



The Importance of Network Topology in Local Contribution Games

Citation

Corbo, Jacomo, Antoni Calvó-Armengol, and David C. Parkes. 2007. The importance of network topology in local contribution games. In *Internet and Network Economics*, ed., X. Deng, F.C. Graham, 388-395. Berlin, Heidelberg: Springer. Previously published in *Lecture Notes in Computer Science* 4858: 388-395.

Published Version

doi:10.1007/978-3-540-77105-0_43

Permanent link

<http://nrs.harvard.edu/urn-3:HUL.InstRepos:4000304>

Terms of Use

This article was downloaded from Harvard University's DASH repository, and is made available under the terms and conditions applicable to Other Posted Material, as set forth at <http://nrs.harvard.edu/urn-3:HUL.InstRepos:dash.current.terms-of-use#LAA>

Share Your Story

The Harvard community has made this article openly available.
Please share how this access benefits you. [Submit a story](#).

[Accessibility](#)

The Importance of Network Topology in Local Contribution Games

Jacomo Corbo*

Antoni Calvó-Armengol[†]

David Parkes*

Abstract

We consider a stylized model of content contribution in a peer-to-peer network. The model is appealing because it allows for linear-quadratic payoff functions and for very general interaction patterns among agents. We ask: How do different link patterns affect contributions? How are contributions affected by whether these goods are strategic complements, substitutes, or something in between? Who contributes and who free rides? How is the computational complexity of optimal play affected by the externalities of the game? What networks optimize social contributions and social welfare? And how does the worst-case equilibrium performance of a network compare to its best-case performance? The analysis finds that Nash equilibria of this game always exist and that they are computable by solving a linear complementarity problem. The equilibrium is unique when goods are strategic complements or weak substitutes and contributions are proportional to a network centrality measure called the Bonacich index. In the case of public goods, the equilibrium is non-unique and for every network there is an equilibrium where some individuals contribute optimally and others completely free-ride. Moreover, these equilibria with complete free-riders and optimal contributors are the only candidate stable equilibria and are characterized by k -order maximal independent sets. We find that the structure of optimal networks is hub-like when the game exhibits strict or weak complements. Under public good scenarios, while hub-like networks remain optimal in the best case, they also yield the worst-performing equilibria. Finally, we discuss two network-based policies for improving the equilibrium performance of networks.

1 Introduction

Peer effects, or the dependence of individual outcomes on group behaviour, is a characterizing feature of peer-to-peer systems (Antoniadis et al. 2005). File-sharing systems, for example, rely on participants to provision the network with content. Participants can experience a marginal increase or decrease in utility from the kind of content contributed by others. We call such goods strategic complements and strategic substitutes, respectively. Complementary goods may refer to missing portions of a file sought by a user or simply to diverse and desirable content. Substitutable goods refer to file contributions that duplicate, if only approximately so, a user's own contributions. Furthermore, in many peer-to-peer systems (as in BitTorrent, Tribler, Kazaa) the proximity of peer-provisioned content impacts the extent to which it is accessible by a given participant. This may be either because of the congestion and resulting latency along multi-hop connections (as in BitTorrent) or because persons have only a partial view of the network (as in Kazaa).

Following Ballester et al. (2005), we adopt a simple model for a contribution game in this paper. The model allows for positive and negative *bilateral influences*, meaning two players can positively or negatively affect each other's payoffs.

¹SEAS, Harvard, Cambridge, MA. E-Mail: {jaco, parkes}@eecs.harvard.edu

²ICREA, Universitat Autònoma de Barcelona, Barcelona, Spain. E-Mail: antoni.calvo@uab.es

A player is modeled with a *linear-quadratic* utility function, that allows for utility-dependence on the contribution by other players. The utility structure provides for an individualized component, reflecting decreasing marginal-returns for a player’s own contribution, in addition to a term that reflects *local interaction* that varies across pairs of players, meaning pairs of players can affect each other differently.

The population-wide pattern of these local complementarities or substitutabilities is well-captured by a network. In the case of a peer-to-peer file sharing system, this network closely resembles the underlying network of connectivity between peers and the local complementarities or substitutabilities reflect the payoff interdependencies rooted in this network structure of player links. In such a setting, equilibrium strategies naturally reflect players’ network embeddedness.

The model is appealing because a unique Nash equilibrium of the contribution game can be readily computed as a metric of network centrality when the network exhibits complementarities. This facilitates a rich empirical analysis of the contribution game under this scenario. First proposed by Bonacich (1987) nearly two decades ago, the Bonacich network centrality measure counts the number of all paths emanating from a given node weighted by a decay factor that decreases with the length of these paths.¹ As demonstrated by Ballester et al. (2005) this turns out to be intimately related to equilibrium behavior in our setting, since the paths capture all possible feedbacks from others’ play. The decay factor represents the extent to which others’ actions enter into a player’s own action’s payoff. We also find that Nash equilibria are always (even when they are non-unique) computable by solving a linear complementarity problem. In the case of public goods, while the equilibria can be non-unique, we find that equilibria where players either contribute completely or free-ride completely are the only candidate-stable equilibria and are characterized by k -order maximal independent sets. In the case of public goods (equivalently strongly substitutable goods), equilibria are non-unique and involve free-riding, i.e. some users not contributing at all to network-available content. In such cases, we find that equilibria where players either contribute optimally or free-ride completely are the only candidate-stable equilibria, and that these are characterized by k -order maximal independent sets.

We consider the problems of designing network that maximize aggregate contribution and welfare, and find that the structure of optimal networks is hub-like when the game exhibits strict or weak complements. Under public good scenarios, while hub-like networks remain optimal in the best case, they also yield the worst-performing equilibria. We discuss two network-based policies aimed at improving the equilibrium performance of networks: the removal of a single *key* player and the reorganization of network links. Removing the key player yields dramatic improvements to the worst-case performance of a graph and can completely reconcile best- and worst-case performance in sufficiently sparse graphs. We show that the efficaciousness of the link-reorganization policy scales as $O(\sqrt{e})$ with respect to the performance of completely random graphs with given v, e , and decreases as graphs become larger and denser.

1.1 Related Work

This paper situates itself in a growing body of literature interested in games where endogenous play is susceptible to externalities manifested across network links. Demange (2006), Kung (2005), and Bramoulle and Kranton (2006) equilibrium profiles in a game with public good (*substitutable good*) provisioning. Our model deals with a different payoff structure and allows for externalities to be either complementary or substitutable. Moreover, Kung’s model deals with cooperative coalition

¹It was originally interpreted as an index of influence or power of the actors of a social network. Katz (1953) is a seminal reference.

formation in the core, whereas our model is non-cooperative, so that we use the Nash equilibrium as our solution concept.

One of the main points of concern in the network games literature is the tension between stability and efficiency. See Jackson (2005) for a survey of the area. Johari et al. (2005) and Roughgarden (2004) investigate the effects of network architecture on the worst-case efficiency (otherwise the *price of anarchy*) of equilibria in routing games. Jackson et al. (2005) and Anshelevich et al. (2004) study the best-case efficiency (otherwise the *price of stability*) of equilibria in network formation games. We provide a partial characterization of equilibria in our game and relate a network's efficiency, in both the best and worst case, to its geometric properties.

Our contribution game is based on work that investigates how individuals provision content in the face of the externalities from other individuals' contribution. Our model is based on a model by Ballester et al. (2005) but where players only face uniformly positive or negative bilateral externalities and marginally decreasing returns in their own contribution. This stylized model captures many natural settings; specifically, it is a good model of content provisioning in peer-to-peer file-sharing systems. Ballester et al. do not focus on the relative efficiency of different network topologies, and much of the analysis in this paper, even the characterization of corner equilibria in terms of k -maximal independent sets, can be adapted to their model.

Finally, Kempe et al. (2003) investigate how the topology of a social network affects the propagation of ideas and formulate the problem of how best to maximize the spread of influence for a given network. Taking the view that social networks are endogenously optimizing some feature, Liben-Nowell and Kleinberg (2003) concern themselves with how to predict which links will form as the network evolves. We study a setting where agents freely interact and study how the pattern of interactions affects the optimality of observed outcomes. We study two network-based intervention policies to improve the equilibrium performance of graphs.

2 The Model

2.1 The game

Let $\mathcal{G}(v, e)$ denote the set of undirected and unweighted connected graphs without loops with v vertices and e edges. Connectedness requires that $e \geq v - 1$, which we assume from now on.

Players are connected by a network $\mathbf{g} \in \mathcal{G}(v, e)$ with adjacency matrix $\mathbf{G} = [g_{ij}]$. This is a zero diagonal and non-negative square matrix, with $g_{ij} \in \{0, 1\}$ for all $i \neq j$.

Each player $i = 1, \dots, n$ selects a contribution $x_i \geq 0$, and gets a payoff $u_i(x_1, \dots, x_n)$. We focus on *bilinear* utility functions of the form:

$$u_i(\mathbf{x}, \mathbf{g}) = x_i - \frac{1}{2}x_i^2 + a \sum_{j=1}^n g_{ij}x_ix_j, \quad (1)$$

where $\mathbf{x} = (x_1, \dots, x_n)$.

Note that this utility is marginal-decreasing in an agent's own contribution, with $\frac{\partial^2 u_i}{\partial x_i^2} = -1$.

The external effect of another agent on the utility of agent i is captured by the cross-derivatives $\frac{\partial^2 u_i}{\partial x_i \partial x_j} = ag_{ij}$, for $i \neq j$. When $a > 0$, the effect on agent i of agent j 's contribution is marginal-increasing if and only if i and j are connected in \mathbf{g} ; when $a < 0$, the effect is marginal-decreasing. The network \mathbf{g} reflects the pattern of existing payoff (relative) complementarities (when a is positive) and substitutabilities (when a is negative) across all pairs of players.

Where convenient, we use Σ to refer to the n -player game with payoffs given by Equation 1 and strategy space, the positive real line.

2.2 The Linear Complementarity Problem

We analyze the set of pure strategy Nash equilibria of the game introduced above. We note that an equilibrium exists if and only if the following holds:

$$\begin{aligned} \frac{\partial u_i}{\partial x_i}(x^*) &= 0, \quad \forall i \in N \text{ such that } x_i^* > 0 \\ \frac{\partial u_i}{\partial x_i}(x^*) &< 0, \quad \forall i \in N \text{ such that } x_i^* = 0 \end{aligned} \tag{2}$$

In matrix notation, the above Nash equilibrium necessary and sufficient conditions in Equations 2 become:

$$\begin{aligned} \mathbf{x}^* &\geq \mathbf{0}, \\ -\mathbf{a} + \Sigma \mathbf{x}^* &\geq \mathbf{0}, \\ -\mathbf{x}^{t*}(-\mathbf{a} + \Sigma \mathbf{x}^*) &= \mathbf{0}. \end{aligned} \tag{3}$$

The problem of finding a vector \mathbf{x}^* such that the above conditions hold is known as the linear complementarity problem $LCP(-\mathbf{a}, -\Sigma)$ (Cottle 1993). We can therefore state the following:

Theorem 2.1. *The set of pure strategy Nash equilibria of the contribution game with parameters α and Σ are given by the set of solutions to $LCP(-\alpha, -\Sigma)$.*

The linear complementarity problem is a well-studied problem and we borrow from this literature to address existence of the Nash equilibrium in our game, as well as in our empirical studies to characterize optimally-designed networks.

In the next sections we study the current model under strict complementarities, when $a > 0$, and under substitutabilities, when $a < 0$. The local interaction graph connecting agents becomes irrelevant when $a = 0$, as the contribution levels of other agents does not impact an agent's utility. In this case, the network-independent optimal contribution level for each agent is 1.

Example 2.2. To fix ideas, Figure 1 illustrates how contribution levels and utilities of all nodes vary with parameter a on an arbitrary graph g . The graph g , is a 12-vertex regular graph depicted in Figure 2 later in the paper. The plot shows the contribution levels of all nodes at all equilibria, which have been exhaustively enumerated. All equilibria are superimposed on the graph once, not counting isomorphisms.

The plots illustrate the multiplicity of equilibria under the pure public good scenario, when a is negative below some threshold less than 0. When $a < 0$, we will characterize this threshold and distinguish between pure public goods, (when a is less than this threshold), and show that hidden complementarities exist above this threshold. The plots also show that above this threshold, when complementarities, hidden or otherwise, exist, the equilibrium is unique and the contribution levels are interior, i.e. $x_i > 0$ for all nodes in the graph. The plots also point to the network-independent optimum contribution level (of 1) persisting for some nodes even when substitutabilities are very large. We will show that this is always the case for all graphs. We also note that intermediate contribution levels, i.e. $0 < x_i < 1$, become less prevalent as a becomes more negative. We will show that equilibria characterized by such contribution levels, which we call partially-corner, are not stable to a perturbed best-response dynamic.




Figure 1: The plots show the individual contributions (left) and utilities (right) of all nodes on the graph \mathbf{g} (shown in Figure 4.2), as a function of a . Along with the graph \mathbf{g} , a determines how the externalities of other nodes' play is felt by each agent.

3 Complementary Goods

We first study the game under local complementarities, when $a > 0$. Before turning to the equilibrium analysis, we define a network centrality measure due to Philip Bonacich (1987) that proves useful for this analysis.

3.1 The Bonacich network centrality measure

Let $\mathbf{g} \in \mathcal{G}(v, e)$ and, as before, let \mathbf{G} denote its adjacency matrix. This is a v -square, symmetric, and zero diagonal $(0, 1)$ -matrix. Denote by $\lambda_1(\mathbf{g})$ its largest eigenvalue, also called the *index* of \mathbf{g} . Note that this index is always well-defined and $\lambda_1(\mathbf{g}) > 0$.

Lemma 3.1. *Let $a > 0$. The matrix $\mathbf{B}(\mathbf{g}, a) = [\mathbf{I} - a\mathbf{G}]^{-1}$ is well-defined and non-negative if and only if $a\lambda_1(\mathbf{G}) < 1$. Then, $\mathbf{B}(\mathbf{g}, a) = \sum_{k=0}^{+\infty} a^k \mathbf{G}^k$, and its coefficients $b_{ij}(\mathbf{g}, a)$ count the number of paths in \mathbf{g} starting at i and ending at j , where paths of length k are weighted by a^k .*

Proof. From Theorem III*, page 601 in Debreu and Herstein (1953). ■

Definition 3.2. Suppose that $a\lambda_1(\mathbf{G}) < 1$. The vector of Bonacich centralities of parameter a in \mathbf{g} is $\mathbf{b}(\mathbf{g}, a) = \mathbf{B}(\mathbf{g}, a) \cdot \mathbf{1}$.

The Bonacich (1987) centrality of node i is $b_i(\mathbf{g}, a) = \sum_j b_{ij}(\mathbf{g}, a)$, and counts the *total* number of weighted paths in \mathbf{g} starting from i .

Theorem 3.3. *For $a\lambda_1(\mathbf{G}) < 1$, the game Σ has a unique Nash equilibrium $\mathbf{x}^*(\Sigma)$, which is interior and given by*

$$\mathbf{x}^*(\Sigma) = \mathbf{b}(\mathbf{g}, a) \tag{4}$$

Proof. From Theorem I in Ballester et al. (2005). ■

Corollary 3.4. *The utility of player i at equilibrium is $u_i(\mathbf{x}^*, \mathbf{g}) = \frac{1}{2}x_i^{*2} = \frac{1}{2}b_i(\mathbf{g}, a)^2$.*

Proof.

$$\begin{aligned} u_i(\mathbf{x}^*, \mathbf{g}) &= x_i^* - \frac{1}{2}x_i^{*2} + a \sum_{j=1}^n g_{ij}x_i^*x_j^* \\ &= x_i^* \left(1 + a \sum_{j=1}^n g_{ij}x_j^*\right) - \frac{1}{2}x_i^{*2} \end{aligned}$$

for all i . By the first order conditions $\partial u_i(\mathbf{x}^*, \mathbf{g})/\partial x_i = 0$ that yield the Nash equilibrium² we have:

$$x_i^* = 1 + a \sum_{j=1}^n g_{ij}x_j^*, \text{ for all } i.$$

Replacing above gives:

$$u_i(\mathbf{x}^*, \mathbf{g}) = x_i^{*2} - \frac{1}{2}x_i^{*2} = \frac{1}{2}x_i^{*2} = \frac{1}{2}b_i(\mathbf{g}, a)^2.$$

■

Then the total equilibrium welfare $\mathbf{u}(\mathbf{x}^*, \mathbf{g}) \cdot \mathbf{1}$ is $\frac{1}{2} \|\mathbf{b}(\mathbf{g}, a)\|^2 = \frac{1}{2} \mathbf{b}(\mathbf{g}, a) \cdot \mathbf{b}(\mathbf{g}, a)$. We remark that the correspondence between the Bonacich centrality indices of a graph and its equilibrium when $a > 0$ establishes uniqueness and interiority of equilibria when $a\lambda_1(\mathbf{G}) < 1$. It also establishes that an equilibrium fails to exist when $a\lambda_1(\mathbf{G}) > 1$. Under the latter conditions, the positive feedback from other agents' contributions is too high and contributions increase without bound.

3.2 Comparative Statics

The previous results relate individual equilibrium outcomes to the Bonacich centrality in the network \mathbf{g} of local complementarities. The next result establishes a positive relationship between the aggregate equilibrium outcome and the pattern of local complementarities. First we remark that the cross effects for the payoff function given in Equation 1 is captured by the matrix $\Sigma = -\mathbf{I} + a\mathbf{G}$. For any two matrices Σ and Σ' , we write $\Sigma' > \Sigma$ if $\sigma'_{ij} > \sigma_{ij}$, for all i, j , with at least one strict inequality.

Theorem 3.5. *Let Σ and Σ' symmetric such that $\Sigma' > \Sigma$. If $a\lambda_1(\mathbf{G}') < 1$ and $a\lambda_1(\mathbf{G}) < 1$ then $\mathbf{x}^*(\Sigma') > \mathbf{x}^*(\Sigma)$.*

Proof. $\Sigma' = \Sigma + aD$, with $d_{ij} \geq 0$ with at least one strict inequality. Theorem 1 holds so $-\Sigma\mathbf{x}^*(\Sigma) = \Sigma'\mathbf{x}^*(\Sigma') = \mathbf{1}$, and $\mathbf{x}^*(\Sigma), \mathbf{x}^*(\Sigma') > \mathbf{0}$. Now observe that $-\mathbf{x}^{*t}(\Sigma')\Sigma\mathbf{x}^*(\Sigma) = \mathbf{x}^{*t}(\Sigma')\mathbf{1} = x^*(\Sigma')$. By symmetry of Σ' , we have: $-\mathbf{x}^{*t}(\Sigma')\Sigma\mathbf{x}^*(\Sigma) = -\mathbf{x}^{*t}(\Sigma')\Sigma'\mathbf{x}^*(\Sigma) + a\mathbf{x}^{*t}(\Sigma')D\mathbf{x}^*(\Sigma) = x(\Sigma) + ax^{*t}(\Sigma')D\mathbf{x}^*(\Sigma)$. ■

Each individual outcome increases with a and with \mathbf{G} . Indeed, we are either increasing the weight to each path, or increasing the number of such paths. Hence, both individual equilibrium outcomes, aggregate equilibrium outcome and welfare increase with a and with \mathbf{G} .

²Recall that this equilibrium is interior, so that these hold with equality.

4 Substitutable Goods

When $a < 0$ we have a substitutability effect between players' contributions, i.e. we have a public good game. We have shown that for all values of a , positive or negative, the equilibria of the contribution game are given by the solutions to an associated LCP problem. We have seen that when $a > 0$ an equilibrium to the game on the graph g fails to exist when $a > 1/\lambda_1(g)$. We remark that an equilibrium always exists in a public good setting, i.e. there exists a Nash equilibrium profile for any network structure. It is easy to show that the best-response function is continuous from the compact convex set $\{x \in \mathcal{R}^n : \forall i, 0 \leq x_i \leq x^*\}$ to itself. The result follows from Brouwer's Fixed Point Theorem.

We now proceed to study the game under two separate conditions: when substitutabilities are weak and when they are strong, i.e. the case of pure public goods.

4.1 Weak Substitutes

We define the complement network $\overline{\mathbf{G}} = \mathbf{J} - \mathbf{I} - \mathbf{G}$, where \mathbf{J} is the all-ones matrix, i.e., $\overline{g}_{ij} = 1 - g_{ij}$, for all $i \neq j$. In words, two vertices are linked in $\overline{\mathbf{G}}$ if and only if they are not linked in \mathbf{G} . We write:

$$\Sigma = (1 + a)\mathbf{I} + a\overline{\mathbf{G}} - a\mathbf{J}.$$

Suppose first that $-1 < a < 0$. Solving for the Nash equilibrium is then equivalent to solving $LCP(-\frac{1}{1+a}\mathbf{e}, \mathbf{I} + \frac{a}{1+a}\overline{\mathbf{G}} - \frac{a}{1+a}\mathbf{J})$.

The solution can be equivalently written in terms of the Bonacich index of nodes on the complement network.

Theorem 4.1. *Consider a game on G where $a < 0$ and let $\overline{\mathbf{G}} = \mathbf{J} - \mathbf{I} - \mathbf{G}$ as before. There exists a unique equilibrium if and only if $-a\lambda_1(\overline{\mathbf{G}}) < 1 + a$. Then, the equilibrium is unique, interior and proportional to Bonacich, that is,*

$$x_i^* = \frac{1}{1 + a + a \sum_{j=1}^n b_j \left(\frac{-a}{1+a}, \overline{\mathbf{G}} \right)} b_i \left(\frac{-a}{1+a}, \overline{\mathbf{G}} \right), \text{ for all } i = 1, \dots, n$$

Recall that we are dealing with the case $-1 < a < 0$. Notice that $-a\lambda_1(\overline{\mathbf{G}}) < 1 + a$ is equivalent to $-\frac{1}{1+\lambda_1(\overline{\mathbf{G}})} < a$. Therefore, the interior unique equilibrium is obtained on $-\frac{1}{1+\lambda_1(\overline{\mathbf{G}})} < a < 0$.

4.2 Public Goods

When substitutabilities are large, i.e. when $a < -\frac{1}{1+\lambda_1(\overline{\mathbf{G}})}$, the above transformations fail to work.

In these circumstances, agents' free-riding on others' contributions is severe enough that some agents do not contribute at all. Equilibria are also no longer unique. In these cases, agents whose equilibrium contribution levels are non-zero either contribute the optimum, (here 1), or some value less than optimum. We call such equilibria corner and partially-corner, respectively.

Example 4.2. To fix ideas, Figure 2 illustrates the contribution levels for 3 different equilibria. The left-most equilibrium is interior and is an equilibrium for $a = -0.49$, the graph in the middle depicts a partially-corner equilibrium for $a = -0.5$, and the right-most graph shows a corner equilibrium for $a = -0.5$. Recall that Figure 1 illustrates how the trajectory of equilibrium profiles change with a and that $a = -0.5$ is the threshold at which equilibria become non-unique and below which equilibria are non-interior. We remark that excluding isomorphisms, there are 9 different




Figure 2: Equilibrium contributions on local interaction graph.

equilibrium profiles that exist on the graph shown for $a = -0.5$ but only one corner profile, namely the one shown on the right.

4.2.1 Partially-Corner Equilibria

Precisely, a partially corner equilibrium profile \mathbf{x}^* on the network $\mathbf{G}(\mathbf{v}, \mathbf{e})$ is one such that there exists some $i, j \in v$ such that $x_i = 0$ and $0 < x_j < 1$. We remark that equilibrium contributions of non-corner agents, i.e. all agents $j \in v$ such that $0 < x_j < 1$, is related to the Bonacich centrality index on the subnetwork joining them:

Lemma 4.3. *Let $a < 0$. Given a partially corner equilibrium profile \mathbf{x} on the network $\mathbf{G}(\mathbf{v}, \mathbf{e})$, the contribution levels of all nodes i such that $0 < x_i < 1$ is given by*

$$x_i^* = \frac{1}{1 + a + a \sum_{j=1}^n b_j \left(\frac{-a}{1+a}, \overline{\mathbf{G}'} \right)} b_i \left(\frac{-a}{1+a}, \overline{\mathbf{G}'} \right), \text{ for all } i = 1, \dots, n$$

Proof. Remove all nodes $i \in v$ (and their adjacent edges) such that $x_i = 0$ since agents whose contributions are zero are effectively absent from the graph. Without loss of generality, consider a single connected component \mathbf{G}' that remains. Note that contribution levels of all nodes in \mathbf{G} remain equilibrium strategies in \mathbf{G}' . The equilibrium is interior, thereby unique, and Theorem 4.1 therefore holds. ■

4.2.2 Corner Equilibria

As mentioned earlier, when local substitutabilities are large, corner equilibria, with agents either free-riding completely or contributing their maximum (1), also exist. For such a situation to be an equilibrium, both free-riders and contributors must gain at doing so. When $a \leq -1$, there is always an equilibrium outcome where some agents exert all the effort and others free-ride. Our proof for this is constructive and provides a method to find all corner equilibria in general graphs. We use a concept from graph theory called maximal independent sets. Given the graph $\mathbf{G} = (v, e)$, let $N_G(v')$ designate the set of neighbours of node $v' \in v$. We find that all corner equilibria for $a \leq -1$ are described by maximal independent sets of contributors.

Definition 4.4. An independent set $S \subseteq v$ for $\mathbf{G}(\mathbf{v}, \mathbf{e})$ is a set of nodes such that no node in the set is adjacent to any other node in the set, i.e. $\forall v' \in S, |N_G(v') \cap S| = \emptyset$.

A maximal independent set is an independent set maximal for the inclusion. Put another way, it is a set S such that every edge of the graph \mathbf{G} has at least one endpoint not in S and every vertex not in S has at least one neighbor in S .

Lemma 4.5. *For $a \leq -1$, there exists a corner equilibrium on the graph \mathbf{G} .*

Proof. We use the following properties of maximal independent sets. Given a maximal independent set S every agent either belongs to S or is connected to an agent who belongs to S . To see this, suppose not. Let S be a maximal independent set, and let i be an agent who does not belong to S and is not connected to any agent who belongs to S . Then the set $S \cup \{i\}$ is an independent set, and hence S is not maximal. Thus, we can partition the population into two disjoint sets of agents: those who belong to maximal independent set S , and those who are linked to an agent in S . For any agent i , there exists a maximal independent set S of the graph G such that i belongs to S . To see this, note that i itself is an independent set. To build a maximal independent set, begin with i and successively add agents not linked to i , then agents not linked to those agents, etc. Now note that free-riders must be connected to at least one (non-zero) contributor to have strict incentives not to contribute at all. Conversely, contributors cannot be connected to another contributor, as otherwise they could substitute effort levels with each other. ■

In the same way, when $-1 < a \leq -\frac{1}{1+\lambda_1(\mathbf{G})}$, complete free-riders need to be connected to more than one contributor, because neighbors' contributions are weighted by $|a| < 1$ and so do not compensate exactly for one's contributions. The following Theorem characterizes the structure of corner equilibria in terms of independent sets with dominating constraints.

Definition 4.6. A set $S \subseteq v$ is called a k -order maximal independent set if and only if it is a maximal independent set such that each node not in the set is connected to at least k nodes in the set, i.e. $\forall v' \notin S, |N_G(v') \cap S| \geq k$.

Theorem 4.7. *Let $\lceil 1/a \rceil$ be the smallest integer higher or equal than $|1/a|$. For $a \leq -\frac{1}{1+\lambda_1(\mathbf{G})}$, a corner profile is a Nash equilibrium if and only if the set of contributing players, i.e. $\{i \in v : x_i = 1\}$, is a $\lceil 1/a \rceil$ -order maximal independent set of the graph \mathbf{G} .*

Maximal independent sets correspond to maximal independent sets of order 1. Every graph has a maximal independent set, therefore there always exists a corner equilibrium for $a \leq -1$. However, for $k \geq 2$, k -order maximal independent sets need not always exist. Therefore, when $-1 < a \leq -\frac{1}{1+\lambda_1(\mathbf{G})}$, we may not have a corner equilibrium, though we may still have a partially corner equilibrium. Recall that an equilibrium is guaranteed to exist for all $a \in \mathbb{R}$.

Consider a star, for example. There are two maximal independent sets, one composed by the hub alone, the other composed by all the spokes. There is a single maximal independent set of order k for $2 \leq k \leq n-1$ composed by all the spokes together.

Notice that there is a sense in which only least-well connected players contribute in a k -maximal independent set, as in the star example above, as free-riders not in the set need to have many connections with nodes in the set, so free riders need to be well-connected.

4.2.3 Stable Equilibria

We now show that corner equilibria are the only candidate stable equilibria to a perturbed best-response procedure, thereby obviating the need to consider partially corner equilibria in our computational results. The correspondence of corner equilibria to maximal independent sets of order k

gives us insight into the computational complexity of computing equilibria under substitutabilities. We report on this later in the paper.

We use a simple notion of stability based on Nash tâtonnement (e.g. Fudenberg 1991). We use this concept because it applies directly to games with continuous action spaces, and because it reduces the number of equilibria in our setting in a natural way. Define $f_i(\mathbf{x})$ as the best-response of i to a profile \mathbf{x} and define \mathbf{f} as the collection of these individual best-responses $\mathbf{f} = (f_1, \dots, f_n)$. An equilibrium \mathbf{x} is stable if and only if there exists a positive number $\rho > 0$ such that for any vector ϵ satisfying $\forall i, |\epsilon_i| \leq \rho$ and $x_i + \epsilon_i \geq 0$ the sequence $\mathbf{x}^{(n)}$ defined by $\mathbf{x}^{(0)} = \mathbf{x} + \epsilon$ and $\mathbf{x}^{(n+1)} = \mathbf{f}(\mathbf{x}^{(n)})$ converges to \mathbf{x} .

This standard notion yields a strong result: only corner equilibria are stable. The result is convenient because it helps to mitigate the problem of multiple equilibria when $a \leq -\frac{1}{1+\lambda_1(\overline{\mathbf{G}})}$, where $\overline{\mathbf{G}}$ is given as the complement of the network \mathbf{G} , as before.

Theorem 4.8. *For any network \mathbf{G} , an equilibrium is stable if and only if it is a corner equilibrium.*

Proof. We first show that interior equilibria are not stable. At an interior equilibrium, at least one player chooses a strictly positive action $x_i > 0$ that is smaller than the individual optimal action denoted x^* (if he chooses a zero action, this is not an interior equilibrium; if all players choose the individual optimal action, they must be all isolated from each other for this action profile to be an equilibrium). Given an equilibrium profile, denote by J the set of all such players, $J \neq \emptyset$. Perturb slightly the action of all the players in J by adding some positive number $\rho > 0$. If the perturbation is small enough, $\rho < \min x^* - x_j : j \in J$, the players in the neighborhood of J all best-respond by decreasing their action. This further triggers an upwards rise of the action by players in J . These best-response dynamics clearly diverge from the original equilibrium profile. We now show that a partially corner equilibria may be stable. Consider a partially corner equilibrium where zero action players are connected to at least two non-zero action players. Perturb the action of all players by some $|\epsilon_i| < x^*/n^2$. Then, direct computations show that the ensuing sequence of best-response dynamics converge back to the original equilibrium. ■

The result is also powerful from a computational perspective, once we observe that the best and worst equilibria, in the region where equilibria are non-unique, are described by maximal independent sets (of order 1).

Lemma 4.9. *For any network \mathbf{G} , and $a \leq -\frac{1}{1+\rho(\overline{\mathbf{G}})}$ such that equilibria are non-unique, the contribution-maximizing and contribution-minimizing stable equilibria are both described by maximal independent sets.*

Proof. All stable equilibria in this region are described by maximal independent sets of order $\lceil 1/a \rceil$. When $a \leq -1$, these equilibria are maximal independent sets. The dominating constraints become more stringent as a increases above -1 , in turn reducing the number of admissible networks over which we conduct the search for contribution-maximizing and contribution-minimizing networks. ■

Together, Theorem 4.8 and Lemma 4.9 provide us with a simplification of how we study the tradeoff between contribution-maximizing and contribution-minimizing (equivalently utility-maximizing and utility-minimizing) networks under substitutabilities, when equilibria are non-unique. While the problems of finding the largest and the smallest maximal independent sets are both \mathcal{NP} -hard, these are well-studied problems and there are algorithms for computing both (e.g. Halldórsson et al. 1998, Gutin

2004). We leverage this to solve for a graph's contribution-maximizing (best-case) and contribution-minimizing (worst-case) equilibria, across all a when equilibria are non-unique, together with the annealing procedure described at the beginning of this section to search over the space of networks.

5 The Complexity of Equilibria

We have shown that in all cases, the equilibria of the contribution game are given by the solutions to an associated LCP problem. This formulation, however, belies the complexity of both deciding whether an equilibrium exists, and of solving for the equilibrium outcome.

For $a > 0$, equilibrium contribution levels are interior, unique, and given by the Bonacich indices of nodes in the network; similarly for $-\frac{1}{1+\rho(\mathbf{G})} < a < 0$, save that the Bonacich index is now on the associated network of non-connections. In both cases, solving for the Nash equilibrium has polynomial-time complexity.

When $a \leq -\frac{1}{1+\rho(\mathbf{G})}$, we lose interiority and uniqueness of equilibria. As described in the last section, when $a \leq -\frac{1}{1+\rho(\mathbf{G})}$ corner equilibria are described by maximal independent sets of order $\lceil 1/a \rceil$ on the graph \mathbf{G} . When $a \leq -1$, these are maximal independent sets, which are guaranteed to exist for all \mathbf{G} . These are polynomial-time solvable³.

Lemma 5.1. *Consider the game Σ with $a \leq -1$. An equilibrium of the game is solvable in polynomial time.*

Proof. Theorem 4.7 establishes that corner equilibria are described by maximal independent sets. Furthermore, the explication of existence of such a set in the Theorem provides an algorithm for its construction that is clearly polynomial in n . ■

For the region where $-1 < a < -\frac{1}{1+\rho(\mathbf{G})}$ corner equilibria may not exist. As before, we focus on these equilibria because they are the only candidate stable equilibria to a perturbed best-response dynamic. Let $k = \lceil 1/a \rceil$. Deciding whether a corner equilibrium exists amounts to proving that an independent sets of order k on \mathbf{G} exists. When $k \geq 2$, we have the following complexity on solving for such equilibria.

Theorem 5.2. *Let $-1 < a < -\frac{1}{1+\rho(\mathbf{G})}$. Deciding if a corner equilibrium of the associated game Σ exists is \mathcal{NP} -complete.*

Proof. We omit the proof for reasons of brevity. The result is an \mathcal{NP} -completeness reduction from an independent set with domination constraints. We show that a maximal independent set of order k corresponds to an independent set of vertices $S \subset v, S \neq \emptyset$ with $|S| \geq \min\{r : r \notin \{0, \dots, k-1\}\}$ such that $\forall v \notin S : |N_G(v) \cap S| \in \{0, \dots, k-1\}$. The classification of the problem of deciding whether the latter sets as \mathcal{NP} -complete is studied in Halldórsson et al. (1998). ■

We remark that when restricted to chordal graphs the k -independent set problem is always solvable in polynomial time.

³Enumerating maximal independent sets or finding the smallest and largest of these remain \mathcal{NP} -hard.

6 Optimal Network Design

A number of edges e and vertices v are given. The problem of optimal network design consists on arranging the v vertices with the e edges in such a way that some objective function is maximized. We identify two different network design problems.

Problem 1 Solve $\max_g \{\mathbf{x}^*(\mathbf{g}, a) : \mathbf{g} \in \mathcal{G}(v, e)\}$.

Problem 2 Solve $\max_g \{\mathbf{u}(\mathbf{x}^*(\mathbf{g}, a)) : \mathbf{g} \in \mathcal{G}(v, e)\}$.

In the first problem, the social planner wants to maximize aggregate activity (or contribution) at equilibrium. In the second problem, the social planner wants to maximize aggregate equilibrium welfare.

We study these problems in our game under complementarities and substitutabilities. When the equilibria are unique, then these problems are well-posed, however when equilibria are non-unique (when substitutabilities become too strong), we consider both the best case and worst case, respectively contribution- or utility-maximizing and contribution- or utility-minimizing, equilibria of networks. We study the relationship between the best-case and worst-case equilibria of a given network, and particularly how contribution- or utility-maximizing networks in the best case perform in the worst case, as well as how contribution- or utility-maximizing networks in the worst case perform in the best case.

We present both analytical and computational results in this section. In our computational experiments, the optimization process is performed using adaptive simulated annealing as described by Allen (1995) and Tsallis et al. (1994). Starting from a given initial network configuration, random rewiring of individual links is performed, the objective function is evaluated and the change is accepted with a certain probability that depends on a temperature coefficient that allows the system to escape local minima in the search space and vanishes to zero as the search results become more certain. When the equilibrium is unique, the equilibrium is solved by computing the associated Bonacich index of nodes, whereas when equilibria are non-unique, the equilibrium is given by algorithms that compute the smallest and largest maximal independent sets. In all cases, our characterization of the equilibria of the game helps us avoid solving the associated LCP.

6.1 Optimizing Under Complementarities

We first consider the case when we have complementarities, i.e $a > 0$. We observe that problems 1 and 2 correspond to maximizing the L_1 and L_2 norms of the Bonacich index vector, i.e. $\max_g \{\mathbf{b}(\mathbf{g}, a) \cdot \mathbf{1} : \mathbf{g} \in \mathcal{G}(v, e)\}$ and $\max_g \{\mathbf{u}(\mathbf{x}^*(\mathbf{g}, a)) : \mathbf{g} \in \mathcal{G}(v, e)\}$, respectively. The correspondence stems from the convenient relationship between utility and contribution quantities and the Bonacich indices in our game, when the equilibrium is unique. The relation between Bonacich centrality and the unique Nash equilibrium of our game under complementarities lends a geometric interpretation to Problem 1: namely, the social welfare-maximizing graph is the graph with maximal weighted paths. For both problems, the relation suggests a way to characterize optimal equilibria using spectral graph theory. We first show that the two network design problems can be reconciled in our game when the equilibrium is interior and unique. Specifically, we show that as network complementarities grow, the two problems become asymptotically the same.

Lemma 6.1. *Let $\mathbf{g} \in \mathcal{G}(v, e)$, and $\lambda_1(\mathbf{G})$ its index. As $a \uparrow \frac{1}{\lambda_1 \mathbf{G}}$, Problems 1 and 2 are equivalent and reduce to:*

$$\max\{\lambda_1(\mathbf{g}) : \mathbf{g} \in \mathcal{G}(v, e)\}. \quad (5)$$

Proof. For a given network \mathbf{g} , $\mathbf{b}(\mathbf{g}, a)$ is only well-defined when $a < 1/\lambda_1(\mathbf{G})$. Therefore, the strict upper bound for a in $\mathcal{G}(v, e)$ is $1/\lambda_1(\mathbf{g}^*)$, where \mathbf{g}^* is the solution. Then,

$$\lim_{a \uparrow 1/\lambda_1(\mathbf{g}^*)} \mathbf{b}(\mathbf{g}^*, a) \cdot \mathbf{1} = \lim_{\mathbf{a} \uparrow 1/\lambda_1(\mathbf{g}^*)} \|\mathbf{b}(\mathbf{g}^*, \mathbf{a})\| = +\infty,$$

whereas

$$\lim_{a \uparrow 1/\lambda_1(\mathbf{g}^*)} \mathbf{b}(\mathbf{g}, a) \cdot \mathbf{1}, \lim_{a \uparrow 1/\lambda_1(\mathbf{g}^*)} \|\mathbf{b}(\mathbf{g}, a)\| \in \mathcal{O}(1), \text{ when } \mathbf{g} \neq \mathbf{g}^*.$$

This drives the equivalence of Problems 1 and 2. \blacksquare

In words, when a is high enough, the networks that maximize aggregate outcome and welfare coincide, and are given by the graph with maximal index.

We use results on graphs with largest maximal eigenvalues to solve the network design problem in this case.

Let $\mathbf{s}_{1,v-1}$ ($v \geq 3$) be the star with vertices $1, \dots, v$ and vertex 1 as hub. For $1 \leq k \leq v-3$, let $\mathbf{h}_{v,k}$ be the graph obtained from $\mathbf{s}_{1,v-1}$ by joining vertex 2 to vertices $3, \dots, k+3$.

Theorem 6.2. *Suppose that a is high enough. Then,*

- (i) *if $e = v - 1$, the unique solution to (3) is $\mathbf{s}_{1,v-1}$;*
- (ii) *if $e = v + k$, $k \in \{0, 1\}$, the unique solution to (3) is $\mathbf{h}_{v,k}$;*
- (iii) *if $e = v + 2$, $\mathbf{h}_{v,k}$ is a solution to (3), but it is not unique;*
- (iv) *if $e = v + k$, $k \geq 2$, there exists $v(k)$ such that, for all $v \geq v(k)$, the unique solution to (3) is $\mathbf{h}_{v,k}$.*

Proof. See Theorem 3.3.2, page 68 in Cvetković et al. (1997). \blacksquare

These asymptotic results reveal a great deal about how the optimal networks (for both objective functions) change with the level of externalities. In the first place, in simulation, we find that L_1 and L_2 -maximizing graphs always coincide, irrespective of a .

As explained earlier, it is the particular payoff complementarities that fix a and induce the network linking players. Consider now the situation reversed: the payoff complementarity between two players i and j is set by two factors: the parameter a and the cumulative distance along all paths in number of node hops between i and j in some graph \mathbf{g} . By this view, we can now understand that the relative network externalities are driving the system to have multiple equilibria and that these increase either by increasing a or by increasing the largest eigenvalue in \mathbf{G} . Moreover, the above results suggest that both objective functions increase without bound as $a \uparrow 1/\lambda_1(\mathbf{G})$. We still treat a as an exogenous parameter that arises from the payoff decomposition seen earlier. So for all a , the optimal network \mathbf{g}^* is the graph on v and e with largest eigenvalue.

The left-hand plot in Figure 3 illustrates precisely this across a large range of a values, from strong substitutabilities to strong complementarities, and for graphs with varying numbers of edges. The graph shows that the first eigenvalue of aggregate contribution- and utility-maximizing graphs (L_1 - and L_2 -maximizing graphs in the case of $a > 0$) corresponds to the largest first eigenvalue possible for graphs with given v, e . In the region where $a < 0$ and equilibria are non-unique, the value attributed to a graph is taken to be its best-performing (best-case) stable equilibrium. With




Figure 3: The left-hand plot shows that the first eigenvalue of aggregate contribution- and utility-maximizing equilibrium graphs corresponds to the largest first eigenvalue possible given v and e across a wide range of a . In the range where equilibria are non-unique, we choose the most favourable, i.e. best case, 'stable' equilibrium performance of the graph. The right-hand figure plots the first eigenvalue of graphs with best worst-case stable equilibrium performance when equilibria are non-unique against the smallest first eigenvalue possible given v and e . Both plots refer to networks with varying number of edges while keeping $v = 12$ and $-3 < a < 0.5$.

v and e fixed, the largest eigenvalue of a graph is a measure of its regularity. A higher eigenvalue corresponds to an irregular hub-like structure, whereas a lower eigenvalue (still for the same number of vertices and edges) refers to a more regular network. We stress that the general trend remains true irrespective of a .

6.2 Optimizing Under Substitutabilities

While interiority holds, as suggested by the left-hand plot in Figure 3, we still find that aggregate contribution- and utility-maximizing graphs coincide and that these networks are maximal index graphs. When substitutabilities are strong enough, i.e. a grows negative enough, we lose interiority and have both partially corner and corner equilibria. As discussed earlier, corner equilibria are the only stable equilibria and so we only consider these. We now consider both the best-case and worst-case equilibrium performance of networks in the case of non-unique equilibria.

Figure 4 illustrates the tension between optimal networks in the best and worst cases. Networks that yield the highest contribution in the best case (among the set of its equilibria) also exhibit the worst worst case performance. We find that these networks are maximal index graphs, which are also optimizing networks when the equilibria is unique. On the other hand, networks that yield the best worst case performance are minimal index graphs, i.e. graphs with the smallest first eigenvalue. This is illustrated in the right-hand plot of Figure 3. Minimal index graphs also exhibit the smallest spread between best and worst case equilibrium performance, as illustrated in the right-hand plot of Figure 4.

6.3 Excluding the Key Player: A Network-Based Policy

We concluded the last section by showing that under substitutabilities and when equilibria are non-unique, the geometry of networks with best equilibrium performance also have worst worst-case

performance, whereas those networks with best worst-case performance are far from optimal in the best case. These networks correspond to maximal and minimal index graphs, respectively.

Now, we briefly investigate a network-based policy aimed at mitigating this discrepancy between best- and worst-case equilibrium performance. We provide a simple geometric criterion to identify the target in the population when the planner wishes to increase optimally the aggregate group outcome. We denote by \mathbf{G}^{-i} (respectively Σ^{-i}) the new adjacency matrix (respectively the matrix of cross-effects), obtained from \mathbf{G} (respectively from Σ) by setting to zero all of its i -th row and column coefficients. The resulting network is \mathbf{g}^{-i} . More formally, we want to solve $\max_{i \in N} \{\mathbf{x}^*(\Sigma^{-i}) - \mathbf{x}^*(\Sigma)\}$, where N is the population of agents. This is equivalently expressed as:

$$\max_{i \in N} \mathbf{x}^*(\Sigma^{-i}) \quad (6)$$

This is a finite optimization problem, that admits at least one solution. We find that a good heuristic for the solution of this problem is the removal of the highest degree node. Recalling that the best-case and worst-case stable equilibria in the range where equilibria are non-unique are described by maximal independent sets, every edge imposes a dominating constraint on the independent set construction, and so a node with highest degree imposes the largest number of constraints. Furthermore, barring the case where multiple nodes share the highest degree, such a node is excluded from the maximal independent set of contributors in the best case equilibria and is included in the set in the worst case. Note that the policy of excluding the key player will improve the network's performance in the worst case, without decreasing its best-case performance⁴.

To fix ideas, consider the star with $n - 1$ spokes. This graph has 2 maximal independent sets. The (best case) equilibrium highest aggregate contribution is described by all spokes contributing (a value of 1) and the hub free-riding completely. The worst case equilibrium is described by the hub being the only contributor. Removing the hub node from the graph leads to no change in aggregate contribution in the best case and completely reconciles the best and worst case outcomes: each node always contributes optimally.

The right-hand plot in Figure 4 illustrates the effectiveness of this policy in reconciling best- and worst- case equilibrium performance. The plot shows that worst-case performance of graphs can be dramatically improved and even match best-case performance when the network graph is sparse. As the graph becomes more dense, this policy becomes less effective.

6.4 Rearranging Network Links: A Network-Based Policy

Armed with an understanding of the network pattern of externalities that lead to the most efficient equilibrium outcomes, the system planner can now prescribe changes to the given graph that maximize its largest eigenvalue. One possible mechanism might be to induce a rearrangement of the network edges, keeping the number of nodes fixed. We explore the limits of this approach for the case where $a > 0$.

Theorem 6.3. *Starting from a sparse random graph $\mathbf{g}(v, p)$, where p is the probability of an edge between any two players i and j , $i \neq j$, the largest eigenvalue can be increased as much as $O(\sqrt{e})$.*

Proof. We omit the full proof for reasons of space. Briefly, consider a connected simple graph \mathbf{g} and let λ_1 be its largest eigenvalue. Let d_i , $i \in V$, denote the degree of a vertex i and let $\Delta = \max_{i \in V} d_i$. (If \mathbf{g} is regular then $d_i = \Delta$ for each $i \in V$, then $\lambda_1 = \Delta$.) If \mathbf{g} is not regular, then there is no graph

⁴Aggregate equilibrium contribution or welfare will only be adversely affected as complementarities increase.




Figure 4: The left-hand plot shows the spread in aggregate equilibrium contribution, i.e. the difference in aggregate contribution for the best- and worst-case equilibria, for graphs with the best best-case and best worst-case performance, fixing a such that equilibria are non-unique, for $v = 15$ (blue/dash), 25 (black/dot-dash), 35 (red/solid), and varying e . The right-hand plot shows the best-case and worst-case equilibrium performance of graphs with best worst-case performance (blue / dot-dash) compared to graphs with best best-case performance (red / solid), for $v = 12$, a such that equilibria are non-unique. Note that best best-case performing graphs have worst worst-case performance. The black / dash curve plots the the worst-case performance of graphs with the *key player* removed, starting from the best best-case performing graphs (red / solid).

such that $\Delta - \frac{1}{2n(n\Delta-1)\Delta^2} \leq \lambda_1 < \Delta$. See Stevanovic (2004) for the proof of this upper bound. In either case, however, $\lambda_1 = O(e)$.

Fixing the number of vertices, it is clear that the greatest difference in first eigenvalue lies in a graph with fewer edges. Consider a sparse random graph, where $\mathbf{g}(v, p)$ and $e = \frac{pv(v-1)}{2}$. It can be shown that the maximum eigenvalue of this graph λ'_1 is almost surely the following: $\lambda'_1 = (1 + O(1))\sqrt{\Delta}$, where the $O(1)$ term tends to zero as $\sqrt{\Delta}$ gets large.

Putting the two results together, we have the maximum increase in first eigenvalue by any manipulation of edges starting from a random graph, $\frac{\lambda_1}{\lambda'_1}$, to be bounded by $O(\sqrt{e})$. ■

Corollary 6.4. *Starting from a random graph, fixing v and e , the gains in aggregate welfare or aggregate contribution from manipulating the graph to achieve an optimal equilibrium (with respect to either objective) is $O(\sqrt{e})$.*

The interpretation of Theorem 6.3 is that the network designer's task of manipulating the graph will induce a change in first eigenvalue of size at most $O(\sqrt{e})$ and that this difference reduces as the network becomes more dense, i.e. edges are added.

As to the gains from such manipulation (Corollary 6.4), consider the example of a small enough so that the optimal graph is the one with largest eigenvalue. At a first order in a , the Bonacich index of i is $b_i = 1 + ad_i$, and so aggregate equilibrium is $v + 2ae$. The square of i 's Bonacich index is $1 + 2ad_i$, and social welfare is $v + 4ae$. The largest eigenvalue increase is of the order of \sqrt{e} . The optimal graph is such that the inverse of the largest eigenvalue equals a . Therefore, both the aggregate equilibrium and the social welfare increase for the optimal graph are roughly of the order of \sqrt{e} . In other words, the gains from random to optimally designed network are marginally decreasing as $O(\sqrt{e})$. Symmetrically, welfare losses in random versus optimal networks are sizeable, and the smaller the larger the network.

That the greatest gains from manipulating the network coincide with smaller networks (where such an approach may be easier to implement), preliminarily suggests the viability of the network designer's task. On the other hand, it also shows that a network designer is very limited in his capacity to affect the outcome on a large graph in this way, i.e. keeping e and v fixed.

Acknowledgements

We are grateful to Jan Arne Telle for his help with some of the complexity analysis in this paper. We also thank Matthew Jackson and Stuart Shieber for their comments.

References

- [1] ALBERT, R., H. JEONG, AND A-L. BARABÁSI (2000): "Error and Attack Tolerance of Complex Networks," *Nature*, 406, 378-382.
- [2] ANSHELEVICH, E., A. DASGUPTA, J. KLEINBERG, E. TARDOS, T. WEXLER, AND T. ROUGH-GARDEN (2004): "The Price of Stability for Network Design with Fair Cost Allocation," *Foundations of Computer Science*.
- [3] ANTONIADIS, P., C. COURCOUBETIS, AND R. MASON (2004): "Comparing Economic Incentives in Peer-to-Peer Networks," *To appear in Special Issue on Network Economics, Computer Networks*.
- [4] BALLESTER, C., A. CALVÓ-ARMENGOL, AND Y. ZENOU (2004) "Who's Who in Crime Networks. Wanted: The Key Player," CEPR Discussion Paper No. 4421.
- [5] BOLLOBÁS, B., AND O. RIORDAN (2003) "Robustness and Vulnerability of Scale-Free Random Graphs," *Internet Mathematics*, 1, 1-35.
- [6] BONACICH, P. (1987): "Power and Centrality: A Family of Measures," *American Journal of Sociology*, 92, 1170-1182.
- [7] CORBO, J., D. PARKES (2005): "The Role of Bilateral Consent in Strategic Network Formation," *Proceedings of the twenty-fourth annual ACM symposium on Principles of distributed computing*, 99 - 107.
- [8] CVETKOVIĆ, D., AND P. ROWLINSON (1990): "The Largest Eigenvalue of a Graph –A Survey," *Linear and Multilinear Algebra*, 28, 3-33.
- [9] CVETKOVIĆ, D., P. ROWLINSON, AND S. SIMIĆ (1997): *Eigenspaces of Graphs*. Cambridge: Cambridge University Press.
- [10] DEBREU, G., AND I.N. HERSTEIN (1953): "Nonnegative Square Matrices," *Econometrica*, 21, 597-607.
- [11] DURLAUF, S.E. (2004): "Neighborhood Effects," in *Handbook of Regional and Urban Economics Vol. 4*, ed. by J.V. Henderson and J-F. Thisse. Amsterdam: Elsevier Science, 2173-2242.
- [12] FUDENBERG, D., J. TIROLE (1991): "Game Theory," in *MIT Press*, Cambridge, 1991, 2325.
- [13] GOYAL, S., AND S. JOSHI (2003): "Networks of Collaboration in Oligopoly," *Games and Economic Behavior*, 43, 57-85.
- [14] GUIMERÀ, R., A. DÍAZ-GUILERA, F. VEGA-REDONDO, A. CABRALES, AND À. ARENAS (2002): "Optimal Network Topologies for Local Search with Congestion," *Physical Review Letters*, 89, 248701, 1-4.
- [15] GUTIN, G. (2004) "Independent Sets and Cliques," *Chapter 5.3 of the Handbook of Graph Theory*, eds. Johnathan Gross and Jay Yellen, CRC Press.

- [16] HALLDORSSON, M. M., J. KRATOCHVIL, AND J. A. TELLE (1998): "Independent Sets with Domination Constraints," *Automata, Languages and Programming*, 176-187.
- [17] JACKSON, M. O. (2005a): "The Stability and Efficiency of Economic and Social Networks," in *Microeconomics*, 0211011, EconWPA.
- [18] JACKSON, M. O. (2005): "A Survey of Models of Network Formation: Stability and Efficiency," in *Group Formation in Economics: Networks, Clubs and Coalitions*, ed. by G. Demange and M. Wooders. Cambridge: University Press.
- [19] , JOHARI, R., J. TSITSIKLIS (2004): "Efficiency Loss in a Network Resource Allocation Game," in *Mathematics of Operations Research*, 29 (3), 407-435.
- [20] KATZ, L. (1953): "A new status index derived from sociometric analysis" in *Psychometrika*, 18(1), 39-43.
- [21] KEARNS M., M. LITTMAN, AND S. SINGH (2001): "Graphical Models for Game Theory," *UAI*.
- [22] KRIVELEVICH M., B. SUDAKOV (2003): "The Largest Eigenvalue of Sparse Random Graphs," in *Combinatorics, Probability and Computing*, 12, 61-72.
- [23] PENNA T. J. P. (1994): "Travelling Salesman Problem and Tsallis Statistics" in *Physical Review E*, 51, R1.
- [24] , ROUGHGARDEN, T., E. TARDOS (2004): "Bounding the Inefficiency of Equilibria in Nonatomic Congestion Games," in *Games and Economic Behavior*, 47 (2), 389-403
- [25] STEVANOVIĆ D. (2004): "The largest eigenvalue of nonregular graphs," *Journal of Combinatorial Theory, Ser. B*, 91(1), 143-146.
- [26] TSALLIS C., D. A. STARIOLO (1995): "Generalized Simulated Annealing," in *Annual Review of Computational Physics II*.
- [27] WASSERMAN, S., AND K. FAUST (1994): *Social Network Analysis. Methods and Applications*. Cambridge: Cambridge University Press.