

Networks Preserving Evolutionary Equilibria and the Power of Randomization

Michael Kearns and Siddharth Suri
Computer and Information Science
University of Pennsylvania

Abstract

We study a natural extension of classical evolutionary game theory to a setting in which pairwise interactions are restricted to the edges of an undirected graph or network. We generalize the definition of an evolutionary stable strategy (ESS), and show a pair of complementary results that exhibit the power of randomization in our setting: subject to minimal edge density conditions, the classical ESS of *any* game are preserved when the graph is chosen randomly and the mutation set is chosen adversarially, or when the graph is chosen adversarially and the mutation set is chosen randomly. We examine natural strengthenings of our generalized ESS definition, and show that similarly strong results are not possible for them.

1 Introduction

In this paper, we introduce and examine a natural extension of classical evolutionary game theory (EGT) to a setting in which pairwise interactions are restricted to the edges of an undirected graph or network. This generalizes the classical setting, in which all pairs of organisms in an infinite population are equally likely to interact. The classical setting can be viewed as the special case in which the underlying network is a clique.

There are many obvious reasons why one would like to examine more general graphs, the primary one being in that many scenarios considered in evolutionary game theory, all interactions are in fact not possible. For example, geographical restrictions may limit interactions to physically proximate pairs of organisms. More generally, as evolutionary game theory has become a plausible model not only for biological interaction, but also economic interaction in which dynamics are more imitative than optimizing, the network constraints may come from similarly more general sources. Evolutionary game theory on networks

has been considered before, but not in the generality we will do so here (see Section 4).

We generalize the definition of an evolutionary stable strategy (ESS) to networks, and show a pair of complementary results that exhibit the power of randomization in our setting: subject to minimal edge density conditions, the classical ESS of *any* game are preserved when the graph is chosen randomly and the mutation set is chosen adversarially, or when the graph is chosen adversarially and the mutation set is chosen randomly. We examine natural strengthenings of our generalized ESS definition, and show that similarly strong results are not possible for them.

The work described here is part of recent efforts examining the relationship between graph topology or structure and properties of equilibrium outcomes. Previous works in this line include studies of the relationship of topology to properties of correlated equilibria in graphical games [KKLO03], and studies of price variation in graph-theoretic market exchange models [KKO⁺04].

2 Classical EGT

The fundamental concept of evolutionary game theory is the evolutionarily stable strategy (ESS). Intuitively, an ESS is a strategy such that if all the members of a population adopt it, then no mutant strategy could invade the population [Smi82]. To make this more precise, we describe the basic model of evolutionary game theory, in which the notion of an ESS resides.

The standard model of evolutionary game theory considers an infinite population of organisms, each of which plays a strategy in a fixed, 2-player, symmetric game. The game is defined by a fitness function F . All pairs of members of the infinite population are equally likely to interact with one another. If two organisms interact, one playing strategy s and the other playing strategy t , the s -player earns a fitness

of $F(s|t)$ while the t -player earns a fitness of $F(t|s)$.

In this infinite population of organisms, suppose there is a $1 - \epsilon$ fraction who play strategy s , and call these organisms *incumbents*; and suppose there is an ϵ fraction who play t , and call these organisms *mutants*. Assume two organisms are chosen uniformly at random to play each other. The strategy s is an ESS if the expected fitness of an organism playing s is higher than that of an organism playing t , for all $t \neq s$ and all sufficiently small ϵ . Since an incumbent will meet another incumbent with probability $1 - \epsilon$ and it will meet a mutant with probability ϵ , we can calculate the expected fitness of an incumbent, which is simply $(1 - \epsilon)F(s|s) + \epsilon F(s|t)$. Similarly, the expected fitness of a mutant is $(1 - \epsilon)F(t|s) + \epsilon F(t|t)$. Thus we come to the formal definition of an ESS [Wei95].

Definition 2.1 *A strategy s is an evolutionarily stable strategy (ESS) for the 2-player, symmetric game given by fitness function F , if for every strategy $t \neq s$, there exists an ϵ_t such that for all $0 < \epsilon < \epsilon_t$, $(1 - \epsilon)F(s|s) + \epsilon F(s|t) > (1 - \epsilon)F(t|s) + \epsilon F(t|t)$.*

A consequence of this definition is that for s to be an ESS, it must be the case that $F(s|s) \geq F(t|s)$, for all strategies t . This inequality means that s must be a best response to itself, and thus any ESS strategy s must also be a Nash equilibrium. In general the notion of ESS is more restrictive than Nash equilibrium, and not all 2-player, symmetric games have an ESS.

In this paper our interest is to examine what kinds of network structure *preserve* the ESS strategies for those games that do have a standard ESS. First we must of course generalize the definition of ESS to a network setting.

3 EGT on Graphs

In our setting, we will no longer assume that two organisms are chosen uniformly at random to interact. Instead, we assume that organisms interact only with those in their local neighborhood, as defined by an undirected graph or network. As in the classical setting (which can be viewed as the special case of the complete network or clique), we shall assume an infinite population, by which we mean we examine limiting behavior in a family of graphs of increasing size.

Before giving formal definitions, some comments are in order on what to expect in moving from the classical to the graph-theoretic setting. In the classical (complete graph) setting, there exist many sym-

metries that may be broken in moving to the the network setting, at both the group and individual level. Indeed, such asymmetries are the primary interest in examining a graph-theoretic generalization. For example, at the group level, in the standard ESS definition, one need not discuss any particular set of mutants of population fraction ϵ . Since all organisms are equally likely to interact, the survival or fate of any specific mutant set is identical to that of any other. In the network setting, this may not be true: some mutant sets may be better able to survive than others due to the specific topologies of their interactions in the network. For instance, foreshadowing some of our analysis, if s is an ESS but $F(t|t)$ is much larger than $F(s|s)$ and $F(s|t)$, a mutant set with a great deal of “internal” interaction (that is, edges between mutants) may be able to survive, whereas one without this may suffer. At the level of individuals, in the classical setting, the assertion that one mutant dies implies that all mutants die, again by symmetry. In the network setting, individual fates may differ within a group all playing a common strategy. These observations imply that in examining ESS on networks we face definitional choices, which we will explain, that were obscured in the classical model.

If G is a graph representing the allowed pairwise interactions between organisms (vertices), and u is a vertex of G playing strategy s_u , then the fitness of u is given by

$$F(u) = \frac{\sum_{v \in \Gamma(u)} F(s_u|s_v)}{|\Gamma(u)|}.$$

Here s_v is the strategy being played by the neighbor v , and $\Gamma(u) = \{v \in V : (u, v) \in E\}$. One can view the fitness of u as the average fitness u would obtain if it played each of its neighbors, or the expected fitness u would obtain if it were assigned to play one of its neighbors chosen uniformly at random.

We now give a definition for a family of mutant vertex sets in an infinite graph family to *contract*.

Definition 3.1 *Let $G = \{G_n\}_{n=0}^\infty$ be an infinite family of graphs, where G_n has n vertices. Let $M = \{M_n\}_{n=0}^\infty$ be any family of subsets of vertices of the G_n such that $|M_n| \geq \epsilon n$ for some constant $\epsilon > 0$. Suppose all the vertices of M_n play a common (mutant) strategy t , and suppose the remaining vertices in G_n play a common (incumbent) strategy s . We say that M_n contracts if for all but $o(n)$ of the $j \in M_n$, j has an incumbent neighbor i such that $F(j) < F(i)$.*

A reasonable alternative would be to ask that the condition above hold for *all* mutants rather than all

but $o(n)$. In addition, we require that a mutant have only *one* incumbent neighbor of higher fitness in order to die. In Sections 6.1 and 6.2 we consider strengthening these conditions and demonstrate that our results can no longer hold.

In order to properly define an ESS for an infinite family of graphs in a way that recovers the classical definition asymptotically in the case of the family of complete graphs, we first must give a definition that restricts attention to families of mutant vertices that are smaller than some invasion threshold $\epsilon'n$, yet remain some constant fraction of the population. This prevents “invasions” that survive merely by constituting a vanishing fraction of the population.

Definition 3.2 *Let $\epsilon' > 0$, and let $G = \{G_n\}_{n=0}^\infty$ be an infinite family of graphs, where G_n has n vertices. Let $M = \{M_n\}_{n=0}^\infty$ be any family of (mutant) vertices in G_n . We say that M is ϵ' -linear if there exists an ϵ , $\epsilon' > \epsilon > 0$, such that for all sufficiently large n , $\epsilon'n > |M_n| > \epsilon n$.*

We can now give our definition for a strategy to be evolutionarily stable when employed by organisms interacting with their neighborhood in a graph.

Definition 3.3 *Let $G = \{G_n\}_{n=0}^\infty$ be an infinite family of graphs, where G_n has n vertices. Let F be any 2-player, symmetric game for which s is a strategy. We say that s is an ESS with respect to F and G if for all mutant strategies $t \neq s$, there exists an $\epsilon_t > 0$ such that for any ϵ_t -linear family of mutant vertices $M = \{M_n\}_{n=0}^\infty$ all playing t , for n sufficiently large, M_n contracts.*

Thus, to *violate* the ESS property for G , one must witness a family of mutations M in which each M_n is a (small but nonzero) constant fraction of the population of G_n , but does not contract (i.e. every mutant set has a subset of linear size that survives all of its incumbent interactions). It can be shown that the definition given coincides with the classical one in the case where G is the family of complete graphs, and in the limit of large n . We note that even in the classical model, small sets of mutants were allowed to have greater fitness than the incumbents, so long as the size of the set was $o(n)$ [Vic87].

In the definition above there are three parameters: the game F , the graph family G and the mutation family M . Our main results will hold for *any* 2-player, symmetric game F . Since classical ESS are only defined for 2-player symmetric games (see Definition 2.1), and we show results that preserve those and only those strategies in a graphical setting, this

is the most general class of games we could possibly include in our model. We will also study two rather general settings for G and M : that in which G is a family of random graphs and M is arbitrary, and that in which G is nearly arbitrary and M is randomly chosen. In both cases, we will see that, subject to rather mild conditions on edge density (essentially forcing connectivity of G but not much more), for *any* 2-player, symmetric game, the ESS of the classical settings, and only those strategies, are always preserved. Thus a common theme of these results is the power of randomization: as long as either the network itself is chosen randomly, or the mutation set is chosen randomly, classical ESS are preserved.

4 Related Work

There has been previous work that analyzes which strategies are resilient to mutant invasions with respect to various types of graphs. What sets our work apart is that the model we consider encompasses a significantly more general class of games and graph topologies. We will briefly survey this literature and point out the differences in the previous models and ours.

In [ESS98], [Blu93] and [Blu95], the authors consider specific families of graphs, such as cycles and lattices, where players play specific games, such as 2×2 -games or $k \times k$ -coordination games. In these papers the authors specify a simple, local dynamic for players to improve their payoffs by changing strategies, and analyze what type of strategies will grow to dominate the population. The model we propose is more general than both of these, as it encompasses a larger class of graphs as well as a richer set of games.

Also related to our work is that of [LHN05], where the authors propose two models. The first assumes organisms interact according to a weighted, undirected graph. However, the fitness of each organism is simply assigned and does not depend on the actions of each organism’s neighborhood. The second model has organisms arranged around a directed cycle, where neighbors play a 2×2 -game. With probability proportional to its fitness, an organism is chosen to reproduce by placing a replica of itself in it’s neighbors position, thereby “killing” the neighbor. We consider more general games than the first model and more general graphs than the second.

Finally, the works most closely related to ours are [Ell93], [Mor00], and [Chw00]. The authors consider 2-action, coordination games played by players in a general undirected graph. In these three works, the authors specify a dynamic for a strategy to repro-

duce, and analyze properties of the graph that allow a strategy to overrun the population. Here again, one can see that our model is more general than these, as it allows for organisms to play any 2-player, symmetric game.

5 Networks Preserving ESS

We now proceed to state and prove two complementary results in the network ESS model defined in Section 3. First, we consider a setting where the graphs are generated via the $G_{n,p}$ model of Erdős and Rényi [Bol01]. In this model, every pair of vertices are joined by an edge independently and with probability p (where p may depend on n). The mutant set, however, will be constructed adversarially (subject to the linear size constraint given by Definition 3.3). For these settings, we show that for any 2-player, symmetric game, s is a classical ESS of that game, if and only if s is an ESS for $\{G_{n,p}\}_{n=0}^{\infty}$, where $p = \Omega(1/n^c)$ and $0 \leq c < 1$, and any mutant family $\{M_n\}_{n=0}^{\infty}$, where each M_n has linear size. We note that under these settings, if we let $c = 1 - \gamma$ for small $\gamma > 0$, the expected number of edges in G_n is $n^{1+\gamma}$ or larger — that is, just superlinear in the number of vertices and potentially far smaller than $O(n^2)$. It is easy to convince oneself that once the graphs have only a linear number of edges, we are flirting with disconnectedness, and there may simply be large mutant sets that can survive in isolation due to the lack of any incumbent interactions in certain games. Thus in some sense we examine the minimum plausible edge density.

The second result is a kind of dual to the first, considering a setting where the graphs are chosen arbitrarily but the mutant sets are chosen randomly. It states that for any 2-player, symmetric game, s is an ESS for that game, if and only if s is an ESS for any $\{G_n = (V_n, E_n)\}_{n=0}^{\infty}$ in which for all $v \in V_n$, $\deg(v) = \Omega(n^\gamma)$ (for any constant $\gamma > 0$), and a family of mutant sets $\{M_n\}_{n=0}^{\infty}$, that is chosen randomly (that is, in which each organism is labeled a mutant with constant probability $\epsilon > 0$). Thus, in this setting we again find that classical ESS are preserved subject to minimal edge density restrictions. Next, since the previous assumption placed on the degree of each vertex is quite strong, we prove another result which only assumes that $|E_n| \geq n^{1+\gamma}$, and shows that there must exist at least 1 mutant with an incumbent neighbor of higher fitness.

5.1 Random Graphs and Adversarial Mutations

Now we state and prove a theorem which shows that if s is a classical ESS, then s will be an ESS for random graphs, where a linear sized set of mutants is chosen by an adversary.

Theorem 5.1 *Let F be any 2-player, symmetric game, and suppose s is a classical ESS of F . Let the infinite graph family $\{G_n\}_{n=0}^{\infty}$ be drawn according to $G_{n,p}$, where $p = \Omega(1/n^c)$ and $0 \leq c < 1$. Then with probability 1, s is an ESS.*

The main idea of this proof is to divide mutants into 2 categories, those with normal fitness and those with abnormal fitness. First, we show all but $o(n)$ of the *organisms* have an incumbent neighbor of *normal* fitness. This will imply that all but $o(n)$ of the mutants of normal fitness have an incumbent neighbor of *higher* fitness. The vehicle for doing so is Theorem 2.15 of [Bol01], which gives an upper bound on the number of vertices not connected to a sufficiently large set. This theorem assumes that the size of this large set is known with equality, which necessitates the union bound argument below. Secondly, we show that there can be at most $o(n)$ mutants with abnormal fitness. Since there are so few of them, even if none of them have an incumbent neighbor of higher fitness, s can still be an ESS with respect to F and G .

Proof: (Sketch) Let $t \neq s$ be the mutant strategy. Since s is a classical ESS, there exists an ϵ_t such that $(1 - \epsilon)F(s|s) + \epsilon F(s|t) > (1 - \epsilon)F(t|s) + \epsilon F(t|t)$, for all $0 < \epsilon < \epsilon_t$. Let M be any mutant family that is ϵ_t -linear. Thus for any fixed value of n that is sufficiently large, there exists an ϵ such that $|M_n| > \epsilon n$ and $\epsilon_t > \epsilon > 0$. Also, let $I_n = V_n \setminus M_n$ and let $I' \subseteq I_n$ be the set of incumbents that have fitness in the range $(1 \pm \tau)[(1 - \epsilon)F(s|s) + \epsilon F(s|t)]$ for some constant τ , $0 < \tau < 1/6$. Lemma 5.1 below shows $(1 - \epsilon)n \geq |I'| \geq (1 - \epsilon)n - \frac{24 \log n}{\tau^2 p}$. Finally, let

$$T_{I'} = \{x \in V \setminus I' : \Gamma(x) \cap I' \neq \emptyset\}.$$

(For the sake of clarity we suppress the subscript n on the sets I' and T .) The union bound gives us

$$\Pr(|T_{I'}| \geq \delta n) \leq \sum_{i=(1-\epsilon)n - \frac{24 \log n}{\tau^2 p}}^{(1-\epsilon)n} \Pr(|T_{I'}| \geq \delta n \text{ and } |I'| = i) \quad (1)$$

Letting $\delta = n^{-\gamma}$ for some $\gamma > 0$ gives $\delta n = o(n)$. We will apply Theorem 2.15 of [Bol01] to the summand on the right hand side of Equation 1. If we let

$\gamma = (1 - c)/2$, and combine this with the fact that $0 \leq c < 1$, all of the requirements of this theorem will be satisfied (details omitted). Now, when we apply this theorem to Equation 1, we get

$$\begin{aligned} \Pr(|T_{I'}| \geq \delta n) &\leq \sum_{i=(1-\epsilon)n - \frac{24 \log n}{\delta^2 p}}^{(1-\epsilon)n} \exp\left(-\frac{1}{6}C\delta n\right) \\ &= o(1) \end{aligned} \quad (2)$$

This is because equation 2 has only $\frac{24 \log n}{\delta^2 p}$ terms, and Theorem 2.15 of [Bol01] gives us that $C \geq (1 - \epsilon)n^{1-c} - \frac{24 \log n}{\delta^2 p}$. Thus, we have shown, with probability tending to 1 as $n \rightarrow \infty$, at most $o(n)$ *organisms* are not attached to an incumbent which has fitness in the range $(1 \pm \tau)[(1 - \epsilon)F(s|s) + \epsilon F(s|t)]$. This implies that the number of mutants of approximately normal fitness, not attached to an incumbent of approximately normal fitness, is also $o(n)$.

Now, those mutants of approximately normal fitness that *are* attached to an incumbent of approximately normal fitness have fitness in the range $(1 \pm \tau)[(1 - \epsilon)F(t|s) + \epsilon F(t|t)]$. The incumbents that they are attached to have fitness in the range $(1 \pm \tau)[(1 - \epsilon)F(s|s) + \epsilon F(s|t)]$. Since s is an ESS of F , we know $(1 - \epsilon)F(s|s) + \epsilon F(s|t) > (1 - \epsilon)F(t|s) + \epsilon F(t|t)$, thus if we choose τ small enough, we can ensure that all but $o(n)$ mutants of normal fitness have a neighboring incumbent of higher fitness.

Finally by Lemma 5.1, we know there are at most $o(n)$ mutants of abnormal fitness. So, even if all of them are more fit than their respective incumbent neighbors, we have shown all but $o(n)$ of the mutants have an incumbent neighbor of higher fitness. \square

We now state and prove the lemma used in the proof above.

Lemma 5.1 *For almost every graph $G_{n,p}$ with $(1 - \epsilon)n$ incumbents, all but $\frac{24 \log n}{\delta^2 p}$ incumbents have fitness in the range $(1 \pm \delta)[(1 - \epsilon)F(s|s) + \epsilon F(s|t)]$, where $p = \Omega(1/n^c)$ and ϵ, δ and c are constants satisfying $0 < \epsilon < 1, 0 < \delta < 1/6, 0 \leq c < 1$. Similarly, under the same assumptions, all but $\frac{24 \log n}{\delta^2 p}$ mutants have fitness in the range $(1 \pm \delta)[(1 - \epsilon)F(t|s) + \epsilon F(t|t)]$.*

Proof: We define the *mutant degree* of a vertex to be the number of mutant neighbors of that vertex, and *incumbent degree* analogously. Observe that the only way for an incumbent to have fitness far from its expected value of $(1 - \epsilon)F(s|s) + \epsilon F(s|t)$ is if it has a fraction of mutant neighbors either much higher or much lower than ϵ . Theorem 2.14 of [Bol01] gives us a bound on the number of such incumbents. It states

that the number of incumbents with mutant degree outside the range $(1 \pm \delta)p|M|$ is at most $\frac{12 \log n}{\delta^2 p}$. By the same theorem, the number of incumbents with incumbent degree outside the range $(1 \pm \delta)p|I|$ is at most $\frac{12 \log n}{\delta^2 p}$. From the linearity of fitness as a function of the fraction of mutant or incumbent neighbors, one can show that for those incumbents with mutant and incumbent degree in the expected range, their fitness is within a constant factor of $(1 - \epsilon)F(s|s) + \epsilon F(s|t)$, where that constant goes to 1 as n tends to infinity and δ tends to 0. The proof for the mutant case is analogous. \square

We note that if in the statement of Theorem 5.1 we $c = 0$, then $p = 1$. This, in turn, makes $G = \{K_n\}_{n=0}^\infty$, where K_n is a clique of n vertices. Then for any K_n all of the incumbents will have identical fitness and all of the mutants will have identical fitness. Furthermore, since s was an ESS for G , the incumbent fitness will be higher than the mutant fitness. Finally, one can show that as $n \rightarrow \infty$, the incumbent fitness converges to $(1 - \epsilon)F(s|s) + \epsilon F(s|t)$, and the mutant fitness converges to $(1 - \epsilon)F(t|s) + \epsilon F(t|t)$. In other words, s must be a classical ESS, providing a converse to Theorem 5.1.

5.2 Adversarial Graphs and Random Mutations

We now move on to our second main result. Here we show that if the graph family, rather than being chosen randomly, is arbitrary subject to a minimum degree requirement, and the mutation sets are randomly chosen, classical ESS are again preserved. A modified notion of ESS allows us to considerably weaken the degree requirement to a minimum edge density requirement.

Theorem 5.2 *Let $G = \{G_n = (V_n, E_n)\}_{n=0}^\infty$ be an infinite family of graphs in which for all $v \in V_n$, $\deg(v) = \Omega(n^\gamma)$ (for any constant $\gamma > 0$). Let F be any 2-player, symmetric game, and suppose s is a classical ESS of F . Let t be any mutant strategy, and let the mutant family $M = \{M_n\}_{n=0}^\infty$ be chosen randomly by labeling each vertex a mutant with constant probability ϵ , where $\epsilon_t > \epsilon > 0$. Then with probability 1, s is an ESS with respect to F, G and M .*

Proof: Let $t \neq s$ be the mutant strategy and let X be the event that every incumbent has fitness within the range $(1 \pm \tau)[(1 - \epsilon)F(s|s) + \epsilon F(s|t)]$, for some constant $\tau > 0$ to be specified later. Similarly, let Y be the event that every mutant has fitness within the range $(1 \pm \tau)[(1 - \epsilon)F(t|s) + \epsilon F(t|t)]$. Since $\Pr(X \cap Y) =$

$1 - \Pr(\neg X \cup \neg Y)$, we proceed by showing $\Pr(\neg X \cup \neg Y) = o(1)$.

$\neg X$ is the event that there exists an incumbent with fitness outside the range $(1 \pm \tau)[(1 - \epsilon)F(s|s) + \epsilon F(s|t)]$. If $\deg_M(v)$ denotes the number of mutant neighbors of v , similarly, $\deg_I(v)$ denotes the number of incumbent neighbors of v , then an incumbent i has fitness $\frac{\deg_I(i)}{\deg(i)}F(s|s) + \frac{\deg_M(i)}{\deg(i)}F(s|t)$. Since $F(s|s)$ and $F(s|t)$ are fixed quantities, the only variation in an incumbents fitness can come from variation in the terms $\frac{\deg_I(i)}{\deg(i)}$ and $\frac{\deg_M(i)}{\deg(i)}$. One can use the Chernoff bound followed by the union bound to show that for any incumbent i ,

$$\Pr(F(i) \notin (1 \pm \tau)[(1 - \epsilon)F(s|s) + \epsilon F(s|t)]) \\ \leq 4 \exp\left(-\frac{\epsilon \deg(i) \tau^2}{3}\right).$$

Next one can use the union bound again to bound the probability of the event $\neg X$,

$$\Pr(\neg X) \leq 4n \exp\left(-\frac{d_i \tau^2}{3}\right)$$

where $d_i = \min_{i \in V \setminus M} \deg(i)$, $0 < \epsilon \leq 1/2$. An analogous argument can be made to show $\Pr(\neg Y) < 4n \exp(-\frac{\epsilon d_j \tau^2}{3})$, where $d_j = \min_{j \in M} \deg(j)$ and $0 < \epsilon \leq 1/2$. Thus, by the union bound,

$$\Pr(\neg X \cup \neg Y) < 8n \exp\left(-\frac{\epsilon d \tau^2}{3}\right)$$

where $d = \min_{v \in V} \deg(v)$, $0 < \epsilon \leq 1/2$. Since $\deg(v) = \Omega(n^\gamma)$, for all $v \in V$, and ϵ , τ and γ are all constants greater than 0,

$$\lim_{n \rightarrow \infty} \frac{8n}{\exp(\epsilon d \tau^2 / 3)} = 0,$$

so $\Pr(\neg X \cup \neg Y) = o(1)$. Thus, we can choose τ small enough such that $(1 + \tau)[(1 - \epsilon)F(t|s) + \epsilon F(t|t)] < (1 - \tau)[(1 - \epsilon)F(s|s) + \epsilon F(s|t)]$, and then choose n large enough such that with probability $1 - o(1)$, every incumbent will have fitness in the range $(1 \pm \tau)[(1 - \epsilon)F(s|s) + \epsilon F(s|t)]$, and every mutant will have fitness in the range $(1 \pm \tau)[(1 - \epsilon)F(t|s) + \epsilon F(t|t)]$. So, with high probability, every incumbent will have a higher fitness than every mutant. \square

By arguments similar to those following the proof of Theorem 5.1, if we let $G = \{K_n\}_{n=0}^\infty$, each incumbent will have the same fitness and each mutant will have the same fitness. Furthermore, since s is an ESS for G , the incumbent fitness must be higher than the

mutant fitness. Here again, one has to show that as $n \rightarrow \infty$, the incumbent fitness converges to $(1 - \epsilon)F(s|s) + \epsilon F(s|t)$, and the mutant fitness converges to $(1 - \epsilon)F(t|s) + \epsilon F(t|t)$. Observe that the exact fraction mutants of V_n is now a random variable. So, to prove this convergence we use an argument similar to one that is used to prove that sequence of random variables that converges in probability also converges in distribution (details omitted). This in turn establishes that s must be a classical ESS, and we thus obtain a converse to Theorem 5.2.

The assumption on the degree of each vertex of Theorem 5.2 is rather strong. The following theorem relaxes this requirement and only necessitates that every graph have $n^{1+\gamma}$ edges, for some constant $\gamma > 0$, in which case it shows there will always be at least 1 mutant with an incumbent neighbor of higher fitness. A strategy that is an ESS in this weakened sense will essentially rule out stable, static sets of mutant invasions, but not more complex invasions. An example of more complex invasions are mutant sets that survive, but only by perpetually ‘‘migrating’’ through the graph under some natural evolutionary dynamics, akin to ‘‘gliders’’ in the well-known Game of Life.

Theorem 5.3 *Let F be any game, and let s be a classical ESS of F , and let $t \neq s$ be a mutant strategy. For any graph family $G = \{G_n = (V_n, E_n)\}_{n=0}^\infty$ in which $|E_n| \geq n^{1+\gamma}$ (for any constant $\gamma > 0$), and any mutant family $M = \{M_n\}_{n=0}^\infty$ which is determined by labeling each vertex a mutant with probability ϵ_t , where $\epsilon_t > \epsilon > 0$, the probability that there exists a mutant with an incumbent neighbor of higher fitness approaches 1 as $n \rightarrow \infty$.*

Proof: (Sketch) The main idea behind the proof is to show that with high probability, over only the choice of mutants, there will be an incumbent-mutant edge in which both vertices have high degree. If their degree is high enough, we can show that close to an ϵ fraction of their neighbors are mutants, and thus their fitnesses are very close to what we expect them to be in the classical case. Since s is an ESS, the fitness of the incumbent will be higher than the mutant.

We call an edge $(i, j) \in E_n$ a $g(n)$ -barbell if $\deg(i) \geq g(n)$ and $\deg(j) \geq g(n)$. Suppose G_n has at most $h(n)$ edges that are $g(n)$ -barbells. This means there are at least $|E_n| - h(n)$ edges in which at least one vertex has degree at most $g(n)$. We call these vertices *light* vertices. Let $\ell(n)$ be the number of light vertices in G_n . Observe that $|E_n| - h(n) \leq \ell(n)g(n)$. This is because each light vertex is incident on at

most $g(n)$ edges. This gives us that

$$|E_n| \leq h(n) + \ell(n)g(n) \leq h(n) + ng(n).$$

So if we choose $h(n)$ and $g(n)$ such that $h(n) + ng(n) = o(n^{1+\gamma})$, then $|E_n| = o(n^{1+\gamma})$. This contradicts the assumption that $|E_n| = \Omega(n^{1+\gamma})$. Thus, subject to the above constraint on $h(n)$ and $g(n)$, G_n must contain at least $h(n)$ edges that are $g(n)$ -barbells.

Now let H_n denote the subgraph induced by the barbell edges of G_n . Note that regardless of the structure of G_n , there is no reason that H_n should be connected. Thus, let m be the number of connected components of H_n , and let c_1, c_2, \dots, c_m be the number of vertices in each of these connected components. Note that since H_n is an edge-induced subgraph we have $c_k \geq 2$ for all components k . Let us choose the mutant set by first flipping the vertices in H_n only. We now show that the probability, with respect to the random mutant set, that *none* of the components of H_n have an incumbent-mutant edge is exponentially small in n . Let A_n be the event that every component of H_n contains only mutants or only incumbents. Then algebraic manipulations can establish that

$$\begin{aligned} \Pr[A_n] &= \prod_{k=1}^m (\epsilon^{c_k} + (1-\epsilon)^{c_k}) \\ &\leq (1-\epsilon)^{(1-\frac{\epsilon\beta^2}{2}) \sum_{k=1}^m c_k} \end{aligned}$$

where β is a constant. Thus for ϵ sufficiently small the bound decreases exponentially with $\sum_{k=1}^m c_k$. Furthermore, since $\sum_{k=1}^m \binom{c_k}{2} \geq h(n)$ (with equality achieved by making each component a clique), one can show that $\sum_{k=1}^m c_k \geq \sqrt{h(n)}$. Thus, as long as $h(n) \rightarrow \infty$ with n , the probability that all components are uniformly labeled will go to 0.

Now, assuming that there exists a non-uniformly labeled component, by construction that component contains an edge (i, j) where i is an incumbent and j is a mutant, that is a $g(n)$ -barbell. We also assume that the $h(n)$ vertices already labeled have been done so arbitrarily, but that the remaining $g(n) - h(n)$ vertices neighboring i and j are labeled mutants independently with probability ϵ . Then via a standard Chernoff bound argument, one can show that with high probability, the fraction of mutants neighboring i and the fraction of mutants neighboring j is in the range $(1 \pm \tau) \frac{(g(n) - h(n))\epsilon}{g(n)}$. Similarly, one can show that the fraction of incumbents neighboring i and the fraction of mutants neighboring j is in the range $1 - (1 \pm \tau) \frac{(g(n) - h(n))\epsilon}{g(n)}$.

Since s is an ESS, there exists a $\zeta > 0$ such that $(1-\epsilon)F(s|s) + \epsilon F(s|t) = (1-\epsilon)F(t|s) + \epsilon F(t|t) + \zeta$. If we choose $g(n) = n^\gamma$, and $h(n) = o(g(n))$, we can

choose n large enough and τ small enough to force $F(i) > F(j)$, as desired. \square

6 Stronger Models and Their Limitations

In this section we show that if one tried to strengthen the model described in Section 3 in two natural ways, one would not be able to prove results as strong as Theorems 5.1 and 5.2, which hold for every 2-player, symmetric game.

6.1 A Stronger Notion of Contraction for the Mutant Set

In Section 3 we alluded to the fact that we made certain design decisions in arriving at Definitions 3.1, 3.2 and 3.3. One such decision was to require that all but $o(n)$ mutants have incumbent neighbors of higher fitness. Instead, we could have required that *all* mutants have an incumbent neighbor of higher fitness. The two theorems in this subsection show that if one were to strengthen our notion of contraction for the mutant set, given by Definition 3.1, in this way, it would be impossible to prove theorems analogous to Theorems 5.1 and 5.3.

Recall that Definition 3.1 gave the notion of contraction for a linear sized subset of mutants. In what follows, we will say an *edge* (i, j) contracts if i is an incumbent, j is a mutant, and $F(i) > F(j)$. Also, recall that Theorem 5.1 stated that if s is a classical ESS, then it is an ESS for random graphs with adversarial mutations. Next, we prove that if we instead required *every* incumbent-mutant edge to contract, this need not be the case.

Theorem 6.1 *Let F be a 2-player, symmetric game that has a classical ESS s for which there exists a mutant strategy $t \neq s$ with $F(t|t) > F(s|s)$ and $F(t|t) > F(s|t)$. Let $G = \{G_n\}_{n=0}^\infty$ be an infinite family of random graphs drawn according to $G_{n,p}$, where $p = \Omega(1/n^c)$ for any constant $0 \leq c < 1$. Then with probability approaching 1 as $n \rightarrow \infty$, there exists a mutant family $M = \{M_n\}_{n=0}^\infty$, where $\epsilon_t n > |M_n| > \epsilon n$ and $\epsilon_t, \epsilon > 0$, in which there is an edge that does not contract.*

Proof: (Sketch) With probability approaching 1 as $n \rightarrow \infty$, there exists a vertex j where $\deg(j)$ is arbitrarily close to ϵn . So, label j mutant, label one of its neighbors incumbent, denoted i , and label the rest of j 's neighborhood mutant. Also, label all of i 's neighbors incumbent, with

the exception of j and j 's neighbors (which were already labeled mutant). In this setting, one can show that $F(j)$ will be arbitrarily close to $F(t|t)$ and $F(i)$ will be a convex combination of $F(s|s)$ and $F(s|t)$, which are both strictly less than $F(t|t)$. \square

Theorem 5.3 stated that if s is a classical ESS, then for graphs where $|E_n| \geq n^{1+\gamma}$, for some $\gamma > 0$, and where each organism is labeled a mutant with probability ϵ , one edge must contract. Below we show that, for certain graphs and certain games, there will always exist one edge that will not contract.

Theorem 6.2 *Let F be a 2-player, symmetric game that has a classical ESS s , such that there exists a mutant strategy $t \neq s$ where $F(t|s) > F(s|t)$. There exists an infinite family of graphs $\{G_n = (V_n, E_n)\}_{n=0}^\infty$, where $|E_n| = \Theta(n^2)$, such that for a mutant family $M = \{M_n\}_{n=0}^\infty$, which is determined by labeling each vertex a mutant with probability $\epsilon > 0$, the probability there exists an edge in E_n that does not contract approaches 1 as $n \rightarrow \infty$.*

Proof: (Sketch) Construct G_n as follows. Pick $n/4$ vertices $u_1, u_2, \dots, u_{n/4}$ and add edges such that they form a clique. Then, for each $u_i, i \in [n/4]$ add edges (u_i, v_i) , (v_i, w_i) and (w_i, x_i) . With probability 1 as $n \rightarrow \infty$, there exists an i such that u_i and w_i are mutants and v_i and x_i are incumbents. Observe that $F(v_i) = F(x_i) = F(s|t)$ and $F(w_i) = F(t|s)$. \square

6.2 A Stronger Notion of Contraction for Individual Mutants

The model of Section 3 requires that for an edge (i, j) to contract, the fitness of i must be greater than the fitness of j . One way to strengthen this notion of contraction would be to require that the *maximum* fitness incumbent in the *neighborhood* of j be more fit than the *maximum* fitness mutant in the *neighborhood* of j . This models the idea that each organism is trying to take over each place in its neighborhood, but only the most fit organism in the neighborhood of a vertex gets the privilege of taking it. If we assume that we adopt this notion of contraction for individual mutants, and require that all incumbent-mutant edges contract, we will next show that Theorems 6.1 and 6.2 still hold, and thus it is still impossible to get results such as Theorems 5.1 and 5.2 which hold for every 2-player, symmetric game.

In the proof of Theorem 6.1 we proved that $F(i)$ is strictly less than $F(j)$. Observe that maximum fitness mutant in the neighborhood of j must have fitness at least $F(j)$. Also observe that there is only

1 incumbent in the neighborhood of j , namely i . So, under this stronger notion of contraction, the edge (i, j) will not contract.

Similarly, in the proof of Theorem 6.2, observe that the only mutant in the neighborhood of w_i is w_i itself, which has fitness $F(t|s)$. Furthermore, the only incumbents in the neighborhood of w_i are v_i and x_i , both of which have fitness $F(s|t)$. By assumption, $F(t|s) > F(s|t)$, thus, under this stronger notion of contraction, neither of the incumbent-mutant edges, (v_i, w_i) and (x_i, w_i) , will contract.

References

- [Blu93] L. E. Blume. The statistical mechanics of strategic interaction. *Games and Economic Behavior*, 5:387–424, 1993.
- [Blu95] L. E. Blume. The statistical mechanics of best-response strategy revision. *Games and Economic Behavior*, 11(2):111–145, November 1995.
- [Bol01] B. Bollobás. *Random Graphs*. Cambridge University Press, 2001.
- [Chw00] Michael Suk-Young Chwe. Communication and coordination in social networks. *Review of Economic Studies*, 67:1–16, 2000.
- [Ell93] Glenn Ellison. Learning, local interaction, and coordination. *Econometrica*, 61(5):1047–1071, Sept. 1993.
- [ESS98] I. Eshel, L. Samuelson, and A. Shaked. Altruists, egoists, and hooligans in a local interaction model. *The American Economic Review*, 88(1), 1998.
- [GS01] Geoffrey R. Grimmett and David R. Stirzaker. *Probability and Random Processes*. Oxford University Press, 3rd edition, 2001.
- [KKLO03] S. Kakade, M. Kearns, J. Langford, and L. Ortiz. Correlated equilibria in graphical games. *ACM Conference on Electronic Commerce*, 2003.
- [KKO⁺04] S. Kakade, M. Kearns, L. Ortiz, R. Pemantle, and S. Suri. Economic properties of social networks. *Neural Information Processing Systems*, 2004.

- [LHN05] E. Lieberman, C. Hauert, and M. A. Nowak. Evolutionary dynamics on graphs. *Nature*, 433:312–316, 2005.
- [Mor00] S. Morris. Contagion. *Review of Economic Studies*, 67(1):57–78, 2000.
- [Smi82] J. M. Smith. *Evolution and the Theory of Games*. Cambridge University Press, 1982.
- [Vic87] William L. Vickery. How to cheat against a simple mixed strategy ess. *Journal of Theoretical Biology*, 127:133–139, 1987.
- [Wei95] Jürgen W. Weibull. *Evolutionary Game Theory*. The MIT Press, 1995.