

The Impact of Network Topology on Pure Nash Equilibria in Graphical Games

Bistra Dilkina and Carla P. Gomes and Ashish Sabharwal

Department of Computer Science
Cornell University, Ithaca, NY 14850, U.S.A.
{bistra,gomes,sabhar}@cs.cornell.edu

Abstract

Graphical games capture some of the key aspects relevant to the study and design of multi-agent systems. It is often of interest to find the conditions under which a game is stable, i.e., the players have reached a consensus on their actions. In this paper, we characterize how different topologies of the interaction network affect the probability of existence of a pure Nash equilibrium in a graphical game with random payoffs. We show that for tree topologies with unbounded diameter the probability of a pure Nash equilibrium vanishes as the number of players grows large. On the positive side, we define several families of graphs for which the probability of a pure Nash equilibrium is at least $1 - 1/e$ even as the number of players goes to infinity. We also empirically show that adding a small number of connection “shortcuts” can increase the probability of pure Nash.

Introduction

In recent years, game theory has moved to the forefront of a number of disciplines: in economics, game theory is used to model complex problems of strategic interaction and confrontation, as found in markets and auctions; in computer science, game theory provides much of the foundation for the design of distributed algorithms and network protocols; in artificial intelligence, game-theoretic models effectively capture key aspects relevant to the study and design of multi-agent systems.

The notion of stability or equilibrium among players is central in game-theoretic settings. While several notions of equilibrium have been proposed, Nash’s concept of equilibrium (Nash 1951) in a non-cooperative setting is arguably the most important and widely used solution concept in game theory. A Nash equilibrium is a profile of strategies in which each player has no incentive to deviate from his strategy, given the other players’ strategies. Nash proved that every game has a mixed or randomized Nash equilibrium in which a player’s strategy is captured by a probability distribution over his actions. On the other hand, when each player has to choose a *pure strategy*, i.e., a single action instead of a randomized mix of actions, a Nash equilibrium is not guaranteed to exist. In this paper, we focus on the study of the

conditions under which pure Nash equilibria (PNEs) exist for graphical games with random payoffs.

Graphical games were proposed by Kearns, Littman, & Singh (2001) as a game-theoretic model for studying large-scale networks of agents with limited interaction. In a graphical game, nodes represent players and undirected edges represent interactions between them, thereby capturing the locality of interactions among players. This is a departure from the standard game theoretic models where each agent potentially interacts with all other agents. Graphical games naturally occur in markets, the Internet, and numerous settings with non-trivial network topologies. The graph can capture many types of interactions. For example, the nodes can represent individuals with edges indicating who knows whom, or different countries with edges representing trade agreements, or different species with edges representing biological interactions. For games with sparsely connected graphs, the graphical game representation gives an exponential space savings compared to the standard game representation as a matrix.

While a mixed Nash equilibrium always exists, deciding whether a pure Nash equilibrium exists is NP-complete for graphical games, even when each player interacts with at most 3 other players (Gottlob, Greco, & Scarcello 2003). Nevertheless, there are several reasons why pure Nash equilibria are more desirable than mixed Nash equilibria. For example, mixed strategies may be too complicated to implement in practice, which has been pointed out in the context of bounded rationality. Additionally, when playing mixed strategies (even if they lead to an equilibrium in expectation), the worst possible outcome can have a much lower payoff than the expected payoff. Consequently, players may want to avoid playing a mixed strategy in games with high stakes or where specific guarantees are needed, e.g., when the games concern critical operations or when the players are accountable for the final outcome, such as to their superiors. In economic contexts, Vega-Redondo (2003) points out that one seldom observes agents resorting to stochastic mechanisms: decision rules used by economic agents may be quite complex but should usually be conceived as deterministic. Yet another reason to study PNEs is that when using best-response learning dynamics to discover optimal strategies, convergence occurs only if there exists a PNE (Vega-Redondo 2003).

There have been several studies related to PNEs. For example, Daskalakis & Papadimitriou (2006) recently analyzed the algorithmic complexity of finding PNEs. Galeotti *et al.* (2006) have characterized the structure of PNEs in so-called network games. In this work, we explore the *existence of PNEs* as a probabilistic property of random-payoff graphical games, in contrast with the algorithmic decision problem or a structural analysis.

Known formal results on the existence of PNEs are for random normal form (or standard matrix) games, where each player’s payoff, chosen uniformly at random, is a function of the selection of actions for *all* players. Normal form games are equivalent to graphical games on complete (clique) graphs. The probability of existence of a PNE for a normal form game with random payoffs converges to $1 - 1/e \approx 0.63$ as the number of players grows (Rinott & Scarsini 2000). A key question concerning PNEs for graphical games is, *do such equilibria exist, given the topology of the underlying interaction graph?* In this paper, we study this question in detail with respect to graphical games with random payoffs.

Preview of Results. We show that for graphical games, different interaction graph topologies lead to radically different behavior: depending on the topology, the probability of having a PNE remains quite high or vanishes completely as the number of players grows. We note that the probability here is with respect to random payoffs; the topologies themselves are structured (except, of course, when we consider random interaction graphs).

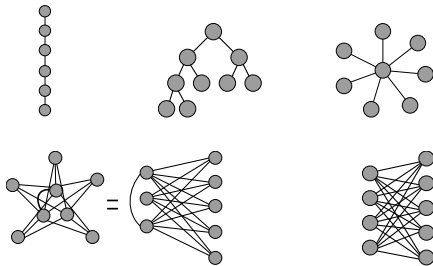


Figure 1: Topologies: path, tree, star, m-star, augmented complete bipartite, and complete bipartite.

Figure 1 depicts the different graph topologies we consider in our theoretical analysis. Our results are with respect to games where each player has 2 actions, unless otherwise specified. Our most basic topology consists of a linear chain or “path” of interactions between players. We prove that as the number of players grows, the probability of having a PNE converges to 0. We then extend our topology to tree structures, and again formally show that the probability of having a PNE converges to zero as the number of players grows. Interestingly, our empirical data reveals that the convergence to zero is much slower on trees than on the linear chain with the same number of nodes. In fact, the higher the branching factor of the tree, the slower the rate of convergence. For example, for a 50 player linear chain, the probability of having a PNE is around 0.04, while for a tree with a

branching factor of 5 and 50 players, this probability is 0.25. Intuitively, while long chains of interactions prevent emergence of PNEs, short paths of interactions between players increases this probability. For instance, when we consider the so-called star topology — one central player interacting with all other players independently — we can show that the probability of having a PNE converges to 0.75 as the number of players grows. Table summarizes our key theoretical findings.

Topology	Prob. of Pure Nash
star	0.75
2-star	0.683
augmented bipartite	$1 - (1 - \frac{1}{2m})^{2m}$
bipartite	$1 - 1/e \approx 0.632$
matrix	$1 - 1/e \approx 0.632$
tree	0
path	0

Table 1: Asymptotic probability of having a PNE in 2-action games with various interaction topologies.

We define a family of graphs that generalize the star graph by allowing the center of a graph of n nodes to have m nodes. The “outer” nodes interact (only) with each of the m center nodes, while the center nodes themselves are connected arbitrarily (see Figure 1). We refer to this topology as an *m-star*, or equivalently an *augmented bipartite graph*. We show that for such graphs and with k possible actions per player, the probability of having a PNE converges to $1 - (1 - 1/k^m)^{km}$ as $n \rightarrow \infty$ whenever m is such that $n/3 - m \rightarrow \infty$ (i.e., the center of the graph is small enough). For 2-action games ($k = 2$), the lower bound for the probability of having a PNE given this topology lies between $1 - 1/e \approx 0.63$ and 0.75.

Another topology we consider is that of complete bipartite graphs. Such interaction graphs occur naturally in several settings such as when modeling the interactions between buyers and sellers or between bidders and auctioneers. We show that as the number of players grows, the probability of having a PNE converges to $1 - 1/e \approx 0.63$. (This is the same value as for standard matrix games, but for a richer underlying topology.)

Finally, we study empirically what happens when we morph a highly regular graph into a random graph, through a sequence of random re-wirings. This process was inspired by the work on Small World Graphs (Watts & Strogatz 1998). Our experiments show that the probability of having a PNE increases significantly with more rewiring. For a large number of players, the regular graph game behaves similarly to the chain and the tree graph games in that the probability of having a PNE goes to zero; however, with only a few random re-wirings, the probability of a PNE dramatically increases. A similar increasing probability phenomenon is observed as more and more edges are added to purely random interaction graphs.

In summary, our results demonstrate that the topology of player interactions can greatly affect the probability of existence of a pure Nash equilibrium. We show that long linear chains of interactions dramatically reduce the chance of

having a PNE. One can reverse this effect by adding a small number of random interactions. Such new interactions also implicitly shorten the longest chains in the network. Interestingly, and perhaps contrary to one’s intuitions, having more interactions between players can actually increase the probability of having a PNE.

Overall, our analysis suggests that one can exploit the topology of the interaction graph when designing networks of interacting agents or components. Ideally, one would choose a topology that maximizes the probability of having a pure Nash equilibrium, which would be beneficial to all players. Our results provide insights for identifying such preferred topologies.

Preliminaries

We begin by formally defining games, strategy profiles, Nash equilibrium, and graphical games. We will then discuss random-payoff games and some topologies of interest.

Definition 1. A *game* is a triple $\Gamma = (P, \{A_p \mid p \in P\}, \{u_p \mid p \in P\})$ where P is a set of players, A_p is the set of possible actions for player p , and $u_p : \prod_{i \in P} A_i \rightarrow \mathbb{R}$ is the payoff function of p .

When $|P| = n$, Γ is referred to as an n -player game. When each $|A_p|$ is finite, we have a *finite game*. Throughout this paper, we will work with n -player finite games.

For $W \subseteq P$, a *pure (partial) strategy profile*, s_W , for Γ is an element of the set $\prod_{p \in W} A_p$, i.e., it is a $|W|$ -tuple that specifies one action for each player in W . For succinctness, we will refer to this simply as a strategy profile for W , and when $W = P$, denote it by s . With slight abuse of notation, we will denote by s_p the action specified for player p by the strategy profile s . Given a (partial) profile s_W , we will use $(s_{W-\{p\}}, a_p)$ to denote the profile where all players but p have the same action as in s_W and p has action a_p .

A pure Nash equilibrium is a strategy profile where each player has no incentive to unilaterally deviate and change his action in order to achieve a better payoff, given the actions chosen by the other players. Formally,

Definition 2. A *Pure Nash equilibrium (PNE)* for $\Gamma = (P, \{A_p\}, \{u_p\})$ is a strategy profile s for Γ such that $\forall p \in P, \forall a \in A_p : u_p(s) \geq u_p(s_{P-\{p\}}, a)$.

Graphical games capture the notion that in many scenarios, a player’s payoff may only depend on the choices of a subset of the other players (his “neighbors”). The interactions between players are represented by an undirected graph (i.e., we work with *symmetric* games) where nodes represent players and edges represent mutual dependencies between players. We will use nodes and players interchangeably. For a graphical game, u_p , the payoff of p , depends only on the actions of players that are adjacent to p in the underlying graph, denoted by $Nbr(p)$.

Definition 3. A *graphical game* is a triple $\Gamma = (G = (V, E), \{A_p \mid p \in V\}, \{u_p \mid p \in V\})$ where G is a (connected) undirected graph, V represents the set of players, A_p is the set of available actions for player p , and $u_p : \prod_{q \in \{p\} \cup Nbr(p)} A_q \rightarrow \mathbb{R}$ is the payoff function of p .

Our focus will be on *random-payoff graphical games*, that is, graphical games where payoff functions are chosen at random. Specifically, let Γ be an n -player graphical game defined over $G = (V, E); |V| = n$. For any player p , the payoff function u_p can be thought of as a *payoff table* U_p with one payoff value for each tuple of possible actions of $\{p\} \cup Nbr(p)$. Our *random distribution* on games over G is defined as follows: for all payoff tables U_p in Γ , choose each payoff value in U_p uniformly and independently at random from the interval $[0, 1]$. (We will shortly replace this uniform random distribution over payoffs with an essentially equivalent but simpler discrete random distribution over “best response tables.”)

Interestingly, in order to determine whether a strategy profile s is a PNE of a graphical game Γ , we only need to consider whether or not for each player p , the action s_p is p ’s best response to the actions $s_{Nbr(p)}$ of his neighbors; it does not matter what the exact payoff values are. This motivates the following definition.

Definition 4. Let Γ be a graphical game. A *best response action* of p w.r.t. a strategy profile s , denoted $BR_p(s_{Nbr(p)})$, is an action that maximizes p ’s payoff with respect to $s_{Nbr(p)}$.

When finitely many payoffs are drawn randomly from continuous distributions (e.g., uniformly from $[0, 1]$), they are all distinct with probability 1. Therefore, every player nearly always has a *unique* best response action for each tuple of actions of the other players. Throughout this paper, we will assume this uniqueness and think in terms of *best response tables* T_p derived naturally from payoff tables U_p : for each strategy profile $s_{Nbr(p)}$, T_p specifies the unique best response action of p , $BR_p(s_{Nbr(p)})$ (which in turn is defined by U_p).

The uniqueness of best response actions simplifies our *random distribution* on games Γ over a graph G to the following: *for all best response tables T_p in Γ , choose each best response action in T_p uniformly and independently at random from A_p .* We will focus on games where $|A_p| = k$ for each p ; k will be 2 unless otherwise stated. In this case, there are $k^{|Nbr(p)|}$ entries or “rows” in T_p for each of which the best response action of p is specified, and the random distribution can be thought of as each T_p being chosen uniformly and independently at random from the set of $k^{|Nbr(p)|}$ possible best response tables.

Example 1. Consider a graphical game on a “path” graph consisting of three players $\{p, q, r\}$ and edges $\{\{p, q\}, \{q, r\}\}$, with action sets $A_p = A_q = A_r = \{0, 1\}$. The payoff function u_p of p is defined over $A_q \times A_p$, and the payoff values are selected uniformly at random from $[0, 1]$. Figure 2 gives an example of the payoff table U_p , and the best response table T_p determined by these payoffs and defined over A_q . Similarly, since $Nbr(q) = \{p, r\}$, the best response table of q is defined over $A_p \times A_r$. Finally, since $Nbr(r) = \{q\}$, the best response table of r is defined over A_q . Figure 2 gives examples of the corresponding best response tables T_q and T_r .

Each best response table can be viewed as a *constraint* that must be satisfied by a strategy profile s

q	p	u_p
0	0	0.3
0	1	0.5
1	0	0.7
1	1	0.2

q	BR_p
0	1
1	0

p	r	BR_q
0	0	0
0	1	1
1	0	1
1	1	0

q	BR_r
0	0
1	0

Figure 2: Best response function of p , and best response tables of p, q , and r for the path graph $\boxed{p}-\boxed{q}-\boxed{r}$

for it to be a PNE. For example, T_p disallows profiles $(000), (001), (110), (111)$ for (p, q, r) and allows $(010), (011), (100), (101)$. Similarly, T_p and T_q together only allow profiles $(011), (101)$ as PNEs, which in particular implies that r must take action 1. However, r 's best response is not 1 when q plays either 0 or 1, and hence there is no PNE in this game. \square

The key question we are interested in is the following: *For an interaction graph G and possible player actions $\{A_p\}$, what is the probability that a graphical game Γ , defined on G and $\{A_p\}$ and with best response tables chosen uniformly at random, has a PNE?* When the underlying graph is implicit, we will often use the phrase “the probability of a PNE” to represent this probability. We will study this problem for various families of graphs, i.e., for different *graph topologies* parameterized by the size of the graph. For each such topology, we will explore how the probability of a PNE behaves as the size of the graph grows.

We briefly describe the key topologies of interest (refer to Figure 1). The simplest of these is the *path graph* P_n , which is a connected graph with $n - 2$ vertices of degree 2 and two “end” vertices of degree 1. A *tree graph* is an acyclic connected graph, without any degree restrictions. A *star graph* S_n is a tree of n nodes with one “center” node and $n - 1$ “outer” nodes. A *clique graph* K_n has n nodes with each pair of nodes connected. For two disjoint sets of vertices X and Y , the *complete bipartite graph* $K(X, Y)$ is the graph with nodes $X \cup Y$ and edges $\{(x, y) \mid x \in X, y \in Y\}$.

We will also consider an extension of complete bipartite graphs: an *augmented complete bipartite graph* $\tilde{K}(X, Y, E_X)$ is a graph where the set of vertices is $X \cup Y$ and the set of edges is $E_X \cup \{(x, y) \mid x \in X, y \in Y\}$, where E_X is an arbitrary set of edges within X . This can alternatively be thought of as an extended star graph with X as the set of arbitrarily connected center vertices and Y as the set of outer vertices, each connected to all of the center vertices.

The diameter of a graph will be a parameter of interest to us. It is defined as the maximum shortest distance over all pairs of vertices. In the case of tree graphs, the diameter is simply the length of the longest path in the tree.

Theoretical Results

In this section, we derive formal results about the probability of a PNE in n -player random-payoff graphical games on sequences of graphs G_n . When the considered sequence of graphs has a tree topology with unbounded diameter and two actions per player, we will show that this probability converges to 0. On the other hand, when the sequence of

graphs has a (augmented) complete bipartite topology (see Preliminaries), we will show that this probability converges to a non-zero quantity, sometimes as high as 0.75.

We begin with the result for tree topologies. For this, we first prove that the probability of a PNE in a random game on a path converges to 0 as $n \rightarrow \infty$. We then show that adding a single vertex with one edge connecting this vertex to the existing graph cannot increase the probability of a PNE. Combining these two building blocks and viewing a tree as being constructed by sequentially adding a set of vertices to a path representing the tree’s diameter, we will have the desired result for unbounded diameter trees. We first give a small example to illustrate the case-based reasoning technique.

Example 2. Consider the 2-action 2-player random payoff game Γ on the path graph P_2 . We will use $0_p, 1_p$ to denote the actions of player $p \in \{1, 2\}$. Let us denote by A the event that the best response of player 1 to both actions of player 2 is the same, i.e. $BR_1(0_2) = BR_1(1_2)$. In a random-payoff game, the probability of A is $1/2$ because of the independence of the best responses. In the case when event A happens, let $x_1 = BR_1(0_2) = BR_1(1_2), x_1 \in \{0_1, 1_1\}$. Then the profile $(x_1, BR_2(x_1))$ is necessarily a PNE since both players are in their best response state, that is, $\Pr[\Gamma \text{ has a PNE} \mid A] = 1$. On the other hand, in the case that event A does not happen, i.e., when $BR_1(0_2) \neq BR_1(1_2)$, let $x_1 = BR_1(0_2)$ and $\bar{x}_1 = BR_1(1_2)$. There are two candidate profiles for having a PNE: $(x_1 0_2)$ and $(\bar{x}_1 1_2)$. Given \bar{A} , the probability that $(x_1 0_2)$ is a PNE is equal to the probability that $BR_2(x_1) = 0_2$ which is $1/2$. Similarly for $(\bar{x}_1 1_2)$. Since $BR_2(x_1)$ and $BR_2(\bar{x}_1)$ are independent events, we obtain that $\Pr[\Gamma \text{ has a PNE} \mid \bar{A}] = 3/4$. Finally, the overall probability of Γ having a PNE is $\frac{1}{2} \times 1 + \frac{1}{2} \times \frac{3}{4} = \frac{7}{8}$. \square

Lemma 1 (Path). *Let Γ_n be a 2-action random-payoff game on the path graph P_n . Then $\Pr[\Gamma_n \text{ has a PNE}] \leq \left(\frac{63}{64}\right)^{n-1}$.*

Proof Sketch. Fix the sets of player actions to $\{0, 1\}$. Let \Pr_n denote the probability that Γ_n does not have a PNE. Let $(v_1, v_2, v_3, \dots, v_n)$ be the nodes of P_n , and consider the end player v_1 . Since the choice of actions of player v_1 depends only on player v_2 , there are four possible best response tables T_{v_1} which occur with equal probability of $1/4$: for $i, j \in \{0, 1\}$, $T_{v_1}^{ij}$ corresponds to the best response of v_1 being i when v_2 plays 0 and j when v_2 plays 1. We first consider $T_{v_1}^{00}$. Denote the rows in v_2 's table where v_1 plays 0 as $T_{v_2 0}$, and where v_1 plays 1 as $T_{v_2 1}$. Both $T_{v_2 0}$ and $T_{v_2 1}$ are equivalent to a best response table for v_2 with respect to his other neighbor v_3 . Since the table $T_{v_1}^{00}$ of v_1 precludes any PNE where v_1 plays 1, we need only consider the probability that there is no PNE when v_1 plays 0. Given $T_{v_1}^{00}$, the probability of choosing $\{T_{v_i}\}_{i=2..n}$ such that there is no PNE is equal to the probability of choosing $\{T_{v_i}\}_{i=3..n} \cup T_{v_2 0}$ such that there is no PNE. Hence, conditioned on $T_{v_1}^{00}$, the probability of not having a PNE in a n -path game is equal to the probability of not having a PNE in an $(n - 1)$ -path game, \Pr_{n-1} . The same holds for $T_{v_1}^{11}$.

When the table of v_1 is either $T_{v_1}^{01}$ or $T_{v_1}^{10}$, v_1 responds with different actions to the play of v_2 . In particular, $T_{v_1}^{01}$

only allows profiles where v_1 and v_2 play the same action. Given $T_{v_1}^{01}$, the probability of not having a PNE is equal to the probability of choosing $\{T_i\}_{i=2..n}$ such that any profile $(0, 0, s_3, \dots, s_n)$ or $(1, 1, s_3, \dots, s_n)$ is disallowed. Let us denote this probability by $\Pr_{n-1}(00, 11)$. The case for $T_{v_1}^{10}$ is symmetric. Considering the four possible tables $T_{v_1}^{ij}$, we have the recursive expression $\Pr_n = \frac{1}{4}\Pr_{n-1} + \frac{1}{4}\Pr_{n-1} + \frac{1}{4}\Pr_{n-1}(00, 11) + \frac{1}{4}\Pr_{n-1}(00, 11)$. Now, we need to expand on $\Pr_{n-1}(00, 11)$. Due to space limitations, the rest of the proof is omitted. We can derive a set of six mutually recursive expressions that explicitly capture the probability of not having a PNE in Γ_n . Using these expressions, we can inductively prove that $\Pr_n \leq \left(\frac{63}{64}\right)^{n-1}$. \square

For a graph $G = (V, E)$, we use the term G -game to denote a random-payoff graphical game on G with uniform distribution over best response tables as described earlier. Let Γ be such a randomly chosen G -game. For a vertex $u \in V$, let U_0 and U_1 denote the events that a Γ has a PNE where u plays 0 and 1, respectively. $\overline{U_0}$ and $\overline{U_1}$ will represent the complements of these two events, respectively. Note that the probability that Γ has a PNE equals, by definition, the probability that at least one of the events U_0 and U_1 happens.

To analyze the probability of a PNE in Γ , we consider the set of best response tables T_W of all players other than u ; $W = V \setminus \{u\}$. Observe that given the choice of T_W , the probability that Γ has a PNE depends only on the best response table of u , T_u . Our proofs will, in general, do a case-based analysis exploiting this fact. We begin with a general property of random-payoff graphical games which will be used in proving Lemma 3.

Lemma 2. *Let Γ be a random-payoff graphical game over $G = (W \cup \{u\}, E)$ and T_W be best response tables. Then*

$$\Pr[\overline{U_0} \cap \overline{U_1} \mid T_W] \leq \Pr[\overline{U_0} \mid T_W] \cdot \Pr[\overline{U_1} \mid T_W].$$

Proof. It is easy to see that if either of $\Pr[\overline{U_0} \mid T_W]$ and $\Pr[\overline{U_1} \mid T_W]$ is 0 or 1, the claimed inequality holds. Assume therefore these probabilities are strictly between 0 and 1. In particular, this means that $\Pr[U_0 \mid T_W] > 0$ and $\Pr[U_1 \mid T_W] > 0$, which we will exploit later. For now, we can rewrite what we need:

$$\Pr[\overline{U_0} \mid T_W] \geq \frac{\Pr[\overline{U_0} \cap \overline{U_1} \mid T_W]}{\Pr[\overline{U_1} \mid T_W]} = \Pr[\overline{U_0} \mid \overline{U_1} \cap T_W]$$

This is equivalent to $\Pr[U_0 \mid T_W] \leq \Pr[U_0 \mid \overline{U_1} \cap T_W]$, i.e., given T_W , the probability of G having a PNE where u plays 0 should be not decrease if we condition on G not having a PNE where u plays 1. Intuitively, this holds because not having a PNE where u plays 1 gives u more ‘‘freedom’’ to play 0 more often in his best response tables. We now prove this formally.

Consider the set S' of strategy profiles allowed by T_W , i.e., strategy profiles s where each of the players in W plays his best response action given the actions of the other players in s . There exists a PNE in Γ given T_W iff at least one $s \in S'$ is allowed by the best response table T_u of u . Since the responses of u only depend on the actions of the players in $Nbr(u)$, we will focus on the profiles over $Nbr(u) \cup \{u\}$ allowed by T_W . Let this set of profiles be S . When projected

over $Nbr(u)$, S is partitioned into three disjoint sets of profiles over $Nbr(u)$: those for which u can only play 0 to be acceptable by T_W , those for which u can only play 1, and those for which u can play either 0 or 1. Let us denote these by S_0, S_1 , and S_{01} , respectively. Note that the profiles in S_0, S_1 , and S_{01} correspond to various rows of T_u . Recall that we are assuming $\Pr[U_0 \mid T_W] > 0$ and $\Pr[U_1 \mid T_W] > 0$, so that $|S_0 \cup S_{01}| > 0$ and $|S_1 \cup S_{01}| > 0$.

Let $\{B_k\}$ denote the set of events that T_u assigns some specific best responses ‘‘k’’ to the rows $S_1 \cup S_{01}$. There are $2^{|S_1 \cup S_{01}|}$ such events that are disjoint and equally likely. Given such an event B_i with the additional property that none of the S_{01} rows of T_u have 0 as the best response for u , the probability that T_u does not allow a PNE with u playing 0 equals the fraction of the possible tables where the best response of u to each S_0 row is 1, which is exactly $2^{-|S_0|}$. It follows that $\Pr[U_0 \mid B_i \cap T_W] = (1 - 2^{-|S_0|})$ for such an event B_i . On the other hand, given an event B_j with the additional property that some S_{01} row of T_u prescribes 0 as the best response for u , $\Pr[U_0 \mid B_j \cap T_W] = 1$. With these two disjoint kinds of events B_i and B_j together, we have that $\Pr[U_0 \mid T_W] = \sum_{k=1}^{2^{|S_1 \cup S_{01}|}} \Pr[U_0 \mid B_k \cap T_W] \cdot \Pr[B_k \mid T_W] = 2^{-|S_1 \cup S_{01}|} \sum_{k=1}^{2^{|S_1 \cup S_{01}|}} \Pr[U_0 \mid B_k \cap T_W] \leq \max_k \Pr[U_0 \mid B_k \cap T_W]$.

On the other hand, the event $\overline{U_1}$ is equivalent to choosing T_u such that the best response to each row in $S_1 \cup S_{01}$ is 0. There is only one such event B' . Notice that when $|S_{01}| = 0$, $\Pr[U_0 \mid B' \cap T_W] = (1 - 2^{-|S_0|}) = \max_k \Pr[U_0 \mid B_k \cap T_W]$, and also when $|S_{01}| > 0$, $\Pr[U_0 \mid B' \cap T_W] = 1 = \max_k \Pr[U_0 \mid B_k \cap T_W]$. Hence, $\Pr[U_0 \mid \overline{U_1} \cap T_W] = \max_k \Pr[U_0 \mid B_k \cap T_W] \geq \Pr[U_0 \mid T_W]$, which is what we needed. \square

We now show that adding a degree 1 vertex v to a graph G cannot increase the probability of having a PNE.

Lemma 3. *Let $G = (V, E)$ with $u \in V$ and $v \notin V$, and $G' = (V \cup \{v\}, E \cup \{(v, u)\})$. Then $\Pr[a G\text{-game has a PNE}] \leq \Pr[a G'\text{-game has a PNE}]$.*

Proof Sketch. Let $V = \{v_1, v_2, \dots, v_n\}$. Without loss of generality, let $u = v_n$ and $W = \{v_1, \dots, v_{n-1}\}$. Games over G' are defined by the best response tables of v , u , and the players in W . Since u is the only neighbor of v , the best response table for v has 2 entries specifying the best response actions of v to the play of u . Let the neighbors of u in G be N_u and let $|N_u| = d_u$. In G , the best response table of u is defined over N_u and has 2^{d_u} rows. In G' , u 's best response table is defined over $N_u \cup \{v\}$ and has 2^{d_u+1} rows. The tables for the rest of the players remain unaffected. To prove the lemma, we will show that the probability of not having a PNE in G' is no less than the probability of not having a PNE in G . Recall that the probability of not having a PNE in a G -game is $\Pr[\overline{U_0} \cap \overline{U_1}]$, where U_0 and U_1 denote events that there is a PNE where u plays 0 and 1, respectively. Similarly, we define U'_0 and U'_1 for a G' -game and consider $\Pr[\overline{U'_0} \cap \overline{U'_1}]$.

We will show that for every instantiation of the tables T_W of the players in W , the probability that G' has no PNE is no less than the probability that G has no PNE. Given T_W , there are 4 possible instantiations of the table T_v of v , each

of which occurs with probability $1/4$. We will show that in each of the four disjoint cases and conditioned on T_v and T_w , the probability that G' has no PNE is no less than the probability that G has no PNE. This will prove the lemma because the probability of G' not having a PNE will be ‘‘pointwise’’ no less than the probability of G not having a PNE, i.e., the desired inequality will hold for every possible instantiation of T_v and T_w .

To this end, observe that given T_v and T_w , the probability of not having a PNE is simply equal to the fraction of the $2^{2^{d_u+1}}$ possible table instantiations for u that do not allow for a PNE (we can talk of fractions of tables because the distribution is uniform). We can compute this probability by considering the four possible tables of v and analyzing the events U'_0 and U'_1 in relation to U_0 and U_1 , using Lemma 2. We omit the fairly involved details for lack of space. \square

Theorem 1 (Trees). *Let Γ_n be a 2-action random-payoff game on a tree graph T_n such that $\text{diameter}(T_n)$ grows without bounds with n . Then $\Pr[\Gamma_n \text{ has a PNE}] \rightarrow 0$ as $n \rightarrow \infty$.*

Proof. One can think of any tree graph T as a sequence of vertex additions starting with a path graph P with $\text{length}(P) = \text{diameter}(T)$. Lemma 3 says that $\Pr[T \text{ has PNE}] \leq \Pr[P \text{ has PNE}]$, and from Lemma 1 we know that the latter goes to zero as $\text{length}(P)$ goes to infinity. We therefore have the desired result whenever $\text{diameter}(T)$ is unbounded. \square

Notice that such tree topologies include as special cases trees with a constant branching factor and also graphs like a star attached to a long path. On the other hand, we show that other topologies are conducive to the existence of a PNE, where the probability of a PNE is bounded from below by a non-zero constant even as the number of players grows.

Theorem 2 (Augmented Complete Bipartite). *Let Γ_n be a random-payoff game on an augmented complete bipartite graph $G_n = \tilde{K}(X_n, Y_n, E_{X_n})$ such that $|X_n \cup Y_n| = n$, $|X_n| = m$, $n/3 - m \rightarrow \infty$ as $n \rightarrow \infty$, and each player in X_n has k actions. Then $\Pr[\Gamma_n \text{ has a PNE}] \rightarrow \left(1 - \left(1 - \frac{1}{k^m}\right)^{k^m}\right)$ as $n \rightarrow \infty$.*

Proof. We will refer to the players in X_n and Y_n as the center and leaf players, respectively. The center players admit k^m different profiles; let c be one such profile. In response to c , each leaf node has a unique best response action which only depends on c . These best responses together form a best response strategy profile $s_{Y_n}^c$. The profiles $S = \{(s_{Y_n}^c, c) \mid c \in \prod_{p \in X_n} A_p\}$ are the only potential global strategy profiles that could be PNEs since in any other profile at least one of the leaf players will have an incentive to deviate. It follows that Γ_n has no PNE iff each potential profile in S is rejected by the center players.

Let c, c' be two profiles of the center players. Let us assume for now that the corresponding profiles $s_{Y_n}^c$ and $s_{Y_n}^{c'}$ of the leaf players are distinct. (We will soon prove that this happens with probability approaching 1.) Now, $(s_{Y_n}^c, c)$ is a PNE iff $\forall x \in X_n, BR_x(s_{Y_n}^c, c) = c_x$, which happens with probability exactly k^{-m} . Further, since $s_{Y_n}^c \neq s_{Y_n}^{c'}$, the two events that $(s_{Y_n}^c, c)$ is a PNE and $(s_{Y_n}^{c'}, c')$ is a PNE are independent.

Extending this argument to all profiles in S and assuming that all $s_{Y_n}^c$ are distinct, the probability that the center players reject all profiles in S is $(1 - k^{-m})^{k^m}$. Hence the probability that at least one of these profiles is a PNE given that the $s_{Y_n}^c$'s are all distinct is $1 - \left(1 - \frac{1}{k^m}\right)^{k^m}$, as claimed.

We now justify the assumption above by proving that with probability approaching 1, all $s_{Y_n}^c$ are distinct for the k^m center profiles c of X_n . Observe that for each of the k^m center profiles, $s_{Y_n}^c$ is chosen from k^{n-m} possible profiles independently and uniformly at random (since the best response table of the leaf players are independent and chosen uniformly at random). Therefore, if k^{n-m} is sufficiently larger than k^m , all these profiles will be distinct. The following calculation shows that it suffices to have $n/3 - m$ be $\omega(1)$ in n .

For $i \in \{1, \dots, k^m\}$, the probability that the i^{th} best response is different from the previous $i - 1$ responses is $\frac{k^{n-m} - (i-1)}{k^{n-m}}$. Hence the probability that all k^m best response choices are distinct is:

$$\begin{aligned} & \left(\frac{k^{n-m}}{k^{n-m}}\right) \left(\frac{k^{n-m} - 1}{k^{n-m}}\right) \cdots \left(\frac{k^{n-m} - (k^m - 1)}{k^{n-m}}\right) \\ & > \frac{(k^{n-m} - k^m)^{k^m}}{(k^{n-m})^{k^m}} = \left(1 - \frac{k^m}{k^{n-m}}\right)^{k^m} > e^{-\frac{2k^m \cdot k^m}{k^{n-m}}} = e^{-\frac{2}{k^{n-3m}}} \end{aligned}$$

This converges to 1 as $n - 3m \rightarrow \infty$. \square

Notice that the restriction on the size of the set X_n of center players is quite flexible. E.g., it allows $m = cn$ for $c < 1/3$, $m = n/3 - \log n$, etc. The probability of a PNE in this family of graphs is maximized when $m = 1$, which corresponds to the star graph S_n . When the center player has 2 actions, the probability of a PNE is considerably higher than a random-payoff matrix game.

Corollary 1 (Star). *Let Γ_n be a random-payoff game on the star graph S_n where the center player has 2 actions. Then $\Pr[\Gamma_n \text{ has a PNE}] \rightarrow 0.75$ as $n \rightarrow \infty$.*

In addition, the proof did not depend in any way on the interactions between the center players. If they are completely disconnected, then the topology corresponds to a complete bipartite graph (see Figure 1). However, we can say something stronger for complete bipartite graphs that is true even when both X and Y have large sizes. The proof of this next theorem is based on viewing complete bipartite topologies as 2-player games with unbounded action sets and using the result of Powers (1986).

Theorem 3 (Complete Bipartite). *Let Γ_n be a k -action random-payoff game on a complete bipartite graph $G_n = K(X_n, Y_n)$ where both $|X_n|$ and $|Y_n|$ grow without bound as n grows, Then $\Pr[\Gamma_n \text{ has a PNE}] \rightarrow 1 - 1/e$ as $n \rightarrow \infty$.*

Proof. Given Γ_n , one can easily construct a 2-player game Γ'_n with the following three properties: (1) the players of Γ'_n in X_n correspond to a single player x_n of Γ'_n with $k^{|X_n|}$ actions $A_{x_n} = \prod_{p \in X_n} A_p$; (2) the players in Y_n correspond to a player y_n of Γ'_n with $k^{|Y_n|}$ actions $A_{y_n} = \prod_{p \in Y_n} A_p$; and (3) Γ'_n has a PNE if and only if Γ_n has a PNE.

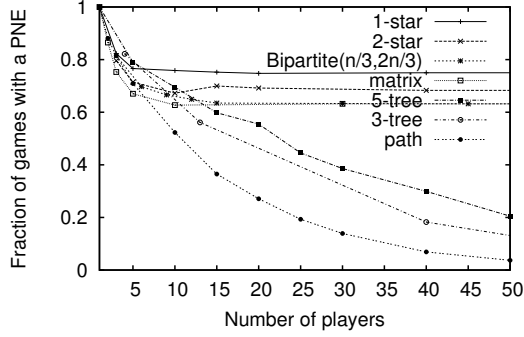


Figure 3: Topologies with analytical results

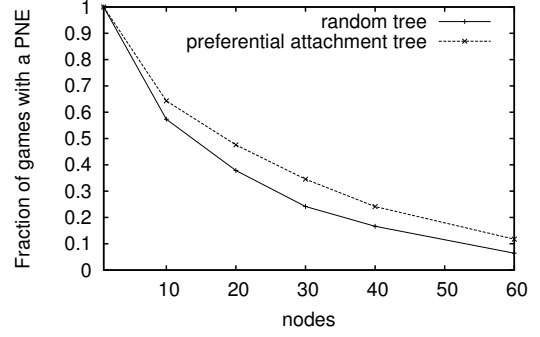


Figure 4: Random trees vs. preferential attachment trees

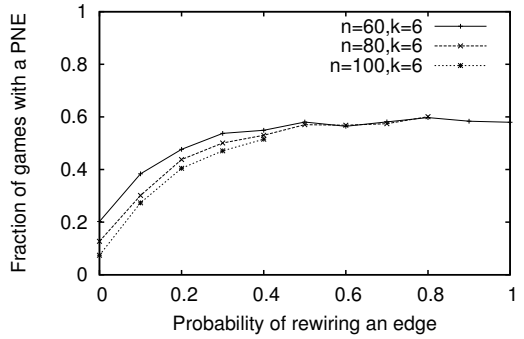


Figure 5: Small world graphs

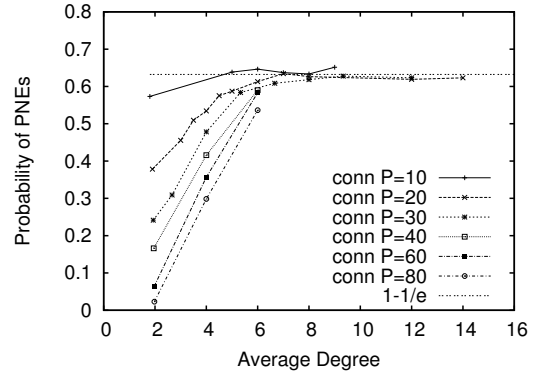


Figure 6: Connected $G(n, m)$ graphs

Since the best responses of the players in X_n are independent of each other (G_n being bipartite), choosing the best responses for X_n uniformly and independently at random in Γ_n corresponds to choosing a best response action for x_n uniformly and independently at random from A_{x_n} in Γ'_n . Similarly, for y_n . Therefore, considering a game over G_n with independent uniform-random payoffs is equivalent to considering a 2-player game with appropriate number of actions and independent uniform-random payoffs. Moreover, since both $|X_n|$ and $|Y_n|$ grow without bound as n grows, $|A_{x_n}|$ and $|A_{y_n}|$ also go to infinity as $n \rightarrow \infty$. Powers (1986) showed that the probability of a PNE in a random-payoff 2-player game converges to $1 - 1/e$ whenever the action sets of the two players grow infinitely large. Resorting to this result completes the proof. \square

Empirical Results

For our experiments, we used the GAMUT library (Nudelman *et al.* 2004), which facilitates the generation of graphical game instances with various topologies. The best response tables were chosen uniformly at random as discussed earlier. Each game instance Γ was translated into an equivalent propositional (SAT) formula F_Γ with the property that F_Γ was satisfiable if and only if Γ had a PNE. F_Γ was created by having two kinds of clauses, the first set making sure that every player chooses an action and the second set making

sure that every strategy profile that is not a PNE (because the action of one of the players is *not* his best response action w.r.t. this strategy profile) is ruled out. We then used the SAT solver *Minisat* (Een & Srensson 2005) to check whether F_Γ was satisfiable, thus determining the existence of a PNE in Γ . In our plots, each data point was obtained by averaging over 4000 game instances.

In Figure 3, we consider various interaction topologies and for each of these, plot the fraction of games with a PNE as a function of the number of players. The probability of having a PNE converges to zero for tree-like and path graphs, and to a non-zero value greater than or equal to $(1 - 1/e)$ for the other topologies. While this is already expected from our theoretical analysis, these plots are interesting because they provide further finer-grained insight into the convergence process, for instance, in terms of the rate of convergence. In particular, we see that the larger the branching factor of the tree-like topologies, the slower the convergence to zero.

Intuitively, a higher branching factor leads to relatively shorter paths between pairs of nodes, thereby increasing the chance of having a PNE. The figure shows this for only a few cases, but we have observed the phenomenon for a variety of tree-like topologies. For example, random trees grown with the so-called preferential attachment model (Albert & Barabasi 2002) have a slower rate of con-

vergence than purely random trees (see Figure 4). A preferential attachment tree can be generated by starting with a single vertex and adding a new node at a time with one edge that is connected to an existing node i with probability $\text{deg}(i)/(\sum_j \text{deg}(j))$. Such a model is characterized by a power-law node degree distribution and therefore it tends to decrease the path length between nodes by creating links to highly connected intermediate nodes.

In addition to highly structured topologies we also considered less structured graph models, such as the random graph model $G(n, m)$ and the Watts-Strogatz small world model (Watts & Strogatz 1998). These models capture some of the characteristics observed in social, biological, and computer networks. A *random graph* $G(n, m)$ is a graph with n vertices and m edges, where each edge is selected uniformly at random from all possible edges without repetition. A *Watts-Strogatz small world graph* $SWG(n, k, p)$ is a random graph with n vertices obtained from a random process that starts with a ring lattice with k edges per vertex; subsequently each of the $nk/2$ edges is rewired to a random vertex with probability $p \in [0, 1]$. Varying the rewiring probability allows us to interpolate between a highly regular graph and a fully random graph.

In our experiments with small world graphs, we start with a highly structured graph, a ring lattice with $k = 6$ edges per vertex when the probability of rewiring is $p = 0$, and end up with a random graph with $3n$ edges when $p = 1$. Figure 5 plots the results for graphs with 60, 80, and 100 players as a function of the probability of rewiring. This plot indicates that the rewiring of edges in a highly regular graph dramatically increases the probability of a PNE. Notice that the probability of a PNE increases monotonically with the increase in rewiring probability, in other words the random graph over the given number of nodes and edges maximizes the probability of PNE. The regular graph game behaves similarly to the chain and the tree graph games in that the probability of having a PNE goes to zero as the number of players increases; however, with only a few random re-wirings, the probability considerably increases. In particular, it is enough to set $p = 0.3$ to come reasonably close to the maximum probability of a PNE, given the number of nodes and edges. This suggests that the process of creating short paths between players as we rewire edges increases the chance of a PNE. A similar phenomenon is observed for random graphs generated using the $G(n, m)$ model in which we create shortcuts between players by increasing the number of edges in the graph. Again, Figure 6 shows that increasing the number of edges between players quickly increases the probability of a PNE. In particular, it seems enough to achieve an average degree of 6 to come reasonably close to the probability of having a PNE in a regular matrix game.

Conclusion

We studied the impact of interaction topology on the existence of pure Nash equilibria in graphical games. We identified a variety of graph structures with diverse asymptotic behavior. Our results for path and tree-like topologies show that with long chains of interactions between players, the probability of a PNE goes to zero as the number of players

grows. One can remedy this situation by creating shortcuts in the interaction graph. The star-like and bipartite structures capture interesting real-world settings with many short interaction paths. We showed that these topologies lead to a surprisingly high asymptotic probability of pure Nash, generally higher than for the standard matrix game scenario. Given the high stability and predictability of (the existence of) pure Nash equilibria, this work suggests ways of designing multi-agent systems that are quite likely to have such equilibria.

Acknowledgments

This work was supported by IISI, Cornell University, AFOSR grant FA-9550-04-1-0151. The authors would like to thank Dexter Kozen, Bart Selman, Eva Tardos, Muthu Venkatasubramaniam, and Tom Wexler for insightful discussions and suggestions.

References

- Albert, R., and Barabasi, A.-L. 2002. Statistical mechanics of complex networks. *Rev. of Modern Physics* 74:47.
- Daskalakis, C., and Papadimitriou, C. H. 2006. Computing pure nash equilibria in graphical games via Markov random fields. In *EC'06*, 91–99.
- Een, N., and Srensson, N. 2005. Minisat - a SAT solver with conflict-clause minimization. In *SAT'05*.
- Galeotti, A.; Goyal, S.; Jackson, M. O.; Vega-Redondo, F.; and Yariv, L. 2006. Network games. Working Paper.
- Gottlob, G.; Greco, G.; and Scarcello, F. 2003. Pure nash equilibria: hard and easy games. In *TARK'03*, 215–230.
- Kearns, M. J.; Littman, M. L.; and Singh, S. P. 2001. Graphical models for game theory. In *UAI '01*, 253–260.
- Nash, J. 1951. Noncooperative games. *Annals of Mathematics* 54:289–295.
- Nudelman, E.; Wortman, J.; Leyton-Brown, K.; and Shoham, Y. 2004. Run the gamut: A comprehensive approach to evaluating game-theoretic algorithms. In *AA-MAS'04*.
- Powers, I. Y. 1986. *Limiting Distributions of the Number of Pure Strategy Nash Equilibria in n-Person Games*. Ph.D. Dissertation, Yale University.
- Rinott, Y., and Scarsini, M. 2000. On the number of pure strategy nash equilibria in random games. *Games and Economic Behavior* 33(2):274–293.
- Vega-Redondo, F. 2003. *Economics and the Theory of Games*. Cambridge University Press.
- Watts, D. J., and Strogatz, S. H. 1998. Collective dynamics of 'small-world' networks. *Nature* 393(6684):440–442.