, and that there is one population of agents matched uniformly. Now $z \in Z$ can be $z = \frac{n_z}{n}$ where $0 \leq n_z \leq n$ is the number of agents who are using the strategy A . The logit model fits very naturally into our framework. We assume that only one agent can change strategy at a time. Let σ be the "distribution" at which the expected payoffs of the two strategies are equal (notice we can have $\sigma < 0$ or $\sigma > 1$ if one of the strategies is dominant), and define $BR(y)$ as the strategy that maximizes an agents payoff at y , then:

$$r(x, y) = \begin{cases} \infty & \text{if } n|x-y| > 1 \\ 0 & \text{else if } \begin{matrix} x \geq y & \text{and } BR(y) = A \\ x \leq y & \text{and } BR(y) = B \end{matrix} \\ n|y - \sigma| & \text{else} \end{cases} \quad (3)$$

and $W(x, y, \beta) = 1(x \geq y) \left(1 - \frac{n_y}{n}\right)^{|x-y|} + 1(x < y) \left(\frac{n_y}{n}\right)^{|x-y|}$, where $1(p) = 1$ if p is true and zero else. In BRM it is simple to characterize the resistance but we will not characterize the weighting function. The resistance function is:

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

The weighting function is complex because we have to consider matched subsets of agents—half that switch to A and half that switch to B , in this model $\sum_{z \in Z} W(z, y, \beta) e^{-\beta r(z, y)} = 1$.

Young (1993a) shows that instead of directly analyzing $\lim_{\beta \rightarrow \infty} \mu(\beta)$ one can analyze directed graphs over Z . There are many ways to characterize graphs. We think the simplest is a matrix of ones and zeros. Thus a graph is a matrix G which has dimensions $|Z| \times |Z|$ where if we transition from y to x then $G_{x,y} = 1$, and $G_{x,y} = 0$ otherwise. This then implies that the resistance of such a graph is simply:

$$\begin{aligned} r(G) &= \sum_{z \in Z} \sum_{\hat{z} \in Z} r(z, \hat{z}) G_{z, \hat{z}} \\ &= 1'_Z G [r]' G 1_Z \end{aligned} \quad (5)$$

where 1_Z is a column vector of $|Z|$ ones and P' is the transpose of P . In this notation graphs are similar to Markov transition matrices. Now a *tree with root* $X \subseteq Z$ is a graph denoted T such that $\sum_y \sum_{x \in X} T_{y,x} = 0$ —or you do not transition out of X ; and for every $y \in Z \setminus X$ there is a finite sequence $\{z_s\}_{s=1}^S$ with $z_1 = y$ and $z_S \in X$ such that $\prod_{s=1}^{S-1} T_{z_{s+1}, z_s} = 1$ —or you transition to some state in X after a finite number of steps. Let the set of these graphs be $T(X)$. Then the *stochastic potential* of X is:

$$sp(X) = \min_{T \in T(X)} 1'_Z T [r]' T 1_Z ; \quad (6)$$

and x is *stochastically stable* if:

$$x \in \arg \min_{z \in Z} sp(z) , \quad (7)$$

and Young (1993a) shows that this implies that $\lim_{\beta \rightarrow \infty} \mu(x, \beta) > 0$. Thus what we want to find are trees with minimal resistance.

Understanding the implications of this representation is easier if we derive the *cost function*. Since we are analyzing the long run we don't care about how many periods it takes to transition between y and x . This motivates us to analyze the *paths* from y to x , denote the set of paths as $P(x, y)$. In a $P \in P(x, y)$ there is a finite sequence $\{z_s\}_{s=1}^S$ with $z_1 = y$ and $z_S = x$ and $\prod_{s=1}^{S-1} P_{z_{s+1}, z_s} = 1$, without loss of generality we assume that unless $(z, \tilde{z}) = (z_{s+1}, z_s)$ for some s then $P_{z, \tilde{z}} = 0$. The *cost* of going from y to x is:

$$c(x, y) = \min_{P \in P(x, y)} 1'_Z P [r]' P 1_Z, \quad (8)$$

and we notice that this has a unique extension to subsets of Z , for $X \subseteq Z$ and $Y \subseteq Z$:

$$c(X, Y) = \min_{x \in X, y \in Y} c(x, y). \quad (9)$$

As $\beta \rightarrow \infty$ the relatively likelihood of a transition from y to x is dominated by $c(x, y)$. Any other possible transition will have a relative likelihood of zero since it has a strictly higher resistance. Notice that since M_β is ergodic $c(x, y) < \infty$ for all $x \in Z$ and $y \in Z$.

Now we can use the cost function to derive an important object of analysis, the *limit sets*.

Definition 1 $\omega \subseteq Z$ is a limit set if for all $y \in \omega$:

1. for all $x \in \omega$ $c(x, y) = 0$
2. for all $z \in Z \setminus \omega$ $c(z, y) > 0$.

We denote the family of these limit sets as Ω . Another way to characterize these is to define:

$$E_{x, y}^0 = \begin{cases} 1 & \text{if } r(x, y) = 0 \\ 0 & \text{else} \end{cases}, \quad (10)$$

E^0 is equivalent to an element of M_∞ , to be precise we can let $M_{\infty, \cdot, y} = E_{\cdot, y}^0 / \sum_{z \in Z} E_{z, y}^0$. Then a limit set is a minimal set such that there is a distribution (ζ) with support ω such that $\zeta(\omega) = E^0 \zeta(\omega)$.

Notice that once you have entered a limit set you will stay there with a probability that is nearly one as $\beta \rightarrow \infty$, and you will leave with a probability on the order of zero. The terminology (limit set, or sometimes recurrence class) is based on analysis of a limiting transition matrix, M_∞ . In that case these will be proper limit sets. In our analysis every state or set of states will be transitory.

In the sequel we will derive a sequence of resistance functions, r^t for $t \in \{1, 2, 3, \dots\}$ and then we will denote the limit sets with regards to these resistance functions as ω^t , and the set of them as Ω^t . To avoid confusion between x raised to the power t and an x^t we write the former as $(x)^t$.

A useful compliment to the concept of a limit set is the *basin of attraction* (Ellison, 2000). These are the states that will converge to a limit set with a strictly positive probability as $\beta \rightarrow \infty$ or:

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

Obviously we are most interested in this for $X \in \Omega$. As Beggs (2005) points out this definition is different from that in Ellison (2000) because we do not require states to reach ω with a probability nearly one. Consider a multi-role normal form game with a best response dynamic. Then from a given state, z , the direction of movement might depend on which role optimizes first. Ellison (2000) would say that such a z is

not in any basin of attraction—our definition might place it in multiple basins. This *inner* basin of attraction ($\hat{D}(Y) = \{y \in \bar{D}(Y) : \forall x \in Z \setminus \bar{D}(Y) \ c(x, y) > 0\}$), may be better in applications but for our purposes the difference is academic and $\bar{D}(Y)$ is simpler. Our genericity assumption implies $\hat{D}(Y) = \bar{D}(Y)$.

We now characterize the probability of the most likely transition out of a set X such that we will not return. We need to be careful about the case where there is only one limit set, and what to do is problematic. It could be that $\hat{D}(X) = \bar{D}(X) = Z$ and since transitions out of Z are impossible the radius must be infinite. On the other hand it could be that $\hat{D}(X) \subset \bar{D}(X) = Z$, in which case the radius would be large but finite. (Consider a game with a cycle in weakly dominated strategies.) For a generic matrix of resistances the former would be the only case to consider, thus define the *radius* (Ellison, 2000) as infinite:

$$\mathcal{R}(X) = \begin{cases} \min_{z \in Z \setminus \bar{D}(X)} c(z, X) & \bar{D}(X) \subset Z \\ \infty & \bar{D}(X) = Z \end{cases} \quad (11)$$

We follow Rozen (2008) in using $\mathcal{R}(\cdot)$ for this function to avoid confusion with the resistance ($r(\cdot, \cdot)$). A simple insight is that the triangle inequality always holds:

Lemma 1 *for all $z \in Z$, $c(x, y) \leq c(x, z) + c(z, y)$*

Proof. *On the right hand side we consider a constrained set of paths and on the left hand side we remove this constraint. ■*

This is not the same as the result in Kandori and Rob (1995). That result which is referred to as a triangle inequality is that you never have to consider two states within the same basin of attraction. This more sophisticated result we refer to as the *direct jump lemma*.

Now we can prove the key insights of Young (1993a) with relative ease.

Lemma 2 *To find the stochastic potential of states you only have to consider graphs over Ω , and furthermore only an $\omega \in \Omega$ can be stochastically stable.*

Proof. *The first result is a transparent implication of the triangle inequality. Clearly $c(\omega, \omega') \leq c(\omega, z) + c(z, \omega')$ implies that you can drop all other states, and simply attach them to a limit set which they reach with zero cost. Thus a graph over Ω can be extended to a graph over Z without increasing its resistance.*

Let $x \in \bar{D}(\omega) \setminus \Omega$ then by root switching we can see that $sp(\omega) \leq sp(x) - \mathcal{R}(\omega)$ and since $\mathcal{R}(\omega) > 0$ $sp(\omega) < sp(x)$ and x can not be stochastically stable. ■

To understand the power of this result consider the classic coordination game. In this game the most parsimonious representation of the set of states has $n + 1$ elements, in contrast only two states (0 and 1) are limit sets. Thus the problem is relatively simple, all one has is the non-trivial task of calculating the likelihood of these transitions. In general if $|\Omega|$ is small this will not be too difficult of a task, but what if $|\Omega|$ is very large? Is it possible to solve these problems? The answer is yes, and if we approach the problem from a global perspective we only need to do it once.

3 The Emergent Seed

We will now find the emergent seed. This is a most likely path for evolution given that we try to avoid cycles but the process does not terminate. Bortz, Kalos and Lebowitz (1975) provides us with this intuitive description, but we want one that will allow analysis of graphs. This definition is:

Definition 3 An emergent seed, E^* , is a minimal cost graph such that:

1. There is a transition from every $\omega \in \Omega$.
2. There exists some $\omega_c \subseteq \Omega$ that are transitioned to from every $\omega \in \Omega$.

Notice that in the implicit zero'th level of the emergent seed (E^0) we have taken care of all states that are not limit sets, thus we focus on the reduced problem. The difficulty in the definition of E^* is the second part. If it was not for this it would be trivial to find the emergent seed, so let us ignore this goal for now.

Definition 4 The first level of the emergent seed, E^1 , is a minimal cost graph such that there is a transition from each $\omega \in \Omega$.

This is quite simple, from each $\omega \in \Omega$ you transition to a $\tilde{\omega}$ that determines its radius. Precisely remember that $P(z, \omega)$ was the set of paths from a state in ω to z , let $P_{\mathcal{R}}(\omega)$ be the solution to:

$$P_{\mathcal{R}}(\omega) = \arg \min \left\{ \sum_{x \in Z} \sum_{y \in Z} P_{xy} : P \in \arg \min_{z \in Z \setminus \bar{D}(\omega)} \min_{P \in P(z, \omega)} 1'_Z P [r]' P 1_Z \right\}. \quad (12)$$

The outer minimization assures us that $P_{\mathcal{R}}(\omega)$ visits only one state in ω , and implies $P_{\mathcal{R}}(\omega)$ is unique. We note that if the matrix of resistances is not generic this is a point where using $\hat{D}(\omega)$ instead of $\bar{D}(\omega)$ would be important. Then define $\text{com max}(P, Q)$ as the component by component maximum, or $\text{com max}(P, Q)_{x,y} = \max(P_{x,y}, Q_{x,y})$ and

$$E^1 = \text{com max}_{\omega \in \Omega} P_{\mathcal{R}}(\omega). \quad (13)$$

If one thinks of a graph as a subset of pairs, one can say that E^1 is the union of the $P_{\mathcal{R}}(\omega)$. We have found only three papers where the emergent seed has more than one level. In the economics literature the only paper where the emergent seed has more than one level is Ben-Shoham, Serrano and Volij (2004). The other two exceptions are from the physics literature—Neves and Schonman (1992) and Arous and Cerf (1996). These articles analyze the speed of evolution on a lattice with the logit model, the former for two dimensions and the latter for three. Hasker (2014) has solved these problems using the emergent seed.

This simple step might fail because of the second criterion, there must be a core. But perhaps it is a good place to start. This is because of all graphs that satisfy the first criterion this one has the lowest cost. So how to proceed? Well we notice that if we make the limit set transitioned to from ω not be the one given by E^1 then this will increase the cost by $c(\tilde{\omega}, \omega)$ but it will also decrease the cost of the new graph by $\mathcal{R}(\omega)$. This gives us the *first difference resistance function*, which Ellison (2000) called the modified cost function:

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

Using this one can derive the first difference cost function like before, denoted $\Delta c(x, y)$, and at this point we should pause.

We should recognize that $\Delta c(x, y)$ is just a cost function. One which has the nice property that $\Delta c(x, y) \leq c(x, y)$ and furthermore $\Delta c(x, y) < c(x, y)$ for $y \in \omega$ and $x \notin \bar{D}(\omega)$. Because of this we know precisely how

to proceed, indeed we just did this. Let $\omega^1 \in \Omega^1$ be a *first order limit set*—notice that $|\Omega^1| \leq |\Omega|/2$ —define $\Delta\bar{D}(\omega^1)$ just like we defined $\bar{D}(\omega)$, $\Delta\mathcal{R}(\omega^1)$ just like $\mathcal{R}(\omega)$ and:

$$P_{\Delta\mathcal{R}}(\omega^1) = \arg \min \left\{ \sum_{x \in Z} \sum_{y \in Z} P_{xy} : P \in \arg \min_{z \in Z \setminus \Delta\bar{D}(\omega^1)} \min_{P \in \mathcal{P}(z, \omega^1)} 1'_Z P [\Delta r]' P 1_Z \right\}. \quad (15)$$

then $E^2 = \text{com} \max_{\omega^1 \in \Omega^1} P_{\Delta\mathcal{R}}(\omega^1)$. Now define the *merge* operation.

$$P \text{ merge } Q = \begin{cases} P \text{ merge } Q_{\cdot, y} = P_{\cdot, y} & \text{if } \sum_{z \in Z} P_{z, y} > 0 \\ P \text{ merge } Q_{\cdot, y} = Q_{\cdot, y} & \text{else} \end{cases}, \quad (16)$$

remembering that what we are merging are graphs over Z , this operator gives precedence to P . If we transition from y to someplace in P then this is written down as our transition. If we do not then we use Q . Then $E^2 \text{ merge } E^1$ will have the second lowest cost of any graph satisfying the first criterion that is closer to having a core.

At this point we simply iterate the process. Our initial condition is $\Delta^{t-1} r(\cdot, \cdot)$, and from this we derive $\Delta^{t-1} c(\cdot, \cdot)$, Ω^{t-1} , $\Delta^{t-1} \bar{D}(\cdot)$, $\Delta^{t-1} \mathcal{R}(\cdot)$, $P_{\Delta^{t-1} \mathcal{R}}(\omega^1)$, and E^t . We then define the t 'th difference resistance function as:

$$\Delta^t r(x, y) = \begin{cases} \Delta^{t-1} c(x, y) - \Delta^{t-1} \mathcal{R}(\omega^{t-1}) & \text{if } y \in \omega^{t-1}, x \notin \Delta^{t-1} \bar{D}(\omega^{t-1}) \\ \Delta^{t-1} c(x, y) & \text{else} \end{cases}. \quad (17)$$

When do we stop? When we are at a T such that $\Delta^T c(\cdot, \cdot)$ has only one limit set. We note that we do need to construct this cost function for future use. Given this the definition of the emergent seed is:

$$E^* = E^T \text{ merge } E^{T-1} \text{ merge } E^{T-2} \dots \text{ merge } E^1. \quad (18)$$

We have proven by construction that this satisfies our criteria.

Proposition 1 *The emergent seed is as defined in 18.*

4 The Representation Theorem

Like above we will prove our theorem by construction and then present the conclusion. Allow us the conventions that $\Delta^0 c(\cdot, \cdot) = c(\cdot, \cdot)$ and $\omega^0 = \omega$, $\Omega^0 = \Omega$. We begin with the cost of the emergent seed,

$$c(E^*) = \sum_{t=0}^{T-1} \Delta^t c(E^{t+1}). \quad (19)$$

It is useful to consider what the equality implies. On the left hand side we have one transition from each limit set, on the right hand side we have many. However most of these transitions cancel each other out. Analytically the right hand side is easier to work with, but it is useful to remember that it is equivalent to the left hand side. For example assume that ω is in the core. Then it is quite clear that the total cost of leaving ω on the right hand side is $\sum_{t=0}^{T-1} \Delta^t \mathcal{R}(\omega)$, but by consideration of the left hand side we can see that this is just the cost of the path out of $\Delta^{T-1} \bar{D}(\omega)$, or one transition in the Ω^{T-1} space.

Now if we take this cost as our fundamental, how do we find the stochastic potential of limit sets? For $\omega \subseteq \omega_c$ it is quite clear that we just subtract the cost of transitioning from ω to some $\omega^{T-1} \in \Omega^{T-1}$, or

$\sum_{t=0}^{T-1} \Delta^t \mathcal{R}(\omega)$. What about other limit sets? Well now all we need to do is solve a shortest path problem. We should start at ω_c and go to x . Thus we have changed a tree minimization problem to a shortest path problem. But we can go further than that, the path from the core to x is already determined.

To go from the core to x we need a sequence $\{x_s\}_{s=-1}^S$ with $x_{-1} = x$ and $x_S \in \omega_c$. The last step requires no analysis, it is $\Delta^T c(x_{S-1}, x_S) = \Delta^T c(x_{S-1}, \omega_c)$, furthermore we can see that by the triangle inequality this should be the only step that uses the T 'th difference cost function.

Thus for each pair of states where $s < S - 1$ there is some $t < T$ such that the direct cost of transitioning from x_s to x_{s-1} is $\Delta^t c(x_{s-1}, x_s)$, but since we are going from x_s to x_{s-1} we do not need to go to the state that we used to go to from x_s , thus we subtract $\Delta^t \mathcal{R}(x_s)$. The net cost is $\Delta^t c(x_{s-1}, x_s) - \Delta^t \mathcal{R}(x_s)$. Now there are two cases to consider, first it could be that $x_{s-1} \in \Delta^t \bar{D}(x_s)$. In this case this is a new function, and like before we can see that without loss of generality we only need one $\{x_{s-1}, x_s\}$ that use the t 'th difference cost function. On the other hand if $x_{s-1} \notin \Delta^t \bar{D}(x_s)$ then $\Delta^t c(x_{s-1}, x_s) - \Delta^t \mathcal{R}(x_s) = \Delta^{t+1} c(x_{s-1}, x_s)$ and we can consider $\Delta^{t+1} c(x_{s-1}, x_s)$. But since this is a cost function the triangle inequality implies that in either case we can write $\{x_s\}_{s=-1}^S$ as $\{x_t\}_{t=-1}^T$ where x_t is the source in the t 'th difference cost function. Let us now define for arbitrary $X \subseteq Z$ the correspondence:

$$\begin{aligned} \omega^{-1}(X) &= X \text{ and for } t \geq 0 \\ \omega^t(X) &= \{\omega^t \in \Omega^t \mid \omega^{t-1}(X) \cap \Delta^t \bar{D}(\omega^t) \neq \emptyset\} . \end{aligned} \quad (20)$$

We are now ready to recognize a very useful and general insight, *strategic delay*. If $t < T - 1$ and $\omega^{t+1}(x_{t-1}) \neq \omega^{t+1}(x_t)$ then $\Delta^{t+1} c(x_{t-1}, x_t) > \Delta^{t+2} c(x_{t-1}, x_t)$, and we should not transition from x_{t-1} to x_t in this step. We should strategically delay the transition. But this implies that $x_t \in \Delta^{t+1} \bar{D}(\omega^{t+1}(x_t))$ and $x_{t-1} \in \Delta^{t+1} \bar{D}(\omega^{t+1}(x_{t-1}))$. Thus we recognize that $\{x_t\}_{t=-1}^T = \{x, \omega, \omega^t(x)\}_{t=1}^T$. If $x \notin \omega$ we do not know the second element of this sequence, but Lemma 1 tells us we do not need to analyze these states. Thus:

Theorem 1 *For $x \in \omega \in \Omega$, if $[r]$ is generic then:*

$$sp(x) = c(E^*) + \sum_{t=0}^{T-1} \Delta^{t+1} c(\omega^t(x), \omega^{t+1}(x)) - \sum_{t=0}^{T-1} \Delta^t \mathcal{R}(\omega^t(x)) \quad (21)$$

where $\omega^t(x)$ is defined in equation 20.

Let us call the *likelihood potential* of ω as:

$$lp(\omega) = \sum_{t=0}^{T-1} \Delta^t \mathcal{R}(\omega^t(\tilde{\omega})) - \sum_{t=0}^{T-1} \Delta^{t+1} c(\omega^t(\tilde{\omega}), \omega^{t+1}(\tilde{\omega})) , \quad (22)$$

being stochastically stable means that you have maximum likelihood potential. We will use this approach in our applications.

If the genericity assumption holds and the cost function is known then this does not require optimization. Of course we must qualify this by pointing out that avoiding deriving the cost function is one attraction of this methodology. But we must respond to that qualification by pointing out we only need it for local transitions—one still does not have to derive it globally.

Because of the genericity assumption the emergent seed will be unique, and this simplifies exposition and analysis. The genericity assumption does not hold in many applications. In many games with more than one role it will fail at states that are not in limit sets. Consider the bargaining game where the strategies are the share of surplus (Young, 1993b), clearly from $(\frac{1}{4}, \frac{1}{4})$ we can go to either $(\frac{3}{4}, \frac{1}{4})$ or $(\frac{1}{4}, \frac{3}{4})$. This will not matter because we only need it to hold for limit sets. But in the volunteer game (Myatt and Wallace, 2008a) it does not hold for limit sets. The strict Nash equilibria in this game are when someone volunteers, and there are two ways to leave one. First someone could offer to replace the volunteer. Second the volunteer could quit, but then anyone would volunteer. If the latter is more likely then we can transition from that limit set to any other. In the gift giving game (Johnson, Levine, and Pesendorfer, 2001) usually some strategies are equal distance from many other strategies. Thus this example is truly not generic. Notice that when the emergent seed is not unique it is the union of cores that matters, and this means core dominance is more likely to hold. A more subtle example allows us to explain a useful principle. When there is a choice of where to go one chooses the target that will result in the most connected emergent seed, the *principle of connectedness*. In the two dimensional lattice (Ellison, 2000) a square of four agents playing the risk dominant strategy is a strict Nash equilibrium, but from here one can either transition up to a state where more agents are playing the risk dominant strategy or down to the state where no agents are. Where should one go? In a sense it does not matter, if we consider the case where the square points down the resulting cycle will have a first order radius of zero, so we should just assume it is up.

When the emergent seed is not generic this can only make it harder to find the stochastic potential of ω 's for which $\omega^t(\omega) \notin \Omega^t$ for some t . We do not know how hard it will be. On one hand since the cost of all emergent seeds is the same it might be that it doesn't matter. Otherwise it might require a global optimization.

Let us consider the intuition of this representation. Obviously $\sum_{t=0}^{T-1} \Delta^{t+1} c(\omega^t(x), \omega^{t+1}(x))$ measures how easy it is to get to x from the core. Note that if $x \in \omega_c$ then $\sum_{t=0}^{T-1} \Delta^{t+1} c(\omega^t(x), \omega^{t+1}(x)) = 0$. Thus it is very easy to get to these states. On the other hand $\sum_{t=0}^{T-1} \Delta^t \mathcal{R}(\omega^t(x))$ is the appropriate measure of how hard it is to leave the given state, and if this is very high a state might be stochastically stable even if the first term is strictly positive.

In the literature it is very common to solve problems using *core dominance* which is formally:

Corollary 1 (Core Dominance) *If there is an $\omega \subseteq \omega_c$ such that $\sum_{t=0}^{T-1} \Delta^t \mathcal{R}(\omega^t(\omega)) \geq \sum_{t=0}^{T-1} \Delta^t \mathcal{R}(\omega^t(\tilde{\omega}))$ for all $\tilde{\omega} \in \Omega$, then ω is stochastically stable and no $\tilde{\omega} \not\subseteq \omega_c$ is stochastically stable.*

Of course this can also be extended to a local statement, but in this case it is better to write it in the negative, where it might rule out many limit sets.

Corollary 2 (Local Core Dominance) *ω can not be stochastically stable if there is a t such that $\sum_{s=0}^t \Delta^s \mathcal{R}(\omega^s(\omega)) < \max_{\tilde{\omega} \subseteq \omega^t(\omega)} \sum_{s=0}^t \Delta^s \mathcal{R}(\omega^s(\tilde{\omega}))$.*

A positive approach would be to find out which state(s) in each basin of attraction might be stochastically

stable at each step. In the first step the candidates are:

$$\omega_{p,1}(\omega^1) = \arg \max_{\omega \in \Omega \cap \Delta \bar{D}(\omega^1)} \mathcal{R}(\omega) - \Delta c(\omega, \omega^1) . \quad (23)$$

This identifies the *primaries* of first order limit sets, or $\Omega_{p,1} = \cup_{\omega^1 \in \Omega^1} \omega_{p,1}(\omega^1)$, and these are the only limit sets we need to consider in future steps. For the next step it is convenient do define:

$$Ca^1(\omega_{p,1}) = \mathcal{R}(\omega_{p,1}) - \Delta c(\omega_{p,1}, \omega^1(\omega_{p,1})) \quad (24)$$

Then:

$$\omega_{p,2}(\omega^2) = \arg \max_{\omega \in \Omega_{p,1} \cap \Delta^2 \bar{D}(\omega^2)} \Delta \mathcal{R}(\omega^1(\omega)) - \Delta^2 c(\omega^1(\omega), \omega^2(\omega)) + Ca^1(\omega) , \quad (25)$$

and $\Omega_{p,2} = \cup_{\omega^2 \in \Omega^2} \omega_{p,2}(\omega^2) \subset \Omega_{p,1} \subset \Omega$. At this point we can proceed by iteration.

$$\begin{aligned} Ca^t(\omega) &= Ca^{t-1}(\omega) + \Delta^{t-1} \mathcal{R}(\omega^{t-1}(\omega)) - \Delta^t c(\omega^{t-1}(\omega), \omega^t(\omega)) \\ \omega_{p,t}(\omega^t) &= \arg \max_{\omega \in \Omega_{p,t-1} \cap \Delta^{t+1} \bar{D}(\omega^{t+1}(\omega^t))} \Delta^t \mathcal{R}(\omega^t(x)) - \Delta^{t+1} c(\omega^t(x), \omega^{t+1}(x)) + Ca^t(\omega) . \end{aligned} \quad (26)$$

And this outlines an iterative approach to finding stochastically stable states.

In stochastic evolution it is not necessary to know how one gets from ω_c to ω , one only needs the cost. The path might be relevant in other analyzes, and analyzing this structure is both easy and insightful. First recognize that we have to add in everything that was in E^t but not in E^* because, for the appropriately chosen states, we will need to make this transition. Then for $\omega^t \in \Delta^{t+1} \bar{D}(\omega^{t+1}(\omega^t)) \setminus \omega^{t+1}(\omega^t)$ we have to find:

$$P_S(\omega^t) = \arg \min_{z \in \omega^{t+1}(\omega^t)} \min_{P \in P(z, \omega^t)} 1'_Z P [\Delta^t r]' P 1_Z . \quad (27)$$

Call this graph the *core attraction star*. The terminology is because (locally) it will always be a (multi-step) star. There will be an origin (ω^{t+1}) that goes to several locations, and then the path terminates. Denote this graph:

$$S^* = \text{com} \max_{t \in \{0,1,2,\dots,T\}, P \in P_S(\omega^t), \omega^t \in \Omega^t \setminus \Omega^{t+1}} (E^t \text{ merge } E^*, P) . \quad (28)$$

We denote the Humblet graph $H^* = \text{com} \max(E^*, S^*)$.

Let us understand what this means by considering a minimal cost transportation network. In many algorithms there is a delayed construction clause. One first constructs the equivalent of E^* and then turns to the issue of determining which roads to build. This is quite intuitive, surely one needs to know the global solution before spending money? The Humblet graph says no, there is no need to look at the global solution. At each step you first find the best place to go to, and then turn the problem around and from the cycles you have just created make sure that there is a path back to the states going to them (these roads are one way). Indeed one could delegate the local problems to local planners—as long as their incentives are right.

5 Two measures of Waiting time

Another important question in stochastic evolution is how long it takes to reach the stochastically stable state(s). Since our representation theorem requires no optimization it can be used to find two measures of

this waiting time. The first, which we call *coheight*, is a precise measure as $\beta \rightarrow \infty$. Since in some analyses one will be interested in a fixed β or will not have a β we derive another measure called the *censored coradius*. This is a generalization of the modified coradius from Ellison (2000).

The objective we want to solve in this section is to find the log waiting time of ω , or:

$$\ln \tau(\omega) = \lim_{\beta \rightarrow \infty} \frac{\ln E_{\beta}(\min t | x_t \in \omega, x_0 \in \Omega \setminus \omega)}{\beta}. \quad (29)$$

As Beggs (2005) points out the solution to this problem is in Freidlin and Wentzell (1998). The *height* of a subset of states is the expected waiting time to exit that subset. Thus the *coheight* of a limit set is the expected waiting time to exit $\Omega \setminus \omega$. From Beggs (2005) we see that:

$$Ch(\omega) = H(\Omega \setminus \omega) = sp(\omega) - \min_{\tilde{\omega} \in \Omega} sp(\{\omega, \tilde{\omega}\}). \quad (30)$$

We have a simple formula for sp_{ω} and it is not hard to derive one for $sp(\{\omega, \tilde{\omega}\})$ and then find this formula.

In order to find the formula for $sp(\{\omega, \tilde{\omega}\})$ let us recognize that for $\tilde{\omega} \neq \omega$ what we are really looking at is a *split-tree*. In this tree every limit set must either go to ω or $\tilde{\omega}$. Thinking in terms of the emergent seed we recognize that one of these two limit sets must have a path from the core to that limit set, but what of the other one? Without loss of generality assume that there is a path from the core to ω . Then this other one only has to go to some $t \leq T - 1$. The cost of the first step in this path will be $\Delta^t c(x_{t-1}, x_t) - \Delta^t \mathcal{R}(x_t)$ just like before. But now we must have $\Delta^t c(x_{t-1}, x_t) - \Delta^t \mathcal{R}(x_t) < 0$ or we should have the sequence only going to the $t - 1$ 'th level. This is only possible if $x_t \in \omega^t(x_{t-1})$, and we derive $x_s \in \omega^s(x_{s-1})$ for $s < t$ by the same strategic delay argument. We also need to consider what will happen if $\omega^t(\omega) = \omega^t(\tilde{\omega})$ for some $t < T$. In this case the path from ω to the core is the same as the path from $\tilde{\omega}$ to the core. Thus define $\bar{t}(\omega, \tilde{\omega}) = \min_t \{\omega^t(\omega) = \omega^t(\tilde{\omega})\} \leq T$. Now we have to be a little careful about how we characterize this cost. We have to make sure to keep the elements of the t 'th step together. But given we are careful about that we can immediately the stochastic potential of a split tree.

Lemma 3 For arbitrary $\{\omega, \tilde{\omega}\} \subseteq \Omega$ let $\{t(\tilde{\omega}), t(\omega)\} \in \{0, 1, 2, \dots, \bar{t}(\omega, \tilde{\omega})\}^2$ such that $\max[t(\tilde{\omega}), t(\omega)] = \bar{t}(\omega, \tilde{\omega}) = \min_t \{\omega^{t+1}(\omega) = \omega^{t+1}(\tilde{\omega})\}$:

$$sp(\{\omega, \tilde{\omega}\}) = c(E^*) + \left(\sum_{t=\bar{t}(\omega, \tilde{\omega})+1}^T \Delta^t c(\omega^{t-1}(\omega), \omega^t(\omega)) - \Delta^t \mathcal{R}(\omega^t(\omega)) \mathbf{1}(t < T) \right) \mathbf{1}(\bar{t}(\omega, \tilde{\omega}) < T) \quad (31)$$

$$\min_{t(\omega), t(\tilde{\omega})} \left\{ \begin{array}{l} \sum_{t=0}^{t(\tilde{\omega})} [\Delta^t c(\omega^{t-1}(\tilde{\omega}), \omega^t(\tilde{\omega})) - \Delta^t \mathcal{R}(\omega^t(\tilde{\omega}))] \mathbf{1}(t < T) \\ \sum_{t=0}^{t(\omega)} [\Delta^t c(\omega^{t-1}(\omega), \omega^t(\omega)) - \Delta^t \mathcal{R}(\omega^t(\omega))] \mathbf{1}(t < T) \\ + \Delta^{\bar{t}(\omega, \tilde{\omega})} \mathcal{R}(\omega^{\bar{t}(\omega, \tilde{\omega})}(\omega)) \mathbf{1}(t(\omega) = t(\tilde{\omega}) = \bar{t}(\omega, \tilde{\omega}) < T) \end{array} \right\}$$

The expression $\mathbf{1}(p)$ is 1 if p is true and 0 otherwise. The expression $\mathbf{1}(t < T)$ is there just to be sure we do not subtract $\Delta^T \mathcal{R}(\omega_c)$. The final line corrects for a potential double counting in the first two lines, and $\mathbf{1}(\bar{t}(\omega, \tilde{\omega}) < T)$ says that term is not relevant in the case where $\bar{t}(\omega, \tilde{\omega}) = T$. Given this we can then immediately derive a formula for coheight:

Proposition 2 For given $\tilde{\omega}$, let $\{t(\tilde{\omega}), t(\omega)\} \in \{0, 1, 2, \dots, \bar{t}(\omega, \tilde{\omega})\}^2$ such that $\max[t(\tilde{\omega}), t(\omega)] = \bar{t}(\omega, \tilde{\omega})$ then:

$$Ch(\omega, \tilde{\omega}) = \max_{t(\omega), t(\tilde{\omega})} \left\{ \begin{array}{l} \sum_{t=0}^{t(\tilde{\omega})} [\Delta^t \mathcal{R}(\omega^t(\tilde{\omega})) \mathbf{1}(t < T) - \Delta^t c(\omega^{t-1}(\tilde{\omega}), \omega^t(\tilde{\omega}))] \\ + \left(\sum_{t=t(\omega)+1}^{\bar{t}(\omega, \tilde{\omega})} [\Delta^t c(\omega^{t-1}(\omega), \omega^t(\omega)) - \Delta^t \mathcal{R}(\omega^t(\omega)) \mathbf{1}(t < T)] \right) \mathbf{1}(t(\omega) < \bar{t}(\omega, \tilde{\omega})) \\ - \Delta^{\bar{t}(\omega, \tilde{\omega})} \mathcal{R}(\omega^{\bar{t}(\omega, \tilde{\omega})}(\omega)) \mathbf{1}(t(\omega) = t(\tilde{\omega}) = \bar{t}(\omega, \tilde{\omega}) < T) \end{array} \right\} \quad (32)$$

and $Ch(\omega) = \max_{\tilde{\omega} \in \Omega \setminus \omega} Ch(\omega, \tilde{\omega})$.

This formula is not simple. Consider the case where $\omega \subseteq \omega_c$; we will immediately have $t(\omega) = \bar{t}(\omega, \tilde{\omega})$ and the maximization over $t(\tilde{\omega})$ is unconstrained. How could this cause a problem? Well assume $T > 1$ and consider a limit set for which $\Delta^{T-1} \mathcal{R}(\omega^{T-1}(\tilde{\omega})) - \Delta^{T-1} c(\omega^{T-2}(\tilde{\omega}), \omega^{T-1}(\tilde{\omega})) - \Delta^T c(\omega^{T-1}(\tilde{\omega}), \omega^T(\tilde{\omega}))$ is negative and very small. In the coheight we do not have to take that steps and thus this might determine the coheight of all limit sets in the core.

In the very common case where $T = 1$ this formula becomes simple for $\omega \subseteq \omega_c$:

Corollary 3 If $T = 1$ and $\omega \subseteq \omega_c$ $Ch(\omega) = \max_{\tilde{\omega} \in \Omega \setminus \omega} \mathcal{R}(\tilde{\omega})$.

Formula (32) shows that a specific intuition that physicists developed in the analysis of phase changes is completely general. The terminology used is motivated by a simple practical experiment. The experiment is to put a bottle of water in a freezer for a specific amount of time. If the timing is just right when one takes the bottle out of the freezer it will have no ice in it, and if one taps the side lightly a little particle of ice will form and the entire bottle will turn to ice. In this experiment a *critical droplet* forms and from there the phase change is instantaneous. Allegorically then the question is how long it takes to form this critical droplet. This leads them to interest in a *critical path* which is the best way to create this droplet. In this analysis there is a Hamiltonian (which is inversely related to the Potential.) The critical path minimizes the maximum value of the Hamiltonian in the transition between $\tilde{\omega}$ and ω , and the critical droplet is the state that achieves this maximum.

The problem (from our point of view) is that this critical droplet can not be a limit set, and the critical path probably passes through limit sets but those will be just some of the states, so why are limit sets interesting? This is why Arous and Cerf (1996) dispenses with finding limit sets; it just finds the level sets of the Hamiltonian. This representation says that in fact both the critical path and the critical droplet can be characterized using limit sets and the cost function. To be precise this representation implies a path, $\omega_s(\omega, \tilde{\omega})$ ($s \in \{0, 1, 2, \dots, S\}$) through the limit sets, and shows that without loss of generality we can further assume that if $s < \bar{t}(\omega, \tilde{\omega})$ then $\omega_s(\omega, \tilde{\omega}) \in \Omega^s$ and if $s \geq \bar{t}(\omega, \tilde{\omega})$ $\omega_s(\omega, \tilde{\omega}) \in \Omega^{\bar{t}(\omega, \tilde{\omega}) + (\bar{t}(\omega, \tilde{\omega}) - s)}$, thus $S = 2\bar{t}(\omega, \tilde{\omega})$. A problem which might have an unbounded number of elements can be characterized as having less than $2T$ elements. And what is the critical droplet? Well if $t(\tilde{\omega}) < \bar{t}(\omega, \tilde{\omega})$ then it is $s_* = t(\tilde{\omega})$, otherwise it is $s_* = \bar{t}(\omega, \tilde{\omega}) + (\bar{t}(\omega, \tilde{\omega}) - t(\omega)) + 1$. (Note that $t(\omega) \geq 1$).

This insight highlights a strange aspect of this estimate of waiting time. One notices that our definition of the critical droplet is equivalent to saying that the cost of all states further along in the path are discarded. Why? Again we can turn to physics for an explanation. For each $\omega_s(\omega, \tilde{\omega})$ there is a log waiting time

for getting to $\hat{\omega} \in \Omega$, call this $\ln \tau(\hat{\omega}, \omega_s(\omega, \tilde{\omega}))$, if $s \geq s_*$ what this means is that $\ln \tau(\omega, \omega_s(\omega, \tilde{\omega})) < \ln \tau(\tilde{\omega}, \omega_s(\omega, \tilde{\omega}))$. This is related to the cost in the standard manner, the optimal path going to ω from $\omega_s(\omega, \tilde{\omega})$ has a lower cost than the optimal path of going to $\tilde{\omega}$. But if $\ln \tau(\omega, \omega_s(\omega, \tilde{\omega})) < \ln \tau(\tilde{\omega}, \omega_s(\omega, \tilde{\omega}))$ then:

$$\lim_{\beta \rightarrow \infty} \frac{e^{\beta \ln \tau(\omega, \omega^s(\omega, \tilde{\omega}))}}{e^{\beta \ln \tau(\tilde{\omega}, \omega^s(\omega, \tilde{\omega}))}} \rightarrow 0, \quad (33)$$

or the relative amount of time it takes to go from $\omega_s(\omega, \tilde{\omega})$ to ω is zero. Thus we ignore the cost of those states because, in short, they were going to ω already.

Now the practical experiment indicates that β can be large enough that this bound is correct. However this discussion might make one think that for some problems it is not. If we are considering a directed minimal cost spanning tree problem then there is no β . What bound might one be interested in? Well we will have an implicit distribution of goods that have to be transferred to ω , $F(\tilde{\omega}|\omega)$, and we can normalize it so that $\int_{\tilde{\omega} \in \Omega} F(\tilde{\omega}|\omega) = 1$. Then what one cares about is $\int_{\tilde{\omega} \in \Omega} c(\omega, \tilde{\omega}) F(\tilde{\omega}|\omega)$. The representation theorem tells us $c(\omega, \tilde{\omega})$, and we might be interested in an upper bound for this expectation which is the *censored coradius*:

$$\int_{\tilde{\omega}} c(\omega, \tilde{\omega}) F(\tilde{\omega}|\omega) \leq \max_{\tilde{\omega} \in \Omega \setminus \omega} c(\omega, \tilde{\omega}) = \overline{CR}(\omega) = \max_{\tilde{\omega} \in \Omega \setminus \omega} \sum_{t=1}^{\bar{t}(\omega, \tilde{\omega})} \Delta^t c(\omega^t(\omega), \omega^{t+1}(\omega)) + \sum_{t=0}^{\bar{t}(\omega, \tilde{\omega})} \Delta^t \mathcal{R}(\omega^t(\tilde{\omega})) \mathbf{1}(t < T). \quad (34)$$

We refer to this as a coradius because it is a generalization of the modified coradius as defined by Ellison (2000), if $T = 1$ then it is the modified coradius. We call it censored because the first term (the cost of going to ω) requires no analysis. It is known, one only has to (essentially) find the maximizer of the second term—the total radius of the alternative limit set. Knowing the emergent seed has reduced the amount of computation. In Ellison (2000) the first term requires analysis for every single $\tilde{\omega}$ —one has to solve $|\Omega \setminus \omega|$ shortest path problems. However since we know the emergent seed this part requires no analysis. Indeed for most limit sets finding the censored coradius will be simple. Define:

$$\Omega_{T\mathcal{R}} = \arg \max_{\tilde{\omega} \in \Omega} \sum_{t=0}^{T-1} \Delta^t \mathcal{R}(\omega^t(\tilde{\omega})) \quad (35)$$

Then it is immediate that:

Lemma 4 *If there is an $\tilde{\omega} \in \Omega_{T\mathcal{R}}$ such that $\omega^{T-1}(\tilde{\omega}) \neq \omega^{T-1}(\omega)$ then the censored coradius of ω is determined by that $\tilde{\omega}$.*

Strictly speaking we do not have to prove that this bound on log waiting times matters. Since we have a precise measure of waiting time all we need to show is that the censored coradius is a bound for the coheight.

Lemma 5 *For all $\omega \in \Omega$, $\overline{CR}(\omega) \geq Ch(\omega)$, and a sufficient condition for them to be equal is $\omega \subseteq \omega_c$ and either $T = 1$ or the limit set that determines the coheight is in the core.*

Proof. *In the formula for $Ch(\omega, \tilde{\omega})$ we first consider dropping all the negative terms,*

$$- \left(\sum_{t=0}^{t(\tilde{\omega})-1} \Delta^t c(\omega^{t-1}(\tilde{\omega}), \omega^t(\tilde{\omega})) \right) \mathbf{1}(t(\tilde{\omega}) > 0) - \left(\sum_{t=t(\omega)}^{\bar{t}(\omega, \tilde{\omega})} \Delta^t \mathcal{R}(\omega^t(\omega)) + \Delta^{t(\omega)} \mathcal{R}(\omega^{t(\omega)}(\omega)) \right) \mathbf{1}(t(\omega) < \bar{t}(\omega, \tilde{\omega})). \quad (36)$$

and we immediately realize that maximization requires $t(\omega) = 0$, $t(\tilde{\omega}) = \bar{t}(\omega, \tilde{\omega})$. But then the formula is the censored coradius for given $\tilde{\omega}$ and thus the censored coradius is an upper bound for the coheight.

The first sufficient condition is obvious. To prove the second we notice that if ω is in the core then $\sum_{t=0}^T \Delta^t c(\omega^{t-1}(\omega), \omega^t(\omega)) = 0$ and $\overline{CR}(\omega) = \max_{\tilde{\omega} \in \Omega \setminus \omega} \sum_{t=0}^{\bar{t}(\omega, \tilde{\omega})} \Delta^t R(\omega^t(\tilde{\omega}))$. In order for this to be equal to the coheight in this case we need that the maximizer of the coheight is also in the core, which implies $\sum_{t=0}^T \Delta^t c(\omega^{t-1}(\tilde{\omega}), \omega^t(\tilde{\omega})) = 0$. ■

It is still worthwhile to prove the censored coradius can also be derived from primitive analysis. In the appendix (Section 9) there is a proof based on Bortz, Kalos, and Lebowitz (1975). We can also point out that a special result in Myatt and Wallace (2008b) is potentially general. We thank the authors for discovering this curiosity.

Lemma 6 *Let $\tilde{\omega}$ determine the censored coradius of ω . If $\tilde{\omega} \subseteq \omega_c$ and $\omega^{T-1}(\omega) \neq \omega^{T-1}(\tilde{\omega})$ then ω is stochastically stable if and only if its total radius is higher than its censored coradius.*

Proof. *The sufficiency is clear by a simple root switching argument, given we know the emergent seed we can now choose the path it implies. The necessity is because for this critical $\tilde{\omega}$ we must have:*

$$\begin{aligned} \sum_{t=0}^{T-1} \Delta^t \mathcal{R}(\omega^t(\tilde{\omega})) &\leq \sum_{t=0}^{T-1} \Delta^t \mathcal{R}(\omega^t(\omega)) - \sum_{t=0}^{T-1} \Delta^{t+1} c(\omega^t(\omega), \omega^{t+1}(\omega)) \\ \overline{CR}(\omega) &= \sum_{t=0}^{T-1} \Delta^{t+1} c(\omega^t(\omega), \omega^{t+1}(\omega)) + \sum_{t=0}^{T-1} \Delta^t \mathcal{R}(\omega^t(\tilde{\omega})) \leq \sum_{t=0}^{T-1} \Delta^t \mathcal{R}(\omega^t(\omega)) . \end{aligned} \quad (37)$$

■

It is important to mention some points in Kreindler and Young (2013) and Levine and Modica (2014b). Kreindler and Young (2013) shows that in the coordination game and the logit model that if $\ln \tau(\omega, \beta)$ is the true waiting time then for small enough β this will be independent of the population size, thus very fast. Levine and Modica (2014b) points out that if one enters an intermediate limit set the most likely event is that you stay there, thus a transition between limit sets that avoids intermediate limit sets may be shorter for high levels of noise (small β). Combined these points might imply that the critical path for small β might not be either of the ones we have identified. These are both limiting paths where, as Ellison (2000) says, there is a slingshot effect from intermediate limit sets.

6 Applications

In our first three applications we will use BRM in a two role strategic form game, thus we will briefly explain this model first. Our analysis of Myatt and Wallace (2008b) must use a different model, which we shall explain in that section. Our next step will to be find the radii of pure strategy Nash equilibria under BRM. We will then turn to each of our applications. After each application we will review papers with similar emergent seeds. This survey will only cover problems with a finite state space (population) and three or more limit sets.

6.1 Model

In the first three applications we will analyze two role strategic form games. The game will be $G = (I, S, u)$ where $I = \{1, 2\}$, $S = S_1 \times S_2$ and $u : S \rightarrow \mathbb{R}^2$, and there will be n agents in each of the two populations. Players will be matched into pairs uniformly. The decision rule shall be that with probability $\rho \in (0, 1)$ agents will use the same action they did last period. With probability $(1 - \rho)$ they will choose a new strategy. When choosing a strategy with probability $1 - e^{-\beta}$ ($\beta > 0$) they will choose a best response to the current distribution of strategies, and with probability $e^{-\beta}$ they will choose a new strategy with equal likelihood. We will analyze the model as $\beta \rightarrow \infty$ for large n .

6.1.1 The Radii of Pure Strategy Nash Equilibria in a Strategic Form Game

First of all, s is a pure strategy Nash equilibrium if there is a $p^* \in [0, 1]$ such that if the distribution of strategies puts weight $1 - p$ on s and $p \leq p^*$ then s is a best response to this distribution, otherwise it is not. We say s is a *strict* Nash equilibrium if $p^* > 0$. Clearly $p^* = \min [p_1^*, p_2^*]$, where p_i^* will be the minimal amount needed to make s_i not a best response, and this p^* will also determine the radius of s . To be precise assume that we put weight $p_1(\tilde{s}_2)$ on some other strategy \tilde{s}_2 . Then there will be different strategy $s_1(s, \tilde{s}_2)$ that will be a best response if $p_1(\tilde{s}_2)$ is large enough or:

$$(1 - p(\tilde{s}_2)) u_1(s) + p(\tilde{s}_2) u_1(s, \tilde{s}_2) \leq (1 - p(s_2)) u_1(s_1(s, \tilde{s}_2), s_2) + p(s_2) u_1(s_1(s, \tilde{s}_2), \tilde{s}_2) \quad (38)$$

and thus the minimal value of $p(\tilde{s}_2)$ where s_1 is not a best response is:

$$p(\tilde{s}_2) = \min \left[\frac{(u_1(s) - u_1(s_1(s, \tilde{s}_2), s_2))}{(u_1(s) - u_1(s_1(s, \tilde{s}_2), s_2)) + (u_1(s_1(s, \tilde{s}_2), \tilde{s}_2) - u_1(s_1, \tilde{s}_2))}, 1 \right], \quad (39)$$

where the bound handles the case that s_1 is a best response to \tilde{s}_2 . The direct jump lemma in Kandori and Rob (1995) allows us to realize that:

$$\mathcal{R}(s) = \left\lceil n \min_{i \in \{1, 2\}} \min_{s_i \in S_i} p(s_i) \right\rceil. \quad (40)$$

Where $\lceil x \rceil$ is the smallest integer greater than or equal to x , and we can ignore this if n is large enough. We can simply define $R(s)/n = \min_{i \in \{1, 2\}} \min_{s_i \in S_i} p(s_i)$ and notice that for large enough n if $R(s)/n \leq R(\hat{s})/n$ then $R(s) \leq R(\hat{s})$. We shall use this normalization with cost functions, and write $c(x, y)/n$ when we do.

This analysis makes clear the difference between static analysis and stochastic evolution. In a static analysis only the best strategies matter—where the definition of best depends on the decision rule. Here there are two other sets of strategies that are important. First are the *best invaders*. These are the strategies that determine the cost function. Denote this set $BI(s)$. The strategies that are best responses at the critical $p(\tilde{s}_2)$ we denote the best response to invaders, $BRI(s)$. The latter set would be rational but there is no reason the first set should be. If one wishes to assume rationality this would restrict the sets of best invaders.

Notice that if s is a strict pure strategy Nash equilibrium of G then it is a limit set because $R(s) > 0$. If it is not a strict equilibrium then it is not a limit set because $R(s) = 0$. One might wonder why we can not include mixed strategy Nash equilibria in this analysis. In principle there would be a set $U(s)$ that $u(s)$ could be in and we would simply also minimize over $u(s) \in U(s)$. Unfortunately in general the set $U(s)$ is unknown.

6.2 Linear Emergent Seeds—The Bargaining Game

We represent a bargaining problem as a pair of concave and strictly increasing utility functions, $u_i(x)$. The actions of the two agents shall be the unit interval, $A_1 = A_2 = [0, 1]$ and a pair (a_1, a_2) is feasible if $a_1 + a_2 \leq 1$. For a feasible pair i gets $u_i(a_i)$, if a pair is unfeasible i gets zero. This is a standard characterization given we have normalized the disagreement point to zero. We shall assume there is an open set of feasible (a_1, a_2) such that $\min_i u_i(a_i) > 0$. We shall want to have a finite number of strategies, thus assume there is some $\lambda > 0$ such that $\frac{1}{\lambda}$ is an integer and $A_i = \{0, \lambda, 2\lambda, \dots, 1\}$.

In the bargaining game (or the Nash Demand game) $S_i = A_i$, or each role simply submits a demand s_i and if (s_1, s_2) is feasible then i gets $u_i(s_i)$, otherwise they receive zero. The set of strict pure strategy Nash equilibria are (s_1, s_2) such that $s_1 + s_2 = 1$ and $\min_i u_i(s_i) > 0$, and it is simple to show these are also the limit sets (see Young, 1993b).

Now we want to determine the radii of the limit sets. If population one is currently demanding $s_1 > 0$ there are two basic strategies for the agents who error to follow. They can demand either more or less. If they demand more then $p_1^+(k)$ is

$$(1 - p_1^+(k)) u_1(s_1) = u_1(s_1 - k\lambda) \quad (41)$$

because if the players in population one reduce their demand they will get their demand from everyone.

$$p_1^+(k) = (u_1(s_1) - u_1(s_1 - k\lambda)) / u_1(s_1) \quad (42)$$

and obviously this is minimized when $k = 1$, thus $p_1^+ = (u_1(s_1) - u_1(s_1 - \lambda)) / u_1(s_1)$. If they demand less then the formula for $p_1^-(s_2)$ is:

$$\begin{aligned} (1 - p_1^-(s_2)) u_1(s_1) &= p_1^-(s_2) u_1(1 - s_2) \\ p_1^-(s_2) &= u_1(s_1) / (u_1(1 - s_2) + u_1(s_1)) \end{aligned} \quad (43)$$

and again it is obvious that $s_2 = 0$, thus $p_1^- = u_1(s_1) / (u_1(1) + u_1(s_1))$. Now as $\lambda \rightarrow 0$ either p_1^+ or p_2^+ or both go to zero, while p_1^- and p_2^- are constant and large. Thus for small enough λ :

$$\mathcal{R}(s) / n = \min \left[\frac{u_1(s_1) - u_1(s_1 - \lambda)}{u_1(s_1)}, \frac{u_2(s_2) - u_2(s_2 - \lambda)}{u_2(s_2)} \right] \quad (44)$$

Now let us write a limit set as $(\gamma, 1 - \gamma)$ and denote it by the share of player one, or γ then:

Lemma 7 *From each limit set γ we can transition to $\tilde{\gamma} = \gamma - \lambda$, if $p_1^+ \leq p_2^+$ and $\tilde{\gamma} = \gamma + \lambda$ if $p_1^+ \geq p_2^+$. Thus the radius is first strictly increasing and then strictly decreasing, and the limit sets with the highest radius are also in the core. For generic u one of them will be stochastically stable.*

In Young (1993b) it is shown that the stochastically stable limit set will be the Nash Bargaining Solution for the finite grid of strategies, or it will be

$$\gamma_{NBS} \in \arg \max_{\gamma \in \{0, \lambda, 2\lambda, \dots, 1\}} u_1(\gamma) u_2(1 - \gamma) . \quad (45)$$

6.2.1 On the Speed and Dynamics of Evolution in the Bargaining Game

One of the advantages of the emergent seed methodology is that knowing the emergent seed the modified coradius becomes the censored coradius. In this game we have core dominance thus we have a simple formula for the censored coradius, which since γ_{NBS} is in the core is also the coheight:

$$\overline{CR}(\gamma_{NBS}) = Ch(\gamma_{NBS}) = CR^*(\gamma_{NBS}) = \max [c(\gamma_{NBS}, \gamma_{NBS} + \lambda), c(\gamma_{NBS}, \gamma_{NBS} - \lambda)] . \quad (46)$$

Now as $\lambda \rightarrow 0$ $Ch(\gamma_{NBS}) \rightarrow 1$, or evolution will be very fast. However when $R(\gamma_{NBS}) = 1$ all limit sets will be stochastically stable. But what of right before this limit? When λ is small but not too small relative to n . It is not that hard to analyze this case. Our analysis implies that the most likely events are that we go to either $\gamma + \lambda$ or $\gamma - \lambda$. This implies moving towards the core is more likely than moving away from the core—and the relative likelihood of going to the core is increasing in the distance from the core. Indeed from Cui and Zhai (2010) we see that the most likely medium run prediction is that one is in the core, near the stochastically stable state. Furthermore the likelihood that one is in the state $\gamma_{NBS} + k\lambda$ or $\gamma_{NBS} - k\lambda$ will be strictly decreasing in k . Thus the γ series will look much like many price series. It will be a random walk with the distribution of error terms being biased towards γ_{NBS} .⁵ The coincidence of this model and empirical observation is deserving of more analysis.

6.2.2 Other papers with Linear Emergent Seeds

Young (1993b) has created a paradigm of applications of the bargaining game. For example Saez-Marti and Weibull (1999) consider what might happen if some agents best respond to other agents' models. Robles (2008) considers agents who have the outside option of waiting until the next period. Agastya (1999) and Newton (2012) extends the model to games in coalition form. Agastya (2004) considers double auctions. In independent research both Ellingsen and Robles (2002) and Troger (2002) analyze bargaining after an a-priori investment. They find that evolution solves the hold up problem. Dawid and Macleod (2008) considers both joint investment and a probability that investments fail. That paper finds the hold up problem returns, and Andreozzi (2012) shows that this is because of the uncertainty of the outcome (or cost) of the investment.

There are other papers that are not bargaining problems where nonetheless the emergent seed is linear. Naidua, Hwang, and Bowles (2010) analyzes the contract game assuming that agents only deviate to contracts that would be better for them. This implies that the best invader demands only a little more than that role is currently receiving, transforming a star emergent seed into a linear emergent seed. Robles (1997) analyzes population games where your payoff depends on your action and an order statistic of the population. The paper follows Young (1993b) by constructing the emergent seed and showing that the radius is increasing as one moves towards the core.

Several of these articles addresses *location problems*. In these problems there are two extreme locations and intermediate ones. For example Jackson and Watts (2002a) analyzes the coordination game where agents choose a network as well as a strategy. That article assumes that a link must be paid for by both agents, and that both must agree to form the link. Two limit sets are complete networks with agents coordinating on either strategy, and then there are a set of two network limit sets where in different networks agents

⁵For large and small γ it is possible that $\gamma + \lambda$ or $\gamma - \lambda$ is not a strict Nash equilibrium, then one can go anywhere from such a $\gamma + \lambda$ or $\gamma - \lambda$.

coordinate on different strategies. The key result is that these intermediate limit sets are only two errors from either complete network. The first error has someone choose the wrong strategy and be isolated, and in the next error someone changes strategy and joins them. Thus from either complete network one goes to the nearest intermediate limit set, and then to one closer to the other complete network. The emergent seed is linear but the model is indeterminate. Anwar (2002) analyzes choosing locations and strategies in the coordination game where the location can be one of two islands with limited capacity. The limit sets are both islands using the same strategy and one island using each strategy. From either of the extremes one moves to the intermediate state, and either the intermediate state or the risk dominant strategy is stochastically stable. Gerber and Bettzüge (2007) considers the evolution of two markets. The limit sets are all using one market or half using each market. All using the same market always goes to half using each market, and that is stochastically stable and Pareto dominant.

6.3 Star Emergent Seeds—The Contract Game

The contract game is similar to the bargaining game, but the difference has a large impact. In the contract game $S_i = A_1 \times A_2$, and if $s_1 = s_2$ then there is agreement, otherwise both parties get zero. A contract is complete; it specifies the payoffs of all parties. In all other respects the model is the same. Notice this is also a structural model of a pure coordination game.

Of course this change in strategies does have a dramatic effect on the set of strict pure strategy Nash equilibria. Now any $s_1 = s_2 = s$ is a strict Nash equilibrium and a limit set as long as $\min [u_1(s), u_2(s)] > 0$. Like before agents who error can either offer more to one party or less, but unlike before there is no option to compromise—an agent can not make an agreement both with the invaders and the people playing s . Thus assume the invaders offer a contract where role 1 is allowed to claim $\tilde{s}_1 \neq s_1$, players are best responding to accept \tilde{s}_1 if:

$$\begin{aligned} (1 - p_1(\tilde{s}_1)) u_1(s_1) &= p_1(\tilde{s}_1) u_1(\tilde{s}_1) \\ p_1(\tilde{s}_1) &= u_1(s_1) / (u_1(s_1) + u_1(\tilde{s}_1)) \end{aligned} \tag{47}$$

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

$$\mathcal{R}(s)/n = \min \left[\frac{u_1(s_1)}{u_1(s_1) + u_1(1)}, \frac{u_2(s_2)}{u_2(s_2) + u_2(1)} \right]. \tag{48}$$

At this point we must consider two separate cases. The simpler one is when either $(0, 1)$ or $(1, 0)$ are not equilibria, and we address this second. We will now analyze the harder case where *disagreement is irrelevant*, or $u_1(0) > 0$ and $u_2(0) > 0$. Under this assumption any contract is better than disagreement.

6.3.1 If Disagreement is Irrelevant

This implies that all contracts are strict pure strategy Nash equilibria, and is important because the core in this model is $(0, 1)$ and $(1, 0)$. Now we need to calculate the cost of going from something in the core to a given s . You are leaving the state where either role one is getting $u_1(0)$ or role two is getting $u_2(0)$, and thus the relevant p is:

$$p_i(s_i) = u_i(0) / (u_i(0) + u_i(s_i)) \tag{49}$$

Thus

$$\Delta c(s, \omega_c)/n = \min \left[\frac{u_1(0)}{u_1(0) + u_1(s_1)} - \frac{u_1(0)}{u_1(0) + u_1(1)}, \frac{u_2(0)}{u_2(0) + u_2(s_2)} - \frac{u_2(0)}{u_2(0) + u_2(1)} \right] \quad (50)$$

Now for limit sets in the core their stochastic stability will be determined by their radii, but they usually will not be stochastically stable. From now on we will assume they are not. Thus for $s \notin \omega_c$ what we care about is the sum of the first difference cost of going from the core to the limit set and it's radius. Or:

$$\begin{aligned} \mathcal{R}(s)/n - \Delta c(s, \omega_c)/n &= \min \left[\frac{u_1(s_1)}{u_1(s_1) + u_1(1)}, \frac{u_2(s_2)}{u_2(s_2) + u_2(1)} \right] \\ &\quad - \min \left[\frac{u_1(0)}{u_1(0) + u_1(s_1)} - \frac{u_1(0)}{u_1(0) + u_1(1)}, \frac{u_2(0)}{u_2(0) + u_2(s_2)} - \frac{u_2(0)}{u_2(0) + u_2(1)} \right]. \end{aligned} \quad (51)$$

It is also clear that $s_1 + s_2 = 1$ —this increases $\mathcal{R}(s)/n$ and decreases $\Delta c(s, \omega_c)/n$ —or that stochastically stable limit sets must be Pareto efficient. Like before let γ be the share of the first player, and we will write a limit set as γ .

Consider if both $\mathcal{R}(\gamma)/n$ and $\Delta c(\gamma, \omega_c)/n$ are both determined by the same agent's utility, for example agent 1, then:

$$\mathcal{R}(\gamma)/n - c(\gamma, \omega_c)/n = \frac{u_1(\gamma)}{u_1(\gamma) + u_1(1)} - \frac{u_1(0)}{u_1(0) + u_1(\gamma)} + \frac{u_1(0)}{u_1(0) + u_1(1)}, \quad (52)$$

and this is strictly increasing in γ . Now we notice that in order for $\mathcal{R}(\gamma)/n$ to be determined by person 1 then $\gamma \leq \gamma_{KS}$ where:

$$\frac{u_1(\gamma_{KS})}{u_1(1)} = \frac{u_2(1 - \gamma_{KS})}{u_2(1)}, \quad (53)$$

this is the Kalai-Smordinsky solution. In order for $\Delta c(s, \omega_c)/n$ to be determined by person 2 we must have $\gamma \geq \gamma_{WB}$ where:

$$\frac{u_1(0)}{u_1(0) + u_1(\gamma_{WB})} - \frac{u_1(0)}{u_1(0) + u_1(1)} = \frac{u_2(0)}{u_2(0) + u_2(1 - \gamma_{WB})} - \frac{u_2(0)}{u_2(0) + u_2(1)}. \quad (54)$$

This results in a general characterization of stochastic stability.

$$\gamma_* \in \arg \max_{\min[\gamma_{KS}, \gamma_{WB}] \leq \gamma \leq \max[\gamma_{KS}, \gamma_{WB}]} \left\{ \begin{array}{ll} \frac{u_1(\gamma)}{u_1(\gamma) + u_1(1)} - \frac{u_2(0)}{u_2(0) + u_2(1 - \gamma)} + \frac{u_2(0)}{u_2(0) + u_2(1)} & \text{if } \gamma_{WB} < \gamma_{KS} \\ \frac{u_2(1 - \gamma)}{u_2(1 - \gamma) + u_2(1)} - \frac{u_1(0)}{u_1(0) + u_1(\gamma)} + \frac{u_1(0)}{u_1(0) + u_1(1)} & \text{if } \gamma_{WB} > \gamma_{KS} \\ \gamma_{KS} & \text{if } \gamma_{WB} = \gamma_{KS} \end{array} \right\}. \quad (55)$$

This characterization will be precise if $[\min[\gamma_{KS}, \gamma_{WB}], \max[\gamma_{KS}, \gamma_{WB}]]$ contains a limit set.⁶ Otherwise it will be one of the limit sets that are closest to this region. This is the Kalai-Smordinsky solution only if $\gamma_{WB} = \gamma_{KS}$. Young (1998a) characterized stochastic stability as the Kalai-Smordinsky solution. There are two reasons for this difference. First our model is slightly different, there contracts are allowed to be from an open rectangle that is bounded by the origin. More importantly that characterization requires the set of contracts to be large. This causes convergence to the case where disagreement is relevant.

This is a good point to clarify the analytic simplicity of using our implementation of Edmonds' algorithm versus the standard one. With the standard one things that are not in a cycle must be considered in the next

⁶One must check this solution has a higher likelihood potential than the limit sets in the core.

iteration, and this is analytically inefficient. To see this notice that in the first iteration one would find the core. But all of the limit sets that were not in the core would be left for the next iteration. Generically the core would only go to one of these limit sets in the next iteration, and so only one would get removed at each stage. Thus only one limit set would be added to the structure at each step, since there are $\frac{1}{2} \left(\frac{1}{\lambda} + 1 \right) \left(\frac{1}{\lambda} + 2 \right)$ limit sets we would need $\frac{1}{2\lambda^2} (3\lambda + 1)$ iterations.

The Censored Coradius and the Coheight Since the stochastically stable limit set is not in the core one should expect the coheight is strictly lower than the censored coradius. The coradius is quite simple:

$$\overline{CR}(\gamma_*) = \Delta c(\gamma_*, \omega_c) + \mathcal{R}(\gamma_{KS}) . \quad (56)$$

Notice that $\overline{CR}(\gamma_*) > \mathcal{R}(\gamma_{KS}) \geq \mathcal{R}(\gamma_*)$ so in this problem the coradius/modified coradius test will never be successful. Now the coheight will be lower because, in essence, we have the choice of whether to have $\tilde{\gamma}$ go to the core or γ_* . The coheight is:

$$Ch(\gamma_*) = \max_{\tilde{\gamma} \neq \gamma_*} \max [\Delta c(\gamma_*, \omega_c) - \Delta c(\tilde{\gamma}, \omega_c), 0] + \mathcal{R}(\tilde{\gamma}) . \quad (57)$$

Now at the critical $\tilde{\gamma}$ if $\Delta c(\gamma_*, \omega_c) - \Delta c(\tilde{\gamma}, \omega_c) \leq 0$ then obviously $\mathcal{R}(\tilde{\gamma}) \leq \mathcal{R}(\gamma_{KS})$ and $Ch(\gamma_*) < \overline{CR}(\gamma_*)$. Otherwise $\mathcal{R}(\tilde{\gamma}) - \Delta c(\tilde{\gamma}, \omega_c) < \mathcal{R}(\gamma_{KS})$ and we still have $Ch(\gamma_*) < \overline{CR}(\gamma_*)$.

Evolution in this model is counter intuitive. Most of the time we will bounce back and forth between $(0, 1)$ and $(1, 0)$ and if we do not then we will exit to some (s_1, s_2) . In fact from Cui and Zhai (2010) we know that the next most likely event is to either go to $(\lambda, 1 - \lambda)$ or $(1 - \lambda, \lambda)$. Society will not necessarily ever be near γ_* , its just that it will either show up more often or stay around longer than all other options.

6.3.2 If Disagreement is Relevant

It is a little counter-intuitive how the model changes if we allow $u_1(0) \leq 0$ or $u_2(0) \leq 0$. For simplicity let us assume that $u_1(0) \leq 0 < u_2(0)$. This has no impact on the best invaders, they will still offer $(0, 1)$ or $(1, 0)$. But then if we transition from s to $(0, 1)$ this is not a strict pure strategy Nash equilibrium. In other words when agents in role one optimize they will be able to choose any other strategy. Thus we transition from $(0, 1)$ to any feasible s . Furthermore, we will always transition to $(0, 1)$ in one or two steps. We might first go to $(1, 0)$ but from there we will go to $(0, 1)$. Thus every limit set is in the core of an emergent seed.

This makes our life simpler. Since every limit set is in the core of an emergent seed all we have to do is maximize the radius of a limit set. Like before it will be Pareto efficient and in fact it will solve:

$$\max_{\gamma \in \{0, \lambda, 2\lambda, \dots, 1\}} \min \left[\frac{u_1(\gamma)}{u_1(1)}, \frac{u_2(\gamma)}{u_2(1)} \right] \quad (58)$$

which is the Kalai-Smordinsky solution on the finite grid.

Obviously this sensitivity of the solution to an—apparently—random parameter is of concern. For every game with $u_1(0) > 0$ there is a nearby game where we increase the disagreement payoff for role one by $u_1(0)$, and in this nearby game the stochastically stable limit set might be different. However we point out that it is the use of the word random that is inappropriate. In stochastic evolution we care not only about best responses but best the best response to best invaders. If two models have different best responses to the best invaders we should expect the results to change.

In static analysis a natural assumption is that the disagreement point is a relevant alternative, or $u_1(0) < 0$ and $u_2(0) < 0$. In this case our results are the same as Young (1998a) except that ours is not a limiting argument.

6.3.3 Star Emergent Seeds

Most star emergent seeds have the *global attractor property*. In other words from every limit set the radius is determined by a transition to the same key limit set, and this limit set has the highest radius. In models of imitation, as pointed out by Alos-Ferrer and Netzer (2012), this is equivalent to having a global evolutionarily stable strategy or GESS (Schaeffer, 1988). This is a strategy that always does better in a population of two strategies. In games with a global attractor the benefit of the emergent seed is minimal, one skips a simple root switching argument at the end of the proof.

Kandori and Rob (1995) were the first to find a star. In symmetric pure coordination game the best invader is always the Pareto dominant Nash equilibrium and this also has the highest radius. Dutta and Prasad (2002) analyzes risk sharing contracts in a moral hazard environment. The stochastically stable contract is half dominant, however the radius of each other contract is determined by going to the stochastically stable contract. Jackson and Watts (2002b) solves the co-author problem. They show that the complete network—which is Pareto dominated—determines the radius of all other networks and has the highest radius.

Levine and Modica (2014a) construct a model of state competition where states have a power coefficient and an amount of land. The resistance is very general, it only requires that in a conflict if the outcome is not pre-determined then the probability a state wins is increasing in their own power and land, and decreasing in the other sides. Despite the generality of the model the emergent seed is always a star. First of all it should be clear that when one state has all of the land (is a hegemony) its radius is determined by a fanatic band, one with an unsustainable level of state power. But then, as stated in the proof, "all roads lead to Rome." A potential state with the maximum sustainable power will pick up the pieces. Thus from every limit set (hegemony) one goes to the stochastically stable hegemony. Of course the radius/coradius test also verifies that states with a maximum sustainable power are stochastically stable.

A great deal of analysis has focused on the imitation dynamic in aggregator games, understanding these is simple once one recognizes that there is a GESS in most of them. For example Vega-Redondo (1997) analyzes the Cournot game, and in this game the Walrasian equilibrium is a GESS. Tanaka (1999) extends this to firms that have low or high costs, and again there is a GESS—and the same strategy is GESS either in prices or quantities. Hehenkamp and Wambach (2010) analyze product differentiation, and minimal differentiation is a GESS. Alos-Ferrer, Ania, and Schenk-Hoppe (2000) analyze imitation in a standard model of Bertrand Competition. Here there is only a GESS if the number of firms is even. It is a price (near) the price where half the firms produce and make zero profits. If the number of firms is odd then there are a range of prices that all are stochastically stable, but the radius of any other price can be determined by converging to a price in this set.

Goyal and Vega-Redondo (2004) analyzes the coordination game where agents choose networks, but unlike Jackson and Watts (2002a) now only the person who forms the link pays the cost. With this minor modification going from one complete network to the other is as low as the cost of going to any limit set with two networks, and from all those other limit sets one goes to one of the complete networks. Thus a linear emergent seed has been transformed into a star. If the cost of forming a link is low the risk dominant

strategy is selected, if it is high then the Pareto dominant strategy is selected.

In Feri (2007) there are many limit sets in the core because they all have the same stochastic potential. That analysis of one sided network formation finds the stochastically stable network is either a star or complete and all star networks have the same stochastic potential. Every stochastically stable network can be reached from all others with only one decision error.

Myatt and Wallace (2008a) analyzes the volunteering game and here there is a weak global attractor. In every limit set there is one volunteer. Exiting one requires either that the most likely other volunteer offers to volunteer or the current volunteer quits—and anyone would volunteer. If person 1 is the agent most likely to offer to volunteer then from every other limit set either person 1 offers or the current volunteer quits, and person 1 can volunteer. Thus person 1 is in the core, but if this person is more likely to quit than anyone else is to offer to volunteer then this agent will usually not be stochastically stable.

6.4 Other Emergent Seeds—Gift Giving and Contributions

Obviously any classification of structures must end up with a grab bag. Which—in a moment of inspiration—we decided to call "other." There is a common structure here, which is the *short step star*. In a short step star there is a global attractor but one approaches it in a small series of steps. In the gift giving game (Johnson, Levine, and Pesendorfer, 2001) the emergent seeds are either stars or short step stars, but there are multiple emergent seeds. In the contribution game (Myatt and Wallace, 2008b) there is again a weak global attractor. Alos-Ferrer (2004) considers imitation in the Cournot game when agents have (one period) of memory. Core dominance holds but the core is characterized by an interesting seesaw pattern. In Ben-Shoham, Serrano and Volij (2004) the emergent seed generally has more than one level, thus it is in this class as well.

We will analyze the first two mentioned in detail. In the gift giving game our contribution is to characterize stochastic stability for all parameters. In the contribution game Myatt and Wallace (2008b) makes an unnecessary assumption. Knowing the emergent seed and being armed with lemma 6 we can precisely characterize the conditions for cooperation to be successful.

6.4.1 The Gift Giving Game

The gift giving game must, fundamentally, have multiple emergent seeds. No matter what the parameters there are always some limit sets from which one can transition to multiple other limit sets at the same cost. Generically the stochastically stable strategy will be either a *selfish* strategy—which never gives the gift—or a *team* strategy. Johnson, Levine and Pesendorfer (2001) also innovated the assumption that all information sets had to be reached with positive probability in an extensive form game.

Agents live for two periods and 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 a gift or not. 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*, either agents are 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 $\tau : \{r, g\} \times \{0, 1\} \rightarrow \{r, g\}$. There are 64 strategies, but many are equivalent. First of all

either red or green could be good. Evolution can not determine the language society will choose, thus we usually assume green is good, or $a(g) \geq a(r)$. This leaves us with 32 strategies. 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. Of the 16 strategies that remain 12 can not be equilibria because either $\tau(g, 1) = r$ or $\tau(g, 1) = \tau(g, 0)$. Thus only 12 strategies can be Nash equilibria, and 8 are selfish. The remaining four are:

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

When we insert noise into the model tit for tat is not an equilibrium. We need to allow for agents to use different strategies. Johnson, Levine, and Pesendorfer (2001) assumes that each agent has 16 social statuses—one for each transition rule, $f \in \{r, g\}^{16}$. An agent using strategy s then uses the appropriate social status.

We will insert noise into the flag process, with probability $2\eta > 0$ a player's f will be replaced with another one at random. Notice that this means that with probability η they will have the wrong social status, naturally η is small. We want people to be able to expect any state may occur, to do this Φ_t^* be the noiseless distribution of social statuses. Agents may deduce this but not the outcome of the randomizations, thus their beliefs are $\hat{\Phi}_t = E(\Phi_t^*)$ and the probability of any f given $\hat{\Phi}_t$ is strictly positive.

Let us briefly explain why none of the other strategies can be Nash equilibria, and indeed that they are in the basin of attraction of the selfish strategy. Obviously the best response to a generous strategy is a selfish strategy, also if either $\tau(g, 1) = r$ or $\tau(g, 1) = \tau(g, 0)$ the best response is selfish because you either will always or never be rewarded for giving the gift. A best response to tit for tat is always a generous strategy, and this is the best response if there is a positive probability of meeting someone with a red flag. Thus all these strategies are in the basin of attraction of a selfish strategy. We will prove the other strategies are strict Nash equilibria of the noisy game for small η .

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 $\Pr(f_s, f_{s'})$ where this is the probability $f_s \in \{r, g\}$ and $f_{s'} \in \{r', g'\}$. The alternative strategy will always be a strict Nash 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."

$$v(s', s') = \Pr(g, g')((1 - \eta)\alpha - 1) + \Pr(r, g')((1 - \eta)\alpha - 1) + \Pr(g, r')\Pr(1|s', r', 0)\alpha + \Pr(r, r')\Pr(1|s', r', 0)\alpha \quad (59)$$

$$v(s, s') = \Pr(g, g')((1 - \eta)\alpha - 1) + \Pr(r, g')\eta\alpha + \Pr(g, r')(\Pr(1|s', r', 1)\alpha - 1) + \Pr(r, r')\Pr(1|s', r', 0)\alpha \quad (60)$$

$$v(\text{selfish}, s') = \Pr(g, g')\eta\alpha + \Pr(r, g')\eta\alpha + \Pr(g, r')\Pr(1|s', r', 0)\alpha + \Pr(r, r')\Pr(1|s', r', 0)\alpha \quad (61)$$

Where $\Pr(1|s', f_{s'}, a) \in \{\eta, 1 - \eta\}$ 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

$(1 - \eta)\alpha - 1 > \eta\alpha$, 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:

$$v(s, s_I) = \Pr(g) v(s, s_I|g) + \Pr(r) v(s, s_I|r) \geq \min[v(s, s_I|g), v(s, s_I|r)] . \quad (62)$$

And if the right hand side is low enough 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:

$$\begin{aligned} v(s, p|g) &= -1 + (1 - p)(1 - \eta)\alpha \leq (1 - p)\eta\alpha = v(\text{selfish}, p|g) . \\ p &= 1 - \frac{1}{\alpha(1 - 2\eta)} . \end{aligned} \quad (63)$$

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 $-$ team, $-$ weak team, or $-$ insider. Where $-s$ is the strategy that treats red as good—the language of the strategy has changed. Next notice that if $\Pr(r, g')$ then both strategies will call for the same action and this can not affect the choice of strategy. Thus what we care about is $v(s, p|r, r')$.

$$\begin{aligned} v(s, p|r, r') &= (1 - p)\Pr(1|s, r, 0)\alpha + p\Pr(1|s', r', 0)\alpha \\ v(s', p|r, r') &= -1 + (1 - p)\Pr(1|s, r, 1)\alpha + p\Pr(1|s', r', 1)\alpha \end{aligned} \quad (64)$$

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

$$\begin{aligned} v(s, p|r, r') &\leq v(s', p|r, r') \\ (1 - p)\Pr(1|s, r, 0)\alpha + p\Pr(1|s', r', 0)\alpha &\leq -1 + (1 - p)\Pr(1|s, r, 1)\alpha + p\Pr(1|s', r', 1)\alpha \\ 1 + (1 - p)\alpha(\Pr(1|s, r, 0) - \Pr(1|s, r, 1)) &\leq p\alpha(\Pr(1|s', r', 1) - \Pr(1|s', r', 0)) \end{aligned} \quad (65)$$

and we see the choice of s' does not matter, for all of them $\Pr(1|s', r', 1) - \Pr(1|s', r', 0) = 1 - 2\eta$. 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:

$$p = \frac{1}{\alpha(1 - 2\eta)} . \quad (66)$$

For the team strategy $\Pr(1|s, r, 0) - \Pr(1|s, r, 1) = 1 - 2\eta$ so:

$$p = \frac{1}{2} \left(1 + \frac{1}{\alpha(1 - 2\eta)} \right) \quad (67)$$

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:

$$\begin{aligned} v(\text{selfish}, p|g') &= p\alpha\eta \leq -1 + p(1-\eta)\alpha = v(s', p|g') \\ p &= \frac{1}{\alpha(1-2\eta)} \end{aligned} \quad (68)$$

We have thus derived the radii of all the limit sets:

$$\begin{aligned} \mathcal{R}(\text{selfish})/n &= \frac{1}{(1-2\eta)\alpha}, \quad \mathcal{R}(\text{team})/n = \min \left[1 - \frac{1}{(1-2\eta)\alpha}, \frac{1}{2} \left(1 + \frac{1}{(1-2\eta)\alpha} \right) \right] \\ \mathcal{R}(\text{insider})/n &= \mathcal{R}(\text{weak team})/n = \min \left[1 - \frac{1}{(1-2\eta)\alpha}, \frac{1}{(1-2\eta)\alpha} \right]. \end{aligned} \quad (69)$$

Lemma 8 *In the emergent seeds we always can transition from the selfish to any other limit set. If $3 \geq (1-2\eta)\alpha$ we either transition back to the selfish strategy or to a team strategy and then the selfish strategy from any initial starting point. If $3 \leq (1-2\eta)\alpha$ we transition from the selfish strategy to any other strategy and then transition among those strategies. In this case the selfish strategy has a very low radius and is not stochastically stable.*

If $(1-2\eta)\alpha < 2$ then a selfish strategy is stochastically stable; if $(1-2\eta)\alpha = 2$ all limit sets are stochastically stable; and if $(1-2\eta)\alpha > 2$ then a team strategy is stochastically stable.

Since the stochastically stable limit set is in the core and there is one level of the emergent seed the censored coradius and the coheight will be the same. However notice that since there is always a language issue it will also be the radius of a stochastically stable limit set.

6.4.2 The Contribution Game

In a contribution game there are $n = |I|$ agents and a public good that requires $1 < m \leq n$ agents to contribute. Each agent is willing to contribute if and only if necessary. Let z be the subsets of agents who are contributing. The strict Nash equilibria are $z_0 = \emptyset$ —where no one contributes, and Z_* —where exactly m agents contribute. Let b_i^β ($b_i < 1$) be the probability that agent i will contribute when it either has no benefit or is not necessary. Let d_i^β ($d_i < 1$) be the probability that agent i will stop contributing when this means the good is no longer provided. Thus if $i \in z \in Z_*$ then $r(z \setminus i, z) = \ln \frac{1}{d_i}$ and if $i \notin z \in Z_*$ or $|z| < m-1$ $r(z \cup i, z) = \ln \frac{1}{b_i}$. We assume genericity of the $\{b_i, d_i\}_{i \in I}$ and without loss of generality $b_s > b_{s+1}$ for $s \in \{1, 2, 3, \dots, n-1\}$. Myatt and Wallace (2008b) makes the assumption that $d_s < d_{s+1}$ as well.

To exit z_0 we need to get $m-1$ agents to contribute when there is no benefit, thus $\mathcal{R}(z_0) = \min_{z:|z|=m-1} \sum_{i \in z} \ln \frac{1}{b_i} = \sum_{i=1}^{m-1} \ln \frac{1}{b_i}$. If we let $z_+ = \{1, 2, 3, \dots, m-1\}$ then we can go from z_0 to any \hat{z}_k where $n \geq k \geq m$ and $\hat{z}_k = z_+ \cup k$. Now for any $z \in Z_*$ $\mathcal{R}(z) = \min \left[\min_{i \in z} \ln \frac{1}{d_i}, \min_{j \in I \setminus z} \ln \frac{1}{b_j} \right]$. And we notice that for $z \neq \hat{z}_m$ $\min_{i \in I \setminus z} \ln \frac{1}{b_i} = \min_{i \in \hat{z}_m \setminus z} \ln \frac{1}{b_i}$ thus from z we either go closer to \hat{z}_m or to z_0 . This implies that the only limit sets that might be in the core are \hat{z}_k 's or z_0 . From every other state the path from that state does not have any cycles.

Thus let us be more precise about the radii of \hat{z}_k . If $k \neq m$ then $\min_{j \in I \setminus \hat{z}_k} \ln \frac{1}{b_j} = \ln \frac{1}{b_m}$, and for \hat{z}_m $\min_{j \in I \setminus \hat{z}_m} \ln \frac{1}{b_j} = \ln \frac{1}{b_{m+1}}$. We also notice that $z_+ \subseteq \hat{z}_k$ thus $\min_{i \in \hat{z}_k} \ln \frac{1}{d_i} = \min \left[\ln \frac{1}{d_+}, \ln \frac{1}{d_k} \right]$ where $\ln \frac{1}{d_+} =$

$\min_{i \leq m-1} \frac{1}{d_i}$. Thus for $k \neq m$ $R(\hat{z}_k) = \min \left[\ln \frac{1}{d_+}, \ln \frac{1}{d_k}, \ln \frac{1}{b_m} \right]$ and $R(\hat{z}_m) = \min \left[\ln \frac{1}{d_+}, \ln \frac{1}{d_m}, \ln \frac{1}{b_{m+1}} \right]$. We will transition from \hat{z}_k to \hat{z}_m or z_0 and from \hat{z}_m to \hat{z}_{m+1} or z_0 .

Now if either \hat{z}_m or \hat{z}_{m+1} goes to z_0 then z_0 and every \hat{z}_k is in a core, thus a necessary and sufficient condition for z_0 to be stochastically stable is:

$$\sum_{i=1}^{m-1} \ln \frac{1}{b_i} \geq \max_{n \geq k \geq m} \mathcal{R}(\hat{z}_k) . \quad (70)$$

In this case there may be multiple emergent seeds and at least one will be a short step star.

The only case in which z_0 is not in the core is when we have a cycle between \hat{z}_m and \hat{z}_{m+1} . In this case we know that $\mathcal{R}(\hat{z}_m) = \ln \frac{1}{b_{m+1}} > \ln \frac{1}{b_m} \geq \mathcal{R}(\hat{z}_k)$ for $k \neq m$, thus if z_0 is not stochastically stable \hat{z}_m must be. At this point we have to find the censored (modified) coradius for z_0 , and given lemma 6 we know that if the radius of z_0 is higher than the censored coradius it will be stochastically stable. We also notice that the radius of \hat{z}_m does not matter, because we subtract it from the first step and add it at the end. Thus if we go directly from \hat{z}_m to z_0 the cost of this method will be $\min \left[\ln \frac{1}{d_+}, \ln \frac{1}{d_m} \right]$. Now what if we take intermediate steps? We can show that we only need one. Notice that for any z in this path we want to keep $\min_{j \in I \setminus z} \ln \frac{1}{b_j}$ as large as possible. Thus we want $z_+ \subseteq z$ or $z = \hat{z}_k$ for some $k \geq m+1$. But now if we take multiple steps we are first adding and then removing the same agent, and thus we only want to take one step. Thus we derive essentially the same formula as Myatt and Wallace (2008b):

$$\sum_{i=1}^{m-1} \ln \frac{1}{b_i} \geq \min \left[\min \left[\ln \frac{1}{d_+}, \ln \frac{1}{d_m} \right], \min_{k \geq m+1} \left(\min \left[\ln \frac{1}{d_+}, \ln \frac{1}{d_k} \right] + \ln \frac{1}{b_k} - \min \left[\ln \frac{1}{d_+}, \ln \frac{1}{d_k}, \ln \frac{1}{b_m} \right] \right) \right] . \quad (71)$$

To reach the formula in the article assume that $\ln \frac{1}{d_+} > \ln \frac{1}{d_m} > \max_{k > m} \ln \frac{1}{d_k}$ in which case \hat{z}_m always has the highest radius among $z \in Z_*$ and the two formulas merge.

When z_0 is not in the core its censored coradius is the right hand side of equation 71. However its coheight is the radius of \hat{z}_m , or $\ln \frac{1}{b_m}$. To see that this is strictly lower one recognizes that the cost of going to k and then to the core is at least $\ln \frac{1}{b_k} > \ln \frac{1}{b_m}$, and that $\min \left[\ln \frac{1}{d_+}, \ln \frac{1}{d_m} \right] > \ln \frac{1}{b_m}$.

6.4.3 Other

As we said most of these models have short step stars, and most are imitation models. Fisher and Vega-Redondo (2003) analyzes the evolution of comparative advantage. Here one firm changes production at a time, thus while the Walrasian equilibrium is a GESS one approaches it in short steps. Kim and Wong (2011) look at an exchange economy and here it is strictly less costly to take small steps. However one moves towards the Walrasian equilibrium in a smooth manner, and it has maximal radius.

Alos-Ferrer and Kirchsteiger (2010) looks at the evolution of trading platforms. There a sequence of single mutations converges to the market clearing (Walrasian) trading platform, but each step only takes you closer. However from the Walrasian trading platform one can go to other—non market clearing—platforms with a sequence of single errors, thus all these limit sets are in the core and stochastically stable.

Ellison (2000) looks at the two dimensional lattice with BRM. In this model the radius of each limit set is determined by one of two intuitive directions. One either goes up to a limit set where more agents are using the risk dominant strategy or down to a limit set where fewer are. From every limit set down has a

weakly higher cost than going up (if both are possible), and thus in the first level of the emergent seed one converges to the limit set where everyone is using the risk dominant strategy. This limit set also has the highest radius thus this is a short step star.

The emergent seed in the two dimensional lattice is constructed in Hasker (2014), which also analyzes higher dimensions. This is an open problem in physics and economics. Physicists (often) use the logit decision rule, Ellison (1993, 2000) uses BRM. That paper finds the emergent seed for both models. This is possible because the same dynamics that structure the first level continue to structure higher levels. From each limit set one either goes up or down. This implies that the critical path is a linear subspace of the set of limit sets, and one can characterize the emergent seed using the critical path.

Perhaps the most interesting emergent seed in the literature is in Alos-Ferrer (2004). That article considers imitation in the Cournot game with short memory. Now any quantity between the Cournot quantity (q_c) and the Walrasian quantity (q_w) has a radius of two. It is easy to show that one can go from q_c to q_w , but can one go back? To construct a root switching argument one needs to find a minimal cost path from q_w to q_c , or to show all these states are in the core. Allow $q_w(Q)$ to be the Walrasian quantity given total output is Q . Have one firm error by producing no output and a second firm error by producing $q_{w1} = q_w((n-1)q_w) > q_w$ —now everyone will imitate the second firm. One further error one leads to $q_{w2} = q_w(nq_{w1}) < q_w$. But if $q_{w2} > q_c$ one has to do it again, seesawing in and out of the key range.

7 Limitations in Application.

We have mentioned that if a problem has an analytic limiting distribution this is better than finding the emergent seed. However there are also cases where either root switching is superior or the emergent seed methodology is not enough.

Ben-Shoham, Serrano, and Volij (2004) solves a house allocation problem using root switching and we can not find the emergent seed. That paper considers two models. The one which is of interest is when the probability of a undesirable trade is increasing in the number of steps down in their rankings the agents involved in the trade make. One can see that with this model two limit sets can be in a cycle. Assume that i and j have the same preferences, and in the limit set ω agent i has their favorite house, and j has their second favorite. Then there is a $\tilde{\omega}$ where j has the favorite house and i has the second favorite, this trade would have a resistance of one and these two limit sets may be in a cycle. This is the only current publication in economics where the emergent seed has more than one level. The article establishes an elegant *reversible paths* property. The optimal path from $\omega \in \Omega$ to $\tilde{\omega} \in \Omega$ is also the optimal path from $\tilde{\omega}$ to ω , and the difference in the costs of the two paths is the difference in the levels of envy. Now root switching can characterize stochastically stable states easily and precisely. Assume that ω is stochastically stable, then one can switch the root with a minimal envy limit set $\tilde{\omega}$ and if you were right about ω then it must also have minimal envy. A simple and elegant proof.

Kandori and Rob (1995) show that in strict supermodular games that the set of limit sets is the set of pure strategy Nash equilibria but can not go further, and neither can we. The best invaders are one of the extreme strategies, but what are the best responses to these best invaders? Without further structure one can not know. In general if there are K limit sets then there are $(K-1)^K$ possibilities for the first level of the emergent seed. Not every problem can be analytically solved using this technique, proponents of this

technique will need to find games with simple emergent seeds. This is but a tool in the toolbox, there are still many problems that can not be easily solved.

8 Conclusion

This paper finds a fundamental underlying structure in stochastic evolution—the emergent seed. Conditional on this structure, the stochastic potential of a limit set requires no optimization, and this makes it simple to derive formulas for two measures of waiting time—the coheight and the censored coradius. Of the two in stochastic evolution the former is precise, but the latter is usually easier to use and a correct measure for minimal cost spanning tree problems. By surveying other articles we have shown that many could have been solved using the emergent seed methodology without (at least) increasing the complexity of arguments.

We hope the reader has been made aware that this methodology is not always the best. The other two methods in the literature—deriving the limiting distribution and root switching—can both be superior. We hope it is also clear that many types of robustness are arguments that the emergent seed does not matter.

The reader might wonder if this is always the best algorithm, and we do to. In the literature on minimal cost spanning trees on directed graphs there is only one algorithm. We refer the reader to Beggs (2005) for another possible algorithm and also to Trygubenko and Wales (2006). The latter article finds a way to simulate waiting times without iteration. Like Bortz, Kalos and Lebowitz (1975) this might be useful for analyzing our problem.

We hope that the this paper has brought some clarity to the study of stochastic evolution. While this is a promising field the methodology used and the rationale for results is often confusing and opaque. The current methodology is to either study problems where the limiting distribution is analytic or to use root switching. Root switching will always be a guess and verify methodology, but the literature is a tribute to how successful it can be. The emergent seed is another methodology, a necessary and sufficient methodology that provides a representation theorem, but its value in new applications has yet to be proven.

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] Alos-Ferrer, C and N Netzer (2012) "Robust stochastic stability" *University of Zurich Working Paper Series*

- [8] Andreozzi, L (2012) "Property rights and investments: An evolutionary approach" *Games and Economic Behavior* 74:1-11.
- [9] Anwar, AW (2002) "On the co-existence of conventions" *Journal of Economic Theory*, 107:145-155.
- [10] Arous, GB; and R Cerf (1996) "Metastability of the three dimensional Ising model on a Torus at very low temperatures" *Electronic Journal of Probability* 1:1-55.
- [11] Beggs, AW (2005) "Waiting times and equilibrium selection" *Economic Theory*, 25:599-628
- [12] Ben-Shoham, A; Serrano, R; and O. Volij (2004) "The evolution of exchange" *Journal of Economic Theory* 114:310-328.
- [13] Binmore, KG; Samuelson, L; and P Young. (2003) "Equilibrium Selection in Bargaining Models." *Games and Economic Behavior*, 45:296-328.
- [14] Blume, LE (1993) "The Statistical Mechanics of Strategic Interaction." *Games and Economic Behavior* 5:387-424
- [15] Blume, LE (1997) "Population games" In: Arthur, W., Durlauf, S., Lane, D. (Eds.), *The Economy as an evolving complex system II. SFI Studies in the science of complexity, Vol. XXVII*, Addison-Wesley, pp. 425-462.
- [16] 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.
- [17] 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.
- [18] Chiu, YJ; and TH Liu (1965) "On the Shortest Arborescence of Directed Graphs." *Scientia Sinica*, 14:1390-1400.
- [19] Cui, ZW; and JA Zhai (2010) "Escape dynamics and equilibria selection by iterative cycle decomposition" *Journal of Mathematical Economics* 46:1015-1029.
- [20] Dawid, H; and WB Macleod (2008) "Hold-up and the evolution of investment and bargaining norms" *Games and Economic Behavior* 62:26-52.
- [21] Dutta, J; and K Prasad (2002) "Stable risk-sharing" *Journal of Mathematical Economics* 38:411-439
- [22] Edmonds, J (1967). "Optimal Branchings." *Journal of Research of the National Bureau of Standards*, 71B: 233-240.
- [23] Ellingsen, T; and J Robles (2002) "Does Evolution Solve the Hold-Up Problem?" *Games and Economic Behavior* 39:28-53.
- [24] Ellison, G (1993) "Learning, Local Interaction, and Coordination," *Econometrica*, 61:1047-1071.
- [25] 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.
- [26] Feri, F (2007) "Stochastic stability in networks with decay" *Journal of Economic Theory* 135:442-457.
- [27] Fisher, EO and F Vega-Redondo (2003) "The Evolution of Comparative Advantage" *working paper*
- [28] Freidlin, M; and A Wentzell (1998) *Random Perturbations of Dynamical Systems* Springer Verlag, New York *second edition*.
- [29] Fudenberg, D; and LA Imhof (2006) "Imitation processes with small mutations" *Journal of Economic Theory* 51:251-262.

- [30] 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
- [31] Gerber, A; and MO Bettzüge (2007) "Evolutionary choice of markets" *Economic Theory* 30:453-472.
- [32] Goyal, S; and F Vega-Redondo (2004) "Network Formation and Social Coordination," *Games and Economic Behavior*, 50:178-207.
- [33] Harsanyi, JC (1973) "Games with randomly disturbed payoffs: a new rationale for mixed-strategy equilibrium points" *International Journal of Game Theory* 2:1-23.
- [34] Hasker, K (2014) "The Speed of Evolution on Lattices" *work in progress*
- [35] Hehenkamp, B; and A Wambach (2010) "Survival at the center-The stability of minimum differentiation" *Journal of Economic Behavior and Organization*, 76:853-858
- [36] Humblet, P (1983) "A distributed algorithm for minimum weighted directed spanning trees", *IEEE Transactions on Communications*, 6:756-762.
- [37] Jackson, MO; and A Watts (2002a) "On the formation of interaction networks in social coordination games," *Games and Economic Behavior*, 41:265-291
- [38] Jackson, MO; and A Watts (2002b) "The Evolution of Social and Economic Networks." *Journal of Economic Theory*, 106, 265-295
- [39] Johnson, P; Levine, D; and W Pesendorfer. (2001): "Evolution and Information in a Gift-Giving Game." *Journal of Economic Theory*, 100:1-21.
- [40] Kandori, M; Mailath, G; and R Rob. (1993): "Learning Mutation, and Long Run Equilibria in Games." *Econometrica*, 61:9-56.
- [41] 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.
- [42] Kandori, M; and R Rob. (1998): "Bandwagon Effects and Long Run Technology Choice." *Games and Economic Behavior*, 22:30-60.
- [43] 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.
- [44] Kim, C; and KC Wong (2011) "Evolution of Walrasian equilibrium in an exchange economy" *Journal of Evolutionary Economics* 21:619-647.
- [45] Kreindler, GE; and HP Young (2012) "Rapid Innovation Diffusion with Local Interaction," *Oxford Economics Discussion Papers Series*.
- [46] Kreindler, GE; and HP Young (2013) "Fast Convergence in Evolutionary Equilibrium Selection," *Games and Economic Behavior*, 80:39-67.
- [47] Levine, DK; and S Modica (2014a) "Conflict, Evolution, Hegemony, and the Power of the State," *working paper*.
- [48] Levine, DK; and S Modica (2014b) "Interim and Long-Run Dynamics in the Evolution of Conventions," *working paper*.
- [49] Monderer, D; and LS Shapley (1996) "Potential Games," *Games and Economic Behavior* 14:124-143.
- [50] Montanari, A; and A Saberi (2010) "The spread of innovations in social networks," *Proceedings of the National Academy of the Sciences*, 107:20196-20201.
- [51] Myatt, DP; and C Wallace (2003) "A multinomial probit model of stochastic evolution" *Journal of Economic Theory* 113:286-301.

- [52] Myatt, DP; and C Wallace (2008a) "An evolutionary analysis of the volunteer's dilemma" *Games and Economic Behavior* 62:67-76.
- [53] 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.
- [54] Naidua, S; Hwang, S; and S Bowles (2010) "Evolutionary bargaining with intentional idiosyncratic play." *Economics Letters* 109:31–33.
- [55] Neves, EJ and RH Schonman (1992) "Behavior of droplets for a class of Glauber dynamics at very low temperature" *Probability Theory and Related Fields* 91:331-354.
- [56] Newton, J (2012) "Recontracting and stochastic stability in cooperative games" *Journal of Economic Theory* 147:364-381.
- [57] Pak, M (2008) "Stochastic stability and time-dependent mutations" *Games and Economic Behavior* 64:650-665.
- [58] Peski, M (2010) "Generalized risk-dominance and asymmetric dynamics" *Journal of Economic Theory* 145:216-248.
- [59] Robles, J (1997) "Evolution and long run equilibria in coordination games with summary statistic payoff technologies" *Journal of Economic Theory* 75:180-193.
- [60] Robles, J (2008) "Evolution, bargaining, and time preferences" *Economic Theory* 35:19-36.
- [61] Robson, AJ; and F Vega-Redondo (1996) "Efficient equilibrium selection in evolutionary games with random matching" *Journal of Economic Theory* 70:65-92.
- [62] Rozen, K (2008) "A 'Dual'-Improved Shortcut to the Long Run" *Yale University Economics Department Working Paper* No. 41
- [63] Saez-Marti, M; and JW Weibull (1999) "Clever Agents in Young's Evolutionary Bargaining Model" *Journal of Economic Theory*, 86:268 279
- [64] Sandholm, WH (2007) "Simple formulas for stationary distributions and stochastically stable states" *Games and Economic Behavior*, 59:154-162.
- [65] Sandholm, WH (2010) "Orders of limits for stationary distributions, stochastic dominance, and stochastic stability" *Theoretical Economics*, 5:1–26.
- [66] Schaeffer, M. (1988): "Evolutionarily Stable strategies for a Finite Population and a Variable Contest Size," *Journal of Theoretical Biology*, 132:469–478.
- [67] 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.
- [68] Troger, T (2002) "Why Sunk Costs Matter for Bargaining Outcomes: An Evolutionary Approach" *Journal of Economic Theory* 102:375-402
- [69] Trygubenko, SA; and DJ Wales (2006) "Graph transformation method for calculating waiting times in Markov chains" *Journal of Chemical Physics*, 124:234110.
- [70] Vega-Redondo, F (1997) "The evolution of Walrasian behavior" *Econometrica*, 65:375-384.
- [71] Young, HP (1993a): "The Evolution of Conventions." *Econometrica*, 61:29-56.
- [72] Young, HP (1993b) "An Evolutionary Model of Bargaining." *Journal of Economic Theory*, 59:145-168.
- [73] Young, HP (1998a) "Conventional Contracts." *Review of Economic Studies*, 65:776-792.
- [74] Young, HP (1998b) *Individual Strategy and Social Structure - An Evolutionary Theory of Institutions*. Princeton University Press, Princeton, NJ
- [75] Young, HP (2011) "The dynamics of social innovation," *Proceedings of the National Academy of the Sciences*, 108:21285–21291.

9 Appendix

Proposition 3 For $x \in X$ $\ln \tau(x) \leq \overline{CR}(x)$

Proof. We begin with the Markov transition matrix, M_β , where $M_\beta(x, y)$ 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 $M_{\beta, c}$, and then we can define a new matrix on Ω by letting $M_{\beta, c}(\omega, \omega') = \sum_{x \in B(\omega)} \sum_{y \in B(\omega')} M_{\beta, c}(x, y)$. Now Bortz, Kalos and Lebowitz (1975) shows we can construct a new transition matrix where $\hat{M}_{\beta, c}(\omega, \omega') = M_{\beta, c}(\omega, \omega') / (1 - M_{\beta, c}(\omega, \omega))$, and from Ellison (2000) we know that $(1 - M_{\beta, c}(\omega, \omega))$ is on the order of $e^{-\beta \mathcal{R}(\omega)}$. Thus since we are interested in large β , without loss of generality we can let $\hat{M}_{\beta, c}(\omega, \omega') = e^{-\beta(c(\omega, \omega') - \mathcal{R}(\omega))}$. At this point we can iterate the process with $\hat{M}_{\beta, c}$ as our initial matrix and $\Delta r(\omega, \omega') = c(\omega, \omega') - \mathcal{R}(\omega)$ as our resistance function.

For a given initial $\tilde{\omega}$ this tells us that we can characterize the path from $\tilde{\omega}$ to ω as two sub-sequences. One goes to the core, $\{\tilde{\omega}^t\}_{t=0}^T$ where $\tilde{\omega}^0 = \tilde{\omega}$ and $\omega^T = \omega_c$, and one comes from the core $\{\omega^t\}_{t=T}^0$, where $\omega^0 = \omega$ and $\omega^T = \omega_c$. The former must have the cost $\sum_{t=0}^{T-1} \Delta^t \mathcal{R}(\omega^t(\tilde{\omega}))$ by likelihood maximization. The latter has the cost $\sum_{t=1}^T \Delta^t c(\omega^{t-1}, \omega^t)$ and by strategic delay $\{\omega^t\}_{t=T}^0 = \{\omega^t(\omega)\}_{t=T}^0$. Thus an estimate for the likelihood for the transition from $\tilde{\omega}$ to ω is the censored coradius conditional on beginning at $\tilde{\omega}$, and as $\beta \rightarrow \infty$ the maximum such waiting time is the critical one. ■