Available online at www.sciencedirect.com
SCIENCE DIRECT®
ELSEVIER
Games and Economic Behavior 56 (2006) 201–224
GAMES and Economic Behavior
www.elsevier.com/locate/geb

A contract-based model for directed network formation ☆

Ramesh Johari [a,*], Shie Mannor [b], John N. Tsitsiklis [c]

[a] *Department of Management Science and Engineering, Stanford University, CA*
[b] *Department of Electrical and Computer Engineering, McGill University, Montréal, PQ*
[c] *Department of Electrical Engineering and Computer Science and the Laboratory for Information and Decision Systems, MIT, Cambridge, MA*

Received 6 November 2003
Available online 19 October 2005

Abstract

We consider a network game where the nodes of the network wish to form a graph to route traffic between themselves. We present a model where costs are incurred for routing traffic, as well as for a lack of network connectivity. We focus on directed links and the *link stability* equilibrium concept, and characterize connected link stable equilibria. The structure of connected link stable networks is analyzed for several special cases.

JEL classification: C62; C72

Keywords: Network formation games; Traffic routing

1. Introduction

Network formation models describe the interaction between a collection of nodes that wish to form a graph. Such models have been introduced and studied in the economics literature; for

☆ Partially supported by the National Science Foundation under grant ECS-0312921 and a Graduate Research Fellowship to the first author, and by the ARO under grant DAAD10-00-1-0466. This research was done while the authors were with the Laboratory for Information and Decision Systems, MIT.
* Corresponding author.
E-mail addresses: ramesh.johari@stanford.edu (R. Johari), shie@ece.mcgill.ca (S. Mannor), jnt@mit.edu (J.N. Tsitsiklis).

0899-8256/$ – see front matter
doi:10.1016/j.geb.2005.08.010

a recent review, see Jackson (2004), who makes a distinction between two categories of formation models. In the first category, the network is formed by a distributed system, but controlled by a single actor with a single objective; in the second category, which is the one considered in this paper, each of the nodes in the network is a different player, and a network is formed through interaction between the players. We are interested in understanding and characterizing the networks that result when individuals interact to choose their connections. In particular, we will focus on the role of bargaining, negotiation, and contracts between the players.

The structure and characteristics of networks where each player selfishly optimizes its own utility were first studied in Aumann and Myerson (1988). That paper considers an extensive form game where links are sequentially formed, and the formation of a link requires the consent of both endpoints of the link. By contrast, Myerson (1991) develops a simultaneous move game, where all players simultaneously announce the set of players to whom they wish to be linked, and a link is formed between two players if both parties desire the link.

More recently, several network formation games have been introduced as models of various competitive situations, including social networks and markets. Social networks are captured by the connections model (Jackson and Wolinsky, 1996) and the spatial connections model (Johnson and Gilles, 1999). In these models the maintenance of a link induces a cost on the endpoints of the link; furthermore, the players receive a positive utility from the direct and indirect connections to other players (with decreasing utility as the connection becomes increasingly indirect). Other contexts for network formation games include market models, such as free-trade networks (Furusawa and Konishi, 2002), market sharing agreements (Belleflamme and Bloch, 2002), and buyer-seller networks (see Jackson, 2004; Kranton and Minehart, 2000; Slikker and van den Nouweland, 2001b, and references therein). Most of this work assumes that the links are bidirectional and that the link serves the two nodes equally; exceptions include buyer-seller networks (Kranton and Minehart, 2000) and networks in labor markets (Jackson and Calvó-Armengol, 2004). In this paper, we consider a fundamentally different model where the links are *directed*. Specifically, a link from node i to node j indicates that node i can use node j (e.g., goods can be sent from i to j), but not necessarily vice versa. We note that formation games of directed networks were considered in Bala and Goyal (2000). But as we describe below, the motivation and the cost structure of their work is quite different from this paper.

The model we develop is motivated by multi-player environments where bilateral negotiation may result in a contractual agreement between two nodes to form a directed link. This link will be used for flow of goods from the upstream node to the downstream node. The upstream node benefits from the improved connectivity; however, the downstream node suffers a cost due to the additional flow sent through it. This tradeoff is resolved by a contract where the upstream node compensates the downstream node. We note that contractual formation mechanisms have been considered previously (see, e.g., Currarini and Morelli, 2000; Slikker and van den Nouweland, 2001a); however, because contracts in our model are motivated by the need to ensure delivery of goods, we must not only specify the formation of a network, but also a routing of flow on that network. This constitutes a significant departure in the current paper from previous work on network formation.

Our work is related to Bala and Goyal (2000) who consider formation games with both directed links and with undirected links. As in our model, the cost of the link is incurred only to the node initiating the link, which leads to a noncooperative game. There are, however, two major differences between the model of Bala and Goyal (2000) and the model we study. First, in our model a node pays other nodes for establishing a link, thus the nodes bargain to create a link. In Bala and Goyal (2000) money is not transferred between nodes, but rather every link has its own predefined price. Second, we assume that there is some underlying traffic of goods on the network, and that each node incurs a cost proportional to the flow of goods through it. Thus, in our model each node is trying to reduce the amount of incoming traffic, or at least require compensation from other nodes for the traffic they send through it. There is no underlying traffic mechanism in Bala and Goyal (2000); rather, a node incurs a direct cost for having incoming links. The two models share the idea of forming connected graphs and incurring loss for being disconnected.

The model considered in this paper can be motivated from various contexts. As an example, consider a "privatized mail distribution" network, as it might arise from the deregulation of postal and delivery services (Sherman, 1991; Cohen, 2000). In such a network, there are several nodes, representing mail office branches; for every node, a link to another node indicates that mail can reach the other node directly. There is a natural requirement that mail should be able to reach any destination, possibly indirectly; this is the "universal service obligation" for postal services (Rawnsley and Lazar, 1999). There is a cost associated with each node, which represents the storage and labor costs of handling the traffic. This setting leads naturally to a network formation model with directed links, where a link from i to j represents a contractual obligation by branch j to deliver or forward any mail sent from i to j.

As another example, consider an information services network (e.g., a distributed storage system or a replication management network), where each node represents a separate facility, containing data storage (i.e., disks) and communications hardware. To ensure business continuity or for purposes of remote backup, each node (facility) wishes to replicate its data at other nodes; see Sidell et al. (1996) and Cox et al. (2002) for some examples. The contractual agreements between the nodes are to carry information, and deliver data to the intended nodes (Geels and Kubiatowicz, 2002; Fuqua et al., 2003). In many cases, e.g., for data migration or asynchronous data replication, the number of hops (i.e., the number of nodes traversed from origin to destination) is of lesser importance, and the primary requirement is that data can reach their destination. There is a cost associated with each node, which represents data storage and communications costs. As before, a model with directed links is natural, where each link represents a directional flow of traffic, and a contractual obligation that the downstream node will carry the traffic coming from an upstream node.

A last, particularly compelling example of a goods delivery network with a universal service obligation is found in the Internet and other communications networks, with nodes representing different providers. Here *transit* contracts are common (Norton, 2000; Gao and Rexford, 2001); such contracts commit a provider to deliver or forward incoming traffic from another provider or customer. The models of this paper can be applied to understand the transit contract formation process.

Overview of the paper

We consider a model where multiple players (nodes) interact to form a graph with directed links. Each node wishes to send a given amount of traffic to some of the other nodes, and only cares whether the traffic eventually arrives at the destination. Furthermore, there is a handling cost at each node, which is proportional to the amount of traffic through that node. In this model, every node prefers to be connected to all other nodes, but also prefers that nobody is connected to it; in some sense, each player prefers to be a leaf in the graph. However, since not all players can be leaves in a connected graph, there exists an inherent tradeoff between the connectivity cost and the traffic cost to a node. This leads us to consider the following bargaining process. Each node i submits to every other node j a separate bid indicating the amount i is willing to pay in order to form a link from i to j. Simultaneously, each node j declares the amount it is willing to accept for each link that terminates at j. A link (i, j) is formed if the amount j is willing to accept is no larger than the bid by node i.

Our game is specified by three parameters: a traffic matrix $\boldsymbol{B}$, a vector of connectivity costs $\boldsymbol{\alpha}$, and a routing mechanism. The entry b_{ij} in the traffic matrix represents the amount of flow node i wishes to send to node j. The connectivity cost α_i represents the cost per unit of traffic that node i fails to send to nodes it cannot reach. (We note that the connectivity cost may be different for different nodes, to help model heterogeneity among agents.) Finally, the routing mechanism specifies the flow sent through each vertex, given the graph structure. To streamline the presentation of the results, we will assume shortest path routing. However, the results of the paper continue to hold for other routing schemes that satisfy certain natural conditions; see Johari et al. (2004) for details.

We define the game of interest in Section 2, by formally completing the description above, and by introducing some assumptions on the way that traffic is routed once a graph is formed. Our main assumption is that if we add a link between two nodes, the flow through either node cannot decrease. We further show that shortest path routing (either with a fixed tie-breaking rule, or with equal splitting between paths of equal length) satisfies this assumption.

In Section 3, we introduce a suitable equilibrium concept. In contrast to the equilibrium concepts of Aumann and Myerson (1988) and Myerson (1991), we use the notion of *pairwise stability* with side payments suggested by Jackson and Wolinsky (1996). A graph of undirected links and "payments" is pairwise stable if the following two conditions hold:

(1) If there is a link between i and j, neither i nor j will be better off by breaking the link.
(2) If there is no link between i and j, there is no possible payment under which the link is formed, and both i and j become better off.

We adopt the same stability concept, but modify it to apply to our contract-based directed graph formation game; we call the resulting equilibrium concept *link stability*, and say that a graph G is link stabilizable if there exists a link stable strategy vector such that G is formed.

Our main interest is in graphs where each node i can reach all desired destinations, i.e., nodes j for which $b_{ij} > 0$; we refer to such graphs as $\boldsymbol{B}$-connected graphs. In Section 3, we provide an explicit characterization of the set of link stable strategies which lead to $\boldsymbol{B}$-connected graphs. We also prove that the graphs which are formed are "minimal," in the sense that no link can be removed without losing $\boldsymbol{B}$-connectivity.

The special case where each node wants to send exactly one unit of goods to each other node is discussed in Section 4. We study in some detail the possible connected link stabilizable graphs. We show that if the cost of being disconnected is sufficiently large, then every link stabilizable graph is connected. On the other hand, if the cost of being disconnected is below a threshold, then every link stabilizable graph is not connected. Moreover, at this threshold, the only connected link stabilizable graph is a directed cycle. We also consider social welfare for this special case. We define "social cost" as the sum of the costs incurred by nodes (both for routing traffic, and for a lack of connectivity). We show that the connected link stabilizable graph with lowest social cost is the bidirectional star (i.e., a graph with a central hub, where all other nodes have links to and from this hub). We also show that the directed cycle has the highest social cost among all connected link stabilizable graphs. Furthermore, we show that the ratio of the social cost of the directed cycle to the social cost of the bidirectional star grows linearly in the number of nodes, and thus may become arbitrarily large. We conclude in Section 5.

2. The game

We are given a finite set of nodes V, each of whom wishes to route a given amount of traffic between themselves; in particular, we assume that each node i has $b_{ij} \geqslant 0$ units of flow to send to each node $j \neq i$. We define $b_{ii} = 0$ for all i and denote the matrix of flows by $\boldsymbol{B} = (b_{ij}, i, j \in V)$. The nodes play a game to form a directed graph on which traffic will be routed. The strategy of node i is a vector $(p_{ij}, q_{ji}, j \neq i)$. The entry $p_{ij} \geqslant 0$ represents a *bid* from node i to node j, to form the link (i, j). The entry $q_{ji} > 0$ represents the minimum *acceptance value*, i.e., the minimum amount of payment node i is willing to accept from j to form the link (j, i). Note that q_{ji} must be strictly positive.

Given the strategies of all the nodes, denoted by the vector $s = (p_{ij}, q_{ji}, i, j \in V, j \neq i)$, a directed graph $G(s) = (V, A(s))$ is formed, where:

$$A(s) = \{(i, j)\colon\ i \neq j,\ p_{ij} \geqslant q_{ij}\}. \tag{1}$$

Thus, the link (i, j) is formed if the bid p_{ij} is larger than or equal to the acceptance value q_{ij}; since $q_{ij} > 0$, this implies that a link can only be formed if the bid p_{ij} is strictly positive. The graph $G(s)$ is then used to route traffic between the nodes.

Given a set of links A, we will assume that the traffic routing results in a vector $(f_i(A), i \in V)$, where $f_i(A)$ is the amount of flow passing through or terminating at node i; as we will discuss below, we define $f_i(A)$ by assuming shortest path routing—see Section 2.1. To capture graph connectivity, we define $w_i(A) \subset V$ to be the set of nodes $j \neq i$ which are unreachable from node i, given the set of links A.

When a set of links A is used to route traffic, we will assume that node i incurs a cost $c_i \geqslant 0$ for each unit of traffic which either passes through or terminates at i. In addition, if destination j is unreachable from node i, we will assume that node i incurs a cost proportional to the traffic b_{ij} that i wishes to send j; we denote the constant of proportionality by $\alpha_i > 0$. Thus, given the graph $G = (V, A)$, the total cost to node i is

$$C_i(A) = c_i f_i(A) + \alpha_i \sum_{j \in w_i(A)} b_{ij}. \tag{2}$$

Notice that as α_i increases, node i has an increasing incentive to insure connectivity to other nodes.

The following example illustrates the calculation of $C_i(A)$.

Example 1. In Fig. 1(a), we have three nodes ($V = \{1, 2, 3\}$) connected with a directed cycle, so $A = \{(1, 2), (2, 3), (3, 1)\}$. Suppose that each node has one unit of traffic to send to each other node, and that the connectivity costs are all equal to α: $\alpha_i = \alpha$, $i = 1, 2, 3$. In this case, $w_i(A) = \emptyset$ for all i, and the cost experienced by each node is $C_i(A) = 2c_i$.

In Fig. 1(b), we have a new set of links $A' = \{(1, 2), (2, 3)\}$. Since link $(3, 1)$ has been removed, we now have reduced the flow traveling through the network; no node sends any traffic to node 1, and node 3 no longer sends any traffic to other nodes. Further, $w_2(A') = \{1\}$, and $w_3(A') = \{1, 2\}$. This leads to the following costs:

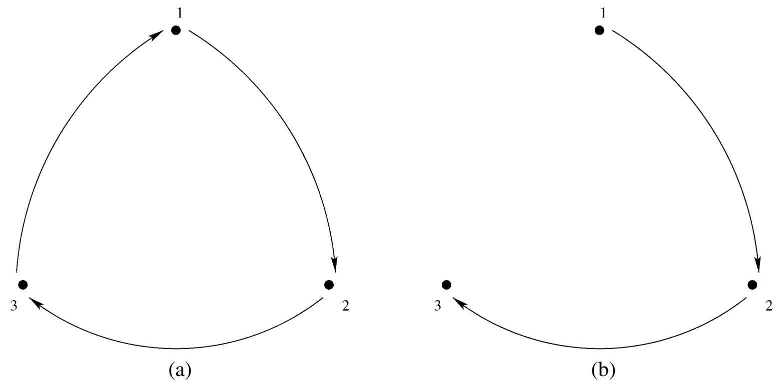

Fig. 1. Cost structure; see Eq. (2).

$$C_1(A') = 0,$$
$$C_2(A') = c_2 + \alpha b_{21} = c_2 + \alpha,$$
$$C_3(A') = 2c_3 + \alpha b_{31} + \alpha b_{32} = 2c_3 + 2\alpha.$$

Given the composite strategy vector s, the payoff to node i is determined entirely by the revenues received from other nodes, the payments made to other nodes, and the cost C_i. The payoff to node i is thus:

$$R_i(s) = \sum_{j:(j,i)\in A(s)} p_{ji} - \sum_{j:(i,j)\in A(s)} p_{ij} - C_i\big(A(s)\big). \tag{3}$$

Notice that only the bids p_{ij} appear in the payoff function; the acceptance values q_{ij} appear indirectly, as they determine whether or not a link (i, j) is formed. This leads to the following remark.

Remark 2. If s and s' are two strategy vectors such that $A(s) = A(s') = A$, and $p_{ij} = p'_{ij}$ for all $(i, j) \in A$, then $R_i(s) = R_i(s')$.

To complete the specification of the model, in the next section we describe a shortest path routing mechanism that characterizes the flow functions $(f_i(A), i \in V)$.

2.1. Routing

Recall that $(f_i(A), i \in V)$ is the amount of traffic entering node i through the various links that end at i. We will assume that given a graph $G = (V, A)$, traffic is routed using *shortest path routing*; in particular, we will assume that there exists a total ordering over all paths, such that shorter paths are preferred to longer paths. This routing then determines the functions $f_i(A)$. We will find that although we make the assumption of shortest path routing, the results we develop depend only on a critical property we refer to as *flow monotonicity*; the results can thus be generalized to any routing mechanism which satisfies flow monotonicity. See Johari et al. (2004) for more details.

We start with a definition and some notation.

Definition 3. A *path* from i to j is a set of links $P = \{(i, i_1), (i_1, i_2), \ldots, (i_{m-2}, i_{m-1}), (i_{m-1}, j)\}$ such that $i_u \neq i_v$ if $u \neq v$, and $i_u \neq i, j$ for all u.

We denote a path $\{(i, i_1), \ldots, (i_{m-1}, j)\}$ by $P = i i_1 i_2 \cdots i_{m-1} j$. By an abuse of notation, if P consists of only the link (i, j) we will write $P = (i, j)$. We will write $|P| = m$ to denote the fact that P contains m links. In addition, for $k \in V$, we will write $k \in P$ if there exists a link $(i, k) \in P$, or a link $(k, j) \in P$.

We use the notation $\mathcal{P}_A(i, j)$ to denote the set of all paths from i to j using the set of links A. We say that a graph (V, A) is *strongly connected* if for every i and j there exists a path from i to j using links in A—i.e., if $\mathcal{P}_A(i, j) \neq \emptyset$ for every i, j. We also let E denote the set of all possible links: $E = \{(i, j)\colon i \neq j\}$. We then define $\mathcal{P}(i, j) \triangleq \mathcal{P}_E(i, j)$, representing all the paths from i to j in the complete graph on the node set V.

Since there may be multiple possible paths between a pair of nodes i and j, we have to specify a mechanism for choosing one of the paths. We assume that there exists a *total ordering on paths*, which uniquely chooses a shortest path from i to j, given a particular set of links A.

Definition 4. A *total ordering on paths* $\prec$ satisfies the following two properties:

(1) For each pair of distinct nodes $i, j \in V$, $\prec$ is a total ordering on $\mathcal{P}(i, j)$.
(2) For $P, P' \in \mathcal{P}(i, j)$, if $|P| < |P'|$, then $P' \prec P$.

We say P is *ranked higher* than P' if $P' \prec P$.

The first condition says that $\prec$ compares all paths between i and j; and the second condition says that shorter paths are always ranked higher than longer paths, coinciding with shortest path routing. Given the set of links A, since the set of paths $\mathcal{P}_A(i, j)$ is finite, there exists a unique highest ranked path, which we denote $P_A(i, j)$; if there are no paths from i to j (i.e., if $\mathcal{P}_A(i, j) = \emptyset$), then we set $P_A(i, j) = \emptyset$ as well. We then assume that on the graph (V, A), all traffic originating at i destined for j flows along $P_A(i, j)$. Thus we have the following definition for $f_i(A)$:

$$f_i(A) = \sum_{j:j \neq i} \sum_k b_{jk} 1_{\{i \in P_A(j,k)\}}. \tag{4}$$

Note that with these definitions, flow is only routed between vertices which are connected. That is, if $P_A(j, k) = \emptyset$, then j cannot reach k, and thus the flow b_{jk} does not contribute to $f_i(A)$ for any i.

We also note that while the absolute information requirements of a total ordering over paths are possibly quite high—a complete specification of the ordering is exponential in the number of nodes—there are simple rules which give rise to a total ordering. For example, if the nodes are numbered $1, \ldots, n$, a reasonable method of breaking ties between shortest paths is to make a choice according to a lexicographic ordering based on node number. Such a tie breaking scheme is succinctly specified, and also gives rise to a total ordering in the precise sense described here.

Many of our proofs depend on the following critical property of the flow functions $f_i(A)$.

Proposition 5. *Suppose that* $(f_i(A), i \in V)$ *is defined according to* (4)*. Then given the set of links A, the following inequalities hold for all pairs* (i, j)*,* $i \neq j$:

$$f_i\big(A \cup \{(i,j)\}\big) \geqslant f_i(A), \tag{5}$$
$$f_j\big(A \cup \{(i,j)\}\big) \geqslant f_j(A). \tag{6}$$

Proof. Fix a set of links A, and a link (i,j). If $(i,j) \in A$, then (5)–(6) are satisfied trivially; so suppose that $(i,j) \notin A$, and let $A' = A \cup \{(i,j)\}$. Consider any pair (k,ℓ) such that $i \in P_A(k,\ell)$. We know that $\mathcal{P}_A(k,\ell) \subseteq \mathcal{P}_{A'}(k,\ell)$, and every path in $\mathcal{P}_{A'}(k,\ell) \setminus \mathcal{P}_A(k,\ell)$ passes through i. It follows that the highest ranked path $P_{A'}(k,\ell)$ must traverse i as well. As a result, any pair (k,ℓ) sending traffic through i in (V,A) also sends traffic through i in (V,A'). This shows that (5) holds. The proof that (6) holds is identical. □

According to the preceding proposition, when a link (i,j) is added to the graph, the total flow through both i and j cannot decrease. This is a natural result: one would expect that adding a link does not decrease the total flow entering the endpoints of that link. We will find that this proposition plays a key role in developing the results of this paper; indeed, our results depend on the routing mechanism only through the conditions (5)–(6). Thus, the results of this paper continue to hold for any routing mechanism for which the result of Proposition 5 holds. For example, it is possible to verify that a routing mechanism which splits flow equally among all available shortest paths also satisfies (5)–(6) (Johari et al., 2004).

3. Link stable equilibria

For the following discussion, recall that E is the set of all possible links: $E = \{(i,j)\colon i \neq j\}$. For $A \subseteq E$, we will use the notation s_A to denote bids and acceptance values of the strategy vector s restricted to the set of links A; i.e.,

$$s_A = \big(p_{ij}, q_{ij}, (i,j) \in A\big). \tag{7}$$

Notice that we have included both bids and acceptance values for all links in A; this means that components of the strategies of all endpoints of links in A are included in s_A. By an abuse of notation, for $(i,j) \in E$, we will write $s_{(i,j)} = s_{\{(i,j)\}} = (p_{ij}, q_{ij})$. Thus $s = (s_{(i,j)}, s_{E\setminus\{(i,j)\}})$.

We have the following definition, which describes the equilibrium concept to be studied in the rest of the paper.

Definition 6. A strategy vector s is *link stable* if for every possible link $(i,j) \in E$, the following conditions hold:

(1) For any $s'_{(i,j)} = (p'_{ij}, q_{ij})$:

$$R_i\big(s'_{(i,j)}, s_{E\setminus\{(i,j)\}}\big) \leqslant R_i(s). \tag{8}$$

(2) For any $s'_{(i,j)} = (p_{ij}, q'_{ij})$:

$$R_j\big(s'_{(i,j)}, s_{E\setminus\{(i,j)\}}\big) \leqslant R_j(s). \tag{9}$$

(3) For any $s'_{(i,j)} = (p'_{ij}, q'_{ij})$, *at least one* of the following holds:

$$R_i\big(s'_{(i,j)}, s_{E\setminus\{(i,j)\}}\big) \leqslant R_i(s), \tag{10}$$
$$R_j\big(s'_{(i,j)}, s_{E\setminus\{(i,j)\}}\big) \leqslant R_j(s). \tag{11}$$

We will also refer to such a strategy vector s as a *link stable equilibrium*.

The conditions in the preceding definition are meant to reflect the different possibilities for deviations between i and j. The condition (8) requires that i has no incentive to unilaterally change the bid p_{ij}, while (9) requires that j has no incentive to unilaterally change the acceptance value q_{ij}. The first two conditions are thus *individual rationality* conditions, as defined by Mas-Colell et al. (1995) and Jackson and Wolinsky (1996). The two conditions (10)–(11) require that no bilateral deviation exists where i and j simultaneously change both the bid and acceptance value for the link (i, j), and at least one of i or j improves its payoff.

We make the following useful remark regarding link stable equilibria.

Remark 7. At a link stable equilibrium we must have $p_{ij} = q_{ij} > 0$ whenever $(i, j) \in A(s)$, since otherwise node i could profitably deviate by reducing the bid p_{ij}. Of course, by definition $0 \leqslant p_{ij} < q_{ij}$ if $(i, j) \notin A(s)$; recall that we had assumed that $q_{ij} > 0$ for all (i, j).

Notice that given a strategy vector s and the graph $G(s)$, breaking a link requires only unilateral deviation (either i reduces p_{ij} or j increases q_{ij}). The following proposition summarizes this point formally: in checking whether a strategy vector s is link stable, Condition 3 of Definition 6 need not be checked for $(i, j) \in A(s)$.

Proposition 8. *Given a strategy vector s, suppose that*:

(1) *Conditions* 1 *and* 2 *of Definition* 6 *hold for every* (i, j); *and*
(2) *Condition* 3 *of Definition* 6 *holds for every* $(i, j) \notin A(s)$.

Then s is link stable.

Proof. Fix a link $(i, j) \in A(s)$. Since both (8) and (9) hold, we must have $p_{ij} = q_{ij} > 0$ (from Remark 7). Fix $s'_{(i,j)} = (p'_{ij}, q'_{ij})$, and let $s' = (s'_{(i,j)}, s_{E \setminus \{(i,j)\}})$. We check that not both (10) and (11) are violated, i.e., that i and j do not both improve their payoff at s'.

Suppose that both (10) and (11) are violated for some link (i, j). If $(i, j) \in A(s')$, then $R_i(s') + R_j(s') = R_i(s) + R_j(s)$; so nodes i and j cannot both improve their payoff under s'. Thus we must have $(i, j) \notin A(s')$. But if breaking the link (i, j) improves the payoff to node i (so that (10) is violated), then node i can achieve this outcome by a unilateral deviation as well, by choosing $p'_{ij} < q_{ij}$; this violates (8). A similar analysis follows for node j; so we conclude that since (8)–(9) have been assumed to hold, that Condition 3 of Definition 6 must hold as well. Thus s is link stable. □

Intuitively, under link stability, if $(i, j) \in A(s)$, then node i does not wish to decrease p_{ij} and node j does not wish to increase q_{ij}. Conversely, if $(i, j) \notin A(s)$, then both nodes could not be made better off by forming the link (at some bid/acceptance value combination). This situation is closely related to the definition of *pairwise stability with side payments* made by Jackson and Wolinsky (1996): unilateral deviations which break links, and bilateral deviations which form links, should not be profitable.

We are interested in characterizing graphs which can possibly be formed as link stable equilibria, captured by the following definition.

Definition 9. A graph G is *link stabilizable* if there exists a link stable strategy vector s such that $G(s) = G$.

In Section 3.1, we characterize link stabilizable graphs that satisfy an additional connectedness property with respect to the traffic matrix $\boldsymbol{B}$, and we show the set of link stable strategy vectors leading to a given graph form a “box” in the strategy space. We specify the boundaries of this box explicitly in terms of the given graph, the traffic matrix $\boldsymbol{B}$, the connectivity cost vector $\boldsymbol{\alpha}$, and the routing mechanism.

3.1. Connected link stabilizable graphs

We first need to characterize a particular class of graphs, captured by the following two definitions.

Definition 10. Given a matrix $\boldsymbol{B}$, a directed graph $G = (V, A)$ is $\boldsymbol{B}$*-connected* if $\mathcal{P}_A(i, j) \neq \emptyset$ for all $i, j \in V$ such that $b_{ij} > 0$.

The preceding definition captures the idea that if $b_{ij} > 0$ then i wishes to have a path available leading to j. A graph which is $\boldsymbol{B}$-connected ensures that no node suffers a connectivity cost, given by the second term in (2).

In general, a $\boldsymbol{B}$-connected path may contain redundant links. Instead, we will show that link stability is closely tied to *minimal* $\boldsymbol{B}$-connected graphs, defined as follows.

Definition 11. Given a matrix $\boldsymbol{B}$, a directed graph $G = (V, A)$ is *minimally* $\boldsymbol{B}$*-connected* if the following two conditions hold:

(1) G is $\boldsymbol{B}$-connected.
(2) For each $(i, j) \in A$, the graph $(V, A \setminus \{(i, j)\})$ is not $\boldsymbol{B}$-connected.

Thus, in a minimal $\boldsymbol{B}$-connected graph, removal of any link destroys $\boldsymbol{B}$-connectivity. Intuitively, given the form of the cost in (2), together with the result of Proposition 5 (monotonicity), nodes have no incentive to form additional links once a graph has become $\boldsymbol{B}$-connected. In light of this insight, we start by relating $\boldsymbol{B}$-connected link stabilizable graphs to minimally $\boldsymbol{B}$-connected graphs.

Proposition 12. *If $G = (V, A)$ is $\boldsymbol{B}$-connected and link stabilizable, then G is minimally $\boldsymbol{B}$-connected.*

Proof. Suppose that s is link stable and $G(s) = (V, A(s)) = G$, but that there exists a link $(i, j) \in A(s)$ which can be removed from $G(s)$ while leaving the graph $\boldsymbol{B}$-connected. We must have $q_{ij} = p_{ij} > 0$, since s is link stable. Consider any unilateral deviation by node i which breaks link (i, j). This does not increase the flow passing through node i (by Proposition 5), and strictly reduces the total payment made by node i. Thus the payoff to node i strictly increases with this deviation, and hence s could not have been link stable; so, if s is link stable, $G(s)$ must be minimally $\boldsymbol{B}$-connected. □

The key result of this section is Theorem 13, which characterizes the set of link stable strategy vectors that result in a given $\boldsymbol{B}$-connected graph. In the discussion to follow, given a graph $G = (V, A)$, we denote the set of strategies under which A is stable by $\mathcal{S}(A)$, that is:

$$\mathcal{S}(A) = \{s\colon\ s \text{ is link stable and } A(s) = A\}.$$

We will investigate two main issues regarding the set $\mathcal{S}(A)$. First, we will determine when the set $\mathcal{S}(A)$ is nonempty; this will yield conditions under which link stable equilibria exist. In addition, we will determine the exact shape of the set $\mathcal{S}(A)$. In particular, we will show that $\mathcal{S}(A)$ can be represented as a "box" in the strategy space, effectively decoupling components of the strategy vector across the links (i, j). Formally, in the next theorem, we show that given a $\boldsymbol{B}$-connected graph $G = (V, A)$, we may calculate the boundaries of the link stable set $\mathcal{S}(A)$. Furthermore, if the inequalities defining this set have empty intersection, then we will conclude that no link stable vector exists.

Given a set of links A, we will need the following operators, defined for each $i \in V$ and $(j, k) \in E$:

$$\begin{aligned}
&F_i^-\big(A, (j,k)\big) = f_i(A) - f_i\big(A \setminus \{(j,k)\}\big),\\
&F_i^+\big(A, (j,k)\big) = f_i\big(A \cup \{(j,k)\}\big) - f_i(A),\\
&W_i^-\big(A, (j,k)\big) = \alpha_i \sum_{j \in w_i(A\setminus\{(j,k)\})} b_{ij} - \alpha_i \sum_{j \in w_i(A)} b_{ij}.
\end{aligned}$$

Thus the operators F_i^- and F_i^+ give the change in flow through node i when a link (j,k) is removed or added, respectively; and W_i^- gives the change in the connectivity cost when link (j,k) is removed from A. Note that W_i^- is nonnegative; by Proposition 5, if $i = j$ or $i = k$, then $F_i^-(A, (j,k))$ and $F_i^+(A, (j,k))$ are nonnegative as well.

We now make the following definitions:

$$m_{ij}(A) = \begin{cases} F_j^-(A, (i,j))c_j, & (i,j) \in A,\\ 0, & (i,j) \notin A, \end{cases} \tag{12}$$

$$M_{ij}(A) = \begin{cases} W_i^-(A, (i,j)) - F_i^-(A, (i,j))c_i, & (i,j) \in A,\\ F_j^+(A, (i,j))c_j, & (i,j) \notin A. \end{cases} \tag{13}$$

The following theorem shows that $m_{ij}(A)$ and $M_{ij}(A)$ exactly characterize $\mathcal{S}(A)$; thus, given a set of links A we can explicitly compute $\mathcal{S}(A)$.

Theorem 13. *Suppose $G = (V, A)$ is $\boldsymbol{B}$-connected. Then, $s \in \mathcal{S}(A)$ if and only if the following conditions hold*:

$$m_{ij}(A) \leqslant p_{ij} \leqslant M_{ij}(A), \quad \text{for all } (i,j), \tag{14}$$

$$0 < q_{ij} = p_{ij}, \quad (i,j) \in A, \tag{15}$$

$$0 \leqslant p_{ij} < q_{ij}, \quad (i,j) \notin A. \tag{16}$$

Furthermore, the following are equivalent:

(1) *G is link stabilizable (i.e., $\mathcal{S}(A)$ is nonempty)*;
(2) *$m_{ij}(A) \leqslant M_{ij}(A)$ for all (i,j)*;
(3) *G is minimally $\boldsymbol{B}$-connected and $m_{ij}(A) \leqslant M_{ij}(A)$ for all $(i,j) \in A$*;

(4) *G is minimally $\boldsymbol{B}$-connected and for all $(i, j) \in A$:*

$$F_i^-\big(A, (i, j)\big)c_i + F_j^-\big(A, (i, j)\big)c_j \leqslant W_i^-\big(A, (i, j)\big). \tag{17}$$

Proof. Let s be a strategy vector such that (15)–(16) hold; then of course $A(s) = A$. We will characterize the conditions that are necessary and sufficient for (8)–(11) to hold.

Step 1: For $(i, j) \in A$, (8) holds if and only if $p_{ij} \leqslant M_{ij}(A)$. We need only consider unilateral deviations where node i breaks link (i, j), as these are the only deviations which can be profitable for node i (since $p_{ij} = q_{ij}$). Breaking the link (i, j) causes a reduction of p_{ij} in the payment by i, and a reduction of flow into i by $F_i^-(A, (i, j))$; but i's connectivity cost increases by $W_i^-(A, (i, j))$. Thus, (8) holds if and only if:

$$p_{ij} + F_i^-\big(A, (i, j)\big)c_i \leqslant W_i^-\big(A, (i, j)\big).$$

This is precisely the condition that $p_{ij} \leqslant M_{ij}(A)$.

Step 2: For $(i, j) \in A$, (9) holds if and only if $m_{ij}(A) \leqslant p_{ij}$. Again, we need only consider unilateral deviations where node j breaks link (i, j), since these are the only deviations which can be profitable for node j. Breaking the link (i, j) causes a reduction of p_{ij} in the revenue of j, and a reduction of flow into j of $F_j^-(A, (i, j))$. Thus (9) holds if and only if:

$$F_j^-\big(A, (i, j)\big)c_j \leqslant p_{ij}.$$

This is precisely the condition $m_{ij}(A) \leqslant p_{ij}$.

Step 3: For $(i, j) \notin A$, (8) and (10) hold. Since G is $\boldsymbol{B}$-connected, if the link (i, j) is formed, then the cost to node i does not decrease (since the connectivity cost is zero, and by Proposition 5, the flow through node i does not decrease). On the other hand, if the link (i, j) is formed then the payment made by node i strictly increases; thus the payoff to node i can only decrease if link (i, j) is formed. Since this is the only deviation with respect to the link (i, j) which changes the payoff of node i, (8) and (10) must hold.

Step 4: For $(i, j) \notin A$, (9) holds if and only if $p_{ij} \leqslant M_{ij}(A)$. We need only consider the unilateral deviation by node j which forms the link (i, j). This leads to an increase of p_{ij} in the revenue to node j, and an increase of $F_j^+(A, (i, j))$ to the flow into j. Thus (9) holds if and only if:

$$p_{ij} \leqslant F_j^+\big(A, (i, j)\big)c_j.$$

This is precisely the condition $p_{ij} \leqslant M_{ij}(A)$.

Step 5: A strategy vector s with $A(s) = A$ is link stable if and only if the conditions (14)–(16) hold. By Steps 1–4 and Proposition 8, we have shown that if s satisfies (14) in addition to (15)–(16), then $s \in \mathcal{S}(A)$. Conversely, if $s \in \mathcal{S}(A)$, then by Remark 7 we know that (15)–(16) hold. Thus from Steps 1–4, together with the fact that $m_{ij}(A) = 0$ for $(i, j) \notin A$, we conclude that (14) holds.

Step 6: Conditions 1–4 are equivalent. We have already shown that $\mathcal{S}(A)$ is nonempty, and G is link stabilizable, if and only if $m_{ij}(A) \leqslant M_{ij}(A)$ for all (i, j). Thus Condition 1 and Condition 2 are equivalent. Condition 3 and Condition 4 are equivalent by definition of $m_{ij}(A)$ and $M_{ij}(A)$. Furthermore, Condition 1 implies Condition 3, by Proposition 12 and Condition 2. For the reverse implication, we only need to check that if G is $\boldsymbol{B}$-connected, then $m_{ij}(A) \leqslant M_{ij}(A)$, for all $(i, j) \notin A$. Indeed, if $(i, j) \notin A$, then $m_{ij}(A) = 0$, and $M_{ij}(A) = F_j^+(A, (i, j))c_j \geqslant 0$ by Proposition 5. □

Remark 14. The proof of the preceding theorem provides an interpretation of the functions $m_{ij}(A)$ and $M_{ij}(A)$ when $G = (V, A)$ is $\boldsymbol{B}$-connected. For $(i, j) \in A$, $m_{ij}(A)$ is the minimum amount that node j is willing to accept for the link (i, j); and $M_{ij}(A)$ is the maximum amount that node i is willing to pay for link (i, j). For $(i, j) \notin A$, $m_{ij}(A)$ is the maximum amount node i is willing to pay for link (i, j), i.e., $m_{ij}(A) = 0$; and $M_{ij}(A)$ is the minimum amount node j is willing to accept for link (i, j).

The previous theorem shows that link stable payments p_{ij} lie in a "box" determined by the functions m_{ij} and M_{ij}, which depend only on the graph. Of course, from Proposition 12, the set $\mathcal{S}(A)$ is empty if G is $\boldsymbol{B}$-connected but not minimally $\boldsymbol{B}$-connected.

We conclude by noting that two extreme points of the "box" defined by (14)–(16) are also the equilibrium points of certain two-stage simultaneous games. In the first game, each node j first sets the acceptance levels q_{ij} for $i \in V$; then, after observing the set of all possible q_{ij}, each node i chooses whether or not to bid $p_{ij} = q_{ij}$ for $j \in V$. Informally, we expect that since nodes choose acceptance levels first, the second stage payments will have to be as *large* as possible, i.e., $p_{ij} = M_{ij}(A)$. In the second game, each node i will first set p_{ij} for $j \in V$; then each node j will decide whether or not to accept the bid p_{ij} for $i \in V$. Informally, we expect that in such a game the payment for each link $(i, j) \in A$ will be as *small* as possible, i.e., $p_{ij} = m_{ij}(A)$. For formal statements of these results, see Johari et al. (2004).

4. Models with uniform all-to-all traffic

In this section, we study the structure of the graphs that might result when each node wishes to be connected to all other nodes. Such a situation may arise when all entries of the traffic matrix $\boldsymbol{B}$ are very large (in particular, if the product of α_i and $\boldsymbol{B}_{ij}$ is much larger than the product of c_i and f_i for every i, j), so that nodes care most about being connected to all other nodes; if we view the nodes as nodes of a goods delivery network, then the model of this section applies when the paramount goal of the nodes is to ensure that goods can be delivered from any origin to any destination. By considering the tradeoff to each node between the connectivity cost and the cost of transit traffic, we are able to investigate which graphs emerge as strongly connected link stabilizable graphs. We note that a model of free trade networks where all nodes are symmetric has been studied by Furusawa and Konishi (2002).

In this section we focus on uniform all-to-all traffic, i.e., we assume that $b_{ij} = 1$ for all $i \neq j$; we also assume that $\alpha_i = \alpha$ for all i, for some positive constant α (which we interpret as the cost to each node per disconnected destination). We assume that $c_i = c$ for all i. Finally, we assume throughout that $V = \{1, \ldots, n\}$, so that in particular $|V| = n$. Note that in this setting, since $b_{ij} = 1$, the flow functions f_i take a particularly simple form: given a set of links A, $f_i(A)$ is simply the number of shortest paths which contain node i.

As in earlier parts of the paper, we will be interested in graphs which are minimally $\boldsymbol{B}$-connected. Directed graphs which are minimally $\boldsymbol{B}$-connected, when the traffic matrix $\boldsymbol{B}$ is defined as in the previous paragraph, are called *minimal strongly connected directed graphs*, or MSDs (Harary, 1968; Cunningham, 1982; Grotschel, 1979); for completeness, we include a formal definition.

Definition 15. A directed graph $G = (V, A)$ is a *minimal strongly connected directed graph* (*MSD*) if the following two conditions hold:

(1) G is strongly connected (i.e., $\mathcal{P}_A(i, j) \neq \emptyset$ for all i, j).
(2) For each $(i, j) \in A$, the graph $(V, A \setminus \{(i, j)\})$ is not strongly connected (i.e., there exist k, ℓ such that $\mathcal{P}_{A\setminus\{(i,j)\}}(k, \ell) = \emptyset$).

We expect that for strongly connected link stable equilibria to result, the ratio α/c must be sufficiently large; that is, the cost of being disconnected must be sufficiently large relative to the cost of carrying additional traffic. To formally explore this relationship, we start with two examples in Section 4.1, and then continue in Section 4.2 to characterize link stabilizable graphs for different values of the ratio α/c. In each of these cases, given an MSD G, we will apply Theorem 13 to find a necessary condition on α/c to guarantee that G is link stabilizable. In Example 16 we consider a directed cycle. We prove in Theorem 19 that the directed cycle is the first link stabilizable graph to emerge as the ratio α/c increases. An example where the traffic is concentrated in bottleneck nodes is described in Example 17. In this example the ratio of α/c required for strong connectivity scales in a quadratic manner with the number of nodes; as we show later, in Theorem 21, this is the worst possible scaling rate.

In Section 4.3, we consider the "social welfare" of link stabilizable graphs. We define social cost as the sum of costs incurred by all nodes; a socially optimal graph is then one which minimizes this aggregate cost. We first show the trivial result that the complete graph on n nodes is socially optimal within the set of strongly connected graphs. We then establish that a directed cycle has the worst social cost among all strongly connected link stabilizable graphs, and that the bidirectional star has the best social cost among all strongly connected link stabilizable graphs. The ratio of social costs between the worst and best strongly connected link stabilizable graphs is shown to increase linearly in n; thus as the number of nodes increases, the welfare loss of link stabilizable graphs may become arbitrarily large.

4.1. Examples

We begin with two examples which highlight the importance of the ratio α/c in determining the set of link stabilizable graphs. The key tool in evaluating whether or not a graph $G = (V, A)$ is link stabilizable will be Condition 4 of Theorem 13; specifically, we will be checking whether the relation (17) holds for each link $(i, j) \in A$. We repeat (17) here for easier reference:

$$F_i^-\big(A, (i, j)\big)c_i + F_j^-\big(A, (i, j)\big)c_j \leqslant W_i^-\big(A, (i, j)\big).$$

Example 16 (*Directed cycle*). Let the set of links be $A = \{(1, 2), (2, 3), \ldots, (n-1, n), (n, 1)\}$; in this case we say the graph $G = (V, A)$ is a *directed cycle*. See Fig. 2 for the case $n = 4$.

The graph $G = (V, A)$ is clearly an MSD. By Theorem 13, and in particular (17), G will be link stabilizable if $F_i^-(A, (i, i+1))c + F_{i+1}^-(A, (i, i+1))c \leqslant W_i^-(A, (i, i+1))$ for all i (where the nodes are numbered modulo n). Since the graph is completely symmetric, it suffices to consider $i = 1$. We first note that $f_1(A) = (n-1) + (n-2) + \cdots + 1 = \frac{n(n-1)}{2}$; that is, node 1 receives $n-1$ units of traffic from node n, $n-2$ units of traffic from node $n-1$, and so on. If the link $(1, 2)$ is cut, then clearly node 2 no longer receives any traffic, so that:

$$F_2^-\big(A, \{(1, 2)\}\big) = \frac{n(n-1)}{2}.$$

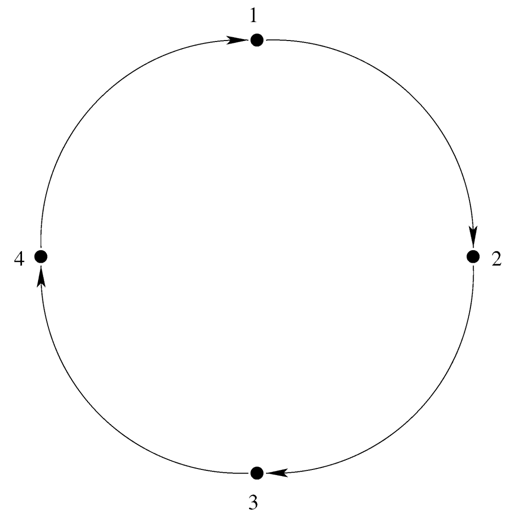

Fig. 2. A cycle on 4 nodes; see Example 16.

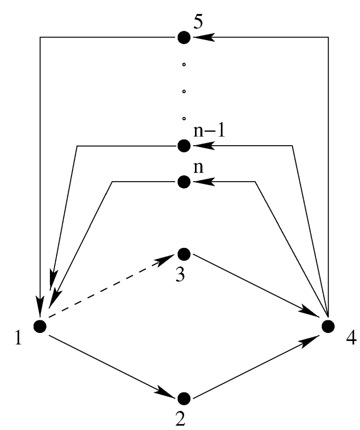

Fig. 3. A graph which requires $\alpha/c \geqslant n^2 - 5n + 10$ for link stability; see Example 17.

Similarly, node 1 only receives the $n-1$ units of traffic destined for it from the other nodes; thus:

$$F_1^-\big(A, \{(1,2)\}\big) = \frac{n(n-1)}{2} - (n-1) = \frac{(n-1)(n-2)}{2}.$$

Finally, cutting the link $(1,2)$ disconnects node 1 from all other nodes, so that $W_1^-(A,(1,2)) = (n-1)\alpha$. Thus (17) is satisfied if and only if $n-1 \leqslant \alpha/c$. Since all links have the same behavior, we have shown the directed cycle is link stabilizable if and only if $\alpha/c \geqslant n-1$.

Example 17 (*Concentrated traffic*). The previous example required that α/c scale linearly in the number of nodes n. In this example, we will demonstrate that there are structures for which α/c must scale quadratically in n.

Consider the example of Fig. 3. The set of links is defined by:

$$A = \{(1,2),(1,3),(2,4),(3,4)\} \cup \{(4,i),(i,1)\colon 5 \leqslant i \leqslant n\}.$$

The graph $G = (V, A)$ is then an MSD.

We will assume that the routing $\prec$ ranks routes passing through node 3 higher than routes passing through node 4; thus, for example, when routing from node 1 to node 4, all traffic flows over route 134 rather than 124. Consider the link $(1, 3)$, drawn with a dashed line in Fig. 3. Removing this link disconnects node 1 from exactly one other node, namely node 3; thus $W_1^-(A, (1, 3)) = \alpha$. Furthermore, it reduces the flow through node 1 by exactly $n - 2$ units, since all other nodes stop sending to node 3, so $F_1^-(A, (1, 3)) = n - 2$.

We now ask: how much is the flow through node 3 reduced when $(1, 3)$ is removed? There are three types of flow which are no longer sent to node 3: (1) traffic destined for node 3, equal to $n - 1$ units; (2) traffic destined for node 4, from all nodes other than 2 and 3, equal to $n - 3$ units; and (3) traffic destined for nodes $5, 6, \dots, n$, from nodes $1, 5, 6, \dots, n$, equal to $(n - 4)^2$ units. Thus the total reduction in flow through node 3 is $F_3^-(A, (1, 3)) = (n - 4)^2 + 2n - 4 = n^2 - 6n + 12$ units.

By (17), for G to be link stabilizable, we must have $\alpha/c \geqslant n^2 - 5n + 10$; in other words, α/c scales at least as fast as $O(n^2)$ as the number of nodes increases.

4.2. Strongly connected link stabilizable graphs

The examples of the previous section give different relationships between α and c, depending on the structure of the graph. In this section, we characterize strongly connected link stabilizable graphs, as a function of the ratio α/c.

Before continuing, we prove a simple bound on the number of links in an MSD. We will say a link (i, j) in the graph $G = (V, A)$ is *bidirectional* if both $(i, j) \in A$ and $(j, i) \in A$, so the link is present in both directions between i and j. A *bidirectional tree* is then an MSD such that every link is bidirectional. It is straightforward to verify that any bidirectional tree may be constructed by first forming an undirected spanning tree on the set of nodes V, and then replacing an undirected link joining i and j by two directed links, (i, j) and (j, i). Note that a bidirectional tree has exactly $2(n - 1)$ links, and a directed cycle has exactly n links. We have the following lemma.

Lemma 18. *Suppose that $G = (V, A)$ is an MSD. Then*:

$$n \leqslant |A| \leqslant 2(n - 1).$$

The lower bound is achieved if and only if G is a directed cycle, and the upper bound is achieved if and only if G is any bidirectional tree.

Proof. Let $G = (V, A)$ be an MSD. Since G is strongly connected, every node has at least one outgoing link; thus $|A| \geqslant n$. We now show that the lower bound is achieved if and only if G is a directed cycle. Suppose that $|A| = n$. Since G is strongly connected, every node has at least one incoming link, and at least one outgoing link; since $|A| = n$, we conclude every node has *exactly* one incoming link and one outgoing link. Thus G is a disjoint collection of cycles. Again applying the fact that G is strongly connected, we conclude G is in fact a single directed cycle.

The remainder of the proof addresses the upper bound. Fix a node i, and consider two sets of links T_1 and T_2. We define $T_1 \subset A$ as the set of links in a directed spanning tree connecting from i to all other nodes in the graph; T_1 contains exactly $n - 1$ links. Similarly, we let $T_2 \subset A$ be the set of links in a directed spanning tree connecting to i from all other nodes in the graph;

again, T_2 contains exactly $n-1$ links. Observe that the graph $(V, T_1 \cup T_2)$ is strongly connected. Since we assumed that G was an MSD, we must in fact have $T_1 \cup T_2 = A$, so that $|A| \leqslant 2(n-1)$.

We now show by induction that the upper bound is achieved if and only if G is a bidirectional tree. Let $G = (V, A)$ be an MSD on n nodes. Then there exists at least one node $i \in V$ with exactly one incoming link (say (j, i)), and exactly one outgoing link (say (i, k), where possibly $j = k$); see Berge (1973, Corollary 3.1) and Grotschel (1979). Let $V' = V \setminus \{i\}$, and define a set of links A' as follows:

$$A' = \begin{cases} A \setminus \{(j,i),(i,k)\}, & \text{if } (V', A \setminus \{(j,i),(i,k)\}) \text{ is strongly connected,} \\ (A \setminus \{(j,i),(i,k)\}) \cup \{(j,k)\}, & \text{otherwise.} \end{cases}$$

Thus we remove the two links (j,i) and (i,k), and add the link (j,k) if necessary to maintain connectivity. With this definition, the new graph $G' = (V', A')$ is an MSD. In the first case in the definition of A', G' is a subgraph of G and thus must be an MSD. In the second case, we have replaced the path jik in A by the link (j,k) in A'; as a result, since G is strongly connected, G' must be strongly connected. Furthermore, if removal of any link other than (j,i) or (i,k) from A disconnects two nodes in G, then the same is true in G'. Finally, removing the link (j,k) from A' leaves a graph which is not strongly connected; thus G' must be an MSD in the second case as well.

Now suppose that $|A| = 2(|V|-1) = 2(n-1)$, so that the upper bound is met. Then again since G' is an MSD, we must have $|A'| \leqslant 2(|V'|-1) = 2(n-2) = |A| - 2$. But this is only possible if $A' = A \setminus \{(j,i),(i,k)\}$, in which case $|A'| = 2(n-2)$. By the inductive hypothesis, G' is in fact a bidirectional tree. We distinguish two cases. If $j = k$, then G is formed by starting with the bidirectional tree G' and adding one "leaf", in which case G is also a bidirectional tree. If on the other hand $j \neq k$, then consider the path from j to k in the bidirectional tree G'; denote the last link on this path by (ℓ, k). We now define a new set of links $\hat{A} = A \setminus \{(\ell, k)\}$; i.e., we remove the link (ℓ, k) from G. The graph $(V, \hat{A})$ remains strongly connected, because from any node we can get to j (by traveling within the bidirectional tree G'), and then follow the links (j,i), (i,k) to get to k. But this would contradict the assumption that G is an MSD. This completes the induction and the proof of the lemma. □

The next theorem reveals that the directed cycle (cf. Example 16) is the first strongly connected link stabilizable graph to emerge as α/c increases.

Theorem 19. *A strongly connected link stabilizable graph exists if and only if $\alpha/c \geqslant n-1$. Furthermore, if $\alpha/c = n-1$, a strongly connected link stabilizable graph must be a directed cycle.*

Proof. We first show that if $\alpha/c < n-1$, no strongly connected link stabilizable graphs exist. By Proposition 12, we may restrict our attention to MSDs. We will use the structure of MSDs to show that a link stable strategy vector can exist only if $\alpha/c \geqslant n-1$.

Suppose that we are given an MSD $G = (V, A)$. By the *ear decomposition* for MSDs (Grotschel (1979); Berge, 1973, Corollary 3.1), there exists a strongly connected subgraph (V', A') of G, together with $a, b \in V'$ with the property that $A \setminus A'$ consists of the links of the path $ai_1i_2 \cdots i_kb$, where $V \setminus V' = \{i_1, \ldots, i_k\}$. (Note that we may have $a = b$.) Here $k \geqslant 1$, and $|V'| = n-k \geqslant 1$. The path $ai_1i_2 \cdots i_kb$ is an "ear," whose addition to the strongly connected subgraph (V', A') produces the entire graph G; see Fig. 4. (Observe that since G is an MSD, the graph (V', A') is also an MSD.) We will focus on the link (a, i_1), and show that stability for this link requires at least $\alpha/c \geqslant n-1$.

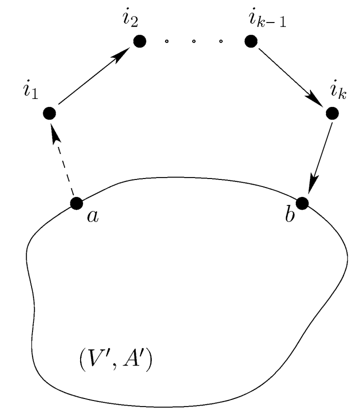

Fig. 4. An ear decomposition of an MSD; see the proof of Theorem 19.

We note that there are two types of traffic across the link (a, i_1): traffic terminating at i_d, for $1 \leqslant d \leqslant k$; and traffic terminating at nodes in V'. We now let X_{ai_1} represent the traffic terminating at nodes in V', which would no longer be sent to a if the link (a, i_1) were removed from G. Since Proposition 5 holds, we have $X_{ai_1} \geqslant 0$. Furthermore, X_{ai_1} is exactly equal to the traffic terminating at nodes in V', which is sent across the link (a, i_1).

We now apply (17) from Theorem 13. If the link (a, i_1) is cut, an amount of traffic equal to X_{ai_1} is no longer sent to a. The traffic from the other $n - k - 1$ nodes of V' to the k nodes $i_1, \ldots, i_k$ is also no longer received by a. Finally, a no longer receives: 1 unit of traffic originating at i_2 and destined for i_1; 2 units of traffic originating at i_3, and destined for i_1, i_2; and so on, up to the $k - 1$ units of traffic originating at i_k, and destined for $i_1, \ldots, i_{k-1}$. This is a total of $k(k - 1)/2$ units of traffic. Thus we have:

$$F_a^-\big(A, (a, i_1)\big) = \left(X_{ai_1} + (n - k - 1)k + \frac{k(k-1)}{2}\right)c.$$

On the other hand, cutting the link (a, i_1) disconnects a from exactly k nodes (since (V', A') is strongly connected). Thus we have:

$$W_a^-\big(A, (a, i_1)\big) = k\alpha.$$

We now consider the amount of traffic no longer sent to i_1 if the link (a, i_1) is cut. We divide this traffic into a portion which travels to V', and a portion which terminates at $i_1, \ldots, i_k$. The first term is exactly X_{ai_1}; and by the same reasoning as for node a, the second term is $(n - k)k + k(k - 1)/2$. (The first term is now $(n - k)k$, because (a, i_1) carries traffic from *all* nodes of V' destined for nodes in $i_1, \ldots, i_k$.) Thus we have:

$$F_{i_1}^-\big(A, (a, i_1)\big) = \left(X_{ai_1} + (n - k)k + \frac{k(k-1)}{2}\right)c.$$

As a result, upon dividing by k and c the relation (17) reduces to:

$$\left(\frac{2X_{ai_1}}{k} + 2(n - k - 1) + k\right) \leqslant \alpha/c.$$

Since $|V'| = n - k \geqslant 1$, we have $k \leqslant n - 1$. Since we also have $X_{ai_1} \geqslant 0$, the left hand side must be at least $n - 1$, achieved with equality if and only if $X_{ai_1} = 0$, and $k = n - 1$; it follows that the only MSD satisfying $X_{ai_1} = 0$ and $k = n - 1$ is a directed cycle. We conclude that a link stable strategy vector exists if and only if $n - 1 \leqslant \alpha/c$; and if $\alpha/c = n - 1$, every link stabilizable MSD is a directed cycle. □

The following theorem shows that bidirectional trees are link stabilizable if and only if $\alpha/c \geqslant 2n - 3$.

Theorem 20. *Suppose that G is a bidirectional tree. Then G is link stabilizable if and only if $\alpha/c \geqslant 2n - 3$.*

Proof. Suppose that $T = (V, A)$ is a bidirectional tree, and fix two nodes i and j which are connected by a bidirectional link in this tree. Suppose we remove the two links $(i, j) \in A$ and $(j, i) \in A$; then T decomposes into two connected components $T_1 = (V_1, A_1)$ and $T_2 = (V_2, A_2)$, such that $i \in V_1$ and $j \in V_2$. Let $|V_1| = k$, so that $|V_2| = n - k$.

If the link (i, j) is broken, then i is disconnected from $n - k$ nodes, so $W_i^-(A, (i, j)) = (n - k)\alpha$. On the other hand, traffic from the $k - 1$ nodes in $V_1 \setminus \{i\}$ destined for V_2 would no longer be routed through i, resulting in a reduction in traffic through i of $(k - 1)(n - k)$; thus $F_i^-(A, (i, j)) = (k - 1)(n - k)$.

Now consider node j. If the link (i, j) is broken, the nodes in V_1 will no longer send traffic destined for V_2 through node j, resulting in a traffic reduction of $k(n - k)$; thus $F_j^-(A, (i, j)) = k(n - k)$.

We conclude that (17) is satisfied if and only if:

$$2k - 1 \leqslant \alpha/c.$$

A symmetric analysis follows for the link (j, i). In any bidirectional tree, the largest possible value for k is $n - 1$; so this requirement is met for all k if and only if $\alpha/c \geqslant 2n - 3$, as desired. □

Finally, the following theorem shows that if $\alpha/c > 2n^2 - 2n$, then only MSDs are link stabilizable; note the relationship to Example 17, where $\alpha/c \approx n^2$ for link stability of a particular MSD to emerge.

Theorem 21. *If $\alpha/c > 2n^2 - 2n$, then G is link stabilizable if and only if G is an MSD.*

Proof. Suppose that we are given a link stabilizable graph $G = (V, A)$. If G is strongly connected, then we know it must be an MSD, by Proposition 12. Suppose, then, that G is not strongly connected, and choose (i, j) such that $\mathcal{P}_A(i, j) = \emptyset$; in particular, $(i, j) \notin A$.

Consider adding the link (i, j) to A. Since the total traffic in the network is $n(n - 1)$, the traffic through nodes i and j can increase by no more than $n(n - 1)$; however, adding the link (i, j) connects i to j, and thus reduces the connectivity cost incurred by i by at least $\alpha > 2n(n - 1)c$. Thus node i would be willing to bid at least $n(n - 1)c$ for the link (i, j), and node j would be willing to accept at least $n(n - 1)c$ as a payment for the link (i, j). We conclude that a profitable deviation exists for the pair i and j, so G could not have been link stabilizable.

Conversely, suppose we are given an MSD $G = (V, A)$. To show G is link stabilizable, it suffices to check that (17) holds for all $(i, j) \in A$, by Theorem 13; this follows by the same reasoning as in the previous paragraph. We conclude, therefore, that G is link stabilizable. □

We summarize the relationship between α/c and link stabilizable graphs in Table 1; in particular, the dependence of α/c on n is the critical determinant of the space of link stabilizable graphs.

Table 1
Link stabilizable MSDs for different values of α/c

Region	Link stabilizable MSDs
$\alpha/c < n-1$	No link stabilizable MSDs
$\alpha/c = n-1$	Directed cycles are the only link stabilizable MSDs
$n-1 < \alpha/c < 2n-3$	Not all link stabilizable graphs are MSDs; none is a bidirectional tree
$2n-3 \leqslant \alpha/c \leqslant 2n^2 - O(n)$	Not all link stabilizable graphs are MSDs; all bidirectional trees are link stabilizable
$\alpha/c > 2n^2 - 2n$	All MSDs are link stabilizable; no other link stabilizable graphs

Recall that a strongly connected link stabilizable graph must be an MSD. Furthermore, if an MSD is link stabilizable, it remains so for larger values of α/c. In the fourth row of the table, the notation $O(n)$ is taken to mean that $O(n)/n$ is bounded below by zero and above by a constant. We note more formally that there exist MSDs for which $\alpha/c \geqslant 2n^2 - 23n + 77$ is required for link stability (Johari et al., 2004).

To complete the picture, we note that the result of Theorem 21 is essentially tight to leading order, in the sense that there exists an MSD which is link stabilizable only if $\alpha/c \geqslant 2n^2 - 23n + 77$; this construction is similar to Example 17, and is developed by Johari et al. (2004).

4.3. Social welfare

In this section we will consider the properties of strongly connected link stabilizable graphs from a "social welfare" point of view. The next definition interprets the social cost of a graph as the total cost incurred by all nodes, both due to routing traffic as well as a lack of connectivity. Similar notions of social welfare have been considered for other network formation models by Fabrikant et al. (2003) and Anshelevich et al. (2003).

Definition 22. Given a graph $G = (V, A)$, the *social cost* $S(G)$ of the graph G is the sum of the costs incurred by all the nodes when the graph G is used to route traffic:

$$S(G) = \sum_{i \in V} C_i(A) = \sum_{i \in V} \bigl(cf_i(A) + \alpha\bigl|w_i(A)\bigr|\bigr). \tag{18}$$

Note that since the sum of payments between all the nodes is zero, $S(G)$ may be alternatively interpreted as the negative of the sum of the payoffs of all the nodes when the graph G is formed. Our interest will be in evaluating the social cost of strongly connected graphs, in which case $\sum_{i \in V} \alpha|w_i(A)| = 0$; thus the social cost reduces to just the cost of routing traffic.

The following simple proposition shows that among all strongly connected graphs, there is a unique choice which minimizes social cost; recall that E is the set of all possible links on the node set V.

Proposition 23. *The complete graph $G_C = (V, E)$ on the node set V has social cost $S(G_C) = n(n-1)c$; and for any other strongly connected graph G on V, $S(G) > S(G_C)$.*

Proof. Let $G = (V, A)$ be any strongly connected graph. Fix a node i; then since G is strongly connected, i receives one unit of traffic from each other node, so that $f_i(A) \geqslant (n-1)c$. Summing over all $i \in V$, we have $S(G) \geqslant n(n-1)c$. The bound is met with equality if and only if $f_i(A) = (n-1)c$ for all $i \in V$; but this is only possible if no node i receives traffic destined for a node other than i. Since G is strongly connected, it must be the complete graph, G_C. □

The previous proposition establishes that among all strongly connected graphs, the lowest possible social cost is $n(n-1)c$. The next proposition considers only link stabilizable strongly connected graphs (which is equivalent to the set of MSDs, by Proposition 12), and shows that the bidirectional star minimizes social cost among all strongly connected link stabilizable graphs. We recall that two graphs G and G' are *isomorphic* if there exists a permutation of the node labels under which they have the same set of links (Harary, 1968).

Proposition 24. *The bidirectional star $G_S = (V, A_S)$, given by $A_S = \{(1,i), (i,1)\colon i = 2, \ldots, n\}$, has social cost $S(G_S) = 2(n-1)^2c$; and for any other strongly connected link stabilizable graph G which is not isomorphic to G_S, $S(G) > S(G_S)$.*

Proof. Since the bidirectional star is an MSD, the only costs are traffic routing costs. Consider any node $i > 1$; then i only receives traffic terminating at i, so $f_i(A_S) = n-1$. Thus $\sum_{i>1} f_i(A) = (n-1)^2$. Node 1 receives, from each other node i, all outgoing traffic from node i; so $f_1(A_S) = (n-1)^2$. Thus $S(G_S) = 2(n-1)^2c$.

Let $G = (V, A)$ be a strongly connected link stabilizable graph; then G is an MSD, by Proposition 12. Furthermore, from Lemma 18, $|A| \leqslant 2(n-1)$. To complete the proof, we will measure the social cost of G by summing the number of nodes traversed on the shortest path between each source-destination pair. Formally, define $h_{ij}(A)$ as the number of nodes traversed on the shortest path from i to j using links in A; we include the terminal node j in computing $h_{ij}(A)$, but exclude the initial node i. Then observe that:

$$S(G) = c\sum_{i\in V} f_i(A) = c\sum_{i\in V}\sum_{j:j\neq i} h_{ij}(A).$$

Now since $|A| \leqslant 2(n-1)$, at most $2(n-1)$ source-destination pairs (i, j) have $h_{ij}(A) = 1$; the remainder must have $h_{ij}(A) \geqslant 2$. Since there are a total of $n(n-1)$ source-destination pairs, we have:

$$\begin{aligned} S(G) &= c\sum_{i\in V}\sum_{j:j\neq i} h_{ij}(A) \\ &\geqslant \big(2(n-1) + 2\big[n(n-1) - 2(n-1)\big]\big)c = 2(n-1)^2c = S(G_S). \end{aligned}$$

Furthermore, since $|A| = 2(n-1)$ if and only if G is a bidirectional tree (from Lemma 18), the only graphs G for which the bound is met with equality are bidirectional trees. However, the bound is met with equality if and only if $h_{ij}(A) \leqslant 2$ for all pairs (i, j), and the only bidirectional tree with this property is the bidirectional star. □

The previous theorem shows that, to leading order, the MSD with lowest social cost has a cost twice as high as that of the social optimum, the complete graph; in particular, note that the social cost of the bidirectional star scales quadratically with the number of nodes n. As the following proposition shows, in the worst case it is in fact possible that the social cost of a strongly connected link stabilizable graph scales as $O(n^3)$ as n increases. As a result, in the worst case the ratio of the social cost of a strongly connected link stabilizable graph to the lowest possible social cost over all strongly connected link stabilizable graphs may become arbitrarily large (and in the worst case, may grow linearly in n).

Proposition 25. *A directed cycle has social cost* $n^2(n-1)c/2$; *and for any other strongly connected link stabilizable graph* G *which is not a directed cycle, the social cost is strictly smaller*: $S(G) < n^2(n-1)c/2$.

Proof. Consider the directed cycle $G_C = (V, A_C)$, where $A_C = \{(1,2),(2,3),\ldots,(n-1,n),(n,1)\}$. From Example 16, $f_1(A_C) = n(n-1)/2$, and by symmetry, $f_i(A_C) = n(n-1)/2$ for all i; thus $S(G_C) = n^2(n-1)c/2$.

Now suppose that we are given a strongly connected link stabilizable graph $G = (V, A)$; then by Proposition 12, G is an MSD. As in the proof of Proposition 24, we define $h_{ij}(A)$ as the number of nodes traversed (including j) on the shortest path from i to j using the links in A, and observe that:

$$S(G) = c\sum_{i\in V} f_i(A) = c\sum_{i\in V}\sum_{j:j\neq i} h_{ij}(A).$$

Fix a node i, and sort and label the remaining nodes as $i_1,\ldots,i_{n-1}$, so that $h_{ii_1}(A) \leqslant h_{ii_2}(A) \leqslant \cdots \leqslant h_{ii_{n-1}}(A)$. Since the kth node in this labeling cannot be more than distance k away from i, it follows that $h_{ii_k}(A) \leqslant k$. We thus have:

$$\sum_{j:j\neq i} h_{ij}(A) \leqslant \sum_{k=1}^{n-1} k = \frac{n(n-1)}{2}.$$

As a result:

$$S(G) = c\sum_{i\in V}\sum_{j:j\neq i} h_{ij}(A) \leqslant \frac{n^2(n-1)c}{2} = S(G_C).$$

Finally, we show the preceding inequality is strict if $G \neq G_C$. We relabel the nodes so that $i_1 = 1$, and $i_k = k$. Then if the preceding inequality holds with equality, we must have $h_{ik}(A) = k$ for all k; observe that this implies each node $1,\ldots,n-1$ has exactly one outgoing link, and the shortest path spanning tree rooted at node 1 must be $\{(1,2),(2,3),\ldots,(n-1,n)\}$. Since G is strongly connected, we must also have $(n,1) \in A$; as a result, G is an MSD and contains a directed cycle, so in fact G must be a directed cycle. This completes the proof. □

The social cost metric considered in this section provides a way of quantifying the overall attractiveness of various strongly connected link stabilizable graphs. We saw that the ratio between the highest and lowest social costs achievable among strongly connected link stabilizable graphs scales linearly in n; this large gap highlights the importance of understanding exactly how game dynamics contribute to selecting among the range of strongly connected link stable equilibria.

5. Conclusion

We considered a network formation game where the different nodes act as players and each node needs to send goods to the other nodes. Our model has two distinctive features: asymmetric contracts between the agents which lead to directed networks, and a cost function which is based on both flow and connectivity. The stability concept we used is similar to the concept of pairwise stability with side payments developed by Jackson and Wolinsky (1996), but our analysis uses graph-theoretic concepts and leads to structural results on link stable equilibria.

Several important factors have been ignored in this formulation. For example, no fixed cost is incurred for forming a link. A more general model that includes such a fixed cost might be appropriate in the context of communication networks and transportation networks. The fixed cost of the link should be split between the two nodes, and the choice of split may even be a part of the contract between the nodes.

Another generalization involves considering the flow as part of the strategy space as well. In our model, while prices are chosen by the nodes, flows are determined by the resulting network structure. Some applications may require a model where the strategy of the nodes includes both routing of the flow and pricing for the links.

Yet another important factor which the cost structure in our model does not take into account is the delay between source and destination. In many contexts, especially in goods delivery and communication networks, it may be natural to include a cost that increases with the delay experienced by a unit of flow. Notice that with this more general cost model, strongly connected link stabilizable graphs need not be minimally connected.

Our model ignores the dynamics of the evolution of networks (e.g., Jackson and Watts, 2002; Bala and Goyal, 2000) as well as the initial process that leads to the formation of the network. As argued by Skyrms and Pemantle (2000), dynamics are essential for making formation games realistic. The dynamics of a network where links are broken and formed over time are of great interest. One advantage of our model in this context is the fact that each player can easily predict (at least in the short term) the effect of breaking or forming a link. Thus, dynamics which focus only on short term best responses of the players might be analytically tractable.

In addition to dynamics, our model does not consider end user response to changes in the network structure (by "end user," we mean any consumer of the service provided by the network nodes). Specifically, our model takes the traffic matrix as fixed. Including the response of end users to the behavior of the network is an important and essential extension to the model, to realistically capture the long term variations in the customer base of network nodes. Indeed, one might expect the response of the end users to act as a stabilizing mechanism which pushes the nodes towards an "efficient" network structure.

References

Anshelevich, E., Dasgupta, A., Tardos, E., Wexler, T., 2003. Near-optimal network design with selfish agents. In: Proceedings of the 35th Annual ACM Symposium on the Theory of Computing, pp. 511–520.

Aumann, R., Myerson, R., 1988. Endogenous formation of links between players and coalitions: An application of the Shapley value. In: Roth, A. (Ed.), The Shapley Value. Cambridge Univ. Press, pp. 175–191.

Bala, V., Goyal, S., 2000. Noncooperative model of network formation. Econometrica 68, 1181–1230.

Belleflamme, P., Bloch, F., 2002. Market sharing agreements and collusive networks. Working paper 443. Department of Economics, Queen Mary University of London.

Berge, C., 1973. Graphs and Hypergraphs. North-Holland, Amsterdam.

Cohen, R., 2000. Consequences of competition. In: Hudgins, E. (Ed.), Mail the Millennium: Will the Postal Service Go Private?. Cato Institute, Washington, DC, pp. 103–113.

Cox, L., Murray, C., Noble, B., 2002. Pastiche: Making backup cheap and easy. In: The Fifth ACM/USENIX Symposium on Operating Systems Design and Implementation.

Cunningham, W.H., 1982. Decomposition of directed graphs. SIAM J. Algebraic Discrete Methods 3, 214–228.

Currarini, S., Morelli, M., 2000. Network formation with sequential demand. Rev. Econ. Design 5, 229–250.

Fabrikant, A., Luthra, A., Maneva, E., Papadimitriou, C., Shenker, S., 2003. On a network creation game. In: Proceedings of the Twenty-Second Annual ACM Symposium on Principles of Distributed Computing, pp. 347–351.

Fuqua, A., Ngan, T., Wallach, D., 2003. Economic behavior of peer-to-peer storage networks. In: Workshop on Economics of Peer-to-Peer Systems.

Furusawa, T., Konishi, H., 2002. Free trade networks. Working paper 458. Department of Economics, Boston College.

Gao, L., Rexford, J., 2001. Stable Internet routing without global coordination. IEEE/ACM Trans. Networking 9 (6), 681–692.

Geels, D., Kubiatowicz, J., 2002. Replica management should be a game. In: Proceedings of the SIGOPS European Workshop.

Grotschel, M., 1979. On minimally strong blocks. J. Graph Theory 3 (3), 213–219.

Harary, F., 1968. Graph Theory. Addison–Wesley.

Jackson, M., 2004. A survey of models of network formation: stability and efficiency. In: Demange, G., Wooders, M. (Eds.), Group Formation in Economics: Networks, Clubs, and Coalitions. Cambridge Univ. Press, Cambridge, UK.

Jackson, M., Calvó-Armengol, T., 2004. Networks in labor markets: Wage and employment dynamics and inequality. J. Econ. Theory. In press.

Jackson, M., Watts, A., 2002. The evolution of social and economic networks. J. Econ. Theory 106 (2), 265–295.

Jackson, M., Wolinsky, A., 1996. A strategic model of social and economic networks. J. Econ. Theory 71 (1), 44–74.

Johari, R., Mannor, S., Tsitsiklis, J.N., 2004. A contract-based model for directed network formation. Publication 2630. MIT Laboratory for Information and Decision Systems.

Johnson, C., Gilles, R., 1999. Spatial social networks. Rev. Econ. Design 5, 273–300.

Kranton, R., Minehart, D., 2000. Competition for goods in buyer–seller networks. Rev. Econ. Design 5, 301–332.

Mas-Colell, A., Whinston, M., Green, J., 1995. Microeconomic Theory. Oxford Univ. Press, Oxford, UK.

Myerson, R.B., 1991. Game Theory: Analysis of Conflict. Harvard Univ. Press.

Norton, W., 2000. Internet service providers and peering. Available on request from: http://www.equinix.com/press/whtppr.htm.

Rawnsley, D., Lazar, N., 1999. Managing the universal service obligation. In: Crew, M., Kleindorfer, P. (Eds.), Emerging Competition in Postal and Delivery Services. Kluwer Academic, Boston, MA, pp. 181–193.

Sherman, R., 1991. Competition in postal service. In: Crew, M., Kleindorfer, P. (Eds.), Competition and Innovation in Postal Services. Kluwer Academic, Boston MA, pp. 191–214.

Sidell, J., Aoki, P., Sah, A., Staelin, C., Stonebraker, M., Yu, A., 1996. Data replication in Mariposa. In: The 12th International Conference on Data Engineering, pp. 485–494.

Skyrms, B., Pemantle, R., 2000. A dynamic model of social network formation. In: 97 Proc. NAS, pp. 9340–9346.

Slikker, M., van den Nouweland, A., 2001a. A one-stage model of link formation and payoff division. Games Econ. Behav. 34 (1), 153–175.

Slikker, M., van den Nouweland, A., 2001b. Social and Economic Networks in Cooperative Game Theory. Kluwer Academic, Boston, MA.