

Balanced Outcomes in Social Exchange Networks *

Jon Kleinberg
Dept. of Computer Science
Cornell University
Ithaca, NY 14853
kleinber@cs.cornell.edu

Éva Tardos
Dept. of Computer Science
Cornell University
Ithaca, NY 14853
eva@cs.cornell.edu

ABSTRACT

The study of bargaining has a long history, but many basic settings are still rich with unresolved questions. In particular, consider a set of agents who engage in bargaining with one another, but instead of pairs of agents interacting in isolation, agents have the opportunity to choose whom they want to negotiate with, along the edges of a graph representing social-network relations. The area of *network exchange theory* in sociology has developed a large body of experimental evidence for the way in which people behave in such network-constrained bargaining situations, and it is a challenging problem to develop models that are both mathematically tractable and in general agreement with the results of these experiments.

We analyze a natural theoretical model arising in network exchange theory, which can be viewed as a direct extension of the well-known Nash bargaining solution to the case of multiple agents interacting on a graph. While this generalized Nash bargaining solution is surprisingly effective at picking up even subtle differences in bargaining power that have been observed experimentally on small examples, it has remained an open question to characterize the values taken by this solution on general graphs, or to find an efficient means to compute it.

Here we resolve these questions, characterizing the possible values of this bargaining solution, and giving an efficient algorithm to compute the set of possible values. Our result exploits connections to the structure of matchings in graphs, including decomposition theorems for graphs with perfect matchings, and also involves the development of new techniques. In particular, the values we are seeking turn out to correspond to a novel combinatorially defined point in the interior of a fractional relaxation of the matching problem.

*This work has been supported in part by NSF grants CCF-0325453, IIS-0329064, CNS-0403340, BCS-0537606, and CCF-0729006, ONR grant N00014-98-1-0589, a Google Research Grant, a Yahoo! Research Alliance Grant, the Institute for the Social Sciences at Cornell, and the John D. and Catherine T. MacArthur Foundation.

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

STOC'08, May 17–20, 2008, Victoria, British Columbia, Canada.
Copyright 2008 ACM 978-1-60558-047-0/08/05 ...\$5.00.

Categories and Subject Descriptors: F.2.2 [Analysis of Algorithms and Problem Complexity]: Non-numerical Algorithms and Problems

General Terms: Algorithms, Economics, Theory

Keywords: social networks, game theory, bargaining, network exchange theory

1. INTRODUCTION

The study of *bargaining* has a long history in economics and sociology, with considerable attention devoted even to cases where just two agents are involved. To take perhaps the most basic example, suppose that two parties are negotiating over how to divide one unit of money, and that each has an *alternate* option — a fallback amount that it can collect in case negotiations break down without a division. Then the prediction of a number of models, including Nash's celebrated *bargaining solution* [13], is that the two agents in such a case will agree on a division that lies halfway between the extremes of their alternate options.

To phrase this principle precisely, suppose agents A and B have alternate options α and β respectively. Then the Nash bargaining solution posits that they split the surplus $s = 1 - \alpha - \beta$ evenly between them: if $s < 0$ then there is no solution that will make them both happy, while if $s \geq 0$, then they will agree on $\alpha + \frac{1}{2}s$ for A and $\beta + \frac{1}{2}s$ for B . Note that *any* division — even a very unbalanced one — in which A and B each get more than α and β respectively is “stable,” in the sense that both agents prefer it to the breakdown of negotiations. Thus, the theory of bargaining in some sense can be viewed as a way to refine the set of possible negotiation outcomes, selecting a particular one that is most “reasonable.” The Nash bargaining solution in particular has a central role in the history of this theory, and has been motivated by axiomatic approaches, optimization approaches, and as the outcome of particular game-theoretic formulations [15].

Despite this history, the problem of bargaining quickly enters unsettled theoretical territory once we consider multiple agents residing at the nodes of a network, with edges representing the pairs that have the opportunity to negotiate. At an experimental level, there has in fact been a long and active history of research in sociology on the way in which people behave in precisely this type of setting; such questions form the foundation of the sociological field of *network exchange theory* [19], which has developed empirical principles for the kinds of positions in networks that confer particular bargaining power. The theoretical challenge in network exchange theory has been to formulate graph-theoretic models

that align with the results of these experiments.

One particularly clean theoretical model arising in network exchange theory, a proposal of Cook and Yamagishi in 1992 [6], can be formulated as a natural generalization of the Nash bargaining solution to arbitrary networks, in which the outside options for each agent arise directly from the network structure, rather than having to be specified externally. It (along with related formalisms) has been found to be surprisingly effective at picking up even subtle differences in bargaining power that have been discovered through experimental work on human subjects [19]. Unfortunately, while it is possible to work out the values predicted by this generalized Nash bargaining solution on small examples, it has remained an open question to characterize these values on general graphs, including the questions of which graphs have a well-defined solution under this notion, and whether such solutions can be computed by an efficient algorithm.

Here we resolve these questions, characterizing the possible values that this generalized Nash bargaining solution can take on arbitrary graphs, and giving an efficient algorithm to compute the set of possible values. Our result exploits connections to the structure of matchings in graphs, via the Edmonds-Gallai structure theorem and the decomposition of a perfectly matchable graph into elementary subgraphs [11]; this in turn provides some further insight into the qualitative observations that arise from experiments in network exchange theory.

We now describe some background on network exchange theory, including the experimental observations that motivate the need for such a model, before turning to the formulation of our results.

Network exchange theory. Network exchange theory originates from the goal of understanding power imbalances in the relationships between pairs of people [8]. This includes relationships in economic or political settings, as well as in social interaction more generally — such as the roles people play in groups of friends, communities, or organizations. A common theme here is to view social relations between two individuals as producing value for both; the extent to which this value is divided between them can be seen as a kind of latent bargaining known as *social exchange*, and *power* is then related to the imbalance in this division — with the powerful party in the relationship getting the majority of the value. We will be deliberately vague in specifying what this value is, since it clearly depends on the type of social relation — in a business partnership, the more powerful party may get better terms, while in a friendship, the more powerful party may be more the focus of attention, may get their way more often, and so forth.

Network exchange theory begins from the hypothesis that these power imbalances may arise in part for *structural* reasons — that is, based on an individual’s position in a social network. How, then, should we characterize a powerful position in a network? For example, in the set of friendships within a group of five people depicted in Figure 1, node v intuitively holds a powerful position: nodes u and w are completely dependent on v for their social relations in the group, while v occupies a central position and has the power to focus his or her social energies on multiple options.

To make such intuitions precise, exchange theorists beginning in the 1970s developed a bargaining-based experimental framework [4, 5] that has continued to be used up to the




Figure 1: A social network on five people, with node v occupying an intuitively powerful position.

present. In a controlled laboratory setting, n human subjects are asked to each play the role of one of the n nodes of a graph such as the one in Figure 1. The value in each social relation is captured by placing a fixed sum of money (say, \$1) on each edge. Nodes now engage in free-form negotiation (say, via instant messaging from different computer terminals) over how to divide the units of money on each edge. By the end of the time limit, each node v is supposed to reach an agreement with *at most one* neighbor w on how to divide the money on their incident edge; if v is able to reach such an agreement, then she gets her agreed-upon share of the money on this edge, and if v is not able to reach an agreement with any neighbor, then she gets 0. A central question is whether certain nodes tend to make a disproportionate amount of money in such an experiment, in a reproducible way across different sets of subjects.

The network in Figure 1 is one that has been studied extensively in these experiments, with predictable results. Since v can only complete one successful transaction, one of nodes u or w in this network will necessarily be left out and hence get 0. Node v uses this power over u and w to obtain close to the full amount on the edge where she does reach an agreement, and so even the node among u and w that completes a transaction gets close to 0. Meanwhile, node x quickly realizes that it’s useless to bargain with v , and so he spends most of his energy bargaining with y on essentially equal terms, reaching a split close to $\frac{1}{2}$ - $\frac{1}{2}$.

Network exchange theory experiments have been carried out under many different conditions — for example, participants can have knowledge of the full graph or only the names of their neighbors; negotiation can be done in free text or through a highly restricted interface that limits communication — and the results are robust, in the sense that the general properties of power imbalances that we discuss here are present essentially regardless of these how these experimental design decisions are made.

Gradations of power. The outcomes of experiments on the network in Figure 1 are relatively simple to understand, and they conceptually arise from the consequences of negotiation on two even simpler graphs: a 3-node path, in which the middle node acquires almost all the value; and a 2-node path, in which the two sides tend to split the value approximately equally. Other small examples show that the results of negotiation cannot be reduced to superficial statements about “central” positions in the network structure. For example, in experiments on the 5-node path in Figure 2(a), nodes v and x acquire almost all the value: although w in-




Figure 2: Three small graphs studied extensively in network exchange theory.

tuitively occupies the “central” position in the network, this does not confer any real bargaining power, since his two potential partners v and x each have alternate, and very weak, neighbors.

The results of the experiments take on further subtleties when we look at other small graphs in which the negotiation outcomes on edges do not gravitate toward the extremes of $\frac{1}{2}$ - $\frac{1}{2}$ or 0-1. On the 4-node path in Figure 2(b), for example, v has the power to exclude u , but it is costly to exercise this power: it requires v to negotiate with w , who has an alternate, weak partner. (The argument for w is symmetric.) Experiments bear out this “weak power”, showing that v tends to obtain between $7/12$ and $2/3$ — but not more — in transactions with u [12, 17]. A similar kind of weak power applies to node v in the graph in Figure 2(c); in experiments here, v tends to obtain strictly between $1/2$ and 1 , and in fact consistently a bit higher than the node labeled v obtains in the 4-node path.

A natural question — pursued by exchange theorists concurrently with the experimental work — has been to find simple graph-theoretic measures that could explain these gradations in bargaining power as a function of position in the network.

Balanced outcomes. We begin by introducing some terminology that will make the discussion easier. We will focus here on the final outcomes experienced by the nodes, rather than the fine-grained dynamics of negotiation that lead to it — that is, with the experimental results providing guidance as to the outcomes that arise in practice, we wish to formalize a prediction for the result of simultaneous negotiation across the edges of a graph G , when each node can engage in at most one transaction.

To begin with, then, we say that an *outcome* of network exchange on a graph G consists of a matching M (indicating the pairs that transact), and a *value* γ_v obtained by each node v , such that each pair of nodes (v, w) that forms an edge in M divides the one unit of value on this edge: $(v, w) \in M$ implies $\gamma_v + \gamma_w = 1$. Nodes v that are unmatched receive $\gamma_v = 0$. Given an outcome, we can define α_v to be the best alternative that v could obtain if it broke its current agreement and formed another one: α_v is the maximum, over all u to which v is connected by an edge not in M , of $1 - \gamma_u$; we set $\alpha_v = 0$ when v has no other alternatives. (That is, v has the option of stealing u away from its current transaction — which yields it γ_u — by offering any value exceeding γ_u by an arbitrarily small amount, and keeping the remainder. Note also that α_v is well-defined even for nodes that are not matched by M .) We will say that an outcome is *stable* if $\alpha_v \leq \gamma_v$ for all nodes v . Stability is an essential but very weak property; there is generally a wide

range of possible (but experimentally implausible) outcomes satisfying it.

To formalize a kind of outcome that agrees with the results of the experiments, Cook and Yamagishi proposed the following natural notion, which can be viewed as a direct extension of the Nash bargaining solution to a network context [6]. One defines a *balanced outcome* to be an outcome in which, for each pair $(v, w) \in M$, the values γ_v and γ_w represent the Nash bargaining solution for v and w , given their alternatives α_v and α_w . Thus, we are simply seeking a set of self-consistent values at the nodes, such that each transaction follows the Nash bargaining solution *given the values in the rest of the network*.

It is striking to consider how the balanced outcomes align with experimental observation. First, in Figure 1 they assign a value of 1 to node v and $\frac{1}{2}$ each to nodes x and y . (While the matching M in this case is not unique, the values are.) The balanced outcomes in Figure 2(a) all give values of 1 to nodes v and x . And perhaps most interestingly, balanced outcomes identify the various gradations of weak power in Figures 2(b) and 2(c): the reader can check that in both cases, the balanced outcome is unique, and it gives $\frac{2}{3}$ to nodes v and w in Figure 2(b), and $\frac{3}{4}$ to node v in Figure 2(c) (i.e. slightly more than on the 4-node path, but still bounded away from 1).

However, the application of this formalism was restricted in the network exchange theory literature to the analysis of small examples, due to the open question mentioned at the outset: Beyond individual examples, it has remained unknown how to characterize and compute the balanced outcomes for an input graph in general. For example, given the fixed-point nature of the definition, it is not even *a priori* clear whether the values in a balanced outcome are necessarily rational or finitely representable.

Our results. In this paper, we resolve this set of questions. We prove that a balanced outcome exists if and only a stable outcome exists, and we provide an algorithm that efficiently computes the set of balanced outcomes for a graph G . The existence of stable outcomes provides a natural condition for the existence of a balanced outcome, and we also find a tractable description of the set of all balanced outcomes in a given graph G . (When balanced outcomes exist for G , at least one outcome — but not necessarily all — consists entirely of rational values.) We also extend the algorithmic and existence results to the case of *weighted* graphs, in which the value being divided across an edge can vary from one edge to another — this too has been an object of experimental work in network exchange theory.

An interesting aspect of the solution is that while balanced outcomes have a very natural definition, they do not

map naturally onto standard values associated with matchings in graphs. It is not hard to see that any stable outcome (and hence any balanced outcome) must involve a maximum matching M , and that the values γ_v must satisfy the constraints for dual variables supporting a fractional relaxation of the maximum matching problem. However, as the examples in Figures 2(b) and 2(c) make clear, a balanced outcome’s node values do not generally occur at or near the extreme points of the dual polytope. While matching theory has tended to focus on questions that are answered by values at these extreme points, balanced outcomes provide a natural instance of a notion whose very definition requires us to consider solutions in the interior of the polytope (or in the relative interior of a face) — a part of the space that has tended to be less well-explored by matching theory.

We begin by partially separating a balanced outcome into its “extreme” and “interior” parts, using the Edmonds-Gallai structure theorem for matchings in graphs [11]. We show that the values in a balanced outcome on certain parts of the decomposition must be 0 or 1, leaving us with a remaining part where the values are generally balanced between these extremes. This remaining part, which always possesses a perfect matching, is where the main source of complexity lies; it is handled by a further decomposition into *elementary subgraphs* [11], and an algorithm that iteratively sets values on edges starting with long (possibly non-simple) paths on which the Nash bargaining solution is as constrained as possible.

Further related work. A number of lines of work in economics have modeled the behavior of buyers and sellers interacting in a network, and this is related in a thematic sense to the special case of our model restricted to bipartite graphs. Indeed, if we designate one side of a bipartite graph as the set of sellers, and the other as the set of buyers, then we can re-cast the model underlying balanced outcomes as an equivalent one in which sellers and buyers bargain about the prices at which various copies of a good — having zero value to sellers and unit value to buyers — will be sold. In this way, our model can be compared at a general level to lines of work in matching markets [16], competitive equilibrium in networks [9, 18], graph-structured auctions [10], and price-determination through bargaining [2, 3, 7, 14].

These economic models differ significantly from our approach in their specifics. Perhaps the most closely related is a model of Corominas-Bosch [7], though it remains quite different as well: it incorporates explicit bargaining dynamics in a graph but leads to a model in which all nodes essentially receive values of 0, $\frac{1}{2}$, or 1. Thus, it deals only with “extreme” outcomes, and is not able to capture the kinds of weak power that are essential both to our model and to the observations of network exchange theory experiments. (We note that follow-up work by Charness, Corominas-Bosch, and Frechette [3] ran human-subject experiments in qualitative agreement with the model of Corominas-Bosch [7]. It is an interesting question to consider the divergence between the experiments of [3] and the wide range of experiments in network exchange theory; while this is not fully settled, it quite possibly involves the rigid synchronization of alternating simultaneous seller proposals and buyer proposals imposed by Charness et al., together with the fact that buyer-seller pairs immediately dropped out of the system as transactions were completed.)

2. BIPARTITE GRAPHS WITH UNIQUE PERFECT MATCHINGS

As a first step toward a general characterization of balanced outcomes, we consider a restricted but illuminating special case: bipartite graphs that possess unique perfect matchings. As the earlier discussion around Figure 2(b) indicates, this special case already contains some of the main sources of subtlety. And for our presentation here, it is useful as a controlled setting in which to introduce some of the arguments; it also has an interesting equivalent formulation as a problem on partially ordered sets.

The connection to posets arises as follows. Let $G = (V, E)$ be a bipartite graph, with color classes X and Y , and suppose that G has a unique perfect matching M . Consider the orientation of G with edges in M oriented from X to Y , and all other edges oriented from Y to X . Since M is the only perfect matching in G , the resulting directed graph must be acyclic. We define a poset on the nodes of X , declaring $v \preceq w$ if and only if w is reachable from v in the directed graph. We also add to the poset two additional elements: an element \perp below all minimal elements, and an element \top above all maximal elements. We write $X^+ = X \cup \{\perp, \top\}$ and refer to this partially ordered set as the poset *induced by M on X^+* .

Now, recall the definitions of stable and balanced outcomes in G from Section 1. We define a corresponding notion of “balance” in posets as follows. Let \mathcal{P} be a poset, and suppose it has a unique minimal element \perp and a unique maximal element \top . We will refer to all elements of \mathcal{P} other than \perp and \top as *internal*. We say that a *consistent labeling* of \mathcal{P} is an assignment of a number $\gamma_a \in [0, 1]$ to each element $a \in \mathcal{P}$, subject to the condition that $\gamma_\perp = 0$, $\gamma_\top = 1$, and $\gamma_a \leq \gamma_b$ whenever $a \preceq b$ in the partial order. (Although it is not important for the discussion here, we note that the set of all labelings satisfying this latter condition, viewing each as a vector $(\gamma_a : a \in \mathcal{P}) \in \mathbf{R}^{|\mathcal{P}|}$, forms the *order polytope* of \mathcal{P} .) For any set $A \subseteq \mathcal{P}$ containing \perp and \top , we can also speak of a consistent labeling of A , by simply considering the partial order restricted to A .

We say that a consistent labeling is *balanced* at an internal element b if the value γ_b is the midpoint of $\max_{a \preceq b} \gamma_a$ and $\min_{c \succeq b} \gamma_c$. More generally, we say that the labeling is balanced at a set of internal elements $S \subseteq \mathcal{P}$ if it is balanced at each $b \in S$.

We connect balanced outcomes in G to balanced poset labelings via the following fact; the translation between these two notions follows from the definition of α_v in Section 1.

PROPOSITION 2.1. *Let $G = (V, E)$ be a bipartite graph with color classes X, Y , and with a unique perfect matching M . Let \mathcal{P} be the poset induced by M on X^+ . Consider an outcome of network exchange G with matching M and values $\{\gamma_v : v \in V\}$, where we also define $\gamma_\perp = 0$ and $\gamma_\top = 1$.*

- (a) *The outcome is stable in G if and only if $\{\gamma_v : v \in X^+\}$ is a consistent labeling of \mathcal{P} .*
- (b) *The outcome is balanced in G if and only if $\{\gamma_v : v \in X^+\}$ is a balanced labeling of \mathcal{P} .*

Constructing a balanced poset labeling. We now prove by an efficient construction that every poset (with unique minimum and maximum) has a unique balanced labeling; in view of Proposition 2.1, this will prove that bipartite graphs

with unique perfect matchings have unique balanced outcomes. In fact, we prove the following slightly more general statement, in which certain poset elements can be pre-set to constant values; this will also be useful in the next section when we consider general bipartite graphs.

THEOREM 2.2. *For every poset \mathcal{P} with unique minimum and maximum, every set $A \subseteq \mathcal{P}$ containing this minimum and maximum, and every consistent labeling γ of A , there exists a unique consistent labeling γ^* extending γ that is balanced at every element of $\mathcal{P} - A$. Moreover, this labeling γ^* can be efficiently computed.*

We will prove this by analyzing an algorithm to construct the labeling γ^* . The algorithm works in stages, as follows. At the start of each stage s , a subset $B_s \supseteq A$ has already been labeled, and the remaining elements are unlabeled. Initially, $B_1 = A$. At the start of a given stage, we say that a chain $C = b_1 \preceq b_2 \preceq \dots \preceq b_r$ in \mathcal{P} is *anchored* if $r > 2$, the endpoints b_1 and b_r of C are labeled, and all its other elements are unlabeled. Notice that as long as $B_s \neq \mathcal{P}$, there is at least one anchored chain at the start of stage s .

In a given stage s , we find an anchored chain $C_s = b_1 \preceq b_2 \preceq \dots \preceq b_r$ that minimizes the following *gap* parameter: $\delta_s = \frac{\gamma_{b_r} - \gamma_{b_1}}{r - 1}$. We assign values to b_2, \dots, b_{r-1} that are equally spaced between γ_{b_1} and γ_{b_r} . This increases the set of labeled elements: we have $B_{s+1} = B_s \cup C_s$. We continue iterating through stages until all elements are labeled.

This is the full algorithm. Clearly it terminates with a labeling of \mathcal{P} ; we now argue that this labeling is balanced at $\mathcal{P} - A$. The first key step in doing this is to show that the gap values $\delta_1, \delta_2, \delta_3, \dots$ found in sequence by the algorithm are non-decreasing.

LEMMA 2.3. *For all $s \geq 1$, we have $\delta_s \leq \delta_{s+1}$.*

Proof. Suppose by way of contradiction that for some s we have $\delta_{s+1} < \delta_s$. Now, if the endpoints of the chain C_{s+1} were already labeled at the start of stage s , then C_{s+1} could have been chosen in stage s , a contradiction. Thus, at least one endpoint of C_{s+1} was labeled during stage s — in other words, at least one endpoint of C_{s+1} lies on C_s . Let us consider the case in which the lower endpoint does but the upper one does not. (The cases in which only the upper one does, or in which both do, are analogous.)

Thus, let $C_s = b_1 \preceq b_2 \preceq \dots \preceq b_r$, let $C_{s+1} = c_1 \preceq c_2 \preceq \dots \preceq c_q$, and suppose that $c_1 = b_p$ for some p . Then since $\delta_{s+1} < \delta_s$, the chain $C'_s = b_1 \preceq b_2 \preceq \dots \preceq b_p \preceq c_2 \preceq \dots \preceq c_q$ is anchored at the start of stage s , and has a strictly smaller gap than C_s . This contradicts our choice of C_s in stage s . ■

An analogous argument, applied to two-element chains with one newly-labeled end, forms the induction step in an argument that after each stage, the labeling is balanced on all elements that have been labeled up through the end of that stage. We omit the proof.

LEMMA 2.4. *For each s , the labeling at the end of stage s is balanced on each element of B_{s+1} .*

Hence, by induction, the labeling at the end of the algorithm is balanced on all elements of $\mathcal{P} - A$, as required.

We now argue that the labeling produced by our algorithm is the unique balanced labeling. Indeed, consider any

consistent labeling β that agrees with γ^* on the elements of A , and is balanced on the elements of $\mathcal{P} - A$. The following fact provides the induction step in proving that β must in fact agree with γ^* everywhere; the proof argues that the smallest gap created by β on elements labeled after stage s must in fact be the gap δ_s .

LEMMA 2.5. *For each $s \geq 1$, if β agrees with γ^* on the elements of B_s , then it must agree with γ^* on the chain C_s as well.*

Proof. Consider the labelings γ^* and β restricted to the elements of B_s , where they are assumed to agree. Among all pairs of elements $a \preceq b$, at least one of which does not belong to B_s , choose a pair for which $\beta_b - \beta_a$ is minimum. Let δ denote this difference.

Suppose $b \notin B_s$. (The case in which $b \in B_s$ is analogous.) Then $b \notin A$, so the labeling β is balanced at b , and hence there is an element b_2 such that $\beta_{b_2} - \beta_b = \delta$. (Note that we have equality here since δ was chosen as the minimum difference between any comparable pair, at least one of which is not in B_s .) If $b_2 \notin A$, then by the same argument there is an element b_3 such that $\beta_{b_3} - \beta_{b_2} = \delta$. We continue this process until we first reach an element $b_r \in B_s$.

Now, if a is also not in B_s , then we can perform an analogous procedure moving downward in the partial order, producing elements $a_q \preceq a_{q-1} \preceq \dots \preceq a_2 \preceq a$ such that the labels of consecutive elements under β differ by exactly δ , and such that a_q is the only element in this sequence belonging to B_s .

Now, if we consider the chain $C = a_q \preceq a_{q-1} \preceq \dots \preceq a_2 \preceq a \preceq b \preceq b_2 \preceq \dots \preceq b_r$, it is anchored at the start of stage s and has gap value δ . Moreover, by the minimality of δ , this must also be the gap $\delta_s = \delta$ of the chain C_s (which may or may not be the same as C).

Finally, suppose that there are two consecutive elements c and c' on C_s for which $\beta_{c'} - \beta_c \neq \delta$. Then there would be two consecutive elements for which this difference is strictly less than δ , again contradicting the minimality of δ . This shows that β agrees with γ^* on the elements of C_s , as required. ■

3. GENERAL BIPARTITE GRAPHS

Next we will extend this solution to the case of bipartite graphs without the assumption that G has a unique perfect matching. Our main result for bipartite graphs is that

THEOREM 3.1. *For any bipartite graph G , and any matching M in G , there is a balanced outcome involving M if and only if M is a maximum matching. In particular, there exists a balanced outcome for every bipartite graph G .*

The proof uses a structure called the *Edmonds-Gallai decomposition* [11], which lets us separate the graph into a portion containing a perfect matching, and a portion on which all stable outcomes take 0-1 values. For the portion of the graph containing a perfect matching, we use a further decomposition into *elementary subgraphs* [11], as we explain below. We will see that the special case considered in the previous section is the case in which each elementary subgraph in the decomposition consists of a single edge; larger elementary subgraphs will give us more flexibility in assigning values in a balanced outcome.

The Edmonds-Gallai decomposition. Our proof of this theorem begins with the following observation, whose proof is straightforward.

PROPOSITION 3.2. *In any stable outcome for a bipartite graph G , the matching M must be maximum; and if v is a node that is not part of every maximum matching, then $\gamma_v = 0$.*

Let D be this set of nodes v that are not part of every maximum matching. The *Edmonds-Gallai decomposition* of a graph $G = (V, E)$ partitions V into three parts: the set D , which in a bipartite graph forms an independent set; a set A consisting of all nodes with a neighbor in D ; and a set $C = V - D - A$ in which there is always a perfect matching [11]. In any stable outcome, we have $\gamma_v = 0$ for all $v \in D$ by Proposition 3.2; we have $\gamma_v = 1$ for all $v \in A$; and hence we can essentially reduce the problem to one in which we need to find values for the subgraph $G[C]$ — i.e. for bipartite graphs containing at least one perfect matching.

Note that the graph in Figure 1 gives a succinct example of this decomposition: in that graph, $D = \{u, w\}$, $A = \{v\}$, and $C = \{x, y\}$; and we have $\gamma_u = \gamma_w = 0$, while $\gamma_v = 1$ and $\gamma_x = \gamma_y = \frac{1}{2}$. Other general principles are reflected in this graph as well: for example, the fact that it is useless for x to negotiate with v is an instance of the more general principle that for any graph G , if a node $x \in C$ has a neighbor $v \in A$, then this neighbor provides x only with $1 - 1 = 0$ as an alternate option.

Bipartite graphs with perfect matchings. Most of the complexity, then, arises for graphs $G[C]$ with perfect matchings. To analyze such a graph, we further decompose C into *elementary subgraphs* [11]: if $E' \subseteq E$ is the set of edges that occur in some perfect matching of C , then the connected components H_1, \dots, H_r of (C, E') are called the elementary subgraphs. We call an elementary subgraph *trivial* if it consists of just two nodes, and *non-trivial* otherwise.

We now connect the structure of the elementary subgraphs to the poset problem defined in the previous section. To do this, we define a notion of reachability on the elementary subgraphs, by constructing a directed graph Γ as follows: We take an arbitrary a perfect matching M , orient the matching edges from X to Y , and orient all other edges from Y to X . Each non-trivial elementary subgraph corresponds to a non-trivial strongly connected component of Γ , and each trivial elementary subgraph consists of two singleton strongly connected components joined by an (oriented) edge in M .

Using the structure of Γ , we can construct the set of all possible balanced outcomes by a reduction to the poset labeling problem. The directed graph Γ defines a poset \mathcal{P} on the components H_i according to reachability. We also add to \mathcal{P} two new elements \perp and \top to serve as the unique minimum and unique maximum. We first consider how to define the labels on each elementary subgraph, and then look at how the labels relate to each other.

If we focus on a non-trivial elementary subgraph H_i as a graph in isolation, its balanced solutions consist of an assignment of a single number $0 \leq \gamma_i \leq 1$ to each node of H_i on the X -side, and the number $1 - \gamma_i$ to each node on the Y -side. If such a labeling is extended to a stable solution on G , the labeling will necessarily be balanced on the subgraph H_i , as $\alpha_v = \gamma_v$ for all nodes $v \in H_i$. Given this observation, any stable outcome on $G[C]$ corresponds to a labeling of the

partially ordered set \mathcal{P} , where the label of component H_i is the value of the labels of the nodes on the X -side of H_i .

Now consider two elementary subgraphs H_i and H_j as subgraphs of G ; we can relate the values a balanced solution takes in each of them as follows. If H_i is an immediate predecessor of H_j in \mathcal{P} , then there is an edge (w, v) from the Y -side of H_i to the X -side of H_j . Then we have $\gamma_w + \gamma_v \geq 1$, which means $(1 - \gamma_i) + \gamma_j \geq 1$, and hence $\gamma_i \leq \gamma_j$. Thus, the values γ_i corresponding to the non-trivial elementary subgraphs form a consistent labeling of \mathcal{P} (adding in $\gamma_\perp = 0$ and $\gamma_\top = 1$).

This is a property that all balanced solutions must have. Conversely, we can use the result of the previous section to show that any consistent labeling γ of the non-trivial elementary subgraphs can be uniquely extended to a balanced solution for all of G . In the notation from Theorem 2.2, we treat the non-trivial elementary subgraphs (together with \perp and \top) as the set A , and we find a consistent labeling γ^* of \mathcal{P} extending γ . We use these values to fill in the balanced solution on the trivial elementary subgraphs using Theorem 2.2. This completely describes all balanced outcomes on G , as we summarize in the next theorem.

THEOREM 3.3. *Stable outcomes on the graph G are in one-to-one correspondence with consistent labelings of the poset \mathcal{P} . A stable outcome is balanced if and only if the labeling of \mathcal{P} is balanced at all elements corresponding to trivial elementary subgraphs. Any set of consistent values at all elements corresponding to non-trivial elementary subgraphs extends uniquely to a set of values associated with a balanced outcome.*

4. GENERAL GRAPHS: FIRST STEPS

From the analysis in the previous section, we see that balanced outcomes exist for every bipartite graph G , and the values in a balanced outcome are unique if and only if $G[C]$ (in the Edmonds-Gallai decomposition) has a unique perfect matching. The situation with non-bipartite graphs becomes more complicated in many directions, including the breakdown of both these facts. This short section highlights some of the issues that arise in the non-bipartite case, to lead into the general solution that will be described in the next section.

To begin with, stable outcomes (and hence balanced outcomes) do not exist for every non-bipartite graph. The 3-node clique (i.e., a triangle) forms a simple example where there is no balanced or even stable outcome: any matching consists of a single edge, and the excluded node would be happy to accept any small value offered. As a more interesting small example with no balanced or stable outcome, consider a 10-node graph G containing a node v such that $G - \{v\}$ consists of three disjoint triangles — that is, the triangles can only reach each other through v .

In fact, the existence of stable and balanced outcomes can be characterized via the Edmonds-Gallai decomposition.

THEOREM 4.1. *In a graph G , the following are equivalent.*

- (a) *There is a balanced outcome.*
- (b) *There is a stable outcome.*
- (c) *The set D in the Edmonds-Gallai decomposition forms an independent set.*

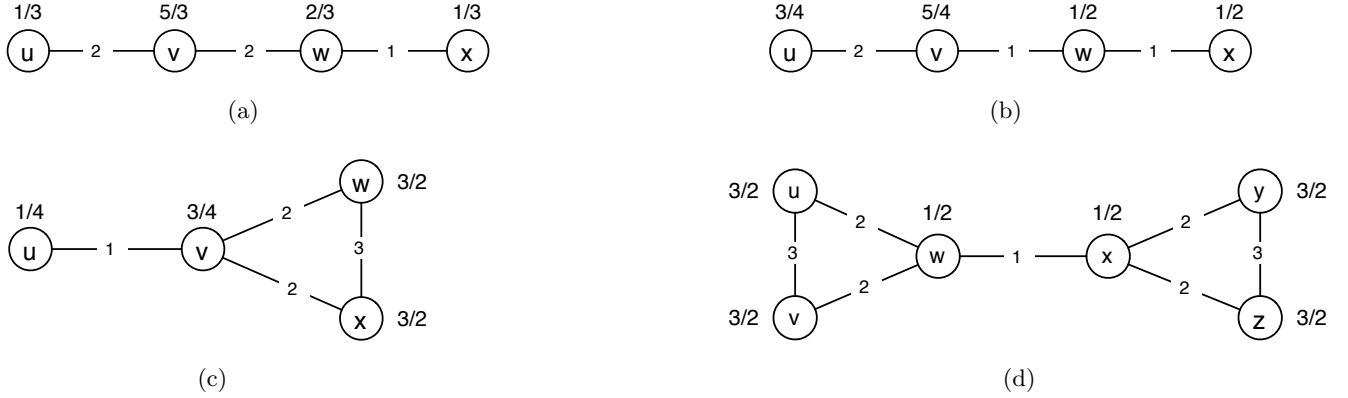


Figure 3: Examples of balanced outcomes in weighted network exchange problems.

The equivalence of (b) and (c) here is straightforward using Proposition 3.2, which applies to all graphs, not just bipartite ones. The non-trivial content of the theorem, that any graph with a stable outcome has a balanced one, will follow from our characterization of balanced outcomes in the next section.

The first step in this characterization is to reduce the problem — as we did for bipartite graphs as well — to the case in which there is a perfect matching. Specifically, if D forms an independent set, we will be able to set $\gamma_v = 0$ for all $v \in D$ and $\gamma_v = 1$ for all $v \in A$, leaving us with the graph $G[C]$ that we again (as in the bipartite case) decompose into elementary subgraphs. For bipartite graphs, recall that in a balanced outcome one could choose any value for the X -side of a non-trivial elementary subgraph; but in general graphs, even these values can be constrained. For example, the 4-node clique is an elementary graph; yet although it has multiple perfect matchings, any balanced outcome assigns value $\gamma_v = \frac{1}{2}$ to each node v .

In the next section — working within the more general framework of weighted graphs — we show how the values in non-trivial elementary subgraphs are constrained, and how a generalization of the procedure to find balanced poset labelings can extend the balanced outcome from the non-trivial elementary subgraphs to the trivial ones as well.

5. BALANCED OUTCOMES IN GENERAL WEIGHTED GRAPHS

Finally, then, we consider the general version of the problem — network exchange on an arbitrary graph, where the value being divided on each edge can be an arbitrary positive number. We first need to make clear how to generalize the notion of a balanced outcome to this weighted case. Then, to construct balanced outcomes, we will no longer be able to use a direct reduction to the poset labeling problem, but we will be able to extend the approach there to our more general setting.

Stable and balanced outcomes in weighted graphs.

To begin with, let $G = (V, E)$ be a graph, with values $\omega_{vw} \geq 0$ on the edges $(v, w) \in E$. As before, each node v represents an agent who is trying to reach an agreement to split the value on at most one of its incident edges. By analogy with

the case of unit-weight edges, an *outcome* of this process is a matching M together with an assignment of a real number $\gamma_v \geq 0$ to each node v , so that $\gamma_v + \gamma_w = \omega_{vw}$ for all edges $(v, w) \in M$, and $\gamma_v = 0$ for unmatched nodes v . We say that the outcome is *stable* if in addition we have $\gamma_v + \gamma_w \geq \omega_{vw}$ for each edge $(v, w) \in E$. To extend the notion of a balanced outcome we need to define for each node v the best alternate value α_v it can obtain by breaking the its current agreement. Formally, this is

$$\alpha_v = \max_{(v,u) \in E: u \neq w} (\omega_{vu} - \gamma_u)^+,$$

where w is the matched pair of v , we take the maximum to be 0 over the empty set, and for a real number α we use $(\alpha)^+$ to denote its positive part $\max(\alpha, 0)$. We say a stable outcome is a *balanced* if, in addition to the stability properties, we have $\gamma_v - \alpha_v = \gamma_w - \alpha_w$ for every edge (v, w) in the matching.

Figure 3 illustrates some sample instances of weighted network exchange; edge weights are drawn inside the edges, and the values γ_v from the (in these cases, unique) balanced outcomes are shown next to the nodes.

It is not hard to show that a matching M in a stable or balanced outcome may no longer be of maximum size; however, it must be of maximum weight.

PROPOSITION 5.1. *The matching M in a stable outcome must be a maximum-weight matching. If an outcome γ is stable with a matching M , then it is also stable with any other maximum-weight matching M' .*

To make the discussion cleaner we introduce the following piece of additional notation: we define the *slack* of an edge (v, w) to be

$$\sigma_{vw} = \gamma_v + \gamma_w - \omega_{vw}.$$

Note that an outcome is stable if and only if $\sigma_{vw} \geq 0$ for all edges $(v, w) \in E$. Further, we will define the *slack of a node* to be $\sigma_v = \gamma_v - \alpha_v$. Using the notion of node slacks, we can say that a stable outcome γ is balanced if and only if for every edge (v, w) in the matching, we have $\sigma_v = \sigma_w$.

The slack of a node is closely related to the slacks of the adjacent edges. To see this connection consider an edge (v, w) in the matching, with values γ_v and $\gamma_w = \omega_{vw} - \gamma_v$, and assume the value α_v is determined by edge (v, u) ; that




Figure 4: Simple examples of the four classes of alternating structures that can cause the solution to a 2-variable-per-inequality system not to exist, and can also serve to limit the maximum slack σ .

is $\alpha_v = \omega_{vu} - \gamma_u$. In this case, $\gamma_v - \alpha_v = \sigma_{uv}$. When $\alpha_v = 0$ we get $\sigma_v = \gamma_v$, so if v is matched to w then

$$\sigma_v = \min(\gamma_v, \min_{(v,u) \in E: u \neq w} \sigma_{uv}),$$

the minimum of the value γ_v and the smallest slack of an unmatched edge adjacent to v .

In three of the four examples of Figure 3, the slack at each internal node is determined by the unique unmatched edge. In Figure 3(b), however, the internal node v has value γ_v exceeding the weight of the adjacent unmatched edge (v, w) , and so we have $\alpha_w = 0$.

Overview of the main result. Our main result shows that a graph has a balanced outcome whenever it has a stable one, and it gives a procedure to compute balanced outcomes, along with a characterization of all balanced outcomes.

THEOREM 5.2. *If a graph G has a stable outcome, then it has a balanced outcome, and the set of all balanced outcomes can be constructed in polynomial time.*

We start by characterizing the graphs that have stable outcomes, and understanding the connection of slack values with the slack variables of a certain linear program. By Proposition 5.1 if an outcome consisting of a matching M and values γ is stable, then this same set of values together with any other maximum-weight matching must also form a stable outcome. So we can use any maximum-weight perfect matching M in our stable outcome without loss of generality. We can think of the stability requirement as a linear program with two variables per inequality:

$$\begin{aligned} \gamma_v &\geq 0 \quad \forall v \in V \\ \gamma_v &= 0 \quad \forall v \in V \text{ unmatched} \\ \gamma_v + \gamma_w &= \omega_{vw} \quad \forall (v, w) \in M \\ \gamma_v + \gamma_w &\geq \omega_{vw} \quad \forall (v, w) \in E \setminus M \end{aligned}$$

Aspvall and Shiloach [1] show how to solve such linear programs via a dynamic programming algorithm, and give a characterization of the existence of the solution using alternating paths and cycles that are possibly not simple, as shown on Figure 4. In particular, the structures relevant to their characterization are simple alternating paths (Figure 4(a)), alternating cycles (Figure 4(b)), and blossoms and bi-cycles: non-simple alternating paths and cycles that contain within them simple odd alternating cycles, as shown in Figures 4(c) and 4(d). We say that an alternating path or blossom is *anchored* at its degree-1 node(s), and in the characterization only consider copies of these structures that are anchored either at unmatched nodes or matching edges.

Aspvall and Shiloach show that an instance of the linear program is infeasible if and only if, for a subgraph of G consisting of one of these structures, the total weight on unmatched edges exceeds the total weight on matched edges. (To do this, we need to count each edge on the path portions of the blossoms and bi-cycles twice, as these structures should be viewed as non-simple walks that traverse this part twice.)

PROPOSITION 5.3. *In a graph G , a stable outcome can be found if one exists by computing a maximum-weight matching M , and then solving the inequality system above to get the values γ_v .*

We can think of the values γ_v in a stable outcome as dual variables for a fractional relaxation of the matching problem in which the (primal) constraints are simply that the fractional degree of each node should be bounded by 1. A matching M and set of values γ form a stable outcome if γ is feasible for the dual of this fractional matching relaxation, and the matching M and dual γ satisfy complementary slackness. In a bipartite graph such values always exist, as the matching polytope is described by the degree constraints. In the non-bipartite case, however, we do not always get such primal-dual solutions: the degree constraints form only a relaxation of the matching polytope, since the odd-subset constraints are missing.

Constructing a balanced outcome in weighted graphs.

We prove Theorem 5.2 via a procedure to find all balanced outcomes. After showing how to construct one balanced outcome we show that our procedure can generate any balanced outcome. The approach is analogous to finding a balanced labeling of a poset from Section 2. We start with a maximum-weight matching, and we set values on the matched edges (v, w) gradually, spreading out the slack through the graph as much as possible, similar to how we spread out the gaps between labels in posets. A key difference, however, is that with posets we spread out the slack by considering only paths (i.e. chains); here, we have to consider the range of structures (paths, cycles, blossoms and bi-cycles) arising in the Aspvall-Shiloach characterization.

At the first iteration we find stable outcome values γ_v such that the slacks of all unmatched edges, and the slacks of all nodes, are at least a parameter σ , for as high a value of σ as possible. For example, if there are any unmatched nodes v in M then $\sigma_v = \gamma_v = 0$ in all outcomes, and hence the first phases will have maximum $\sigma = 0$, and we will set $\gamma_v = 0$ for all unmatched nodes v . In a general step, we have a subset A of the nodes where the outcome values γ_v for all $v \in A$ have already been set, and we seek to maximize the slack on nodes and edges where values have not been set

by A . The process will set the small slack values first. In general, it is useful to consider analogies between the process as we specify it here, and the simpler process in Section 2 for producing a balanced labeling of a poset; there too, we committed gradually to the values at the elements, starting with small gaps (the analogues of slacks for poset elements).

To describe our process we need a few definitions. Let $G = (V, E)$ be a graph, and let M together with values γ_v be a stable outcome of network exchange on G . Let A be a subset of the nodes. We will consider the values γ_v “set” for $v \in A$, but we’ll maintain values for all nodes $v \in V$. (The values on nodes $V - A$ are viewed as still changeable.) We will only work with sets A having the property that edges in the matching M are either contained in A or are disjoint from A ; the set A will also have the property that it contains all unmatched nodes, with $\gamma_v = 0$ for each unmatched node v .

For running the process to find a balanced outcome, we will let σ_v be defined only by the values that are “set”; that is, we let $\sigma_v(A)$ for a node $v \in A$ be

$$\sigma_v(A) = \min(\gamma_v, \min_{(v,u) \in E: u \neq w, u \in A} \sigma_{uw}),$$

the minimum of the value γ_v and the smallest slack of an unmatched edge adjacent to v inside A . We say that the outcome is *balanced in A* if for all matching edges (v, w) with $v, w \in A$, we have $\sigma_v(A) = \sigma_w(A)$; and we will assume that for all nodes $v \in A$, we have $\sigma_v(A) \leq \sigma$, while all slacks of nodes and edges not contained in A are at least σ .

To prove that this process finds a balanced outcome we will use the following strengthening of Theorem 5.2.

THEOREM 5.4. *Consider a graph $G = (V, E)$, with a matching M and values γ_v constituting a stable outcome of network exchange on G . Let A be a subset of the nodes such that $\sigma_v \leq \sigma$ for all nodes $v \in A$, and $\sigma_v \geq \sigma$ for all nodes $v \notin A$, and suppose γ_v is balanced in A . Then there is a balanced outcome γ' such that $\gamma_v = \gamma'_v$ on all nodes $v \in A$.*

Consider the following linear program seeking to maximize the smallest slack of all edges and nodes not in A . Let $E(A)$ be the set of edges contained in A .

$$\begin{aligned} \max \sigma' \\ \gamma'_v &\geq \sigma' \quad \forall v \in V \setminus A \\ \gamma'_v + \gamma'_w &= \omega_{vw} \quad \forall (v, w) \in M \\ \gamma'_v + \gamma'_w &\geq \omega_{vw} + \sigma' \quad \forall (v, w) \in E \setminus (M \cup E(A)) \\ \gamma'_v &= \gamma_v \quad \forall v \in A. \end{aligned}$$

We will proceed using the Aspvall-Shiloach characterization of feasibility mentioned above. We find a structure C from their theorem that limits the maximum σ' in the linear program, and we “spread out” this available slack σ' across C . Recall that an alternating path or blossom is *anchored* at the end(s) of the path; because values γ_v are already set for nodes in A , we consider only alternating paths and blossoms anchored either at a matching edge or at a node in A .

To do this, we need to define the available *slack* $\sigma(C)$ for any of the structures arising from the Aspvall-Shiloach theorem. This available slack will be built from a simpler quantity $\phi(C)$, defined for all these structures as the difference in weight between matched and unmatched edges, counting the path portions of blossoms and bi-cycles twice. Also, the extent to which we can spread out the slack depends on a

quantity $n(C)$ that captures the number of distinct slack values over which it can be spread out.

We now define the slack $\sigma(C)$, as well as $n(C)$, in several cases, depending on the structure of C . First, for a cycle, or a path anchored in matching edges, we simply define $\sigma(C) = \phi(C)$. If $\sigma(C) < 0$ then augmenting the matching along C results in a matching of larger weight. If $\sigma(C) \geq 0$ then in any stable solution the sum of the slacks of the unmatched edges along C (plus the slacks of the end nodes) sum to $\sigma(C)$, and hence $\sigma(C)/n(C)$ is a bound on the maximum slack σ' , where $n(C)$ is the number of unmatched edges along C for a cycle, and the number of unmatched edges plus two for a path, so as to also count the slack variables of the end nodes.

This observation also extends to a bi-cycle or blossom C anchored at matching edges: again we define $\sigma(C) = \phi(C)$. We define $n(C)$ by counting the edges (and the end node) along the path portion twice. It is not hard to see that if $\sigma(C) < 0$, then no stable outcome can exist (even though the matching may be of maximum weight). Also, since $\sigma(C)$ is the sum of the slacks of the unmatched edges (and end node) along the structure, $\sigma(C)/n(C)$ bounds the maximum slack.

For a blossom C anchored at a node $v \in A$, we define $\sigma(C) = \phi(C) - 2\gamma_v$; for a path P anchored at nodes $v, w \in A$, we define $\sigma(P) = \phi(P) - \gamma_v - \gamma_w$; and finally, if a path P is anchored at one end at a node $v \in A$, and at a matched edge at the other end, we define $\sigma(P) = \phi(P) - \gamma_v$. Note that the value depends only on the values γ_v for the nodes $v \in A$ and the structure, and doesn’t depend on any values not yet “set”. Let $n(C)$ for such structures denote the number of slack variables not yet set along the structure, i.e., not counting slacks for the ends anchored in A . As before we have that $\sigma(C)$ is the sum of the slacks of the unmatched edges along the structure plus the ends not anchored in A , and hence $\sigma(C)/n(C)$ is a bound on the maximum slack.

From the Aspvall-Shiloach characterization of feasibility [1] we get that the highest possible value σ' is limited by the smallest ratio over the structures C .

PROPOSITION 5.5. *The maximum value of σ' of the linear program above is equal to the smallest value of $\sigma(C)/n(C)$, where C ranges over all possible alternating cycles, paths, blossoms, and bi-cycles.*

Proof of Theorem 5.4. First we put each unmatched node v into A and set $\gamma_v = 0$ for such nodes. We then proceed by induction on the number of nodes in $V \setminus A$. We use the Aspvall-Shiloach dynamic programming algorithm [1] with σ' as a parameter to solve this linear program. Let γ^* be a solution with value σ^* . Recall that by assumption we have $\sigma^* \geq \sigma$.

By Proposition 5.5 the maximum of σ^* is limited by a path, cycle, blossom or bi-cycle C . We find such a limiting structure C , and let $A^* = A \cup C$. Note that the slacks σ_{vw} for unmatched edges in C must equal σ^* as well as $\gamma_v = \sigma^*$ if C is anchored at a node $v \notin A$, as these values average to σ^* .

We claim γ^* , σ^* and A^* satisfy the theorem’s conditions, which will finish the proof by induction on the size of $V \setminus A$.

First note that the nodes $v \in A$ remain balanced in A . This is true as all slacks not fully contained in A are at least $\sigma^* \geq \sigma$ and hence for nodes in A the minimum slack value is attained as before: σ_v does not change for $v \in A$. Second, matching edges do not leave A^* . This follows as the limiting

structure C contains the matched edge adjacent to each of its nodes. Finally, note that all nodes newly added to A^* are balanced, as all nodes in $A^* \setminus A$ have slack σ^* . ■

Representing all balanced outcomes. Finally, we want to argue that *all* balanced outcomes can be obtained by this procedure. Before turning to the proof of this statement, we consider to what extent the execution of the procedure was uniquely determined. If C is a path or blossom, then the values γ_v^* are fully determined by σ^* and the values of the anchored ends (if anchored in A), and hence the values set in this iteration are unique. In this case, we can show that all balanced solutions γ' that agree with γ on A (i.e., have $\gamma'_v = \gamma_v$ for all $v \in A$) must also have $\gamma'_v = \gamma_v^*$ for all $v \in C$. However, when the limiting structure C is a cycle or bi-cycle, then the values set in this iteration may not be fully determined just by σ^* and the structure. Setting a value γ_v^* for a node $v \in C$ fully determines the values at the other nodes, but there may be a range of values possible for γ_v^* . (Note that this is a more complex analogue of the freedom we had to set values on nodes in the non-trivial elementary subgraphs of a unit-edge-weight bipartite graph.) The interval of possible values for such a node v can also be determined by the Aspvall-Shiloach [1] dynamic program.

THEOREM 5.6. *Consider an exchange network $G = (V, E)$ with stable outcome values γ_v , and a value σ , and let A be a subset of nodes v such that $\sigma_v \leq \sigma$ for all nodes $v \in A$, and $\sigma_v \geq \sigma$ for all nodes $v \notin A$. Suppose all matching edges (v, w) internal to A are balanced (i.e., have $\sigma_v = \sigma_w$).*

If γ^ is a set of balanced outcome values that satisfies $\gamma_v^* = \gamma_v$ for all nodes $v \in A$, then it is an optimal solution for the above linear program.*

Proof. Let γ^* be a balanced outcome that agrees with γ on all nodes in A . Let σ^* be the smallest slack of any unmatched edge or node not contained in A . We need to argue that $\sigma^* \geq \sigma'$. We will find a structure (alternating path, cycle, blossom or bi-cycle) containing v with slack equal to this minimum value σ^* . The average slack $\sigma(C)/n(C)$ for this structure is now also σ^* and hence $\sigma^* \geq \sigma'$ as claimed.

We find the structure by following an alternating walk on nodes and edges with slack equal to σ^* . We start with the edge or node where the smallest slack occurs outside of A . Now suppose we have followed such a walk for some number of nodes, and consider an alternating path P found so far; we describe how to extend P in both directions. Let v be the last node of the path. Assume first that P so far ends with an unmatched edge (u, v) (and hence $\sigma_{uv}^* = \sigma^*$). If $v \in A$ this will anchor the path. Otherwise we have $v \notin A$, and now consider v 's matched pair w . The outcome γ^* is balanced, so $\sigma_v^* = \sigma_w^* = \sigma^*$, and we add the matched edge (v, w) to the path P . If P so far ends with a matched edge, let w be the last node. If $\sigma_w^* = \gamma_w^*$ then node w will anchor our path, else we extend the path with the unmatched edge (w, x) that is defining the slack $\sigma_w^* = \sigma_{wx}^*$.

Following the walk P in both directions, P can either get anchored, or close a cycle. If the cycle C closed is even, then we have found the structure C , as claimed. If the cycle C is odd, we continue the walk P . Note that the walk may now visit nodes and edges for the second time, but these visits will occur in the opposite direction. This process will result in finding the claimed structure: we either close an

even alternating cycle; or find two odd cycles, closing a bi-cycle; or we get both ends of the walk anchored, closing an alternating path or blossom. ■

Acknowledgments. We thank Larry Blume, David Easley, and Ankur Moitra for many discussions and insights, and Michael Macy, Arnout van de Rijt, and Shane Thye for discussions on the sociological context.

6. REFERENCES

- [1] B. Aspvall, Y. Shiloach. A polynomial time algorithm for solving systems of linear inequalities with two variables per inequality. *SIAM J. Computing*, 9(1980).
- [2] A. Calvó-Armengol. Bargaining power in communication networks. *Math. Soc. Sci.*, 41(2001).
- [3] G. Charness, M. Corominas-Bosch, and G. R. Frechette. Bargaining and network structure: An experiment. *Econometric Society North Am. Summer Mtgs.*, 2004.
- [4] K. S. Cook and R. M. Emerson. Power, equity, and commitment in exchange networks. *American Sociological Review*, 43(5):721–739, Oct. 1978.
- [5] K. S. Cook, R. M. Emerson, M. R. Gillmore, and T. Yamagishi. The distribution of power in exchange networks: Theory and experimental results. *American Journal of Sociology*, 89:275–305, 1983.
- [6] K. S. Cook and T. Yamagishi. Power in exchange networks: A power-dependence formulation. *Social Networks*, 14:245–265, 1992.
- [7] M. Corominas-Bosch. Bargaining in a network of buyers and sellers. *J. Economic Theory*, 115(2004).
- [8] R. M. Emerson. Power-dependence relations. *American Sociological Review*, 27:31–40, 1962.
- [9] S. M. Kakade, M. J. Kearns, L. E. Ortiz, R. Pemantle, and S. Suri. Economic properties of social networks. In *Proc. 17th Adv. Neural Info. Proc. Sys (NIPS)*, 2004.
- [10] R. Kranton and D. Minehart. A theory of buyer-seller networks. *American Economic Review*, 91(3), 2001.
- [11] L. Lovász and M. Plummer. *Matching Theory*. North-Holland, 1986.
- [12] J. Lucas, C. Younts, M. Lovaglia, B. Markovsky. Lines of power in exchange networks. *Soc. Forces*, 80(2001).
- [13] J. Nash. The bargaining problem. *Econometrica*, 18:155–162, 1950.
- [14] N. Navarro, A. Perea. Bargaining in networks and the Myerson value. Working paper, Univ. Carlos III, 2001.
- [15] M. J. Osborne. *An Introduction to Game Theory*. Oxford University Press, 2003.
- [16] A. E. Roth and M. Sotomayor. *Two-sided matching: A study in game-theoretic modeling and analysis*. Cambridge University Press, 1990.
- [17] J. Skvoretz and D. Willer. Exclusion and power: A test of four theories of power in exchange networks. *American Sociological Review*, 58:801–818, 1993.
- [18] V. V. Vazirani. Combinatorial algorithms for market equilibria. In N. Nisan, T. Roughgarden, É. Tardos, and V. Vazirani, editors, *Algorithmic Game Theory*, pages 103–134. Cambridge University Press, 2007.
- [19] D. Willer (editor). *Network Exchange Theory*. Praeger, 1999.