

Strategic Network Formation with Structural Holes

Jon Kleinberg
Dept. of Computer Science
Cornell University, Ithaca NY
kleinber@cs.cornell.edu

Siddharth Suri
Dept. of Computer Science
Cornell University, Ithaca NY
ssuri@cs.cornell.edu

Éva Tardos
Dept. of Computer Science
Cornell University, Ithaca NY
eva@cs.cornell.edu

Tom Wexler
Dept. of Mathematics and
Computer Science
Denison University, Granville
OH
wexlert@denison.edu

ABSTRACT

A fundamental principle in social network research is that individuals can benefit from serving as intermediaries between others who are not directly connected. Through such intermediation, they potentially can broker the flow of information and synthesize ideas arising in different parts of the network. These principles form the underpinning for the theory of *structural holes*, which studies the ways in which individuals, particularly in organizational settings, fill the “holes” between people or groups that are not otherwise interacting.

We apply a game-theoretic approach to this notion, studying the structures that evolve when individuals in a social network have incentives to form links that bridge otherwise disconnected parties. We model payoffs as a trade-off between the benefits of connecting non-neighboring nodes, and the cost, in effort, to maintain links — including settings where the costs are non-uniform to reflect the increased difficulty in spanning different parts of a hierarchical organization.

We find, both through theoretical results and computational experiments, that the equilibrium networks in this model have rich combinatorial structure, and capture qualitative observations arising in the study of structural holes. In particular, even in completely symmetric settings, individuals will differentiate themselves in equilibrium, occupying different social strata and receiving correspondingly different payoffs.

This work has been supported in part by NSF grants CCF-0729006, CCF-0325453, CNS-0403340, BCS-0537606, and IIS-0705774, ONR grant N00014-98-1-0589, a Google Research Grant, a Yahoo! Research Alliance Grant, the Institute for the Social Sciences at Cornell, and the John D. and Catherine T. MacArthur Foundation.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

Copyright 200X ACM X-XXXXX-XX-X/XX/XX ...\$5.00.

Categories and Subject Descriptors

J.4 [Social and Behavioral Sciences]: Economics

General Terms

Economics, Theory

Keywords

algorithmic game theory, social networks, structural holes

1. INTRODUCTION

An important line of work in sociology argues that individuals benefit when they serve as intermediaries or “bridges” between others who are not directly connected [2, 8]. In particular, Burt’s theory of *structural holes* views organizations not just in terms of the tight connections within their social structures, but also in terms of the “holes” where connections have failed to form. Such holes provide opportunities for socially entrepreneurial individuals to play bridging roles in the organization, linking to disparate people or groups who are otherwise not interacting with each other. The argument is that individuals whose connections bridge such holes tend to do particularly well professionally, and this argument has been supported by empirical studies of managers in large corporations, correlating an individual’s success with this type of bridging activity [2].

At its heart, then, the theory of structural holes is about the advantages that accrue to people who occupy such bridging positions, like the node A in the social network of Figure 1. A number of basic reasons have been advanced for the success of nodes like A in this figure. By serving as an intermediary, A gains power by brokering the flow of information between different parts of the network. Moreover, A ’s network position is potentially an amplifier for innovation, since A is well-positioned to synthesize ideas arising from different groups [3].

Part of the premise of structural holes is that people appreciate when such advantages are possible, and implicitly seek out opportunities to realize them. As a result, there is a natural strategic aspect to the theory, in which the network is shaped by link formation decisions made by individuals who are trying to create bridging structures that benefit themselves [5, 7]. However, very little modeling work has been devoted to these strategic questions, despite their importance to the whole framework: while much work has focused on the opportunities that arise from bridging holes, there is relatively little understanding of the kinds of networks that arise




Figure 1: According to the theory of structural holes, node A 's bridging role in the social network of an organization gives it structural advantages relative to node D that sits in the middle of a tightly-knit group.

when the inhabitants of a social network are all trying to act on these opportunities.

We address this question here, building a general model of the payoffs that arise from bridging structural holes in a social network, and then studying the game-theoretic consequences of this model for social network structure. We find that the networks resulting from strategic behavior have a rich combinatorial structure, and correspond in interesting ways to some of the more qualitative observations that have arisen in work on structural holes.

The Model. We assume a setting in which a set V of n distinct agents create links to one another to form a social network. (Following one of the primary motivations for the theory of structural holes, we can imagine these agents to be a set of managers within a large company, or people working within a large organization.) We model this as a game in the following way. Each agent's strategy is a choice of other nodes to whom it will build links. Links are undirected, and hence if u builds a link to v , this link becomes usable in both directions.

The payoffs to an agent arise as a difference of costs and benefits. First, there is a specified cost matrix $\{c_{uv} : u, v \in V\}$: if an agent u creates a link to an agent v , then u incurs a given cost of c_{uv} , reflecting the effort required to maintain the link. The benefits to an agent u are based both on the nodes to whom it is directly linked, as well as on the ways in which its links allow it to bridge structural holes between others. We will define the benefits from direct linkage very simply by saying that a node derives a fixed benefit of α_0 from each of its neighbors in the network.

Bridging benefits are more subtle to define; they arise from lying on paths between pairs of nodes v and w who are not directly connected. Qualitatively, it is clear that these benefits should decrease with the length of the v - w path and also with the number of other nodes who are also playing a bridging role with respect to v and w . Here we formalize bridging benefits as follows. There is empirical evidence that this type of benefit is realized primarily when a node u forms the middle of a two-step path between two non-neighboring nodes v and w ; there appears to be much less

measurable benefit to u if it is an internal node on a path between two nodes at graph distance greater than two [2, 4], as such interactions are in general too "long-range" to confer significant benefit on their intermediaries. Thus, for example, in Figure 1, the argument is that node A benefits from sitting at the middle of many length-two paths; but no node benefits significantly from lying on a B - F path, since all such paths have length at least three. Moreover, the bridging benefit to u is greatest when it forms the unique length-2 path between v and w , and it decreases in the number of other length-2 paths between v and w .

Accordingly, we assume there is a decreasing function β such that a node u obtains a benefit of $\beta(r)$ whenever it lies on a length-2 path between a pair of non-neighboring nodes that, in total, have r length-2 paths between them. (For example, in Figure 1, node A receives a benefit of $7\beta(1)$ from the seven pairs of nodes for which it lies on the unique length-2 path, and it receives an additional benefit of $\beta(4)$ from lying on one of the four length-2 paths between nodes E and H .) We will refer to $\beta(r)$ as the *intermediary benefit* or *middleman benefit* received by each of the nodes that serves as one of the r intermediaries.

In summary, then, the game is as follows.

- Each node u constructs links to a set of nodes $L(u)$.
- Collectively, all constructed links form an undirected graph G . We let $N(u)$ denote the set of neighbors of u in G ; note that $N(u)$ may be a proper superset of $L(u)$, since it can also include nodes that have constructed links to u . (Recall that a constructed link can be used in both directions.) For nodes v and w , we let $r_{vw} = 0$ if they are directly connected, and otherwise we define r_{vw} to be the number of length-2 paths between them.
- The payoff to u is

$$\alpha_0 |N(u)| + \sum_{v,w \in N(u)} \beta(r_{vw}) - \sum_{v \in L(u)} c_{uv}.$$

Some of our results will apply to the model in its full generality, while for others we will make assumptions about the forms of c and β . Primarily, we will focus on cost matrices arising either from

- the uniform metric, in which all costs are equal to 1, or
- *hierarchical metrics* — again following the organizational motivation — in which costs arise from shortest paths in a rooted tree that represents an organizational structure.

For intermediary benefits β , there is a very natural special case that can be defined as follows. If nodes v and w connect directly, they each obtain a benefit of α_0 , for a total shared benefit of $2\alpha_0$. One natural assumption is that if v and w are not connected, but have r_{vw} length-2 paths between them, then this total benefit of $2\alpha_0$ is shared equally between the r_{vw} middle nodes of these paths: each gets $\beta(r_{vw}) = 2\alpha_0/r_{vw}$. (If v and w are at graph distance greater than three, we assume that the benefit of $2\alpha_0$ is lost completely.) We will call this special form for β the *harmonic* special case. Some of our results will concern this special case, and it will also guide certain milder restrictions that one can put on the form of β : for certain results, we will assume that β is an arbitrary decreasing function subject only to conditions that $r\beta(r)$ is either upper-bounded or lower-bounded by a constant as r grows. (Note that $r\beta(r)$ is the total benefit to all r intermediaries between a pair of nodes, and it is equal to a constant independent of r for the harmonic function.)

Our Results. We focus on understanding properties of the networks that arise as possible equilibria in this model. To begin with, agents have large combinatorial strategy sets, since they can choose to link to any subset of other nodes. Despite this, we first show that there is an efficient algorithm for an agent to determine its best response in a given graph, via a reduction to network flow computations. In addition to the establishing a certain tractability of the model, it also makes it feasible to run computational experiments in some of the model’s more complicated variants, as we discuss below.

For the uniform metric, we prove the existence of equilibrium networks. It is possible for best-response dynamics to cycle indefinitely in this game, but we show that best-response dynamics will in fact terminate starting from an empty graph, and we characterize the equilibrium networks that are reached this way — complete multipartite graphs in which each node picks a particular “stratum” and links to all nodes in higher strata. We also find, via computational experiments, that there is a wide range of possible equilibria, some with fairly complex structures; we prove that all equilibria are dense graphs, in the sense that they have a quadratic number of edges (for constants depending on the parameters of the payoff functions).

For hierarchical metrics, when nodes reside in the shortest-path metric on the leaves of a complete b -ary tree (e.g. an organizational hierarchy), we find equilibrium networks through computational experiments that are even more complex, with a small number of “high-status” nodes connected to almost all others, while most nodes connect to a small number of neighbors that are nearby in the hierarchy. We prove that any equilibrium network must in fact exhibit the properties we observe in the computational experiments: there exist “high-status” nodes of linear degree, while the average degree is asymptotically much smaller.

The structure of these equilibria, in both kinds of metrics, illustrate at a concrete level a number of interesting qualitative observations that have been made in the literature on structural holes. In particular, although they are both based on symmetric structures — metrics which look the same from every node’s perspective — nodes nevertheless differentiate in equilibrium, with the symmetry-breaking putting different nodes at different social levels in which they obtain different payoffs. In the uniform metric, nodes stratify into layers, while in the hierarchical metric, nodes appear to approximately settle on a “level” of the hierarchy that they will dominate, with different nodes settling on different levels. And this provides, in turn, a more qualitative explanation for how non-trivial bridging behavior can still arise even when all nodes have incentive to try engaging in it: since the advantages of bridging decreases rapidly with the number of people doing it, there are “first-mover” advantages that translate into broken symmetry in equilibrium.

In this way, our focus here is on structural and dynamic questions surrounding the equilibrium networks in this model. We do not address questions related to the price of anarchy or the price of stability here; while it is clearly interesting to understand how equilibrium networks perform relative to a social optimum, defining the social optimum carefully would bring in an essentially orthogonal set of formalisms that seek to capture the quality of an organization’s global performance as a function of its network structure (see e.g. [13]). We would argue that simply adding the payoffs of all individuals — as usually serves for the social optimum in price-of-anarchy analyses — would not necessarily be a good model of global performance in our case; there is no evidence that the ideal outcome from the overall organization’s perspective is one in which the total of individuals’ bridging payoffs are as high as possible, and indeed it is an interesting open question to understand how

this relates to organizational performance. As a result, we leave the modeling of global performance, and the consequent price-of-anarchy analysis that would be enabled by it, as a direction for future work.

Related Work. Our work is related to a growing body of research on strategic network formation [9]. There is a basic methodological choice required when modeling network construction as a game, arising from the fact that Nash equilibrium studies the possibility of strategy revision made by a single agent, while edges in an undirected graph have not one but two endpoints. Two approaches to this difficulty have been pursued in the literature: modifying the equilibrium concept to *pairwise stability* or one of its variants (e.g. [6, 10]), in which agents must jointly agree to edge formation, or assuming strategies in which agents unilaterally construct edges that can then be used in either direction (e.g. [1]), as we do here.

As noted above, there has been very little work on strategic network formation that explicitly incorporates the notion of structural holes. The two other papers we are aware of to pursue this are Goyal and Vega-Redondo [7] and recent independent work of Buskens and van de Rijt [5]. Each is motivated by similar concerns to ours, though they frame the problem differently. Goyal and Vega-Redondo consider a model in which a node u potentially benefits from serving as an intermediary between nodes v and w even when it resides on an arbitrarily long v - w path, and with this formulation they obtain the star network as essentially the unique most robust equilibrium over a broad range of model parameters. This makes it difficult to study more complex kinds of networks and bridging structures. Buskens and van de Rijt focus, as we do, on benefits that arise from completing two-step paths, although they consider a formulation based on Burt’s notion of *network constraint*, which does not capture bridging benefits directly but rather the “cost” to nodes of wasting effort on redundant length-2 paths. They also use a more complex notion of equilibrium that they call *unilateral stability*. Their paper restricts attention only to uniform metrics, and interestingly, in this context, they also find complete multipartite graphs as some of the key examples of stable networks.

2. PRELIMINARIES

We begin by considering the behavior of best-response dynamics in our game. As noted earlier, the strategy set of each player is exponentially large: each node can choose to construct edges to any subset of the other nodes. Hence, it is not *a priori* clear how to run best-response dynamics efficiently, but we show here how to compute it in polynomial-time via a reduction to a network flow problem. This efficient algorithm forms the basis for our computational experiments to determine Nash equilibria that accompany the theoretical results in the next two sections.

The convergence behavior of best-response dynamics in this game is a delicate issue. We show in this section that best-response dynamics can cycle, so it can not be used in a straightforward way to prove the existence of an equilibrium. In the next section, however, we show that with appropriate scheduling and initial conditions, best-response dynamics converges to an equilibrium in the uniform metric, and we characterize the structure of this equilibrium. More generally, we find through computational experiments that best-response dynamics often converges for choices of initial conditions and scheduling rules, and this forms the basis for our computed equilibria.

2.1 Efficient Best-Response Computation

To implement best-response dynamics in our game, consider a player v , and consider the graph $G = (V, E)$ of edges bought by

other players. As before, $N(v)$ denotes the set of neighbors of v in G , and we let A be the subset of nodes that v is not connected to, that is $A = V \setminus (N(v) \cup \{v\})$. Clearly player v will not want to buy edges to nodes outside of this set A . We show that v 's best response can be found in polynomial time by reducing it the problem of finding a maximum-weight ideal in a partially ordered set.

DEFINITION 2.1. *Given a partially ordered set \mathcal{O} , an ideal $I \subset \mathcal{O}$ is a subset of elements such that for any nodes $v \preceq w$ in the partially ordered set, $w \in I$ implies we must also have $v \in I$. The maximum-weight ideal problem is the following: given weights w_i on the elements of a partially ordered set \mathcal{O} , select an ideal I of maximum total weight $\sum_{i \in I} w_i$.*

The maximum-weight ideal problem can be solved in polynomial time via a reduction to a minimum-cut problem [11].

THEOREM 2.2. *The best response of a player v can be computed in polynomial time.*

Proof. We reduce the problem to the maximum-weight ideal problem. For all pairs of nodes (u, v) , let $r_{u,v}$ denote the number of nodes in G that lie on length-2 paths from u to v . (This is the number of pairs that receive benefits as intermediaries.)

Consider an edge (v, w) for some node $w \in A$. There is a direct benefit of α_0 associated with this edge, and a cost of c_{vw} . In addition to this cost and direct benefit, buying edges also results in potential benefit as an intermediary. Some of this intermediary benefit can also be directly associated with buying a single edge (v, w) , assuming the edge (u, v) has been bought by u .

The cost and benefit discussed thus far is directly associated with buying a single edge, to a node $w \in A$. We will call this the *weight* of the node $w \in A$:

$$weight_w = \alpha_0 - c_{vw} + \sum_{u \notin A, (u,v) \in E, (u,w) \notin A} \beta(r_{u,w} + 1).$$

We also have intermediary benefits associated with pairs of nodes in A not connected by an edge. To represent these we add additional elements to the partially ordered set.

$$\mathcal{O}(v) = A \cup \{(u, w) \notin E, u, w \in A\}.$$

The weight of such a pair of nodes is $weight_{u,w} = \beta(r_{u,w} + 1)$. In the partially ordered set $\mathcal{O}(v)$ we will have the relation that for every pair of nodes $(u, w) \in \mathcal{O}(v)$, both nodes u and w must be included in an ideal that contains (u, w) ; that is, $u, w \preceq (u, w)$ for every pair $(u, w) \in \mathcal{O}(v)$.

Note that the weight of the elements of \mathcal{O} derived from node-pairs is always positive, while elements of \mathcal{O} derived from single nodes may have negative weight. As a result, a maximum-weight ideal will contain a subset $A' \subset A$ and all the pairs in A' not connected by an edge, and the weight of this ideal is exactly v 's benefit of buying edges to every node in A' . Finding an ideal of maximum weight therefore determines v 's best response. ■

2.2 Cycling of Best-Response Dynamics

We now show, via an example that is sketched in Figure 2, that best-response dynamics in our game can cycle. The example in this figure uses the uniform metric, with $\alpha_0 = 0.9$ and the harmonic form of the intermediary benefits β . Figure 2a shows the initial configuration. If player 1 moves, it must compute which of nodes 2, 4 and 5 to connect to (3 already connects to 1). Player 1's current strategy (connecting to 4 and 5) yields a benefit of 1.3. Player 1's




Figure 2: An example of best response dynamics cycling. See Section 2.2 for more details.

unique best response is to buy a link to only 4, which yields a payoff of 1.4. This results in Figure 2b.

When player 2 moves, it must decide which of nodes 1 and 5 to connect to (3 and 4 both already connect to 2). Connecting to neither yields a benefit of 2.4. Connecting to either one but not both yields 2.3. Connecting to both yields 2.8, and thus this is player 2's best response, yielding the graph shown in Figure 2c. Observe that if we permute the player labels according to $(1, 2, 3, 4, 5) \rightarrow (4, 5, 1, 2, 3)$, we have recreated our initial configuration. If we now schedule the best response of players repeatedly in order, after 10 moves we will be back to the initial configuration: in other words, best-response dynamics will have cycled.

It is interesting to note that the cycling of best-response dynamics shows that we do not have an ordinal or an exact potential game [12], a sub-class of games with natural equilibrium existence proofs that have been heavily studied in recent work on algorithmic game theory.

3. UNIFORM METRIC

In this section we prove that for the case of the uniform metric our game always has a pure Nash equilibrium. We will define an infinite family of multipartite graphs $G_{n,k}$ and show that for any choice of parameters, one of the $G_{n,k}$ is an equilibrium. In addition, we will also show that best-response dynamics starting from an empty graph actually converges to such an equilibrium.

3.1 Existence of Equilibria

A key tool in proving the existence of equilibria is the following lemma. It concerns a graph G with an independent set I , such that all nodes $v \notin I$ are connected to every node in I . The first part of the lemma describes how much value a node would get from connecting to all the nodes in I . The second part describes when a node seeking to maximize its payoff would do so.

LEMMA 3.1. *The change in benefit for a node $v \notin I$ from deleting all the edges connecting v to I is $B(n, k) = k(\alpha_0 - 1) + \binom{k}{2}\beta(n - k)$, where $k = |I|$. Node v will keep all the edges connecting v to I if $B(n, k) \geq 0$, and will drop all the edges connecting v to I if $B(n, k) \leq 0$.*

Proof. The cost of buying an edge to every node in I is k , the direct benefit from doing so is $k\alpha_0$. Moreover, the middleman benefit from buying edges to a single pair of nodes in I is $\beta(n-k)$. Since all nodes in $V \setminus I$ are connected to all nodes in I there is no other middlemen benefit associated with these edges. Thus the total benefit of attaching to all nodes in I is, $k(\alpha_0 - 1) + \binom{k}{2}\beta(n-k)$ as claimed.

To prove the second part of the statement, we need to consider the benefit v would obtain if it kept edges to a subset $A \subseteq I$ of $k' \leq k$ nodes. Again, since all edges between I and $V \setminus I$ are included in the graph, all benefit for attaching to any nodes in A comes from direct benefit, and middlemen benefits for pairs of nodes in A . Connecting only to the k' nodes in A results in $k'(\alpha_0 - 1)$ from the direct benefit minus costs, and $\binom{k'}{2}\beta(n-k)$ from middlemen benefits. For fixed k and n this benefit as a function of k' is a convex quadratic function of k' and hence in the range $0 \leq k' \leq k$ its maximized either at $k' = 0$ or at $k' = k$. This implies that v either will want to keep all edges to I or none, which implies the statement. ■

Now we are ready to prove that an equilibrium graph exists. The equilibrium graphs will be complete, multipartite graphs in which there are $q = \lfloor \frac{n}{k} \rfloor$ parts each consisting of a k -vertex independent set, and there is one part consisting of an ℓ -vertex independent set, where $\ell = n \bmod k$. Nodes from higher indexed sets buy all the edges to nodes in lower indexed sets.

DEFINITION 3.2. Let $G_{n,k}$ denote a directed, multipartite graph of the following form. Let $V = V_1 \cup \dots \cup V_q \cup V_{q+1}$ where $V_i \cap V_j = \emptyset$ for $i \neq j$, $|V_1| = \dots = |V_q| = k$, and $|V_{q+1}| = n \bmod k$. For each $u \in V_i$, let $L(u) = \cup_{j=1}^{i-1} V_j$.

THEOREM 3.3. For an appropriate choice of k the graph $G_{n,k}$ is at equilibrium.

Proof. Choose k such that $B(n, k) \geq 0$ and $B(n, k-1) \leq 0$ (we prove that such a k exists below). We claim that for this choice of k , $G_{n,k}$ is at equilibrium. Since $G_{n,k}$ is complete and multipartite, there is no middleman benefit to be gained from attaching to a node in V_i and a node in V_j where $i \neq j$. There is, however, middleman benefit to be gained from attaching to nodes in the same independent set. Thus for any node $v \in V_i$, the costs and benefits associated with buying edges to different sets V_j are independent of each other, so node v needs to make independent decisions about which of the edges to each set V_j it wants to have. We apply Lemma 3.1 for the independent sets V_j for $j \neq i$. As $B(n, k) \geq 0$ node v will keep all its connections to every node in V_j for all $j \leq i$. Then we apply Lemma 3.1 for the independent set $V_i - \{v\}$. As $B(n, k-1) \leq 0$ the node v does not want to buy connections to the other nodes in V_i .

To finish the proof we have to argue that there is a value k with $B(n, k) \geq 0$ and $B(n, k-1) \leq 0$. Consider $B(n, 1) = (\alpha_0 - 1)$. If $B(n, 1) > 0$ than the direct benefit associated with an edge is at least as large as the cost, and hence the complete graph $G_{n,1}$ is an equilibrium. If $B(n, n-1) \leq 0$ then the empty graph $G_{n,n}$ is an equilibrium by Lemma 3.1. If $B(n, 1) < 0$ and $B(n, n-1) > 0$ then there must be an integer k with $B(n, k) \geq 0$ and $B(n, k-1) \leq 0$ and hence $G_{n,k}$ is an equilibrium graph. ■

3.2 Convergence to Equilibrium

Having now seen that equilibria exist, we ask whether they are reachable by a natural sequence of best-response updates. The example in Section 2, showing that arbitrary sequences of best-response updates can cycle, indicates some of the difficulty in addressing this question. However, as we show here, the equilibria

$G_{n,k}$ are in fact reachable via a very natural sequence of best-response updates: starting from a graph with no edges, we number the nodes arbitrarily $1, 2, 3, \dots, n$, and have them perform best-response updates in this order, cycling through the order indefinitely until an equilibrium is reached. (We call this the *round-robin* schedule, since nodes simply take turns performing updates according to a repeating fixed order.)

In particular, we show that best-response dynamics, executed in round-robin order starting from empty graph, converges to $G_{n,k}$ where k is the largest integer such that $B(n, k) \geq 0$ and $B(n, k-1) \leq 0$. The proof requires that we keep track of a number of inductive invariants as the updates proceed, and to give the main ideas in a somewhat simpler setting, we first prove a special case of the result, in Theorem 3.6, for the case in which $k \geq \lceil \frac{n}{2} \rceil$. Following this, we discuss how to extend the proof to any value of k ; the full details are given as Theorem 6.5 in the Appendix.

The proof of Theorem 3.6 relies on two lemmas. The first lemma describes the result of the first round of best response dynamics starting from the empty graph. Intuitively, as each node makes its move it sees an independent set on the nodes who have yet to move. Thus they buy all of the edges to the nodes yet to move to become a middleman between all pairs of the independent set. The first $n-k$ nodes do this resulting in a clique of size $n-k$ where all of the nodes in it are also attached to an independent set of size k .

LEMMA 3.4. If k is the largest integer such that $B(n, k) \geq 0$ and $B(n, k-1) \leq 0$, then at the end of the first round of best response dynamics, executed in round robin order, node i will have bought edges to nodes $i+1, \dots, n$ where $i \in \{n-k\}$.

Proof. The proof is by downward induction on k . If $k = n$ then by Lemma 3.1, node 1 will not buy any edges. Next, let's assume the claim holds for $k = k' + 1$ and show it holds for $k = k'$. By the induction hypothesis, $B(n, n-1), \dots, B(n, k'+1) \geq 0$, and thus by Lemma 3.1 nodes 1 through $n-k'-1$ buy edges to all higher-indexed nodes. Thus when it is the turn of node $n-k'$, the only nodes from which node $n-k'$ could earn middleman benefit from will be nodes $n-k'+1, n-k'+2, \dots, n$, which form an independent set of size k' . Since $B(n, k') \geq 0$, by Lemma 3.1, node $n-k'$ will buy edges to all of the nodes in the independent set. After this move by node $n-k'$, the largest independent set in the graph has size $k'-1$. Since $B(n, k'-1) \leq 0$ and by Lemma 3.1, all subsequent nodes to move will not buy any edges. ■

The next lemma bounds how much a node could earn if it attaches to both a clique and an independent set which are disjoint. Doing so would earn middleman benefit between pairs of the independent set, and between nodes of the clique and the independent set. Intuitively, since a node gets a quadratic increase for each new node it becomes a middleman for, while only paying a linear cost, a node will either attach to all of the nodes in clique and the independent set, or none of them. Here we will use the notation $N(S) = \cup_{v \in S} N(v)$.

LEMMA 3.5. Let $G = (V, E)$ be a graph with an independent set I of size i , and a clique C of size c , where I and C are disjoint and $i + c \leq k$, such that $|N(C) \cap N(I)| \geq n - k$.

1. If $B(n, k) \geq 0$ then the most a node would gain from attaching to $I \cup C$ would be $B(n, k)$. Furthermore, the maximum gain would result from attaching to all nodes in $I \cup C$.
2. If $B(n, k) \leq 0$ then the most a node would gain from attaching to $I \cup C$ would be 0 which would result from not attaching to any nodes in $I \cup C$.

Proof. Say a node v attached to $i' \leq i$ of the nodes in the independent set and $c' \leq c$ of the nodes of the clique. This would cost $i' + c'$ and earn a direct benefit of $\alpha_0(i' + c')$. Observe that v would earn $\binom{i'}{2}$ in middleman benefit from attaching to the nodes in the independent set, but this would be shared with at least $n - k$ other nodes, since $|N(I)| \geq n - k$. Similarly, v would earn $i'c'$ in middleman benefit for attaching the nodes in the clique with nodes in the independent set. This would also be shared with $n - k$ other nodes, since $|N(C) \cap N(I)| \geq n - k$. Thus the marginal payoff of v is:

$$\text{mp}_v(i', c') \leq (\alpha_0 - 1)(i' + c') + \left[\binom{i'}{2} + i'c' \right] \beta(n - k)$$

Observe that this is a convex quadratic function of i' . Thus $\text{mp}_v(i', c')$ is maximized either when $i' = 0$ or when $i' = i$. Next we consider both of these cases. The first case is when $i' = 0$, then $\text{mp}_v(i', c') = 0$. Furthermore, since there is no middleman benefit to be gained from attaching to nodes in C , $c' = 0$. Thus v will not attach to any nodes in $I \cup C$. The second case is when $i' = i$, and since $c' \leq c$ and $i + c \leq k$ we get:

$$\begin{aligned} \text{mp}_v(i', c') &\leq (\alpha_0 - 1)(k) + \binom{k}{2} \beta(n - k) \\ &= B(n, k) \end{aligned}$$

Thus if $B(n, k) \leq 0$, v will not attach to any nodes in $I \cup C$. If $B(n, k) \geq 0$, then v would achieve the maximum gain by attaching to all nodes in $I \cup C$, resulting in a gain of at most $B(n, k)$. ■

Next we consider the special case when two rounds of best-response dynamics converges to $G_{n,k}$. Lemma 3.4 shows that after the first round, the graph will look like a clique of size $n - k$ with each node also attached to an independent set of size k . Then we will use Lemma 3.5 to show that during round two, the nodes in the clique drop edges to each other to form another independent set, yielding a complete bipartite graph overall.

THEOREM 3.6. *If k is the largest integer such that $B(n, k) \geq 0$, $B(n, k - 1) \leq 0$ and $k \geq \lceil \frac{n}{2} \rceil$, then best response dynamics, executed in round robin order, will converge to $G_{n,k}$ in 2 rounds.*

Proof. Lemma 3.4 shows that after the first round of best response dynamics nodes every $j \in \{1, 2, \dots, n - k\}$ will have bought edges to all nodes in $\{j + 1, j + 2, \dots, n\}$. Furthermore, only those edges will have been bought.

To discuss the results of round 2, we first introduce some notation. Let $V_2 = \{1, 2, \dots, n - k\}$ and let $V_1 = \{n - k + 1, \dots, n\}$. Since $|V_1| = k$ and $B(n, k) \geq 0$, Lemma 3.1 implies that every node in V_2 will keep all of their edges to the nodes in V_1 . This implies there is no middleman benefit to be gained from buying an edge to a node in V_1 and a node in V_2 . Thus we can compute benefit of attaching to nodes in V_2 independently of attaching to nodes in V_1 .

During round 2, node 1 will drop all of its edges to V_2 since V_2 forms a clique and node 1 earns a benefit of $(\alpha_0 - 1)|V_2| < 0$ for those edges. Inductively assume that the first i nodes of V_2 drop their edges to the other nodes in V_2 . This means that at the time of node $i + 1$'s move, there is an independent set of size i , and a clique of size c where $c = |V_2| - i - 1$. Furthermore, since $k \geq \lceil \frac{n}{2} \rceil$, and $|V_2| = n - k$, we have that $i + c \leq k - 1$. Thus Lemma 3.5 combined with $B(n, k - 1) \leq 0$ implies that node $i + 1$ will drop all edges to the nodes in V_2 . After the nodes in V_2 do this, the graph is exactly $G_{n,k}$ which by Theorem 3.3 is stable. ■




Figure 3: The density of the equilibria as a function of α_0 . The density is computed as $|E|/\binom{n}{2}$. $G_{n,k}$ are represented by crosses. The alternate equilibria, found by simulations, are represented by circles.

In a relatively straightforward way, one can generalize the above theorem to a schedule where there exists a time t such that before t all nodes perform at least one update, and after t all nodes perform at least one update.

For values $\lfloor \frac{n}{2} \rfloor \geq k > 0$ the process of the best response dynamic of Theorem 3.6 is more complex. The construction of $G_{n,k}$ proceeds in rounds, where round i results in the formation of the i th level in $G_{n,k}$, V_i . Lemma 3.4 is again used to show that round 1 results in a clique of $n - k$ nodes attached to an independent set V_1 of the last k nodes. Round i will result in the formation of the independent set V_i , but surprisingly the set do not come from consecutive k nodes in the round-robin order. Despite this fact, we use Lemma 3.5 to show that in round i a subset of k nodes forms a new level V_i . Furthermore, we show that once a node forms part of a new level, it is “frozen” in the sense that it neither buys additional edges nor drops edges from then on.

3.3 Properties of Equilibria

In the last two sections we considered the multipartite graphs $G_{n,k}$ extensively: We showed that for any α_0 and any function β one of these graphs always forms an equilibrium, and we showed that best-response dynamics starting from the empty graph, and with round-robin scheduling of players, always converges to one of these multipartite graphs.

A natural question is now to consider other equilibrium graphs — asking whether equilibria other than $G_{n,k}$ exist, to what extent $G_{n,k}$ has properties representative of these other equilibria, and whether it is possible to prove non-trivial properties that hold for all equilibria. In this direction, we first discuss the results of computational experiments that uncover a range of other equilibrium graph structures. We then prove that all equilibria in the uniform metric are dense graphs.

To find other equilibrium graphs, we run best response starting from random initial configurations. We have seen that best-response dynamics can cycle, but in many runs it terminates with new equilibria. One of the key properties we observe in our computed equilibria is that they are all dense, with a constant fraction of all node pairs connected by edges (this property is easily checked for the multipartite graphs $G_{n,k}$ for all but very large k). Indeed, as Figure 3 shows, the other equilibria we find are even more dense

than $G_{n,k}$. In our experiments we use the harmonic intermediary benefit function $\beta(r) = \alpha_0/r$, so we have only a single parameter α_0 . In Figure 3, we see that the other equilibria found are denser than the multipartite equilibrium corresponding to the given value of α_0 .

We now prove that what we are observing computationally in fact holds for all equilibria: all equilibrium networks are dense, in that they all have $\Omega(n^2)$ edges, assuming that $r\beta(r) \geq \gamma$ for a constant $\gamma > 0$. Note that $r\beta(r)$ is the total benefit when k intermediaries connect a pair of nodes, and so as noted in the Introduction, we are making the mild assumption — generalizing the harmonic function — that this benefit is non-vanishing (i.e. it remains lower-bounded by some constant) as r increases.

To prove the theorem we will need two technical lemmas.

LEMMA 3.7. *For $a, b, r > 0$, the maximum value of $\sum_{i=1}^r x_i^2$, over all choices of nonnegative real numbers (x_1, \dots, x_r) with $\sum_i x_i \leq a$ and $x_i \leq b$ for all i , occurs on vectors in which all but one coordinate is equal either to 0 or b .*

Proof. If there an optimal vector (x_1, \dots, x_r) in which $\varepsilon \leq x_i \leq x_j \leq b - \varepsilon$ for some $\varepsilon > 0$, then the objective function would be increased by replacing them with $x_i - \varepsilon$ and $x_j + \varepsilon$. ■

LEMMA 3.8. *For $a, r > 0$, the minimum value of $\sum_{i=1}^r \frac{1}{1+x_i}$, over all choices of nonnegative real numbers (x_1, \dots, x_r) with $\sum_i x_i \leq a$, occurs when all x_i are equal, in which case it is $\frac{r^2}{a+r}$.*

Proof. First, if $\sum_i x_i$ is not equal to a , then uniformly scaling up all x_i will only reduce the objective function. Next, suppose that there are i and j with $x_i < x_j$, and let $\varepsilon = \frac{1}{3}(x_j - x_i)$. Then we could decrease the objective function by replacing x_i and x_j with $x_i + \varepsilon$ and $x_j - \varepsilon$, since

$$\begin{aligned} & \frac{1}{1+x_i+\varepsilon} + \frac{1}{1+x_j-\varepsilon} \\ &= \frac{2+x_i+x_j}{(1+x_i)(1+x_j)+\varepsilon(x_j-x_i)-\varepsilon^2} \\ &< \frac{2+x_i+x_j}{(1+x_i)(1+x_j)} = \frac{1}{1+x_i} + \frac{1}{1+x_j}. \end{aligned}$$

Thus the optimum has $\sum_i x_i = a$ and all x_i equal, and so it is $r \cdot \frac{1}{1+(a/r)} = \frac{r^2}{a+r}$. ■

THEOREM 3.9. *For every choice of $\alpha_0 > 0$ and function β , subject to $\beta(r) \geq \frac{\gamma}{r}$ for a constant $\gamma > 0$, every equilibrium graph has at least $\Omega(n^2)$ edges.*

Proof. We will prove that there exists $\delta > 0$ and $n_0 > 1$ such that any equilibrium graph on at least n_0 nodes has at least δn^2 edges. We choose $\delta < \frac{1}{162} \min(1, \gamma)$, and $n_0 \geq \delta^{-1}$. Now, suppose there were an equilibrium graph G on at least n_0 nodes with strictly fewer than δn^2 edges; we will show this leads to a contradiction.

Let d_v denote the degree of node v , and for each pair of nodes v, w , let $r_{v,w}$ be the number of middlemen (i.e. common neighbors) of v and w . We have

$$\sum_{v,w} r_{v,w} \leq \sum_v \binom{d_v}{2} \leq \sum_v d_v^2.$$

Since the degree sequence is a vector of numbers that satisfies $\sum_v d_v \leq 2\delta n^2$ and also $d_v \leq n$ for all v , Lemma 3.7 implies that

$$\sum_{v,w} r_{v,w} \leq (2\delta n + 1)n^2 \leq 3\delta n^3.$$




Figure 4: A 32-node equilibrium network in a hierarchical metric.

Now, consider a node u of minimum degree; let Y be the set of all nodes that are not neighbors of u , and let Γ be the collection of two-element subsets of Y consisting of nodes not connected by edges. Since the number of edges is less than δn^2 , we have $|Y| > \frac{1}{2}n$. If u creates edges to all nodes in Y , then its middleman benefit just among the nodes in Y is

$$\sum_{\{v,w\} \in \Gamma} \beta(1+r_{v,w}) \geq \gamma \sum_{\{v,w\} \in \Gamma} \frac{1}{1+r_{v,w}}.$$

Since $|Y| > \frac{1}{2}n$ and the graph has only δn^2 edges, we have $|\Gamma| \geq (\frac{1}{8} - \delta)n^2 \geq \frac{1}{9}n^2$. Thus, by Claim 3.8, the middleman benefit is at least $\frac{\gamma r^2}{a+r}$ with $r = |\Gamma| \geq \frac{1}{9}n^2$ and $a = 3\delta n^3$. Hence it is at least

$$\frac{\frac{1}{81}\gamma n^4}{3\delta n^3 + n^2} \geq \frac{\frac{1}{81}\gamma n^4}{4\delta n^3} = \frac{1}{324}\gamma \delta^{-1}n > n,$$

where the first inequality follows since $\delta n \geq 1$ for $n \geq n_0$, and the final inequality follows since $\delta < \frac{1}{324}\gamma$. Since the cost for u to create edges to all nodes of Y is at most n , this represents a profitable deviation from the current graph, contradicting our assumption that the graph was in equilibrium. ■

4. HIERARCHICAL METRICS

We have now seen that equilibria can have rich properties even in a uniform metric; but the model and the efficiency of best-response dynamics apply to any metric. Here we look at properties that arise in another important class of metrics.

For modeling the structure of an organization, metrics based on hierarchies are very natural [13]: we put n points at the leaves of a (rooted) tree with edge lengths, and we define the distance between two points to be the total length of the unique simple path between them in the tree. Such a structure thus models the organizational hierarchy of a large company or other organization, with the effort required for two nodes to connect across different parts of the company growing in the tree distance between them. In this section, we will focus for the sake of concreteness on the case in which the n points are placed at the leaves of an n -node rooted complete bi-

nary tree, and the distance between two points whose least common ancestor has height h is equal to 2^{h-1} . However, our results can be extended to metrics derived more generally from exponentially weighted trees.

In computational experiments, we find that equilibria in this metric have a rich structure that aligns in interesting ways with the motivation from organizational structure; Figure 4 shows a characteristic example on 32 nodes. As the figure indicates, in these equilibria there are typically a small number of nodes that form connections to almost everyone, a few other nodes form connections to everyone within moderate-sized subtrees, and most nodes build a very small number of connections to nearby other nodes. Thus, as with the uniform metric, a type of symmetry-breaking occurs in which certain nodes stake out central positions, but it is not profitable in equilibrium for too many to do so — a phenomenon that is closely consistent with the observation that in real organization, a relatively small number of people tend to seek out and bridge structural holes.

It remains an open question to prove the existence of equilibria in this metric for arbitrarily large n (although the computational experiments clearly provide strong evidence). However, we can prove properties that must hold in any equilibrium; specifically, we show that two of the basic features we observe in the computational experiments — very-high-degree nodes and lower average degree — hold in all equilibria. The first result applies to our general model of payoffs, subject to the assumption that α_0 is a constant independent of n , and $\beta(1) > 4/3$.

THEOREM 4.1. *Let α_0 and $\beta(1)$ be constants, with $\beta(1) > 4/3$. In any equilibrium graph G in the hierarchical metric, there is a node of degree $\Omega(n)$.*

Proof. Let $\gamma = \max(\alpha_0, \beta(1))$; note that γ is a constant independent of n . We first observe that if a node v of degree d purchases an edge of cost c to a node w , then $d \geq c/\gamma$. Indeed, the benefit v derives from its edge to w is only through its direct connection to w and through acting as an intermediary between w and other neighbors of v . So if $d < c/\gamma$, then v would lose a benefit of at most $(d-1)\beta(1) + \alpha_0 \leq d\gamma < c$ by dropping the edge to w . As this is less than the cost of the edge, dropping this edge would represent a profitable deviation for v , contradicting the assumption that G is in equilibrium.

Next, we claim that some node v in G has a purchased edge to a node w , such that the least common ancestor of v and w is the root of the tree — in other words, v has a purchased edge of cost $n/2$. For suppose not, and let v be any node. If v purchases edges to all nodes in G , it becomes the unique intermediary between each node in the left sub-tree of the root and each node in the right sub-tree of the root. (It is the unique such intermediary since any other intermediary would have to be a node that purchased an edge of cost $n/2$.) This would provide v with a benefit of $(n^2/4)\beta(1)$. The total cost of all these edges is $(n/2)(n/2) + (n/4)(n/4) + (n/8)(n/8) + \dots < n^2/3$. Since $\beta(1) > 4/3$, purchasing all these edges would be a profitable deviation, contradicting the assumption that G is in equilibrium.

Thus, some node purchases an edge of cost $n/2$; by our earlier argument, this node must have degree at least $n/(2\gamma) = \Omega(n)$. ■

We now show that the average degree is asymptotically much smaller than linear. For this, we clearly need some restriction on the form of the payoff function in terms of the function $\beta(\cdot)$, since it is easy to create payoffs high enough that a clique will form. We therefore impose a restriction analogous to the ones earlier, generalizing the harmonic payoff, that $\max_r(r\beta(r))$ is a constant independent of n . (Recall that in the harmonic case, by definition, we have $r\beta(r) = 2\alpha_0$ for all r .)




Figure 5: The fraction of nodes in computed equilibrium networks with at least each given saturation level. All equilibria were computed via best response dynamics executed in round robin order with $n = 512$ and $\alpha_0 = 6.5$. The equilibrium represented by the black bars was computed starting from the empty graph. The two equilibria represented by the light and dark gray bars were computed starting from two random graphs.

THEOREM 4.2. *Let α_0 and $\max_r(r\beta(r))$ be constants. In any equilibrium graph G in the hierarchical metric, the average node degree is $O(\sqrt{n})$.*

Proof. We write β^* for $\max_r r\beta(r)$, and $\gamma = \max(\alpha_0, \beta^*)$. Since no node has an incentive to change its purchased edges, each node in G has the property that the benefit it receives is at least as large as the cost it incurs. Summing this inequality over all nodes, we see that the total benefit to all nodes is at least as large as the total cost paid by all nodes. The total benefit is at most a constant γ per node pair, and hence the total benefit to all nodes is at most $\gamma \binom{n}{2}$. The total cost incurred by all nodes, which is simply the total cost of all edges in G , is thus at most $\gamma \binom{n}{2}$.

Now suppose that G contains more than $m = 2^{j+1} \cdot n/2$ edges, where we choose j as small as possible so that $m \geq \gamma^{1/2} n^{3/2}$. The total cost of all edges in G is at least as large as the cost of the m cheapest node pairs in the metric. There are only $(n/2) \cdot (1 + 2 + \dots + 2^{j-1}) < (n/2) \cdot 2^j = m/2$ pairs of nodes of cost less than 2^j ; hence among these m cheapest pairs, at least half have cost at least $2^j = m/n$. Thus the total cost of the m edges is at least $(m/2)(m/n) = m^2/(2n) = \frac{1}{2}\gamma n^2 > \gamma \binom{n}{2}$. This contradicts our assumption that the total benefit is at least as large as the total cost. ■

Further structure through computational experiments. These two theorems suggest some of interesting structure shared by all equilibrium networks in the hierarchical metric, with nodes occupying different social positions. Through computational experiments with best-response dynamics, we can quantify this notion of a “social position” more fully, and identify more subtle structure in the equilibria. In particular, we find through experiments that most nodes approximately “saturate” some sub-tree that they belong to; that is, a node v tends to construct links to most nodes in some rooted subtree T_v containing v , and almost no links to nodes outside T_v .

Motivated by this observation, we define the ε -saturation level of a node v to be the maximum h such that v constructs links to at least a $(1 - \varepsilon)$ fraction of all nodes in the height- h subtree containing it. In Figure 5, we show the fraction of nodes of saturation level $\geq h$ for each h , in three different computed equilibrium networks on 512 nodes, with the harmonic model of payoffs and a fixed choice of α_0 . (Thus, these are hierarchical metrics in which the maximum subtree height is 9.) We see that many nodes achieve a saturation level of 3, a smaller number achieve 4 and 5, and a few achieve a saturation level of 9; this latter group constitutes the nodes described in Theorem 4.1 that construct links to almost all other nodes.

As this curve essentially characterizes the relative sizes of different “status levels” in the network, a very interesting open question is to try formally characterizing in more detail how stable the shape of this curve is across the full set of possible equilibria, and what its asymptotic shape looks like as n goes to infinity.

Acknowledgments. We thank the members of the Institute for the Social Sciences Networks group for valuable discussions on the topic of structural holes.

5. REFERENCES

- [1] Venkatesh Bala and Sanjeev Goyal. A non-cooperative model of network formation. *Econometrica*, 68:1181–1229, September 2000.
- [2] Ronald S. Burt. *Structural Holes: The Social Structure of Competition*. Harvard University Press, 1992.
- [3] Ronald S. Burt. Structural holes and good ideas. *American Journal of Sociology*, 110(2):349–99, September 2004.
- [4] Ronald S. Burt. Second-hand brokerage: Evidence on the importance of local structure for managers, bankers, and analysts. *Academy of Management Journal*, to appear.
- [5] Vincent Buskens and Arnout van de Rijt. Dynamics of networks if everyone strives for structural holes, 2007. Working paper, Cornell University Department of Sociology.
- [6] Jacomo Corbo and David C. Parkes. The price of selfish behavior in bilateral network formation. In *Proc. 24th ACM Symposium on Principles of Distributed Computing*, pages 99–107, 2005.
- [7] Sanjeev Goyal and Fernando Vega-Redondo. Structural holes in social networks. *Journal of Economic Theory*, to appear.
- [8] Mark Granovetter. The strength of weak ties. *American Journal of Sociology*, 78:1360–1380, 1973.
- [9] Matthew Jackson. The study of social networks in economics. In Joel Podolny and James E. Rauch, editors, *The Missing Links: Formation and Decay of Economic Networks*. Russell Sage Foundation, 2007.
- [10] Matthew Jackson and Asher Wolinsky. A strategic model of social and economic networks. *Journal of Economic Theory*, 71(1):44–74, 1996.
- [11] Jon Kleinberg and Éva Tardos. *Algorithm Design*. Addison Wesley, 2006.
- [12] Dov Monderer and Lloyd S. Shapley. Potential games. *Games and Economic Behavior*, 14:124–143, 1996.
- [13] Timothy Van Zandt. Decentralized information processing in the theory of organizations. In Murat Sertel, editor, *Contemporary Economic Issues*, volume 4 (Economic Design and Behavior), pages 125–160. MacMillan, 1999.

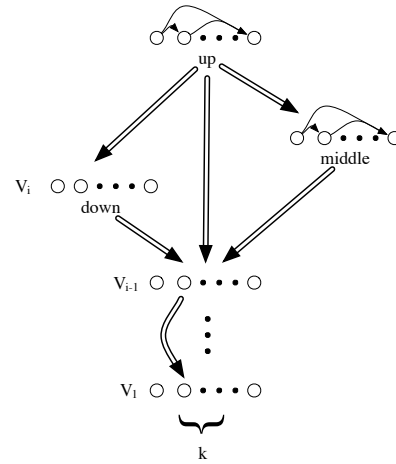


Figure 6: During round i , nodes that have gone down form V_i . Nodes that have gone up will form subsequent levels. Middle nodes have yet to make their move during round i . Single arrows represent edges between nodes. The double arrows represent that each node referred to by the base of the double arrow is connected to each node that is referred to by the arrowhead. (Not pictured here, for the sake of clarity, are double arrows connecting all up, down, and middle nodes to V_1, \dots, V_{i-1} .)

6. APPENDIX: CONVERGENCE OF BEST-RESPONSE DYNAMICS

In this section, we provide the details of the proof deferred from Section 3.2, that in the uniform metric, best-response dynamics with round-robin scheduling starting from the empty graph will converge to $G_{n,k}$. (Recall that in Section 3.2 we proved this in the special case when convergence was to $G_{n,k}$ for $k \geq \lceil \frac{n}{2} \rceil$.)

The proof will track several properties of the graph as it evolves through best-response dynamics, and to formulate this precisely, we need the following preliminary definitions. The intuition behind these definitions is as follows (also see Figure 6). During round i , k nodes form level V_i of the graph. When a node does this, it is said to “go down” to level V_i . If a node chooses not to join V_i during round i it is said to “go up”. If a node has not chosen which level to join, and its turn has not yet come during round i , it is called a “middle node”. T_i denotes all the nodes that have joined level i during round i (gone down), all the nodes that are not yet part of a level and have yet to move during round i (middle nodes), and all the nodes who have not yet committed to a level and have moved during round i (gone up). Since the nodes that have already chosen a level are frozen, T_i is where all the action takes place during round i .

DEFINITION 6.1. After node u ’s turn during round i , u is said to go down if it retains edges to all $v \in \cup_{j=1}^{i-1} V_j$ and only those nodes.

DEFINITION 6.2. $T_i = V \setminus (\cup_{j=1}^{i-1} V_j)$.

DEFINITION 6.3. Node u is said to be a middle node if during round i , u ’s turn has not yet come and $u \in T_i$.

DEFINITION 6.4. After node u ’s turn during round i , u is said to go up if u keeps its edges to all $v \in \cup_{j=1}^{i-1} V_j$ and all middle nodes, and u buys edges to all nodes that have already gone down during round i .

With these definitions in hand, we are now ready to prove the theorem. The proof is by induction which maintains certain invariants. Part 1 of the induction hypothesis shows that each node maintains edges to all nodes below it. Part 3 shows that once a node becomes part of a level it does not buy any edges for the rest of time. These two parts of the induction hypothesis combine to show that once a node becomes part of a level it is essentially “frozen”. Part 2 maintains that, during round i , each middle node either forms part of V_i (goes down) or goes up. Part 4 shows that each level has k nodes in it, with perhaps the exception of the topmost level. These invariants are pictorially represented in Figure 6.

THEOREM 6.5. *If k is the largest integer such that $B(n, k) \geq 0$ and $B(n, k - 1) \leq 0$, then starting from the empty graph, best response dynamics executed in round robin order will converge to $G_{n,k}$ in $\lceil \frac{n}{k} \rceil$ rounds.*

Proof. The proof is by induction on k . Lemma 3.4 shows that the first $n - k$ nodes go up during round 1 and the last k nodes go down. The induction hypothesis has four parts:

1. After v 's turn during round i ,
 - (a) if $v \in V_j$, where $1 \leq j < i$, v keeps edges to all nodes in V_1, \dots, V_{j-1} .
 - (b) if $v \in T_i$, v keeps edges to all nodes in V_1, \dots, V_{i-1} .
2. During round i , nodes in T_i will either go up or down.
3. During round i , if $v \in V_j$ where $1 \leq j < i$, then v will not buy any edges.
4. (a) If $|T_i| \geq k$ at the beginning of round i , then $|V_i| = k$,
(b) If $|T_i| < k$ at the beginning of round i , then $|V_i| = |T_i|$ and i is the last round.

We assume the induction hypothesis is true for rounds 1 through $i - 1$ and prove it is true for round i . To prove part 1a of the inductive step, consider the amount of benefit $v \in V_j$, $1 \leq j < i$, gets from attaching to all the nodes in $V_{j'}$ where $1 \leq j' < j$. By part 4 of the induction hypothesis, we have that $|V_{j'}| = k$. Also, part 2 of the induction hypothesis shows that the nodes that went down to form level j' are not connected to each other. There will be at most $n - k$ nodes attached to the nodes in $V_{j'}$. Thus v will earn at least $B(n, k)$ from attaching to the nodes in $V_{j'}$. Since $B(n, k) \geq 0$, Lemma 3.1 implies v will maintain its edges to $V_{j'}$ during round i . Part 1b of the inductive step follows by a similar argument.

Inductively assume that during round i , all the nodes in T_i will either go down or go up. Consider a node v which is a middle node whose turn has come. Let m denote the number of middle nodes other than v , and let d denote the number of nodes that have gone down. The marginal payoff that v would get from connecting to d' of the down nodes and m' of the middle nodes (denoted $\text{mp}_v(d', m')$) is:

$$\text{mp}_v(d', m') = (\alpha_0 - 1)(d' + m') + \left[\binom{d'}{2} + d'm' \right] \beta(n - d - m) \quad (1)$$

The first term comes from the fact that there is an α_0 direct benefit to attaching to the middle and down nodes and a cost of $d' + m'$ for doing so. By part 1 of the inductive step, and since the nodes that have gone down are not attached to each other, the second term describes the middleman benefit from attaching to d' of the down nodes, and the middleman benefit from becoming the middleman between d' down nodes and m' middle nodes. By part 2 of the induction hypothesis, and the definition of a node going up, all the

middle nodes are attached to each other and thus there is no middleman benefit for attaching to two of them.

Equation 1 is quadratic in d' with a positive coefficient on the quadratic term where $d' \in \{0, 1, \dots, d\}$. Thus it is either maximized at $d' = 0$ or $d' = d$. First, assume $\text{mp}_v(d', m')$ is maximized at $d' = 0$. In this case Equation 1 simplifies to,

$$\text{mp}_v(0, m') = m'(\alpha_0 - 1) \leq 0, \quad (2)$$

since $0 \leq \alpha_0 \leq 1$. Thus, if v chooses not to buy edges to any of the down nodes, it will also choose not to buy edges to any of the middle nodes. Next, assume that $\text{mp}_v(d', m')$ is maximized at $d' = d$. In this case mp_v must strictly increase when $d' = d + 1$ ¹. Thus $\text{mp}_v(d' + 1, 0) > \text{mp}_v(d', 0)$. Next, if we evaluate $\text{mp}_v(d', m')$ at $m' = 1$ we get:

$$\text{mp}_v(d', 1) = (\alpha_0 - 1)(d' + 1) + \left(\frac{d' + 1}{2} \right) \beta(n - d - m).$$

This is also equal to $\text{mp}_v(d' + 1, 0)$, and thus we have shown $\text{mp}_v(d', 1) = \text{mp}_v(d' + 1, 0)$.

Assume for the sake of contradiction that when $d' = d$, mp_v is maximized at $m' = 0$. Thus, $\text{mp}_v(d', 0) \geq \text{mp}_v(d', 1) = \text{mp}_v(d' + 1, 0)$ which provides a contradiction. Since mp_v is a linear function of m' where $m' \in \{0, 1, \dots, d\}$, it will either be maximized at $m' = 0$ or $m' = m$. So we can conclude that when mp_v is maximized at $d' = d$, then $m' = m$. Thus, v will either buy edges to all middle and down nodes, or none of them depending on whether Equation 1 is positive or negative respectively. This combined with part 1 of the inductive step imply that v will either go up or down.

Next we prove part 3 of the inductive step. Let $v \in V_j$ where $1 \leq j < i$. By parts 1 and 2 of the inductive step, the only nodes that do not have links either to, or from v are those nodes in V_j . Thus v would earn at most $B(n, k - 1) \leq 0$ by attaching to the other nodes in V_j . By Lemma 3.1, v would not buy these edges.

Next, we prove part 4a of the inductive step in two parts. First we show that if at the beginning of round i , $|T_i| \geq k$ then $|V_i| \leq k$. Observe that k nodes going down during round i results in a k -element independent set V_i . Furthermore, parts 1 and 2 of the inductive step show that if $v \in V_i$, $|N(v)| = n - k$. Thus each middle node yet to move would earn $B(n, k)$ from attaching to all k of these nodes. Since $B(n, k) \geq 0$, by Lemma 3.1, these middle nodes would all go up.

Now assume for the sake of contradiction that at the beginning of round i , $|T_i| \geq k$, yet $|V_i| < k$. Let v be the last node to go up during round i . Let d be the number of nodes that already went down during round i ; they form an independent set of size d . Let m be the number of middle nodes; they form a clique of size m . Since $|V_i| < k$, $d + m \leq k - 1$. Part 1 of the inductive step shows that all down and middle nodes nodes keep their edges to V_1, \dots, V_{i-1} . Similarly, part 2 of the inductive step shows that all nodes that went up keep their edges to the middle and down nodes. Thus, by Lemma 3.5, v would earn at most $B(n, k - 1) \leq 0$ from going up. Thus v would have have earned at least as high a marginal payoff if it would have went down, which contradicts the fact that v is rational and went up. A similar argument to the above shows that if at the beginning of round i , $|T_i| < k$, then all middle nodes would go down, thus $|V_i| = |T_i|$. Then parts 2 and 3 of the induction hypothesis show that no edges would be bought or dropped after round i . ■

¹Here we are analyzing mp_v strictly as a function of two variables and ignoring the fact that in this model the meaningful range of d' is $\{0, 1, \dots, d\}$.