

Optimizing Social Welfare for Network Bargaining Games in the Face of Unstability, Greed and Spite

T.-H. Hubert Chan, Fei Chen, and Li Ning

Department of Computer Science, The University of Hong Kong
{hubert,fchen,lning}@cs.hku.hk

Abstract. Stable and balanced outcomes of network bargaining games have been investigated recently, but the existence of such outcomes requires that the linear program relaxation of a certain maximum matching problem has integral optimal solution.

We propose an alternative model for network bargaining games in which each edge acts as a player, who proposes how to split the weight of the edge among the two incident nodes. Based on the proposals made by all edges, a selection process will return a set of accepted proposals, subject to node capacities. An edge receives a commission if its proposal is accepted. The social welfare can be measured by the weight of the matching returned.

The node users, as opposed to being rational players as in previous works, exhibit two characteristics of human nature: greed and spite. We define these notions formally and show that the distributed protocol by Kanoria et. al can be modified to be run by the edge players such that the configuration of proposals will converge to a pure Nash Equilibrium, without the LP integrality gap assumption. Moreover, after the nodes have made their greedy and spiteful choices, the remaining ambiguous choices can be resolved in a way such that there exists a Nash Equilibrium that will not hurt the social welfare too much.

1 Introduction

Bargaining games have been studied with a long history, early in economics [19] and sociology, and recently in computer science, there has been a lot of attention on bargaining games in social exchange networks [16,1,4,15], in which users are modeled as nodes in an undirected graph $G = (V, E)$, whose edges are weighted. An edge $\{i, j\} \in E$ with weight $w_{ij} > 0$ means that users i and j can potentially form a contract with each other and split a profit of w_{ij} . A capacity vector $b \in \mathbb{Z}_+^V$ limits the maximum number b_i of contracts node i can form with its neighbors, and the set M of executed contracts form a b -matching in G .

In previous works, the nodes bargain with one another to form an *outcome* which consists of the set M of executed contracts and how the profit in each contract is distributed among the two participating nodes. The outside option of a node is the maximum profit the node can get from another node with whom there

is no current contract. An outcome is *stable* if for every contract a node makes, the profit the node gets from that contract is at least its outside option. Hence, under a stable outcome, no node has motivation to break its current contract to form another one. Extending the notion of Nash bargaining solution [19], Cook and Yamagishi [11] introduced the notion of balanced outcome. An outcome is balanced if, in addition to stability, for every contract made, after each participating node gets its outside option, the surplus is divided equally between the two nodes involved. For more notions of solutions, the reader can refer to [7].

Although stability is considered to be an essential property, as remarked in [4,15], a stable outcome exists *iff* the linear program (LP) relaxation (given in Section 4) for the b -matching problem on the given graph G has integrality gap 1. Hence, even for very simple graphs like a triangle with unit node capacities and unit edge weights, there does not exist a stable outcome. Previous works simply assumed that the LP has integrality gap 1 [15,8] or considered restriction to bipartite graphs [16,4], for which the LP always has integrality gap 1.

We think the integrality gap condition is a limitation to the applicability of such framework in practice. We would like to consider an alternative model for network bargaining games, and investigate different notions of equilibrium, whose existence does not require the integrality gap condition.

Our Contribution and Results. In this work, we let the edges take over the role of the “rational” players from the nodes. Each edge $e = \{i, j\} \in E$ corresponds to an agent, who proposes a way to divide up the potential profit w_{ij} among the two nodes. Formally, each edge $\{i, j\}$ has the action set $A_{ij} := \{(x, y) : x \geq 0, y \geq 0, x + y \leq w_e\}$, where a proposal (x, y) means that node i gets amount x and j gets amount y .¹ Based on the configuration $m \in A_E := \times_{e \in E} A_e$ of proposals made by all the agents, a selection process (which can be randomized) will choose a b -matching M , which is the set of contracts formed. An agent e will receive a commission if his proposal is selected; his payoff $u_e(m)$ is the probability that edge e is in the matching M returned.² Observe that once the payoff function u is defined, the notion of (pure or mixed) Nash Equilibrium is also well-defined. We measure the social welfare $\mathcal{S}(m)$ by the (expected) weight $w(M)$ of the matching M returned, which reflects the volume of transaction.

We have yet to describe the selection process, which will determine the payoff function to each agent, and hence will affect the corresponding Nash Equilibrium. We mention earlier that the rational players in our framework will be the edges, as opposed to the nodes in previous works; in fact, in the selection process we assume the node users will exhibit two characteristics of human nature: greed and spite.

Greedy Users. For a node i with capacity b_i , user i will definitely want an offer that is strictly better than his $(b_i + 1)$ -st best offer. If this happens for both users forming an edge, then the edge will definitely be selected. We also say the resulting payoff function is greedy.

¹ In case $x + y < w_{ij}$ the remaining amount is lost and not gained by anyone.

² The actual gain of an agent could be scaled according to the weight w_e , but this will not affect the Nash Equilibrium.

Spiteful Users. Spite is an emotion that describes the situation that once a person has seen a better offer, he would not settle for anything less, even if the original better offer is no longer available. If user i with capacity b_i sees that an offer is strictly worse than his b_i -th best offer, then the corresponding edge will definitely be rejected. We also say the resulting payoff function is spiteful.

One can argue that greed is a rational behavior (hence the regime of greedy algorithms), but spite is clearly not always rational. In fact, we shall see in Section 2 that there exist a spiteful payoff function and a configuration of agent proposals that is a pure Nash Equilibrium, in which all proposals are rejected by the users out of spite, even though no single agent can change the situation by unilaterally offering a different proposal. The important question is that: can the agents follow some protocol that can avoid such bad Nash Equilibrium? In other words, can they collaboratively find a Nash Equilibrium that achieves good social welfare?

We answer the above question in the affirmative. We modify the distributed protocol of Kanoria et. al [15] to be run by edge players and allow general node capacities b . As before, the protocol is iterative and the configuration of proposals returned will converge to a fixed point m of some non-expansive function \mathcal{T} . In Section 3, we show that provided the payoff function u is greedy and spiteful, then any fixed point m of \mathcal{T} is in the corresponding set \mathcal{N}_u of pure Nash Equilibria.

In Section 4, we analyze the social welfare through the linear program (LP) relaxation of the maximum b -matching problem. As in [15], we investigate the close relationship between a fixed point of \mathcal{T} and (LP). However, we go beyond previous analysis and do not need the integrality gap assumption, i.e., (LP) might not have an integral optimum. We show that when greedy users choose an edge, then all (LP) optimal solutions must set the value of that edge to 1; on the other hand, when users reject an edge out of spite, then all (LP) optimal solutions will set the value of that edge to 0. We do need some technical assumptions in order for our results to hold: either (1) (LP) has a unique optimum, or (2) the given graph G has no even cycle such that the sum of the weights of the odd edges equals that of the even edges; neither assumption implies the other, but both can be achieved by perturbing slightly the edge weights of the given graph. Unlike the case for simple 1-matching, we show (in the full version) that assumption (2) is necessary for general b -matching, which indicates that there is some fundamental difference between the two cases.

The greedy behavior states that some edges must be selected and the spiteful behavior requires that some edges must be rejected. However, there is still some freedom to deal with the remaining *ambiguous* edges.³ Observe that a fixed point will remain a Nash Equilibrium (for the edge players) no matter how the ambiguous edges are handled, so it might make sense at this point to maximize the total number of extra contracts made from the ambiguous edges. However, optimizing the cardinality of a matching can be arbitrarily bad in terms of weight, but a maximum weight matching is a 2-approximation in terms of cardinality. Therefore, in Section 5, we consider a greedy and spiteful payoff function u that

³ As a side note, we remark that our results implies that under the unique integral (LP) optimum assumption, there will be no ambiguous edges left.

corresponds to selecting a maximum weight matching (approximate or exact) among the ambiguous edges (subject to remaining node capacities b'); in reality, we can imagine this corresponds to a centralized clearing process or a collective effort performed by the users. We show that if a $(1+c)$ -approximation algorithm for maximum weight matching is used for the ambiguous edges, then the social welfare is at least $\frac{2}{3(1+c)}$ fraction of the social optimum, i.e., the price of stability is $1.5(1+c)$. Finally, observe that the iterative protocol we mention will converge to a fixed point, but might never get there exactly; hence, we relax the notions of greed and spite in Section 5.1 and show that the same guarantee on the price of stability can be achieved eventually (and quickly).

We remark that if the topology of the given graph and the edge weights naturally indicate that certain edges should be selected while some should be rejected (both from the perspectives of social welfare and selfish behavior), then our framework of greed and spite can detect these edges. However, we do not claim that our framework is a silver bullet to all issues; in particular, for the triangle example given above, all edges will be ambiguous and our framework simply implies that one node will be left unmatched, but does not specify how this node is chosen. We leave as future research direction to develop notions of fairness in such situation.

Related Work. Kleinberg and Tardos [16] recently started the study of network bargaining games in the computer science community; they showed that a stable outcome exists *iff* a balanced outcome exists, and both can be computed in polynomial time, if they exist. Chakraborty et. al [9,10] explored equilibrium concepts and experimental results for bipartite graphs. Celis et. al [8] gave a tight polynomial bound on the rate of convergence for unweighted bipartite graphs with a unique balanced outcome. Kanoria [14] considered *unequal division* (UD) solutions for bargaining games, in which stability is still guaranteed while the surplus is split with ratio $r : 1 - r$, where $r \in (0, 1)$. They provided an FPTAS for the UD solutions assuming the existence of such solutions.

Azar et. al [1] considered a local dynamics that converges to a balanced outcome provided that it exists. Assuming that the LP relaxation for matching has a unique integral optimum, Kanoria et. al [15] designed a local dynamics that converges in polynomial time. Our distributed protocol is based on [15], but is generalized to general node capacities, run by edges and does not require the integrality condition on (LP).

Batani et. al [4] also considered general node capacities; moreover, they showed that the network bargaining problem can be recast as an instance of the well-studied cooperative game [12]. In particular, a stable outcome is equivalent to a point in the core of a cooperative game, while a balanced outcome is equivalent to a point in the core and the prekernel. Azar et. al [2] also studied bargaining games from the perspective of cooperative games, and proved some monotonicity property for several widely considered solutions.

In our selection process, we assume that the maximum weight b' -matching problem is solved on the ambiguous edges. This problem is well-studied and can be solved exactly in polynomial time [22][Section 33.4]; moreover, the problem

can be solved by a distributed algorithm [5], and $(1 + c)$ -approximation for any $c > 0$ can be achieved in poly-logarithmic time [18,20,17].

2 Notation and Preliminaries

Consider an undirected graph $G = (V, E)$, with vertex set V and edge set E . Each node $i \in V$ corresponds to a *user* i (vertex player), and each edge $e \in E$ corresponds to an *agent* e (edge player). Agents arrange contracts to be formed between users where each agent $e = \{i, j\}$ gains a commission when users i and j form a contract. Each edge $e = \{i, j\} \in E$ has weight $w_e = w_{ij} > 0$, which is the maximum profit that can be shared between users i and j if a contract is made between them. Given a node i , denoted by $N(i) := \{j \in V : \{i, j\} \in E\}$ the set of its neighbors in G , there exists a capacity vector $b \in \mathbb{Z}_+^V$ such that each node i can make at most b_i contracts with its neighbors in $N(i)$, where at most one contract can be made between a pair of users; hence, the set M of edges on which contracts are made is a b -matching in G .

Agent Proposal. For each $e = \{i, j\} \in E$, agent e makes a proposal of the form $(m_{j \rightarrow i}, m_{i \rightarrow j})$ from an action set A_e to users i and j , where $A_e := \{(x, y) : x \geq 0, y \geq 0, x + y \leq w_{ij}\}$, such that if users i and j accepts the proposal and form a contract with each other, user i will receive $m_{j \rightarrow i}$ and user j will receive $m_{i \rightarrow j}$ from this contract.

Selection Procedure and Payoff Function u . Given a configuration $m \in A_E := \times_{e \in E} A_e$ of all agent's proposals, some selection procedure is run on m to return a b -matching M , where an edge $e = \{i, j\} \in M$ means that a contract is made between i and j . The procedure can be (1) deterministic or randomized, (2) centralized or (more preferably) distributed.

If i and j are *matched* in M , i.e., $e = \{i, j\} \in M$, agent e will receive a commission, which can either be fixed or a certain percentage of w_e ; since an agent either gains the commission or not, we can assume that its payoff is 1 when a contract is made and 0 otherwise. Hence, the selection procedure defines a payoff function $u = \{u_e : A_E \rightarrow [0, 1] | e \in E\}$, such that for each $e \in E$, $u_e(m)$ is the probability that the edge e is in the b -matching M returned when the procedure is run on $m \in A_E$. We shall consider different selection procedures, which will lead to different payoff functions u . However, the selection procedure should satisfy several natural properties, which we relate to the human nature of the users as follows.

We use $\max^{(b)}$ to denote the b -th maximum value among a finite set of numbers (by convention it is 0 if there are less than b numbers). Given $m \in A_E$, we define $\widehat{m}_i = \max_{j \in N(i)}^{(b_i)} m_{j \rightarrow i}$ and $\overline{m}_i = \max_{j \in N(i)}^{(b_i+1)} m_{j \rightarrow i}$.

Greedy Users. If both users i and j see that they cannot get anything better from someone else, then they will definitely make a contract with each other. Formally, we say that the payoff function u is *greedy* (or the users are greedy), if for each $e = \{i, j\} \in E$ and $m \in A_E$, if $m_{j \rightarrow i} > \overline{m}_i$ and $m_{i \rightarrow j} > \overline{m}_j$, then $u_e(m) = 1$.

Spiteful Users. It is human nature that once a person has seen the best, they will not settle for anything less. We try to capture this behavior formally. We say that the payoff function u is *spiteful* (or the users are spiteful) if for each $e = \{i, j\} \in E$ and $m \in A_E$, if $m_{j \rightarrow i} < \widehat{m}_i$, then $u_e(m) = 0$, i.e., if user i cannot get the b_i -th best offer from j , then no contract will be formed between i and j .

Game Theory and Social Welfare. We have described a game between the agents, in which agent e has the action set A_e , and has payoff function u (determined by the selection procedure). In this paper, we consider pure strategies and pure Nash Equilibria. A configuration $m \in A_E$ of actions is a Nash equilibrium if no single player can increase its payoff by unilaterally changing its action.

Given a payoff function u , we denote by $\mathcal{N}_u \subset A_E$ the set of Nash Equilibria. Given a configuration $m \in A_E$ of proposals and a payoff function u , we measure social welfare by $\mathcal{S}_u(m) := \sum_{e \in E} w_e \cdot u_e(m)$, which is the expected weight of the b -matching returned. When there is no ambiguity, the subscript u is dropped. The optimal social welfare $\mathcal{S}^* := \max_{m \in A_E} \mathcal{S}(m)$ is the maximum weight b -matching; to achieve the social optimum, given a maximum weight b -matching M , every agent $e \in M$ proposes $(\frac{w_e}{2}, \frac{w_e}{2})$, while other agents proposes $(0, 0)$. The weight of the b -matching can be an indicator of the volume of transactions or how active the market is. The Price of Anarchy (PoA) is defined as $\frac{\mathcal{S}^*}{\min_{m \in \mathcal{N}} \mathcal{S}(m)}$ and the Price of Stability (PoS) is defined as $\frac{\mathcal{S}^*}{\max_{m \in \mathcal{N}} \mathcal{S}(m)}$.

Proposition 1 (Infinite Price of Anarchy). *There exists an instance of the game such that when the users are spiteful, there exists a Nash Equilibrium $m \in A_E$ under which no contracts are made.*

We defer the proof of Proposition 1 to the full version.

3 A Distributed Protocol for Agents

We describe a distributed protocol for the agents to update their actions in each iteration. The protocol is based on the one by Kanoria et. al [15], which is run by nodes and designed for (1-)matchings. The protocol can easily be generalized to be run by edges and for general b -matchings. In each iteration, two agents only need to communicate if their corresponding edges share a node. Given a real number $r \in \mathbb{R}$, we denote $(r)_+ := \max\{r, 0\}$. Moreover, as described in [15,3] a damping factor $\kappa \in (0, 1)$ is used in the update; we can think of $\kappa = \frac{1}{2}$.

Although later on we will also consider the LP relaxation of b -matching, unlike previous works [21,6,15], we do not require the assumption that the LP relaxation has a unique integral optimum.

In Algorithm 1, auxiliary variables $\alpha^{(t)} \in \mathbb{R}_+^{2|E|}$ are maintained. Intuitively, the parameter $\alpha_{i \setminus j}$ is meant to represent the b_i -th best offer user i can receive if user j is removed. Suppose $W := \max_{e \in E} w_e$ and we define a function $\mathcal{T} : [0, W]^{2|E|} \rightarrow [0, W]^{2|E|}$ as follows.

Given $\alpha \in [0, W]^{2|E|}$, for each $\{i, j\} \in E$, define the following quantities.

$$S_{ij}(\alpha) = w_{ij} - \alpha_{i \setminus j} - \alpha_{j \setminus i} \quad (1)$$

```

Input:  $G = (V, E, w)$ 
Initialization: For each  $e = \{i, j\} \in E$ , agent  $e$  picks arbitrary
 $(m_{j \rightarrow i}^{(0)}, m_{i \rightarrow j}^{(0)}) \in A_e$ .
for agent  $e = \{i, j\} \in E$  do
     $\alpha_{i \setminus j}^{(1)} := \max_{k \in N(i) \setminus j}^{(b_i)} m_{k \rightarrow i}^{(0)}$ ;  $\alpha_{j \setminus i}^{(1)} := \max_{k \in N(j) \setminus i}^{(b_j)} m_{k \rightarrow j}^{(0)}$ 
end
for  $t = 1, 2, 3, \dots$  do
    for agent  $e = \{i, j\} \in E$  do
         $S_{ij} := w_{ij} - \alpha_{i \setminus j}^{(t)} - \alpha_{j \setminus i}^{(t)}$ 
         $m_{j \rightarrow i}^{(t)} := (w_{ij} - \alpha_{j \setminus i}^{(t)})_+ - \frac{1}{2}(S_{ij})_+$ ;  $m_{i \rightarrow j}^{(t)} := (w_{ij} - \alpha_{i \setminus j}^{(t)})_+ - \frac{1}{2}(S_{ij})_+$ 
    end
    for agent  $e = \{i, j\} \in E$  do
         $\alpha_{i \setminus j}^{(t+1)} = (1 - \kappa) \cdot \alpha_{i \setminus j}^{(t)} + \kappa \cdot \max_{k \in N(i) \setminus j}^{(b_i)} m_{k \rightarrow i}^{(t)}$ ;
         $\alpha_{j \setminus i}^{(t+1)} = (1 - \kappa) \cdot \alpha_{j \setminus i}^{(t)} + \kappa \cdot \max_{k \in N(j) \setminus i}^{(b_j)} m_{k \rightarrow j}^{(t)}$ 
    end
end
    
```

Algorithm 1. A Distributed Protocol for Agents. For each time t , agent $e = \{i, j\}$ computes its action $(m_{j \rightarrow i}^{(t)}, m_{i \rightarrow j}^{(t)}) \in A_e$; the first value is sent to other edges incident on i and the second to edges incident on j .

$$m_{j \rightarrow i}(\alpha) = (w_{ij} - \alpha_{j \setminus i})_+ - \frac{1}{2}(S_{ij}(\alpha))_+ \quad (2)$$

Then, we define $\mathcal{T}(\alpha) \in [0, W]^{2|E|}$ by $(\mathcal{T}(\alpha))_{i \setminus j} := \max_{k \in N(i) \setminus j}^{(b_i)} m_{k \rightarrow i}(\alpha)$. It follows that Algorithm 1 defines the sequence $\{\alpha^{(t)}\}_{t \geq 1}$ by $\alpha^{(t+1)} := (1 - \kappa)\alpha^{(t)} + \kappa\mathcal{T}(\alpha^{(t)})$.

Given a vector space D , a function $T : D \rightarrow D$ is *non-expansive* under norm $\|\cdot\|$ if for all $x, y \in D$, $\|T(x) - T(y)\| \leq \|x - y\|$; a point $\alpha \in D$ is a fixed point of T if $T(\alpha) = \alpha$. As in [15], it can be proved that the function \mathcal{T} is non-expansive, and using a result by Ishikawa [13] (we defer the details to the full version), the next theorem follows.

Theorem 1 (Convergence to a Fixed Point). *The distributed protocol shown in Figure 1 maintains the sequence $\{\alpha^{(t)}\}$ which converges to a fixed point of the function \mathcal{T} under the ℓ_∞ norm.*

Properties of a Fixed Point. Given a fixed point α of the function \mathcal{T} , the quantities $S \in \mathbb{R}^{|E|}$ and $m \in A_E$ are defined according to Equations (1) and (2). We also say that (m, α, S) , (m, α) or m is a fixed point (of \mathcal{T}). Similar to [15], we give several important properties of a fixed point, whose details are given in the full version. Theorems 1 and 2 imply that as long as the payoff function is greedy and spiteful, the game defined between the agents (edge players) always has a pure Nash Equilibrium. We defer the proof of Theorem 2 to the full version.

Theorem 2 (Fixed Point is NE). *Suppose the payoff function u is greedy and spiteful. Then, any fixed point $m \in A_E$ of \mathcal{T} is a Nash Equilibrium in \mathcal{N}_u .*

4 Analyzing Social Welfare via LP Relaxation

Theorem 2 states that a fixed point (m, α) of the function \mathcal{T} is a Nash Equilibrium in \mathcal{N}_u , as long as the underlying payoff function is greedy and spiteful. Our goal is to show that there exists some greedy and spiteful u such that the fixed point m also achieves good social welfare $\mathcal{S}_u(m) = \sum_{e \in E} w_e \cdot u_e(m)$.

As observed by Kanoria et. al [15], the network bargain game is closely related to the linear program (LP) relaxation of the b -matching problem, which has the form $\mathcal{S}_{LP} := \max_{x \in \mathcal{L}} w(x)$, where $w(x) := \sum_{\{i,j\} \in E} x_{ij} w_{ij}$ and $\mathcal{L} := \{x \in [0, 1]^E : \forall i \in V, \sum_{j: \{i,j\} \in E} x_{ij} \leq b_i\}$ is the set of feasible fractional solutions. Given $x \in \mathcal{L}$, we say a node i is *saturated* under x if $\sum_{j: \{i,j\} \in E} x_{ij} = b_i$, and otherwise *unsaturated*.

They showed that when the LP relaxation has a unique integral maximum, a fixed point (m, α, S) corresponds naturally to the unique maximum (1-)matching. However, their analysis cannot cover the case when the optimal solution is fractional or when the maximum matching is not unique.

In this section, we fully exploit the relationship between a fixed point and the LP relaxation, from which we show that good social welfare can be achieved. Note that we do not require the unique integral optimum assumption. On the other hand, we assume that either (1) the LP has a unique optimum or (2) the following technical assumption.

No Cycle with Equal Alternating Weight. We say that a cycle has *equal alternating weight* if it is even, and the sum of the odd edges equals that of the even edges. We assume that the given weighted graph G has no such cycle. The weights of any given graph can be perturbed slightly such that this condition holds. Observe that the optimum of (LP) might not be unique even with this assumption. For the 1-matching case, the conclusions listed in Theorem 3 have been derived by Kanoria et. al [15] without the “no cycle with equal alternating weight” assumption. However, for general b -matching, this assumption in Theorem 3 is necessary for cases (b) and (c). The reason is shown in the full version.

Theorem 3 (Fixed Point and LP). *Suppose (LP) has a unique optimum or the graph G has no cycle with equal alternating weight, and (m, α, S) is a fixed point of \mathcal{T} . Then, for any edge $\{i, j\} \in E$, the following holds.*

- (a) *Suppose (LP) has a unique integral optimum corresponding to the maximum b -matching M^* . Then, $S_{ij} \geq 0$ implies that $\{i, j\} \in M^*$.*
- (b) *Suppose $S_{ij} > 0$. Then, any optimal solution x to (LP) must satisfy $x_{ij} = 1$.*
- (c) *Suppose $S_{ij} < 0$. Then, any optimal solution x to (LP) must satisfy $x_{ij} = 0$.*

Although the three statements in Theorem 3 look quite different, they can be implied by the three similar-looking corresponding statements in the following lemma.

Lemma 1 (Fixed Point and LP). *Suppose (m, α, S) is a fixed point of \mathcal{T} , and x is a feasible solution to (LP). Then, for each $\{i, j\} \in E$, the following properties hold.*

- (a) If $S_{ij} \geq 0$ and $x_{ij} = 0$, then there is $\hat{x} \in \mathcal{L}$ such that $\hat{x} \neq x$ and $w(\hat{x}) \geq w(x)$.
 (b) If $S_{ij} > 0$ and $x_{ij} < 1$, then there is $\hat{x} \in \mathcal{L}$ such that $\hat{x} \neq x$ and $w(\hat{x}) \geq w(x)$.
 (c) If $S_{ij} < 0$ and $x_{ij} > 0$, then there is $\hat{x} \in \mathcal{L}$ such that $\hat{x} \neq x$ and $w(\hat{x}) \geq w(x)$.

Moreover, strict inequality holds for (b) and (c), if in addition the graph G has no cycle with equal alternating weight.

4.1 Finding Alternative Feasible Solution via Alternating Traversal

Lemma 1 shows the existence of alternative feasible solutions under various conditions. We use the unifying framework of the *alternating traversal* to show its existence.

Alternating Traversal. Given a fixed point (m, α, S) of \mathcal{T} and a feasible solution $x \in \mathcal{L}$, we define a structure called *alternating traversal* as follows.

(1) An alternating traversal \mathcal{Q} (with respect to (m, α, S) and x) is a path or circuit (not necessarily simple and might contain repeated edges), which alternates between two disjoint edge sets \mathcal{Q}^+ and \mathcal{Q}^- (hence \mathcal{Q} can be viewed as a multiset which is the disjoint union of \mathcal{Q}^+ and \mathcal{Q}^-) such that $\mathcal{Q}^+ \subset S^+$ and $\mathcal{Q}^- \subset S^-$, where $S^+ := \{e \in E : S_e \geq 0\}$ and $S^- := \{e \in E : S_e \leq 0\}$.

The alternating traversal is called *feasible* if in addition $\mathcal{Q}^+ \subset E^+$ and $\mathcal{Q}^- \subset E^-$, where $E^+ := \{e \in S^+ : x_e < 1\}$ and $E^- := \{e \in S^- : x_e > 0\}$.

An edge e is called **critical** if e is in exactly one of E^+ and E^- , and is called **strict** if $S_e \neq 0$. Given an edge $e \in E$, we denote by $r_{\mathcal{Q}}(e)$ the number of times e appears in \mathcal{Q} , and by $\text{sgn}_{\mathcal{Q}}(e)$ to be $+1$ if $e \in \mathcal{Q}^+$, -1 if $e \in \mathcal{Q}^-$ and 0 otherwise. Given a multiset \mathcal{U} of edges, we denote by $w(\mathcal{U}) := \sum_{e \in \mathcal{U}} r_{\mathcal{U}}(e)w_e$ the sum of the weights of the edges in \mathcal{U} in accordance with each edge's multiplicity.

(2) The following additional properties must be satisfied if the traversal \mathcal{Q} is a path. If one end of the path has edge $\{i, j\} \in \mathcal{Q}^+$ and end node i , then i is unsaturated under x , i.e., $\sum_{e:i \in e} x_e < b_i$; if the end has edge $\{i, j\} \in \mathcal{Q}^-$ and end node i , then $\alpha_{i \setminus j} = 0$. Observe that there is a special case where the path starts and ends at the same node i ; we still consider this as the path case as long as the end node conditions are satisfied for both end edges (which could be the same).

Lemma 2 (Alternative Feasible Solution.) *Suppose \mathcal{Q} is a feasible alternating traversal with respect to some feasible $x \in \mathcal{L}$. Then, there exists feasible $\hat{x} \neq x$ such that $w(\hat{x}) - w(x)$ has the same sign ($\{-1, 0, +1\}$) as $w(\mathcal{Q}^+) - w(\mathcal{Q}^-)$.*

Proof. Suppose \mathcal{Q} is a feasible alternating traversal. Then, for some $\lambda > 0$, we can define an alternative feasible solution $\hat{x} \neq x$ by $\hat{x}_e := x_e + \lambda \cdot \text{sgn}_{\mathcal{Q}}(e) \cdot r_{\mathcal{Q}}(e)$. Moreover, $w(\hat{x}) - w(x) = \lambda(w(\mathcal{Q}^+) - w(\mathcal{Q}^-))$. \square

Lemma 3 (Alternating Traversal Weight). *Suppose \mathcal{Q} is an alternating traversal. Then, the following holds.*

- (a) We have $w(\mathcal{Q}^+) \geq w(\mathcal{Q}^-)$, where strict inequality holds if \mathcal{Q} contains a strict edge.

(b) If \mathcal{Q} is a simple cycle with no strict edges, then $w(\mathcal{Q}^+) = w(\mathcal{Q}^-)$, i.e., \mathcal{Q} is a cycle with equal alternating weight; in particular, with the “no cycle with alternating weight” assumption, any alternating traversal that is an even cycle must contain a strict edge.

Lemma 4 (Growing Feasible Alternating Traversal). *Suppose a fixed point (m, α, S) and a feasible $x \in \mathcal{L}$ are given as above.*

1. *Suppose $\{i, j\} \in E^+$ and node j is saturated (we stop if j is unsaturated). Then, there exists some node $k \in N(j) \setminus i$ such that $\{j, k\} \in E^-$.*
2. *Suppose $\{j, k\} \in E^-$ and $\alpha_{k \setminus j} > 0$ (we stop if $\alpha_{k \setminus j} = 0$). Then, there exists some node $l \in N(k) \setminus j$ such that $\{k, l\} \in E^+$.*

The proofs of Lemmas 3 and 4 are deferred to the full version.

Lemma 5 (Unifying Structural Lemma). *Suppose edge $e \in E$ is critical (with respect to some fixed point (m, α) and feasible $x \in \mathcal{L}$). Then, there exists a feasible alternating traversal \mathcal{Q} ; if in addition e is strict and there is no cycle with equal alternating weight, then \mathcal{Q} contains a strict edge.*

Proof. To find a feasible alternating traversal \mathcal{Q} , we apply a growing procedure (described in the full version) that starts from the critical edge $e = \{i, j\}$. Moreover, if \mathcal{Q} is a simple even cycle, then by Lemma 3(b), \mathcal{Q} contains a strict edge under the “no cycle with equal alternating weight” assumption; otherwise, \mathcal{Q} contains the edge e , in which case e being strict implies that \mathcal{Q} contains a strict edge. \square

Proof of Lemma 1: It suffices to check the given edge $\{i, j\}$ is critical in each of the three cases. Then, Lemma 4 promises the existence of a feasible alternating traversal, which contains a strict edge where appropriate. Then, Lemmas 3 and 2 guarantee the existence of feasible $\hat{x} \neq x$ such that $w(\hat{x}) \geq w(x)$, where strict inequality holds where appropriate. \square

5 Achieving Social Welfare with Greedy and Spiteful Users

We saw in Proposition 1 that a Nash Equilibrium m can result in zero social welfare if users are spiteful. In this section, we investigate under what conditions can a fixed point (m, α, S) of \mathcal{T} achieve good social welfare, even if the underlying payoff function u is greedy and spiteful. Given $m \in A_E$, recall that for each node i , \hat{m}_i is the b_i -th best offer to i and \bar{m}_i is the $(b_i + 1)$ -st best offer to i . Observe that each edge $e = \{i, j\} \in E$ falls into exactly one of the following three categories.

1. **Greedy Edges:** $m_{j \rightarrow i} > \bar{m}_i$ and $m_{i \rightarrow j} > \bar{m}_j$. Edge e will be selected and $u_e(m) = 1$.
2. **Spiteful Edges:** $m_{j \rightarrow i} < \hat{m}_i$ or $m_{i \rightarrow j} < \hat{m}_j$. Edge e will be rejected and $u_e(m) = 0$.
3. **Ambiguous Edges:** these are the remaining edges that are neither greedy nor spiteful.

Given a fixed point (m, α, S) , by propositions shown in the full version, the category of an edge $e \in E$ can be determined by the sign of S_e : greedy (+1),

ambiguous (0). Observe that after the greedy edges are selected and the spiteful edges are rejected, even if ambiguous edges are chosen arbitrarily (deterministic or randomized) to form a b -matching, the resulting payoff function is still greedy and spiteful. Since no agent (edge player) has motivation to unilaterally change his action for fixed point m , and any contract made for an ambiguous edge will be within the best b_i offers for a node i (i.e., if $\{i, j\} \in E$ is ambiguous, then $m_{j \rightarrow i} = \widehat{m}_i$ and $m_{i \rightarrow j} = \widehat{m}_j$), we can optimize the following, subject to remaining node capacity constraints b' (after greedy edges are selected).

- Find a maximum cardinality b' -matching among the ambiguous edges, hence optimizing the number of contracts made.
- Find a maximum weight b' -matching among the ambiguous edges, hence optimizing the social welfare.

Choosing Maximum Weight Matching among Ambiguous Edges. Observe that a maximum cardinality matching can be arbitrarily bad in terms of weight, but a maximum weight matching must be maximal and so is a 2-approximation for maximum cardinality. Hence, we argue that it makes sense to find a maximum weight b' -matching among the ambiguous edges. This step can be performed centrally or as a collective decision by the users. We give the main result in the following theorem and leave the details in the full version.

Theorem 4 (Price of Stability). *Suppose the given graph has no cycle with equal alternating weight or (LP) has a unique optimum. Then, there exists a greedy and spiteful payoff function u such that any fixed point m of \mathcal{T} is a Nash Equilibrium in \mathcal{N}_u ; moreover, the social welfare $\mathcal{S}_u(m) \geq \frac{2}{3}\mathcal{S}_{LP} \geq \frac{2}{3}\max_{m' \in A_E} \mathcal{S}_u(m')$, showing that the Price of Stability is at most 1.5*

5.1 Rate of Convergence: ϵ -Greedy and ϵ -Spiteful Users

Although the iterative protocol described in Figure 1 will converge to some fixed point (m, α, S) , it is possible that a fixed point will never be exactly reached. However, results in Section 4 and 5 can be extended if we relax the notions of greedy and spiteful users, and relax also the technical assumption on no cycle with equal alternating weight. The details are given in the full version and we state the convergence result.

Theorem 5 (Rate of Convergence). *Suppose $\epsilon > 0$, and the given graph has maximum edge weight W and has no cycle with ϵ -equal alternating weight. Then, there exists an ϵ -greedy and ϵ -spiteful payoff function u such that the following holds. For any sequence $\{(m^{(t)}, \alpha^{(t)})\}$ produced by the iterative protocol in Figure 1, and for all $t \geq \Theta(\frac{W^2|V|^4}{\epsilon^2})$, the social welfare $\mathcal{S}_u(m^{(t)}) \geq \frac{2}{3}\mathcal{S}_{LP}$.*

References

1. Azar, Y., Birnbaum, B., Celis, L.E., Devanur, N.R., Peres, Y.: Convergence of local dynamics to balanced outcomes in exchange networks. In: FOCS 2009, pp. 293–302 (2009)

2. Azar, Y., Devanur, N.R., Jain, K., Rabani, Y.: Monotonicity in bargaining networks. In: SODA 2010, pp. 817–826 (2010)
3. Baillon, J., Bruck, R.E.: The rate of asymptotic regularity is $O(1/\sqrt{n})$. In: Theory and Applications of Nonlinear Operators of Accretive and Monotone Type. Lecture Notes in Pure and Appl. Math., pp. 51–81 (1996)
4. Bateni, M., Hajiaghayi, M., Immorlica, N., Mahini, H.: The Cooperative Game Theory Foundations of Network Bargaining Games. In: Abramsky, S., Gavoille, C., Kirchner, C., Meyer auf der Heide, F., Spirakis, P.G. (eds.) ICALP 2010. LNCS, vol. 6198, pp. 67–78. Springer, Heidelberg (2010)
5. Bayati, M., Borgs, C., Chayes, J., Zecchina, R.: Belief-propagation for weighted b-matchings on arbitrary graphs and its relation to linear programs with integer solutions. *SIAM J. Discrete Math.* 25, 989 (2011)
6. Bayati, M., Shah, D., Sharma, M.: Max-product for maximum weight matching: Convergence, correctness, and lp duality. *IEEE Transactions on Information Theory* 54(3), 1241–1251 (2008)
7. Binmore, K.: *Game Theory and the Social Contract*, vol. 2. MIT Press (1998)
8. Celis, L.E., Devanur, N.R., Peres, Y.: Local Dynamics in Bargaining Networks via Random-turn Games. In: Saberi, A. (ed.) WINE 2010. LNCS, vol. 6484, pp. 133–144. Springer, Heidelberg (2010)
9. Chakraborty, T., Kearns, M.: Bargaining Solutions in a Social Network. In: Papadimitriou, C., Zhang, S. (eds.) WINE 2008. LNCS, vol. 5385, pp. 548–555. Springer, Heidelberg (2008)
10. Chakraborty, T., Kearns, M., Khanna, S.: Network bargaining: algorithms and structural results. In: EC 2009, pp. 159–168 (2009)
11. Cook, K.S., Yamagishi, T.: Power in exchange networks: A power-dependence formulation. *Social Networks* 14, 245–265 (1992)
12. Driessen, T.S.H.: *Cooperative Games: Solutions and Applications*. Kluwer Academic Publishers (1988)
13. Ishikawa, S.: Fixed points and iteration of a nonexpansive mapping in a banach space. *Proceedings of the American Mathematical Society* 59(1), 65–71 (1976)
14. Kanoria, Y.: An FPTAS for Bargaining Networks with Unequal Bargaining Powers. In: Saberi, A. (ed.) WINE 2010. LNCS, vol. 6484, pp. 282–293. Springer, Heidelberg (2010)
15. Kanoria, Y., Bayati, M., Borgs, C., Chayes, J.T., Montanari, A.: Fast convergence of natural bargaining dynamics in exchange networks. In: SODA 2011, pp. 1518–1537 (2011)
16. Kleinberg, J., Tardos, É.: Balanced outcomes in social exchange networks. In: STOC 2008, pp. 295–304 (2008)
17. Koufogiannakis, C., Young, N.E.: Distributed Fractional Packing and Maximum Weighted b-Matching via Tail-Recursive Duality. In: Keidar, I. (ed.) DISC 2009. LNCS, vol. 5805, pp. 221–238. Springer, Heidelberg (2009)
18. Lotker, Z., Patt-Shamir, B., Pettie, S.: Improved distributed approximate matching. In: SPAA 2008, pp. 129–136. ACM (2008)
19. Nash, J.: The bargaining problem. *Econometrica* 18, 155–162 (1950)
20. Nieberg, T.: Local, distributed weighted matching on general and wireless topologies. In: Proceedings of the Fifth International Workshop on Foundations of Mobile Computing, DIALM-POMC 2008, pp. 87–92. ACM, New York (2008)
21. Sanghavi, S., Malioutov, D.M., Willsky, A.S.: Linear programming analysis of loopy belief propagation for weighted matching. In: NIPS 2007 (2007)
22. Schrijver: *Combinatorial Optimizatin*, vol. A-C. Springer (2004)