Theory of Computing Systems
https://doi.org/10.1007/s00224-019-09958-4

Optimizing Social Welfare for Network Bargaining Games in the Face of Instability, Greed and Idealism

T.-H. Hubert Chan[1] · Fei Chen[2] · Li Ning[3]

Published online: 04 December 2019

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 have 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 exhibit two characteristics of human nature: greed and idealism. 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 integrality gap assumption. Moreover, after the nodes have made their greedy and idealistic 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.

Keywords Network bargaining game · Nash equilibrium · Optimizing social welfare · Unstable outcome · Greed and idealism

1 Introduction

Bargaining games have been studied with a long history, early in economics [20] and sociology, and recently in computer science, there has been a lot of attention

The conference version of the paper has appeared in the European Symposium on Algorithms 2012.
This research was partially funded by Hong Kong RGC under the contract 17202715.

This work was done while the author was at the University of Hong Kong.

✉ T.-H. Hubert Chan
hubert@cs.hku.hk

Extended author information available on the last page of the article.

on bargaining games in social exchange networks [1, 4, 6, 16, 17], in which users are modeled as nodes in an undirected simple 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 work, 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 [20], Cook and Yamagishi [12] 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 [8].

Although stability is considered to be an essential property, as remarked in [4, 6, 16], 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 work simply assumed that the LP has integrality gap 1 [9, 16] or considered restriction to bipartite graphs [4, 17], for which the LP always has integrality gap 1.

We consider the integrality gap condition as 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 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 all transactions.

[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.

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 players in our framework will be the edges, as opposed to the nodes in previous work; in fact, in the selection process we assume the node users will exhibit two characteristics of human nature: greed and idealism.

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.

Idealistic Users. Idealism captures 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 idealistic.

One can argue that greed is a natural behavior (hence the regime of greedy algorithms), but idealism is clearly not always rational. In fact, we shall see in Section 2 that there exist a idealistic 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 idealism, 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. [6, 16] 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 idealistic, 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 [6, 16], 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 idealism, 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 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 that assumption (2) is necessary for general b-matching, showing that there is some fundamental difference between the two cases.

The greedy behavior states that some edges must be selected and the idealistic 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 idealistic payoff function u that corresponds to selecting a maximum 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 idealism in Section 6 to analyze the properties of a near fixed point. Specifically, user i with capacity b_i is ϵ-greedy, if he will definitely want an offer that is strictly ϵ better than his $(b_i + 1)$-st best offer. Similarly, user i is ϵ-idealistic, if he will definitely reject an offer that is strictly ϵ worse than his b_i-th best offer. We show that the same guarantee on the price of stability can be achieved eventually (and quickly). For ease of understanding, we encourage the reader to first read Sections 4 and 5 for the special case $\epsilon = 0$ in order to understand our approach. Section 6 describes how our arguments need to be augmented for the general case $\epsilon > 0$.

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 idealism 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.

Last but not least, this work is by no means an attempt to model realistic human behavior (which is of course a daunting, if not impossible, task). What we have shown here is that given the rules of greed and idealism, agents can exploit rules to achieve desirable Nash equilibria.

Related Work Kleinberg and Tardos [17] 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. [10, 11] explored equilibrium concepts and experimental results for bipartite graphs. Celis et al. [9] gave a tight polynomial bound on the rate of convergence for unweighted bipartite graphs with a unique

[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.

balanced outcome. Kanoria [15] 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. [6, 16] designed a local dynamics that converges in polynomial time. Our distributed protocol is based on [6, 16], but is generalized to general node capacities, run by edges and does not require the integrality condition on LP.

Bateni 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 [13]. 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 [23][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, 19, 21].

2 Notation and Preliminaries

Consider an undirected simple 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, denoting 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 $(m_{j \to i}, m_{i \to j})$ 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 $m_{j \to i}$ and user j will receive $m_{i \to j}$ from this contract.

Selection Procedure and Payoff Function *u* Given a configuration $m \in A_E := \times_{e \in E} A_e$ of all agents' proposals, a 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 $u = \{u_e : A_E \to [0, 1] | e \in E\}$, 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 $\widehat{m}_i = \max^{(b_i)}_{j \in N(i)} m_{j \to i}$ and $\overline{m}_i = \max^{(b_i+1)}_{j \in N(i)} m_{j \to i}$.

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 $m_{j \to i} > \overline{m}_i$ and $m_{i \to j} > \overline{m}_j$, then $u_e(m) = 1$.

Idealistic 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 *idealistic* (or the users are idealistic) if for each $e = \{i, j\} \in E$ and $m \in A_E$, if $m_{j \to i} < \widehat{m}_i$, 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 idealistic, there exists a Nash Equilibrium $m \in A_E$ under which no contracts are made.*

Proof We take G to be the complete graph K_5 on five nodes, where each edge has unit weight, and each node has unit capacity. It is straight forward to construct a configuration $m \in A_E$ of proposals that has the following properties (Fig. 1).

(a) Each agent splits the weight into (0.4, 0.6).
(b) Each user gets two offers with profit 0.4 and two offers with profit 0.6.

Observe that for idealistic users, no contract will be accepted, because for each contract there will be a user getting 0.4, which is worse than his best choice 0.6. Hence, $\mathcal{S}(m) = 0$.

We next show that m is a Nash Equilibrium. Consider an edge $e = \{i, j\}$ where under m, user i gets 0.4 and user j gets 0.6. Since currently the best offer i receives is 0.6, in order for edge e to have any chance to be considered by user i, agent e must offer at least 0.6 to i, which means there is only at most 0.4 to be offered to user j, who will definitely reject e because user j still has another offer with 0.6. Hence, there is no way for any agent to change his strategy unilaterally to increase his payoff. □

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. [6, 16], 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 $r \in \mathbb{R}$,

Fig. 1 K_5

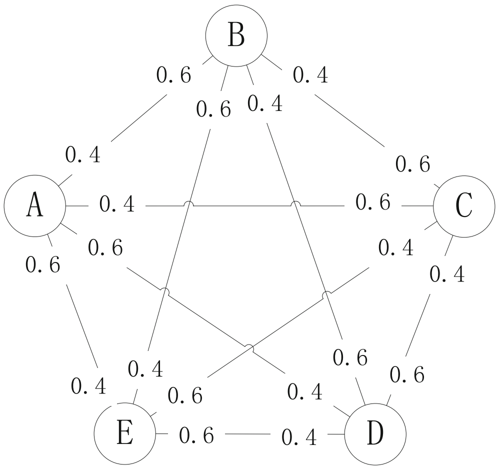

we denote $(r)_+ := \max\{r, 0\}$. Moreover, as described in [3, 16] a damping factor $\kappa \in (0, 1)$ is used in the update; we can think of $\kappa = \frac{1}{2}$.

Although later on we will also consider the LP relaxation of b-matching, unlike previous works [7, 16, 22], we do not require the assumption that the LP relaxation has a unique integral optimum.

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

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

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 $\mathcal{T} : [0, W]^{2|E|} \to [0, W]^{2|E|}$ as follows.

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

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

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

Then, we define $\mathcal{T}(\alpha) \in [0, W]^{2|E|}$ by $(\mathcal{T}(\alpha))_{i\setminus j} := \max^{(b_i)}_{k\in N(i)\setminus j} m_{k\to i}(\alpha)$. 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 $T : D \to D$ 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 [6, 16], it can be proved that the function $\mathcal{T}$ is non-expansive. The following result by Ishikawa [14] shows that by applying a non-expansive function repeatedly, convergence to a fixed point can be obtained.

Fact 1 ([14]) Suppose $T : D \to D$ is a non-expansive function under norm $||\cdot||$ and for some $\kappa \in (0, 1)$ and some initial $\alpha^{(1)} \in D$, the sequence $\{\alpha^{(t)}\}$ is defined by $\alpha^{(t+1)} := (1-\kappa)\cdot\alpha^{(t)} + \kappa \cdot T(\alpha^{(t)})$. Suppose further that the sequence $\{\alpha^{(t)}\}$ is bounded under norm $||\cdot||$. Then, the sequence $\{\alpha^{(t)}\}$ converges (under norm $||\cdot||$) to some fixed point of T.

Claim 1 Given a weighted graph $G = (V, E)$ with maximum weight W, the function $\mathcal{T} : [0, W]^{2|E|} \to [0, W]^{2|E|}$ is non-expansive under the ℓ_∞ norm.

Proof The claim can be proved in a similar way as in [16]. We observe the following facts.

– The '$\max^{(b_i)}$' in the mapping is non-expansive. To prove this, it suffices to show that given two vectors x and y with the same dimension d larger than b, the following holds

$$|\max_{i\in[d]}{}^{(b)} x_i - \max_{i\in[d]}{}^{(b)} y_i| \le \max_{i\in[d]} |x_i - y_i|.$$

Let $b_1 := \arg\max^{(b)}_{i\in[d]} x_i$ and $b_2 := \arg\max^{(b)}_{i\in[d]} y_i$. Without loss of generality, assume $x_{b_1} \le y_{b_2}$. If $y_{b_1} \ge y_{b_2}$, then $|x_{b_1} - y_{b_2}| \le |x_{b_1} - y_{b_1}|$. If $y_{b_1} < y_{b_2}$, there exists $k \in [d]$ satisfying $y_k \ge y_{b_2}$ such that $x_k \le x_{b_1}$. Then $|x_{b_1} - y_{b_2}| \le |x_k - y_k|$. Therefore, $|x_{b_1} - y_{b_2}| \le \max_{i\in[d]} |x_i - y_i|$.

– The variable $m = m(\alpha)$ is non-expansive according to its definition. Let $m_{i\to j} = f(\alpha_{i\backslash j}, \alpha_{j\backslash i})$, where $f(x, y)$ is given by

$$f(x, y) = \begin{cases} \frac{w_{ij}-x+y}{2} & x + y \le w_{ij}, \\ (w_{ij} - x)_+ & \text{otherwise.} \end{cases}$$

It can be checked that f is continuous in $\mathbb{R}^2_+$. Also, it is differentiable except in $\{(x, y) \in \mathbb{R}^2_+ : x + y = w_{ij} \text{ or } x = w_{ij}\}$, and satisfies $|\frac{\partial f}{\partial x}| + |\frac{\partial f}{\partial y}| \le 1$. Therefore, f is Lipschitz continuous in the ℓ_∞ norm with Lipschitz constant 1 and is non-expansive. □

Fact 1 and Claim 1 give Theorem 1.

Theorem 1 (Convergence to a Fixed Point) *The distributed protocol shown in Algorithm 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 (1) and (2). We also say that (m, α, S) or (m, α) or m is a fixed point (of $\mathcal{T}$). Similar to [16], we give several important properties of a fixed point. In particular, we give the following propositions on which

Theorem 2 is based. Recall that $\widehat{m}_i := \max^{(b_i)}_{k\in N(i)} m_{k\to i}$ and $\overline{m}_i := \max^{(b_i+1)}_{k\in N(i)} m_{k\to i}$, and in addition to (2), a fixed point (m, α) also satisfies $\alpha_{i\setminus j} = \max^{(b_i)}_{k\in N(i)\setminus j} m_{k\to i}$.

Proposition 2 (**Outside option** α) *Suppose for each* $\{i, j\} \in E$*, we have* $\alpha_{i\setminus j} = \max^{(b_i)}_{k\in N(i)\setminus j} m_{k\to i}$*. Then, the following properties hold.*

(a) *If* $m_{j\to i} < \widehat{m}_i$*, then* $\alpha_{i\setminus j} = \widehat{m}_i$ *; if* $m_{j\to i} \geq \widehat{m}_i$*, then* $\alpha_{i\setminus j} = \overline{m}_i$*.*
(b) $m_{j\to i} \geq \widehat{m}_i$ *iff* $m_{j\to i} \geq \alpha_{i\setminus j}$*;* $m_{j\to i} \leq \overline{m}_i$ *iff* $m_{j\to i} \leq \alpha_{i\setminus j}$*;*

Proof (a) If $m_{j\to i} < \widehat{m}_i$, that is, the offer that user i gets from j is not one of its best b_i offers, then $\alpha_{i\setminus j} = \max^{(b_i)}_{k\in N(i)\setminus j} m_{k\to i} = \max^{(b_i)}_{k\in N(i)} m_{k\to i} = \widehat{m}_i$.

If $m_{j\to i} \geq \widehat{m}_i$, that is, the offer that user i gets from j is at least as good as its b_i-th best offer, then $\alpha_{i\setminus j} = \max^{(b_i)}_{k\in N(i)\setminus j} m_{k\to i} = \max^{(b_i+1)}_{k\in N(i)} m_{k\to i} = \overline{m}_i$.

Note that the equalities hold even if $\widehat{m}_i = \overline{m}_i$.

(b) Proposition 2(a) and its proof imply the following

- If $m_{j\to i} \geq \widehat{m}_i$, then $\alpha_{i\setminus j} = \overline{m}_i \leq \widehat{m}_i \leq m_{j\to i}$. If $m_{j\to i} < \widehat{m}_i$, then $\alpha_{i\setminus j} = \widehat{m}_i > m_{j\to i}$.
- If $m_{j\to i} \leq \overline{m}_i$, then $\alpha_{i\setminus j} = \widehat{m}_i \geq \overline{m}_i \geq m_{j\to i}$. If $m_{j\to i} > \overline{m}_i$, then $m_{j\to i} \geq \widehat{m}_i$ and thus $\alpha_{i\setminus j} = \overline{m}_i < m_{j\to i}$. □

Proposition 3 (α defining (S, m)) *Suppose given* $\alpha \in [0, W]^{2|E|}$*, the quantities* $S \in \mathbb{R}^{|E|}$ *and* $m \in A_E$ *are defined as in* (1) *and* (2)*. Then, the following properties hold for each* $\{i, j\} \in E$*.*

(a) $S_{ij} > 0$ *iff* $m_{j\to i} > \alpha_{i\setminus j}$*;*
(b) $S_{ij} = 0$ *iff* $m_{j\to i} = \alpha_{i\setminus j}$ *and* $m_{i\to j} = \alpha_{j\setminus i}$*.*

Proof (a) Recall that $S_{ij} = w_{ij} - \alpha_{i\setminus j} - \alpha_{j\setminus i}$ and $m_{j\to i} = (w_{ij} - \alpha_{j\setminus i})_+ - \frac{1}{2}(S_{ij})_+$. If $S_{ij} > 0$, then $w_{ij} - \alpha_{j\setminus i} = S_{ij} + \alpha_{i\setminus j} > 0$. Therefore $m_{j\to i} = (S_{ij} + \alpha_{i\setminus j})_+ - \frac{1}{2}(S_{ij})_+ = \alpha_{i\setminus j} + \frac{1}{2}S_{ij} > \alpha_{i\setminus j}$.

On the other hand, if $S_{ij} \leq 0$, then $m_{j\to i} = (S_{ij} + \alpha_{i\setminus j})_+ = \max(S_{ij} + \alpha_{i\setminus j}, 0) \leq \alpha_{i\setminus j}$.

(b) If $S_{ij} = 0$, then $m_{j\to i} = (S_{ij} + \alpha_{i\setminus j})_+ - \frac{1}{2}(S_{ij})_+ = \alpha_{i\setminus j}$ and similarly $m_{i\to j} = \alpha_{j\setminus i}$.

If $m_{j\to i} = \alpha_{i\setminus j}$ and $m_{i\to j} = \alpha_{j\setminus i}$, then from Proposition 3(a) we have $S_{ij} \leq 0$. Therefore $m_{j\to i} + m_{i\to j} = \alpha_{i\setminus j} + \alpha_{j\setminus i} = w_{ij} - S_{ij} \geq w_{ij}$. Since $m_{j\to i} + m_{i\to j} \leq w_{ij}$, we have $m_{j\to i} + m_{i\to j} = w_{ij}$ and therefore $S_{ij} = 0$. □

Proposition 4 (Fixed Point (m, α)) *Suppose* (m, α) *is a fixed point of* $\mathcal{T}$*. Then, for each* $\{i, j\} \in E$*, the following properties hold.*

(a) *If* $m_{j\to i} > 0$ *and* $m_{j\to i} \geq \widehat{m}_i$*, then* $S_{ij} \geq 0$*; if* $S_{ij} \geq 0$*, then* $m_{j\to i} \geq \widehat{m}_i$*.*
(b) $m_{j\to i} \leq \overline{m}_i$ *iff* $S_{ij} \leq 0$*.*

Proof (a) Suppose $m_{j\to i} > 0$ and $m_{j\to i} \geq \widehat{m}_i$. Then from Proposition 2(b), we have $m_{j\to i} \geq \alpha_{i\setminus j}$. Consider the following two cases:

– $m_{j\to i} > \alpha_{i\backslash j}$. Then, from Proposition 3(a) we have $S_{ij} > 0$.
– $m_{j\to i} = \alpha_{i\backslash j}$. Then, from Proposition 3(a) we have $S_{ij} \leq 0$. Since $m_{j\to i} > 0$, it follows that $\alpha_{i\backslash j} = m_{j\to i} = (w_{ij} - \alpha_{j\backslash i})_+ - \frac{1}{2}(S_{ij})_+ = w_{ij} - \alpha_{j\backslash i}$ and so $S_{ij} = 0$.

Therefore, $S_{ij} \geq 0$.

For the converse, $S_{ij} \geq 0$ implies from Proposition 3 that $m_{j\to i} \geq \alpha_{i\backslash j}$, which implies from Proposition 2(b) that $m_{j\to i} \geq \widehat{m}_i$.

(b) We have by Proposition 2(b) that $m_{j\to i} \leq \overline{m}_i$ *iff* $m_{j\to i} \leq \alpha_{i\backslash j}$, which is equivalent to $S_{ij} \leq 0$ by Proposition 3(a). □

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

Proof Let (m, α, S) be a fixed point of $\mathcal{T}$. We show that for each $e = \{i, j\} \in E$, agent e cannot increase $u_e(m)$ by changing its action m_e unilaterally.

If $S_{ij}(m) < 0$, i.e. $\alpha_{i\backslash j} + \alpha_{j\backslash i} > w_{ij}$, then any proposal $(m'_{j\to i}, m'_{i\to j}) \in A_e$ must satisfy $m'_{j\to i} + m'_{i\to j} \leq w_{ij} < \alpha_{i\backslash j} + \alpha_{j\backslash i}$, which implies that $m'_{i\to j} < \alpha_{j\backslash i}$ or $m'_{j\to i} < \alpha_{i\backslash j}$. Since the payoff function is idealistic, it follows that i and j cannot be matched. Hence, $u_e = 0$ if other agents maintain their actions.

If $S_{ij}(m) > 0$, then by Proposition 4(b), $m_{j\to i} > \overline{m}_i$ and $m_{i\to j} > \overline{m}_j$. Since the payoff function u is greedy, we already have $u_e(m) = 1$ and there is no more room for improvement.

If $S_{ij}(m) = 0$, then by Proposition 3(b), $m_{j\to i} = \alpha_{i\backslash j}$ and $m_{i\to j} = \alpha_{j\backslash i}$. This also implies that $m_{i\to j} + m_{j\to i} = w_{i,j}$. Then, any change of $(m_{i\to j}, m_{j\to i})$ will lead to either $m'_{j\to i} < \alpha_{i\backslash j}$ and $m'_{i\to j} < \alpha_{j\backslash i}$, which means there is no chance for $\{i, j\}$ to be matched for idealistic payoff function u. □

Theorems 1 and 2 imply that as long as the payoff function is greedy and idealistic, the game defined between the agents (edge players) always has a pure Nash Equilibrium.

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 idealistic. Our goal is to show that there exists some greedy and idealistic 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. [16], the network bargain game is closely related to the linear program LP relaxation of the b-matching problem, which has the form $\mathcal{S}_{\mathsf{LP}} := \max_{x\in\mathcal{L}} w(x)$, 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.

The main technical properties are as follows.

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 $x_{ij} = 0$.*

For the 1-matching case, the conclusions listed in Theorem 3 have been derived by Kanoria et al. [16] without the no-alternating-cycle assumption. However, for general b-matching, this assumption in Theorem 3 is necessary for cases (b) and (c). We show that without this technical assumption, there is a counter example for Theorem 3(b).

Consider the graph given in Fig. 2, which contains a cycle with equal alternating weight, e.g. $\{C, D, E, F, C\}$. Each note has capacity 2. The offers specified on the edges form a fixed point of $\mathcal{T}$. It can be checked that $S_{AC} > 0$, while every optimal solution of the corresponding LP has $x_{AC} = 0$. Thus, case (b) of Theorem 3 is incorrect for this given graph. By slight modification of the given graph, a counter example for Theorem 3(c) can be constructed.

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

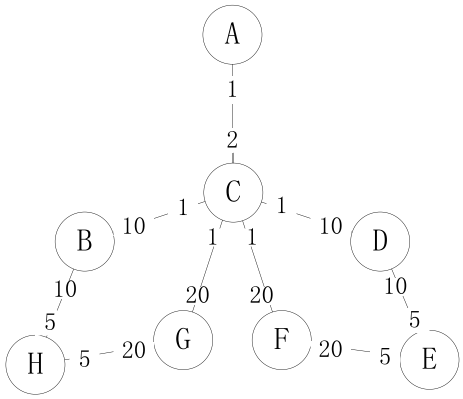

Fig. 2 Graph with Equal Alternating Weight: The numbers on the edges indicate how the agents split the weights, e.g. C gets 2 and A gets 1 from edge $\{A, C\}$

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 *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 $\mathrm{sgn}_{\mathcal{Q}}(e)$ 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 $\widehat{x} \neq x$ such that $w(\widehat{x}) - w(x)$ 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 $\widehat{x} \neq x$ by $\widehat{x}_e := x_e + \lambda \cdot \mathrm{sgn}_{\mathcal{Q}}(e) \cdot r_{\mathcal{Q}}(e)$. Moreover, $w(\widehat{x}) - w(x) = \lambda(w(\mathcal{Q}^+) - w(\mathcal{Q}^-))$. □

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.*

Proof Consider consecutive edges $\{i, j\} \in \mathcal{Q}^+$ and $\{j, k\} \in \mathcal{Q}^-$ in the alternating traversal.

Since $S_{ij} \geq 0$, by Proposition 4(a), we have $m_{i\to j} \geq \widehat{m}_j$. Since $\alpha_{j\setminus k}$ is either $\widehat{m}_j$ or $\overline{m}_j$, we have $\alpha_{j\setminus k} \leq \widehat{m}_j$. Therefore, $m_{i\to j} \geq \alpha_{j\setminus k}$.

Suppose $\{j, k\}$ is strict, i.e., $S_{jk} < 0$. Then, by definition, we have $\alpha_{j\setminus k} + \alpha_{k\setminus j} > w_{jk}$.

Suppose $\{i, j\}$ is strict, i.e., $S_{ij} > 0$. Then we show that either $m_{i\to j} > \alpha_{j\setminus k}$ or $\alpha_{j\setminus k} + \alpha_{k\setminus j} > w_{jk}$ holds. From Proposition 3(a), we have $m_{i\to j} > \alpha_{i\setminus j}$, which implies by Proposition 2(b) that $m_{i\to j} > \overline{m}_j$. If $\widehat{m}_j = \overline{m}_j$, then $m_{i\to j} > \widehat{m}_j \geq \alpha_{j\setminus k}$. If $\widehat{m}_j > \overline{m}_j$, then $S_{jk} \leq 0$ together with Proposition 4(b) implies $m_{k\to j} \leq \overline{m}_j < \widehat{m}_j$. Then from Proposition 2(a) we have $\alpha_{j\setminus k} = \widehat{m}_j$ and thus $m_{k\to j} < \alpha_{j\setminus k}$. It follows from Proposition 3 that $S_{jk} < 0$, that is, $\alpha_{j\setminus k} + \alpha_{k\setminus j} > w_{jk}$.

We next show that $w(\mathcal{Q}^+) \geq w(\mathcal{Q}^-)$ by a "paying" argument. Observe that for $S_{ij} \geq 0$, we have $w_{ij} = m_{j\to i} + m_{i\to j}$, and for $S_{jk} \leq 0$, we have $\alpha_{j\setminus k} + \alpha_{k\setminus j} \geq w_{jk}$. Since we have $m_{i\to j} \geq \alpha_{j\setminus k}$, it follows that we can split the weight of every edge into 2 parts, such that each part of the weight from an edge in $\mathcal{Q}^+$ can be used to pay for a part of the weight in a neighboring edge in $\mathcal{Q}^-$. Observe that if $\{j, k\} \in \mathcal{Q}^-$ and node k is an end point of the traversal, then $\alpha_{k\setminus j} = 0$ and so there is no need for k to have a neighboring edge in $\mathcal{Q}^+$ to pay for this part. Hence, we have $w(\mathcal{Q}^+) \geq w(\mathcal{Q}^-)$. Observe that if $\mathcal{Q}$ contains a strict edge, then at least one of the inequalities $m_{i\to j} \geq \alpha_{j\setminus k}$ and $\alpha_{j\setminus k} + \alpha_{k\setminus j} \geq w_{jk}$ becomes strict and hence we have $w(\mathcal{Q}^+) > w(\mathcal{Q}^-)$.

Finally, if $\mathcal{Q}$ is a simple cycle with no strict edges, then all edges $\{i, j\} \in \mathcal{Q}$ satisfy $S_{ij} = 0$. Hence, the roles of $\mathcal{Q}^+$ and $\mathcal{Q}^-$ can be exchanged and so $w(\mathcal{Q}^+) = w(\mathcal{Q}^-)$ follows. □

Growing Procedure Using Lemma 4, we can grow an alternating traversal with edges alternating between E^+ and E^-. We start the procedure from a critical edge $e = \{i, j\}$, i.e., in exactly one of E^+ and E^-.

1. If the growth process stops at both ends eventually without revisiting nodes, then we have a simple generalized alternating traversal path containing the edge e.
2. Suppose the growth process revisits a node starting from one end, say for the one from node j. Then, at some point a node is revisited, and suppose k is the first node to be revisited and we have a simple cycle C. If C is even, then C forms a generalized alternating traversal $\widehat{\mathcal{Q}}$ that is an even cycle.
3. If C is odd and $k \neq i$, then we continue to grow from k; we next consider the case $k = i$, and suppose l is the node before i is revisited again. Observe that since C is odd, the edges $\{i, j\}$ and $\{i, l\}$ are either both in $\widehat{E}^+$ or both in $\widehat{E}^-$. Suppose we continue to grow from i with respect to the edge $\{l, i\}$ according to $\mathcal{R}$ and the next node is h; since $\{i, j\}$ is in exactly one of $\widehat{E}^+$ and $\widehat{E}^-$, it follows that $h \neq j$, and so we can continue the growth process. At this point, the partial traversal forms a "lollipop" graph with the odd cycle C, and we are growing the stem.
4. If the growth process at the stem stops at some node g without revisiting nodes, then we have a generalized alternating traversal $\widehat{\mathcal{Q}}$ that is considered to be a path starting at g, traveling along the stem to k, then along the cycle C back to k, and finally returning to g along the stem again. Note that the traversal $\widehat{\mathcal{Q}}$ contains edge e.
5. Suppose the growth process at the stem revisit some node g. If g is a node in the cycle C other than k, then since C is odd, we must have formed a generalized alternating traversal that is an even cycle.
6. Finally, suppose the revisited node g is on the stem (including k) forming another cycle C'. If C' is even, then C' forms a generalized alternating traversal; otherwise, we have two odd cycles C and C' connected by the stem between g and k. In this case, we have a generalized alternating traversal $\widehat{\mathcal{Q}}$ that is a circuit going through each cycle once and the stem back and forth, which contains edge e.

Note that a generalized alternating traversal obtained from the above growing procedure has one of the following forms. Apart from the first case of simple even cycle, the edge e from which the procedure starts is contained in the generalized alternating traversal.

(a) Simple Even Cycle.
(b) Simple Path.
(c) Lollipop with Odd Cycle. The traversal is a path starting from the end of the stem, traveling along the stem to some node, then along an odd cycle back to this node, and then along the stem, finally returning to the end of the stem.
(d) Dumbbell - Two Odd Cycles connected with a path.

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* $\{k, l\} \in E^+$.

Proof 1. Suppose $\{i, j\} \in E^+$ and node j is saturated. Since $\sum_{k \in N(j)} x_{jk} = b_j$ and $x_{ij} < 1$, there are at least b_j nodes k in $N(j) \setminus i$ such that $x_{jk} > 0$. We pick the k such that $m_{k \to j}$ is the smallest. Since $S_{ij} \geq 0$, we conclude from Proposition 4(a) that $m_{i \to j} \geq \widehat{m}_j$. It follows that $m_{k \to j}$ is at most as large as the minimum offer to j among $b_j + 1$ offers. Hence, $m_{k \to j} \leq \overline{m}$, which implies from Proposition 4(b) that $S_{jk} \leq 0$. Hence, $\{j, k\} \in E^-$.
2. Suppose $\{j, k\} \in E^-$ and $\alpha_{k \setminus j} > 0$. By Proposition 4(b), $S_{jk} \leq 0$ implies that $m_{j \to k} \leq \overline{m}_k$, i.e., node j's offer to k is as worst as the $(b_k + 1)$-st offer and so $\widehat{m}_k = \alpha_{k \setminus j} > 0$. Moreover $x \in \mathcal{L}$ and $x_{jk} > 0$ implies that there are at most $b_k - 1$ neighbors $i \in N(k) \setminus j$ such that $x_{ik} = 1$. Suppose $l \in N(k) \setminus j$ such that $x_{kl} < 1$ and $m_{l \to k}$ is the largest. It follows node l's offer to k is at least as good as the b_k-th offer and hence $m_{l \to k} \geq \widehat{m}_k > 0$, which implies that $S_{kl} \geq 0$, by Proposition 4(a). Hence, we have $\{k, l\} \in E^+$. □

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 the growing procedure 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. □

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 $\widehat{x} \neq x$ such that $w(\widehat{x}) \geq w(x)$, where strict inequality holds where appropriate. □

5 Achieving Social Welfare with Greedy and Idealistic Users

We saw in Proposition 1 that a Nash Equilibrium m can result in zero social welfare if users are idealistic. 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 idealistic. Given $m \in A_E$, recall that for each node i, $\widehat{m}_i$ is the b_i-th best offer to i and $\overline{m}_i$ 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:** $m_{j \to i} > \overline{m}_i$ and $m_{i \to j} > \overline{m}_j$. Edge e will be selected and $u_e(m) = 1$.
2. **Idealistic Edges:** $m_{j \to i} < \widehat{m}_i$ or $m_{i \to j} < \widehat{m}_j$. Edge e will be rejected and $u_e(m) = 0$.
3. **Ambiguous Edges:** these are the remaining edges that are neither greedy nor idealistic.

Given a fixed point (m, α, S), by Propositions 2 and 3, the category of an edge $e \in E$ can be determined by the sign of S_e: greedy (+1), idealistic (-1), ambiguous (0). Observe that after the greedy edges are selected and the idealistic 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 idealistic. We first recover the result in [16] for the special case where the LP has a unique integral optimum, for which we shall see that the ambiguous edges are handled trivially.

Given a fixed point (m, α, S), the following result implies that there are no ambiguous edges in this case. Hence, after all the greedy edges are selected, the optimal social welfare is achieved automatically, and so the price of stability is 1.

Theorem 4 (Fixed Point and Integral LP) *Suppose (m, α, S) is a fixed point of $\mathcal{T}$ and LP has a unique integral optimum x corresponding to the maximum b-matching M^*. Then, for each edge $\{i, j\} \in E$, $S_{ij} > 0$ if and only if $x_{ij} = 1$.*

Proof For each edge $\{i, j\} \in E$, if $S_{ij} > 0$, then it follows from Theorem 3 (a) that $x_{ij} = 1$.

If $x_{ij} = 1$, then from Theorem 3 (c) we have $S_{ij} \geq 0$. Suppose $S_{ij} = 0$. Then, from Proposition 3 (b) we have $m_{j\to i} = \alpha_{i\backslash j}$ and $m_{i\to j} = \alpha_{j\backslash i}$. Since $m_{j\to i} + m_{i\to j} = \alpha_{i\backslash j} + \alpha_{j\backslash i} = w_{ij} > 0$, at least one of $m_{j\to i}$ and $m_{i\to j}$ is positive. Assume $m_{j\to i} > 0$ without loss of generality. Then, $\widehat{m}_i = \overline{m}_i > 0$. Let node k be a neighbor of i such that $m_{k\to i} = \alpha_{i\backslash j}$. Then, we have $m_{k\to i} = m_{j\to i} = \alpha_{i\backslash k} > 0$ and thus $S_{ik} = 0$. Then, from Theorem 3 (a) we have $x_{ik} = 1$. Since from Proposition 4(a) any $l \in N(i)$ such that $m_{l\to i} \geq \widehat{m}_i > 0$ satisfies $S_{il} \geq 0$, we know that there are more than b_i edges incident to node i appear in M^*, which is a contradiction. Therefore, S_{ij} must be positive. □

5.1 Handling Ambiguous Edges

In general, given $m \in A_E$, there will be ambiguous edges, and how the ambiguous edges are handled will affect the payoff function u and the social welfare. However, observe that a fixed point m of $\mathcal{T}$ will remain a Nash Equilibrium no matter how the ambiguous edges are handled. For the rest of the section, we analyze the social welfare of a fixed point m.

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 $m_{j\to i} = \widehat{m}_i$ and $m_{i\to j} = \widehat{m}_j$), 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. We can imagine this step to be performed centrally or as a collective decision by the users. From now on, we consider a payoff function u that results from this selection procedure and analyze the social welfare $\mathcal{S}_u(m)$ for a fixed point m of $\mathcal{T}$. Recall that we assume at least one of the following holds: (1) LP has unique optimum, (2) the graph has no cycle with equal alternating weight.

Analyzing $\mathcal{S}(m)$ for Fixed Point m Suppose (m, α, S) is a fixed point and let $\widehat{E} := \{e \in E : S_e > 0\}$ and $\overline{E} := \{e \in E : S_e = 0\}$. As observed before, $\widehat{E}$ is the set of greedy edges and $\overline{E}$ is the set of ambiguous edges. Suppose $H \subseteq \overline{E}$ is the maximum weight b'-matching that is chosen in $\overline{E}$ by the selection process. Then, $\mathcal{S}(m) = \sum_{e \in \widehat{E}} w_e + \sum_{f \in H} w_f$, which we compare with the value of an optimal LP solution, which we can assume is half-integral from a standard fact. We include its proof for completeness.

Fact 2 (Half-Integral LP Optimum) There exists a half-integral optimal solution x to LP, i.e., for all $e \in E$, $x_e \in \{0, \frac{1}{2}, 1\}$.

Proof Let A be the incident matrix of G. It suffices to show that the polytope $P = \{x : Ax \leq b, x \geq 0 \text{ and } x \leq 1\}$ is half-integral, that is, the vertices of P are half-integral. Note that P is the set of points x satisfying

$$\begin{bmatrix} A \\ I \\ -I \end{bmatrix} x \leq \begin{bmatrix} b \\ 1 \\ 0 \end{bmatrix}.$$

Let x be a vertex of P, and without loss of generality assume $x = (x_f, x_g)$, where x_f of size l and x_g of size $(m - l)$ are vectors consisting of fractional and integral entries of x, respectively. Let A' be the submatrix of A consisting of the first l columns of A and $b' = b - A \cdot (0, x_g)$. Then the inequality $A' x_f \leq b'$ holds and there is a subsystem (A'', b'') of (A', b') where A'' is nonsingular satisfying $A'' x_f = b''$. Note that A' is formed from A by deleting columns in A with integral x values. Therefore, A' is the incident matrix of a subgraph of G which consists of all edges with fractional x values. Suppose the subgraph has k connected components $C_1, C_2, \cdots, C_k$, and $A_1, A_2, \cdots, A_k$ are their incident matrices, respectively.

In the following proof, when we intend to argue that x is not a vertex of P, we show that for arbitrarily small $\delta > 0$, there is x'_f of size l satisfying $0 < \|x'_f - x_f\|_\infty \leq \delta$ such that $A'' x'_f = b''$ and $A' x'_f \leq b'$. We say that x'_f satisfies the *closeness condition*. For each edge e with $x_e \in x_f$, define $g_e := \min(x_e, 1 - x_e)$.

Consider the connected component C_1 consisting of n_1 nodes and m_1 edges, together with the corresponding fractional vector x_1 of size m_1 and subsystem (A_1, b_1). Recall that (A'', b'') is a subsystem of (A', b') such that A'' is nonsingular and $A'' x_f = b''$. Then there is a subsystem (A''_1, b''_1) of (A_1, b_1) such that A''_1 is nonsingular and $A''_1 x_1 = b''_1$. Note that A_1 is of dimension $n_1 \times m_1$. If $n_1 < m_1$, then there exists a non-zero y_1 such that $A_1''(x_1 + y_1) = b_1''$ and $A_1''(x_1 - y_1) = b_1''$ while both $x_1 + y_1$ and $x_1 - y_1$ are in the interval (0, 1). Extending y_1 to y of size m with other entries being 0, we see that both $x + y$ and $x - y$ are in P, which contradicts the fact that x is a vertex of P. Therefore we have $n_1 \geq m_1$. On the other hand, since C_1 is connected, we have $m_1 \geq n_1 - 1$. Then $m_1 = n_1 - 1$ or $m_1 = n_1$. That is, C_1 is a single node, a path or a graph consisting of one cycle. We define $\delta_m := \min_{e \in E(C_1)} g_e$.

We show that C_1 cannot contain an even cycle. Assume C_1 contains an even cycle with consecutive edges $e_1, e_2, \cdots, e_{2j}$. Then for any positive $\delta < \delta_m$ we can get a vector x_1' defined as

$$(x_1')_e = \begin{cases} (x_1)_e + \delta & \text{if } e = e_i, \text{ where } i \in [2j] \text{ and } i \text{ is odd,} \\ (x_1)_e - \delta & \text{if } e = e_i, \text{ where } i \in [2j] \text{ and } i \text{ is even,} \\ (x_1)_e & \text{otherwise.} \end{cases}$$

Then we know that $A_1 x_1' \leq b_1$ and $A_1'' x_1' = b_1''$. Let x_f' be the vector obtained from x_f by replacing x_1 with x_1', then x_f' satisfies the closeness condition and thus x is not a vertex of P, which is a contradiction.

Next we show that if C_1 is not a single node, each node in C_1 is of degree at least 2. Suppose, on the contrary, that C_1 contains at least one node with degree 1. Below we consider two cases.

If there are at least two nodes with degree 1, let $D = \{q, e_1, v_2, e_2, \cdots, w\}$ be the simple path between any of such two nodes q and w. Since each of q and w is incident to only one edge with fractional value and b_1 is integral, their corresponding inequalities strictly hold. Given any positive $\delta < \delta_m$ define vector x_1' of size m_1 as follows:

$$(x_1')_e = \begin{cases} (x_1)_e + \delta & \text{if } e = e_i \in E(D) \text{ and } i \text{ is odd,} \\ (x_1)_e - \delta & \text{if } e = e_i \in E(D) \text{ and } i \text{ is even.} \end{cases}$$

Then we can get a vector x_f' satisfying the closeness condition and leading to a contradiction.

If C_1 has exactly one node with degree 1, then it must be an odd cycle $L = \{v_1, e_1^L, v_2, e_2^L, \cdots, v_{2j+1}, e_{2j+1}^L, v_1\}$ connected to a simple path $D = \{w_1, e_1^D, w_2, e_2^D, \cdots\}$. Assume $w_1 = v_1$ without loss of generality. Given any positive $\delta < \frac{\delta_m}{2}$, define vector x_1' of size m_1 as follows:

$$(x_1')_e = \begin{cases} (x_1)_e + \delta & \text{if } e = e_i^L \text{ and } i \text{ is odd,} \\ (x_1)_e - \delta & \text{if } e = e_i^L \text{ and } i \text{ is even,} \\ (x_1)_e - 2\delta & \text{if } e = e_i^D \text{ and } i \text{ is odd,} \\ (x_1)_e + 2\delta & \text{if } e = e_i^D \text{ and } i \text{ is even.} \end{cases}$$

Again we can get a vector x_f' satisfying the closeness condition where a contradiction occurs.

Recall that C_1 can not contain any even cycle. Now we know that if C_1 is not a single node, then it must be an odd cycle. In the latter case, the incident square matrix A_1 is nonsingular. Therefore we have $A_1 x_1 = b_1$. It can be shown by induction that the determinant $|A_1| = 2$. Since b_1 is integral, x_1 is half-integral by Cramer's rule.

Generally, for each connected component C_i, where $i \in [k]$, the corresponding edge vector is either empty or half-integral. In conclusion, x_f must be half-integral and thus x is half-integral. $\square$

Propositions 3(b) and (c) state that any optimal LP solution x must set $x_e = 1$ for a greedy edge $e \in \widehat{E}$, and set $x_e = 0$ for a idealistic edge e. Hence, we analyze the contribution of the ambiguous edges to the optimal value.

Lemma 6 (**Integrality Gap**) *Suppose x is a half-integral solution to* LP *(with node capacity vector b') that takes non-zero values on the edge set E' . Then, there exists a b'-matching H in E' such that $w(H) \geq \frac{2}{3} \sum_{e \in E'} w_e \cdot x_e$.*

Proof Since x is half-integral, we can first include all edges $e \in E'$ such that $x_e = 1$ in H; this ensures that for edges e such that $x_e = 1$, they contribute the same to $w(H)$ and $w(x)$. We next transform the solution x, if necessary, such that the set J of $\frac{1}{2}$-edges form vertex-disjoint odd cycles.

Observe that if there is a node i such that its degree in J is odd, then there must exist a path in J from i to another odd degree node j; moreover, both i and j are unsaturated in x. Hence, using standard alternating path argument, we can transform the solution x without decreasing $w(x)$ such that all edges on the path has value either 0 or 1.

We can now assume that all degrees of nodes in J are even. Hence, each connected component in J has an Euler circuit. If the circuit is even, then again we can use alternating circuit argument to choose a solution (without decreasing its value) such that all edges on the circuit is either 0 or 1. If the circuit is odd but not simple, then there must be an even circuit, which can be eliminated again; hence, any remaining edges in J form vertex-disjoint odd cycles.

Consider each odd cycle C. Observe that the edges in C can be partitioned into three sets C_1, C_2 and C_3 such that each set forms a (1-)matching; moreover, the contribution of C to the value of $w(x)$ solution is $\frac{1}{2} w(C) = \sum_{e \in C} w_e x_e$. Hence, if we pick the C_r with the largest weight and include it in H, then we have $w(C_r) \geq \frac{1}{3} w(C) = \frac{2}{3} \sum_{e \in C} w_e x_e$.

Since H gets the same contribution as $w(x)$ from integral edges and at least $\frac{2}{3}$ fraction from the fractional edges, it follows that $w(H)$ is at least $\frac{2}{3} w(x)$, as required. $\square$

We summarize the main result of this section in the following theorem.

Theorem 5 (**Price of Stability**) *Suppose the given graph has no cycle with equal alternating weight or* LP *has unique optimum. Then, there exists a greedy and idealistic payoff function u such that any fixed point m of $\mathcal{T}$ is a Nash Equilibrium in $\mathcal{N}_u$; moreover, the social welfare $\mathcal{S}_u(m) \geq \frac{2}{3} \mathcal{S}_{\mathsf{LP}} \geq \frac{2}{3} \max_{m' \in A_E} \mathcal{S}_u(m')$, showing that the Price of Stability is at most 1.5*

Proof The statements about fixed point and Nash Equilibrium follow from Theorem 2. We focus on analyzing the social welfare of a fixed point m under the payoff function u that results from choosing a maximum weight matching H from the set of ambiguous edges $\overline{E}$. Suppose x is an optimal solution to LP. Since greedy edges $\widehat{E}$ will be chosen under m, we have $\mathcal{S}(m) = \sum_{e \in \widehat{E}} w_e + \sum_{f \in H} w_f$.

Suppose $x \in \mathcal{L}$ is an optimal solution to LP, i.e., $w(x) = \mathcal{S}_{\mathsf{LP}}$. Observe that Theorem 3(b) implies that for $e \in \widehat{E}$, $x_e = 1$, and Theorem 3(c) implies that for $e \in E$ such that $x_e > 0$, $e \in \widehat{E} \cup \overline{E}$. Hence, it follows that $\mathcal{S}_{\mathsf{LP}} = \sum_{e \in \widehat{E}} w_e + \sum_{f \in \overline{E}} x_f \cdot w_f$.

Observe that x restricted to $\overline{E}$ is still an optimal solution to the LP restricted to $\overline{E}$ with remaining node capacity vector b' (after accounting for the greedy edges). Then, from Lemma 6, we have $\sum_{f \in H} w_f \geq \frac{2}{3} \sum_{f \in \overline{E}} x_f \cdot w_f$, and hence $\mathcal{S}(m) \geq \frac{2}{3} \mathcal{S}_{\mathsf{LP}}$, as required. □

Remark 1 We remark that if we use a distributed algorithm such as [18, 19, 21] to find a $(1 + c)$-approximate maximum matching among the ambiguous edges, then we can show that the resulting price of stability is at most $1.5(1 + c)$.

6 Rate of Convergence: ϵ-Greedy and ϵ-Idealistic Users

Although the iterative protocol described in Algorithm 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 Sections 4 and 5 can be extended if we relax the notions of greedy and idealistic users.

Suppose $\epsilon \geq 0$. We say the payoff function u is ϵ-*greedy* (or the users are ϵ-greedy), if for each $e = \{i, j\} \in E$ and $m \in A_E$, if $m_{j \to i} > \overline{m}_i + \epsilon$ and $m_{i \to j} > \overline{m}_j + \epsilon$, then $u_e(m) = 1$. We say the payoff function u is ϵ-*idealistic*, if for each $e = \{i, j\} \in E$ and $m \in A_E$, if $m_{j \to i} < \widehat{m}_i - \epsilon$, then $u_e(m) = 0$. Given $m \in A_E$, we can place each edge $e = \{i, j\} \in E$ in the following ϵ-categories (with respect to m).

1. **ϵ-Greedy Edges:** $m_{j \to i} > \overline{m}_i + \epsilon$ and $m_{i \to j} > \overline{m}_j + \epsilon$. If u is ϵ-greedy, then $u_e(m) = 1$.
2. **ϵ-Idealistic Edges:** $m_{j \to i} < \widehat{m}_i - \epsilon$ or $m_{i \to j} < \widehat{m}_j - \epsilon$. If u is ϵ-idealistic, then $u_e(m) = 0$.
3. **ϵ-Ambiguous Edges:** these are the remaining edges that are neither ϵ-greedy nor ϵ-idealistic.

As in Section 5, we consider an ϵ-greedy and ϵ-idealistic payoff function u that corresponds to the selection process of finding a maximum weight matching among the ϵ-ambiguous edges, after accepting the ϵ-greedy edges and rejecting the ϵ-idealistic edges.

Recall that given $\alpha \in [0, W]^{2|E|}$, $m = m(\alpha) \in A_E$ is defined by (2). We shall use the following fact about convergence.

Fact 3 (**Convergence**) Suppose the sequence $\{\alpha^{(t)}\}$ converges to $\alpha \in [0, W]^{2|E|}$ under the ℓ_∞ norm, where each $\alpha^{(t)}$ and α define $m^{(t)}$ and m respectively. Then, for all $\epsilon > 0$, there exists $T > 0$ such that for all $t \geq T$, the following holds.

1. For all $\{i, j\} \in E$, $|m_{j \to i} - m^{(t)}_{j \to i}| \leq \epsilon$ and $|\alpha_{i \backslash j} - \alpha^{(t)}_{i \backslash j}| \leq \epsilon$.

2. For all $i \in V$, $|\widehat{m}_i - \widehat{m}_i^{(t)}| \leq \epsilon$ and $|\overline{m}_i - \overline{m}_i^{(t)}| \leq \epsilon$.

Theorem 6 (**Convergence of Social Welfare**) *Suppose the given graph has no cycle with equal alternating weight or* LP *has a unique optimum. For any $\epsilon > 0$, there exists an ϵ-greedy and ϵ-idealistic payoff function u such that the following holds. For any sequence $\{(m^{(t)}, \alpha^{(t)})\}$ produced by the iterative protocol in Algorithm 1, there exists $T > 0$ such that for all $t \geq T$, the social welfare $\mathcal{S}_u(m^{(t)}) \geq \frac{2}{3}\mathcal{S}_{\mathsf{LP}}$.*

Proof Fix $\epsilon > 0$. By Theorem 1, the sequence $\{(m^{(t)}, \alpha^{(t)})\}$ converges to (m, α).

Consider the ϵ-greedy and ϵ-idealistic payoff function u, that when acting on $m^{(t)}$, corresponds to selecting a maximum matching among the ϵ-ambiguous edges $\overline{E}^{(t)}$, after selecting the ϵ-greedy edges $\widehat{E}^{(t)}$ (both with respect to $m^{(t)}$).

Suppose $\widehat{E}$ is the set of greedy edges and $\overline{E}$ is the set of ambiguous edges, both with respect to m. Observe that as long as the payoff function u is conservative in selecting greedy edges (i.e., $\widehat{E}^{(t)} \subseteq \widehat{E}$) and rejecting idealistic edges (i.e., $\widehat{E} \cup \overline{E} \subseteq \widehat{E}^{(t)} \cup \overline{E}^{(t)}$), then $\mathcal{S}_u(m^{(t)}) \geq \mathcal{S}_{\widehat{u}}(m) \geq \frac{2}{3}\mathcal{S}_{\mathsf{LP}}$, where $\widehat{u}$ is the greedy and idealistic payoff function as in Theorem 5.

Hence, it suffices to show that for large enough t, both $\widehat{E}^{(t)} \subseteq \widehat{E}$ and $\widehat{E} \cup \overline{E} \subseteq \widehat{E}^{(t)} \cup \overline{E}^{(t)}$ hold.

Let $T > 0$ be large enough such that the upper bounds in Fact 3 hold with $\frac{\epsilon}{2}$. Consider $t \geq T$.

To prove $\widehat{E}^{(t)} \subseteq \widehat{E}$, it suffices to show that any ϵ-greed edge $e = \{i, j\}$ with respect to $m^{(t)}$ is greedy with respect to m. Observe that $m_{j\to i}^{(t)} > \overline{m}_i^{(t)} + \epsilon$, $|m_{j\to i}^{(t)} - m_{j\to i}| \leq \frac{\epsilon}{2}$ and $|\overline{m}_i^{(t)} - \overline{m}_i| \leq \frac{\epsilon}{2}$ imply that $m_{j\to i} > \overline{m}_i$.

Similarly, to prove that $\widehat{E} \cup \overline{E} \subseteq \widehat{E}^{(t)} \cup \overline{E}^{(t)}$, it suffices to show that any ϵ-idealistic edge $e = \{i, j\}$ with respect to $m^{(t)}$ is idealistic with respect to m. One also observes that $m_{j\to i}^{(t)} < \widehat{m}_i^{(t)} - \epsilon$, $|m_{j\to i}^{(t)} - m_{j\to i}| \leq \frac{\epsilon}{2}$ and $|\widehat{m}_i^{(t)} - \widehat{m}_i| \leq \frac{\epsilon}{2}$ imply that $m_{j\to i} < \widehat{m}_i$. □

6.1 Rate of Convergence

Although the result in [14] shows that the configurations given by the iterative protocol will converge to a fixed point, it does not give the rate of convergence. However, a result by Baillon and Bruck [3] tells us how fast the protocol can arrive at an approximate fixed point.

Given $\delta \geq 0$, a δ-fixed point α for a function $\mathcal{T}$ satisfies $||\alpha - \mathcal{T}(\alpha)||_\infty \leq \delta$.

Proposition 5 ([3]) *Suppose the maximum edge weight is W. Then, after iteration t of the distributed protocol in Algorithm 1, the $\alpha^{(t)} \in [0, W]^{2|E|}$ returned is a $O(\frac{W}{\sqrt{t}})$-fixed point of $\mathcal{T}$.*

We also need a modified technical assumption on the network topology. Given $\lambda \geq 0$, a cycle is said to have λ-*equal alternating weight* if it is even, and the sum of the even edge weights differs from the sum of the odd edge weights by at most λ. We remark that all our arguments in Sections 4 and 5 can be extended to δ-fixed points, and ϵ-greedy and ϵ-idealistic payoff functions in a straightforward manner. We state the following result and defer the proof to the next section.

Theorem 7 (δ-Fixed Point, ϵ-Greedy and ϵ-Idealistic, and LP Optimum) *Suppose $\epsilon > 0$ and the given graph has no cycle with ϵ-equal alternating weight. Suppose further that (m, α, S) is a δ-fixed point of $\mathcal{T}$, where $\delta = O(\frac{\epsilon}{|V|^2})$, and $x \in \mathcal{L}$ is an optimal solution to* LP. *Then, the following holds.*

1. *If an edge e is ϵ-greedy with respect to m, then $x_e = 1$.*
2. *If an edge e is ϵ-idealistic with respect to m, then $x_e = 0$.*

Applying Proposition 5 and Theorem 7 to previous analysis can give the following theorem.

Theorem 8 (**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-idealistic payoff function u such that the following holds. For any sequence $\{(m^{(t)}, \alpha^{(t)})\}$ produced by the iterative protocol in Algorithm 1, and for all $t \geq \Theta(\frac{W^2|V|^4}{\epsilon^2})$, the social welfare $\mathcal{S}_u(m^{(t)}) \geq \frac{2}{3}\mathcal{S}_{\mathsf{LP}}$.*

6.2 δ-Fixed Point and LP Optimum

In this section we prove Theorem 7. Note that Proposition 3 holds for a δ-fixed point of $\mathcal{T}$. We further state the following properties for a δ-fixed point.

Proposition 6 (δ-Fixed Point) *Suppose (m, α, S) is a δ-fixed point of $\mathcal{T}$. Then for each $\{i, j\} \in E$ and $c \geq 0$, the following properties hold.*

(a) *$\alpha_{i\setminus j} \leq \widehat{m}_i + \delta$ and $\alpha_{i\setminus j} \geq \overline{m}_i - \delta$.*
(b) *If $m_{j\to i} \geq \alpha_{i\setminus j} - c$, then $m_{j\to i} \geq \widehat{m}_i - c - \delta$.*
(c) *If $m_{j\to i} \leq \alpha_{i\setminus j} + c$, then $m_{j\to i} \leq \overline{m}_i + c + \delta$.*
(d) *If $m_{j\to i} > 0$ and $m_{j\to i} \geq \widehat{m}_i - c$, then $m_{i\to j} \geq \widehat{m}_j - c - 2\delta$.*
(e) *If $m_{j\to i} > 0$ and $m_{j\to i} \geq \widehat{m}_i - c$, then $\widehat{m}_i + \widehat{m}_j \leq w_{ij} + 2c + 2\delta$.*
(f) *If $m_{j\to i} \leq \overline{m}_i + c$, then $m_{i\to j} \leq \overline{m}_j + c + 2\delta$.*
(g) *If $m_{j\to i} \leq \overline{m}_i + c$, then $\widehat{m}_i + \widehat{m}_j \geq w_{ij} - 2c - 2\delta$.*
(h) *Suppose further that $c \geq \delta$. If $m_{j\to i} < \widehat{m}_i - c$, then $\widehat{m}_i + \widehat{m}_j > w_{ij} + c - \delta$.*

Proof (a) This follows directly from the definition of α and δ-fixed point, observing that $\overline{m}_i \leq \mathcal{T}(\alpha)_{i\setminus j} \leq \widehat{m}_i$.

(b) The result is true, if $m_{j\to i} \geq \widehat{m}_i$; hence we assume $m_{j\to i} < \widehat{m}_i$, which implies $\mathcal{T}(\alpha)_{i\setminus j} = \widehat{m}_i$. Therefore, by the definition of δ-fixed point, we have $|\alpha_{i\setminus j} - \widehat{m}_i| \leq \delta$ and thus $m_{j\to i} \geq \alpha_{i\setminus j} - c \geq \widehat{m}_i - \delta - c$. Therefore, we always have $m_{j\to i} \geq \widehat{m}_i - c - \delta$.

(c) The result is true, if $m_{j\to i} \leq \overline{m}_i$; hence we assume $m_{j\to i} > \overline{m}_i$, which implies $\mathcal{T}(\alpha)_{i\setminus j} = \overline{m}_i$. Therefore, by the definition of δ-fixed point, we have $|\alpha_{i\setminus j} -$

$\overline{m}_i| \le \delta$ and thus $m_{j\to i} \le \alpha_{i\setminus j} + c \le \overline{m}_i + \delta + c$. Therefore, we always have $m_{j\to i} \le \overline{m}_i + c + \delta$.

(d) Consider the following cases.

- If $m_{j\to i} > \alpha_{i\setminus j}$, then from Proposition 3 we have $S_{ij} > 0$ and thus $m_{i\to j} > \alpha_{j\setminus i}$. From (b) we have $m_{i\to j} \ge \widehat{m}_j - \delta$.
- If $m_{j\to i} \le \alpha_{i\setminus j}$, then from Proposition 3 we have $S_{ij} \le 0$ and thus $m_{j\to i} = w_{ij} - \alpha_{j\setminus i} = \alpha_{i\setminus j} + S_{ij}$. We also have $m_{j\to i} \ge \widehat{m}_i - c \ge \alpha_{i\setminus j} - \delta - c$ from (a). So $S_{ij} \ge -c-\delta$. Then $m_{i\to j} = (w_{ij} - \alpha_{i\setminus j})_+ \ge \alpha_{j\setminus i} + S_{ij} \ge \alpha_{j\setminus i} - c - \delta$. From (b), we have $m_{i\to j} \ge \widehat{m}_j - c - 2\delta$.

(e) From (d) we have $\widehat{m}_i + \widehat{m}_j \le m_{j\to i} + m_{i\to j} + 2\delta + 2c \le w_{ij} + 2\delta + 2c$.

(f) Consider the following cases.

- If $m_{j\to i} > \alpha_{i\setminus j}$, then from Proposition 3 we have $S_{ij} > 0$ and thus $m_{j\to i} = \alpha_{i\setminus j} + \frac{1}{2}S_{ij}$. We also have $m_{j\to i} \le \overline{m}_i + c \le \alpha_{i\setminus j} + \delta + c$, from (a). Hence, $\frac{1}{2}S_{ij} \le c + \delta$. Then $m_{i\to j} = \alpha_{j\setminus i} + \frac{1}{2}S_{ij} \le \alpha_{j\setminus i} + c + \delta$. From (c), we have $m_{i\to j} \le \overline{m}_j + c + 2\delta$.
- If $m_{j\to i} \le \alpha_{i\setminus j}$, then from Proposition 3 we have $S_{ij} \le 0$ and thus $m_{i\to j} \le \alpha_{j\setminus i}$. From (c) we have $m_{i\to j} \le \overline{m}_j + \delta$.

(g) Consider the following cases.

- If $S_{ij} > 0$, then from (f) we have $\widehat{m}_i + \widehat{m}_j \ge \overline{m}_i + \overline{m}_j \ge m_{j\to i} + m_{i\to j} - 2c - 2\delta = w_{ij} - 2c - 2\delta$.
- If $S_{ij} \le 0$, then from (a) we have $\widehat{m}_i + \widehat{m}_j \ge \alpha_{i\setminus j} + \alpha_{j\setminus i} + 2\delta \ge w_{ij} + 2\delta$.

(h) Since $m_{j\to i} < \widehat{m}_i - c$ implies that $\mathcal{T}(\alpha)_{i\setminus j} = \widehat{m}_i$, then by the definition of δ-fixed point we have $|\alpha_{i\setminus j} - \widehat{m}_i| \le \delta$. Since $c \ge \delta$, we have $m_{j\to i} < \widehat{m}_i - c \le \alpha_{i\setminus j} + \delta - c \le \alpha_{i\setminus j}$, that is $m_{j\to i} < \alpha_{i\setminus j}$. From Proposition 3, we have $S_{ij} < 0$ and thus $m_{j\to i} = (w_{ij} - \alpha_{j\setminus i})_+$. Observe that $\widehat{m}_j \ge \alpha_{j\setminus i} - \delta$, from (a). Therefore, we have $\widehat{m}_i + \widehat{m}_j > (m_{j\to i} + c) + (\alpha_{j\setminus i} - \delta) = (m_{j\to i} + \alpha_{j\setminus i}) + c - \delta = ((w_{ij} - \alpha_{j\setminus i})_+ + \alpha_{j\setminus i}) + c - \delta \ge w_{ij} + c - \delta$. □

Observe that Theorem 7 can be implied by the following lemma.

Lemma 7 (δ-Fixed Point and LP Optimum) *Suppose $\epsilon > 0$ and the given graph has no cycle with ϵ-equal alternating weight and further that x is a feasible solution of* LP. *Then there exists a δ-fixed point (m, α, S) of $\mathcal{T}$, where $\delta = O(\frac{\epsilon}{|V|^2})$, such that for each $\{i, j\} \in E$, the following properties hold.*

(a) *If there exists an edge e such that e is ϵ-greedy and $x_e < 1$, then there exists $\widehat{x} \in \mathcal{L}$ such that $\widehat{x} \ne x$ and $w(\widehat{x}) > w(x)$.*
(b) *If there exists an edge e such that e is ϵ-idealistic and $x_e > 0$, then there exists $\widehat{x} \in \mathcal{L}$ such that $\widehat{x} \ne x$ and $w(\widehat{x}) > w(x)$.*

In the rest of this paper we assume ϵ is sufficiently larger than δ, say $\epsilon = \Omega(n^2\delta)$. Similarly to the argument used in Section 4, we use a unifying framework of (ϵ, δ)-alternating traversal.

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

1. A (ϵ, δ)-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). $\mathcal{Q}$ 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 $S^+ := \{\{i, j\} \in E : m_{j\to i} \geq \widehat{m}_i - \epsilon$ and $m_{i \to j} \geq \widehat{m}_j - \epsilon\}$ is the set of edges that are **not** ϵ-idealistic, and $S^- := \{\{i, j\} \in E : m_{j \to i} \leq \overline{m}_i + \epsilon$ or $m_{i \to j} \leq \overline{m}_j + \epsilon\}$ is the set of edges that are **not** ϵ-greedy. We require that an edge appearing for multiple times in $\mathcal{Q}$ cannot appear both in $\mathcal{Q}^+$ and in $\mathcal{Q}^-$.
   The (ϵ, δ)-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 ϵ-***strict*** if e is either ϵ-greedy or ϵ-idealistic. Hence, the edges in statements (a) and (b) of Lemma 7 are both critical and ϵ-strict.
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 $(\mathcal{T}(\alpha))_{i \setminus j} = 0$.
   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).
3. As described in Section 4, the alternating traversal is obtained from the growing procedure starting from some **seed** edge, which in this case is both critical and ϵ-strict. Observe that the alternating traversal might not contain the seed edge.

Lemma 8 ((ϵ, δ)-Alternative Feasible Solution.) *Suppose $\mathcal{Q}$ is a feasible (ϵ, δ)-alternating traversal with respect to some feasible $x \in \mathcal{L}$. Then, there exists feasible $\widehat{x} \neq x$ such that $w(\widehat{x}) - w(x)$ has the same sign ($\{-1, 0, +1\}$) as $w(\mathcal{Q}^+) - w(\mathcal{Q}^-)$.*

Proof The proof is exactly the same as that of Lemma 2. □

In view of Lemma 8, it suffices to show that we can find a feasible (ϵ, δ)-traversal $\mathcal{Q}$ such that $w(\mathcal{Q}^+) - w(\mathcal{Q}^-) > 0$. We make use of the following slack variables for each edge to analyze the weights of the edges.

Slack Variables For each edge $e \in \mathcal{Q}$, there exists some $c_e \geq 0$ such that

(a) for $e = \{i, j\} \in \mathcal{Q}^+$ that is grown from i, $m_{j \to i} \geq \widehat{m}_i - c_e$; define $d_e := \widehat{m}_i + \widehat{m}_j - w_{ij}$.
(b) for $e = \{j, k\} \in \mathcal{Q}^-$ that is grown from j, $m_{k \to j} \leq \overline{m}_j + c_e$; define $d_e := w_{jk} - (\widehat{m}_j + \widehat{m}_k)$.

Observe that by Proposition 6(e) and (g), we have $d_e \leq 2c_e + 2\delta$.

We show that there is some way to grow the alternating traversal such that the slack variables can be kept small. Given an edge e, its *hop number* is the number of steps away from the seed edge in the growing procedure. For instance, the seed edge has hop number 0 and the next edge grown adjacent to the seed edge has hop number 1, and so on.

Lemma 9 (Growing Feasible (ϵ, δ)-Alternating Traversal) *Suppose a δ-fixed point (m, α, S) and a feasible $x \in \mathcal{L}$ are given as above. Assume that $c \geq 0$ is a constant.*

1. *Suppose $\{i, j\} \in E^+$ such that $m_{j\to i} > 0$ and $m_{j\to i} \geq \widehat{m}_i - c$. Suppose further $c \leq \epsilon - 2\delta$ and that node j is saturated (we stop if j is unsaturated). Then, there exists some node $k \in N(j) \setminus i$ such that $\{j, k\} \in E^-$ and $m_{k\to j} \leq \overline{m}_j + c + 2\delta$.*
2. *Suppose $\{j, k\} \in E^-$ such that $m_{k\to j} \leq \overline{m}_j + c$. Suppose further $c \leq \epsilon - 4\delta$ and $(\mathcal{T}(\alpha))_{k\setminus j} > 0$ (we stop if $(\mathcal{T}(\alpha))_{k\setminus j} = 0$). Then, there exists some node $l \in N(k) \setminus j$ such that $\{k, l\} \in E^+$, $m_{l\to k} > 0$ and $m_{l\to k} \geq \widehat{m}_k - c - 2\delta$.*

In particular, it follows that an edge e with hop number t has $c_e \leq (2t + 2)\delta$.

Proof 1. Suppose $\{i, j\} \in E^+$ such that $m_{j\to i} > 0$ and $m_{j\to i} \geq \widehat{m}_i - c$ and node j is saturated. Since $x_{ij} < 1$ and $\sum_{k\in N(j)} x_{jk} = b_j$, there are at least b_j nodes k in $N(j) \setminus i$ such that $x_{jk} > 0$. We pick the k such that $m_{k\to j}$ is the smallest. Then it follows that $m_{k\to j} \leq (\mathcal{T}(\alpha))_{j\setminus i}$. Since $m_{j\to i} > 0$ and $m_{j\to i} \geq \widehat{m}_i - c$, from Proposition 6(d) we have $m_{i\to j} \geq \widehat{m}_j - c - 2\delta$. If $m_{i\to j} \geq \widehat{m}_j$, then $m_{k\to j} \leq (\mathcal{T}(\alpha))_{j\setminus i} = \overline{m}_j$. Otherwise, $m_{i\to j} \leq \overline{m}_j$ and thus $\widehat{m}_j - \overline{m}_j \leq c + 2\delta$. Then $m_{k\to j} \leq (\mathcal{T}(\alpha))_{j\setminus i} = \widehat{m}_j \leq \overline{m}_j + c + 2\delta$.

   To show $\{j, k\} \in E^-$, it suffices to prove that $\{j, k\} \in \mathcal{Q}^-$. With the condition $c \leq \epsilon - 2\delta$, we have $m_{k\to j} \leq \overline{m}_j + \epsilon$. Therefore $\{j, k\} \in \mathcal{Q}^-$.
2. Suppose $\{j, k\} \in E^-$ such that $m_{k\to j} \leq \overline{m}_j + c$ and $(\mathcal{T}(\alpha))_{k\setminus j} > 0$. Then $x \in \mathcal{L}$ and $x_{jk} > 0$ implies that there are at most $b_k - 1$ neighbors $i \in N(k) \setminus j$ such that $x_{ik} = 1$. Suppose $l \in N(k) \setminus j$ such that $x_{kl} < 1$ and $m_{l\to k}$ is the largest. Then it follows that $m_{l\to k} \geq (\mathcal{T}(\alpha))_{k\setminus j} > 0$. Since $m_{k\to j} \leq \overline{m}_j + c$, from Proposition 6(f) we have $m_{j\to k} \leq \overline{m}_k + c + 2\delta$. If $m_{j\to k} \leq \overline{m}_k$, then $m_{l\to k} \geq (\mathcal{T}(\alpha))_{k\setminus j} = \widehat{m}_k$. Otherwise, $m_{j\to k} \geq \widehat{m}_k$ and thus $\widehat{m}_k - \overline{m}_k \leq c + 2\delta$. Then $m_{l\to k} \geq (\mathcal{T}(\alpha))_{k\setminus j} = \overline{m}_k \geq \widehat{m}_k - c - 2\delta$.

To show $\{k, l\} \in E^+$, it suffices to prove that $\{k, l\} \in \mathcal{Q}^+$. With the condition $c \leq \epsilon - 2\delta$ and Proposition 6(d), we have $m_{l\to k} \geq \widehat{m}_k - c - 2\delta \geq \widehat{m}_k - \epsilon$ and $m_{k\to l} \geq \widehat{m}_l - c - 4\delta \geq \widehat{m}_l - \epsilon$. Therefore $\{k, l\} \in \mathcal{Q}^+$.

Moreover, if $e_0 := \{x, y\}$ is the seed edge corresponding to some (ϵ, δ)-alternating traversal, then either i) e_0 is ϵ-greedy, i.e., $m_{y\to x} > \overline{m}_x + \epsilon$ and $m_{x\to y} > \overline{m}_y + \epsilon$, which implies $m_{y\to x} \geq \widehat{m}_x$ and $m_{x\to y} \geq \widehat{m}_y$; or ii) e_0 is ϵ-idealistic, i.e., $m_{y\to x} < \widehat{m}_x - \epsilon$ or $m_{x\to y} < \widehat{m}_y - \epsilon$, which implies $m_{y\to x} \leq \overline{m}_x$ or $m_{x\to y} \leq \overline{m}_y$, indicating that $m_{y\to x} \leq \overline{m}_x + 2\delta$ and $m_{x\to y} \leq \overline{m}_y + 2\delta$ by Proposition 6(f). Hence, we have $c_{e_0} \leq 2\delta$. By induction, it is obvious that an edge e with hop number t satisfies $c_e \leq (2t + 2)\delta$. □

Observe that each edge appears at most twice in the traversal and has hop number at most $2n - 1$ from the seed edge. Hence, from Lemma 9, for all edges $e \in \mathcal{Q}$, we have $c_e \leq 4n\delta$ and $d_e \leq (8n + 2)\delta$.

Lemma 10 (Property of End Node) *Suppose an (ϵ, δ)-alternating traversal $\mathcal{Q}$ is a path and i is an end node with $\{i, j\} \in \mathcal{Q}^-$. If $m_{j\to i} \leq \overline{m}_i + c$, then $\widehat{m}_i \leq c$, where $c \geq 0$ is a constant.*

Proof From the definition of (ϵ, δ)-alternating traversal, we have $(\mathcal{T}(\alpha))_{i\backslash j} = 0$. If $m_{j\to i} < \widehat{m}_i$, then $\widehat{m}_i = (\mathcal{T}(\alpha))_{i\backslash j} = 0 \leq c$. Suppose $m_{j\to i} \geq \widehat{m}_i$. Then $\overline{m}_i = (\mathcal{T}(\alpha))_{i\backslash j} = 0$. Therefore $\widehat{m}_i \leq m_{j\to i} \leq \overline{m}_i + c = c$. □

Lemma 11 (Analyzing Weight with Slack Variables) *Given an (ϵ, δ)-alternating traversal $\mathcal{Q}$, we have $w(\mathcal{Q}^+) - w(\mathcal{Q}^-) \geq -\sum_{e\in\mathcal{Q}} r_{\mathcal{Q}}(e) \cdot d_e - 8n\delta$, where $r_{\mathcal{Q}}(e)$ is the number of times e appears in $\mathcal{Q}$.*

Proof Let $V_{\mathcal{Q}}$ be the set of nodes that appear in the traversal $\mathcal{Q}$. For all $v \in V_{\mathcal{Q}}$, define $E_v := \{e \in \mathcal{Q} : e \text{ is incident to } v\}$. Consider the following two cases.

(1) $\mathcal{Q}$ is a circuit. Then for all $v \in V_{\mathcal{Q}}$, we have $|\mathcal{Q}^+ \cap E_v| = |\mathcal{Q}^- \cap E_v|$. By the definition of the slack variable d, we have $\sum_{e_1\in\mathcal{Q}^+} r_{\mathcal{Q}^+}(e_1)(w_{e_1} + d_{e_1}) = \sum_{e_2\in\mathcal{Q}^-} r_{\mathcal{Q}^-}(e_2)(w_{e_2} - d_{e_2})$. Rearranging the equality implies $w(\mathcal{Q}^+) - w(\mathcal{Q}^-) = -\sum_{e\in\mathcal{Q}} r_{\mathcal{Q}}(e) d_e \geq -\sum_{e\in\mathcal{Q}} r_{\mathcal{Q}}(e) d_e - 8n\delta$.

(2) $\mathcal{Q}$ is a path. Define $V_{\mathcal{Q}^+} := \{x \in V_{\mathcal{Q}} : x \text{ is an end point of } \mathcal{Q} \text{ with } \{x, y\} \in \mathcal{Q}^+\}$ and $V_{\mathcal{Q}^-} := \{x \in V_{\mathcal{Q}} : x \text{ is an end point of } \mathcal{Q} \text{ with } \{x, y\} \in \mathcal{Q}^-\}$. Then, for any $v \in V_{\mathcal{Q}}$, we have the following claims.

   (a) If $v \in V_{\mathcal{Q}^+}$, then $|\mathcal{Q}^+ \cap E_v| - |\mathcal{Q}^- \cap E_v| = 1$.
   (b) If $v \in V_{\mathcal{Q}^-}$, then $|\mathcal{Q}^- \cap E_v| - |\mathcal{Q}^+ \cap E_v| = 1$.
   (c) if $v \in V_{\mathcal{Q}} \setminus (V_{\mathcal{Q}^+} \cup V_{\mathcal{Q}^-})$, then $|\mathcal{Q}^+ \cap E_v| = |\mathcal{Q}^- \cap E_v|$.

   Note that for all $e \in \mathcal{Q}$, we have $c_e \leq 4n\delta$. Then from Lemma 10, for any $v \in V_{\mathcal{Q}^-}$ we have $\widehat{m}_v \leq 4n\delta$. Also observe that $|V_{\mathcal{Q}^-}| \leq 2$. Then, by the definition of the slack variable d, we have $w(\mathcal{Q}^+) - w(\mathcal{Q}^-) = -\sum_{e\in\mathcal{Q}} r_{\mathcal{Q}}(e) d_e + \sum_{v_1\in V_{\mathcal{Q}^+}} \widehat{m}_{v_1} - \sum_{v_2\in V_{\mathcal{Q}^-}} \widehat{m}_{v_2} \geq -\sum_{e\in\mathcal{Q}} r_{\mathcal{Q}}(e) d_e + 0 - 2 \cdot 4n\delta = -\sum_{e\in\mathcal{Q}} r_{\mathcal{Q}}(e) d_e - 8n\delta$. □

We can obtain an (ϵ, δ)-alternating traversal $\mathcal{Q}$ by applying the growing procedure as described in Section 4, which in this case starts from a seed edge e_0 that is critical and ϵ-strict, and goes with the rules set $\mathcal{R}$ indicated in Lemma 9. Observe that either $\mathcal{Q}$ is a simple even cycle or it contains the seed edge e_0.

Lemma 12 ((ϵ, δ)-Alternating Traversal Weight) *Suppose that we have a δ-fixed point (m, α, S) and a feasible solution x to* LP. *Suppose further $\delta = \frac{\epsilon}{22n^2}$ and there is no cycle with ϵ-equal alternating weight. Then, the growing procedure gives an (ϵ, δ)-alternating traversal $\mathcal{Q}$ such that $w(\mathcal{Q}^+) > w(\mathcal{Q}^-)$.*

Proof We consider the following two cases.

(1) The traversal $\mathcal{Q}$ is a simple even cycle. By the no cycle with ϵ-equal alternating weight assumption, we have $|w(\mathcal{Q}^+) - w(\mathcal{Q}^-)| > \epsilon$. Observe that any $e \in \mathcal{Q}$ satisfies $d_e \leq (8n + 2)\delta$. Also note that $|\mathcal{Q}| \leq 2n$. Then from Lemma 11, we

have $w(\mathcal{Q}^+) - w(\mathcal{Q}^-) \geq -\sum_{e \in \mathcal{Q}} r_{\mathcal{Q}}(e) \cdot d_e \geq -2n(8n+2)\delta \geq -\epsilon$. Therefore, $w(\mathcal{Q}^+) - w(\mathcal{Q}^-) \geq \epsilon > 0$.

(2) The traversal $\mathcal{Q}$ is a simple path, lollipop with odd cycle or dumbbell. In this case $e_0 := \{x, y\} \in \mathcal{Q}$. Then one of the following two cases happens.

(a) The seed edge e_0 is ϵ-idealistic. Then either $m_{y \to x} < \widehat{m}_x - \epsilon$ or $m_{x \to y} < \widehat{m}_y - \epsilon$. From Proposition 6(h) we have $\widehat{m}_x + \widehat{m}_y > w_{xy} + \epsilon - \delta$. Since $e_0 \in \mathcal{Q}^-$, we have $d_{e_0} < -(\epsilon - \delta) < 0$. Therefore, from Lemma 11 we have $w(\mathcal{Q}^+) - w(\mathcal{Q}^-) \geq -\sum_{e \in \mathcal{Q}} r_{\mathcal{Q}}(e) \cdot d_e - 8n\delta \geq -\sum_{e \in \mathcal{Q} \setminus \{e_0\}} r_{\mathcal{Q}}(e) \cdot d_e - d_{e_0} - 8n\delta > -(2n-1)(8n+2)\delta + (\epsilon - \delta) - 8n\delta \geq 0$.

(b) The seed edge e_0 is ϵ-greedy. Then $m_{y \to x} > \overline{m}_x + \epsilon$ and $m_{x \to y} > \overline{m}_y + \epsilon$, which implies $m_{y \to x} \geq \widehat{m}_x$ and $m_{x \to y} \geq \widehat{m}_y$. Note that $e_0 \in \mathcal{Q}^+$. Let $e_1 := \{y, z\} \in \mathcal{Q}^-$ be the edge next to e_0 in the traversal $\mathcal{Q}$. If we know that either $d_{e_0} < -(\epsilon - 4\delta)$ or $d_{e_0} + d_{e_1} < -(\epsilon - 4\delta)$, then from Lemma 11, we have $w(\mathcal{Q}^+) - w(\mathcal{Q}^-) \geq -\sum_{e \in \mathcal{Q}} r_{\mathcal{Q}}(e) \cdot d_e - 8n\delta > -(2n-1)(8n+2)\delta + (\epsilon - 4\delta) - 8n\delta \geq 0$.

To finish the proof, we only need to show that $d_{e_0} < -(\epsilon - 4\delta)$ or $d_{e_0} + d_{e_1} < -(\epsilon - 4\delta)$.

Note that $(\mathcal{T}(\alpha))_{y \backslash x} = \overline{m}_y$ and thus $|\alpha_{y \backslash x} - \overline{m}_y| \leq \delta$. Then $m_{x \to y} > \overline{m}_y + \epsilon > \overline{m}_y + \delta \geq \alpha_{y \backslash x}$, which implies $S_{xy} > 0$, from Proposition 3. Hence, $m_{x \to y} + m_{y \to x} = w_{xy}$. Define $g_y := \widehat{m}_y - \overline{m}_y \geq 0$. Then $d_{e_0} = \widehat{m}_x + \widehat{m}_y - w_{xy} \leq m_{y \to x} + (g + \overline{m}_y) - w_{xy} < m_{y \to x} + (g + m_{x \to y} - \epsilon) - w_{xy} = -(\epsilon - g)$.

- If $g \leq 4\delta$, then $d_{e_0} < -(\epsilon - g) \leq -(\epsilon - 4\delta)$.
- If $g > 4\delta$, then from the growing procedure we know that $m_{z \to y} \leq \overline{m}_y + 2\delta = \widehat{m}_y - g + 2\delta < \widehat{m}_y - (g - 3\delta)$. Since $g - 3\delta \geq \delta$, then from Proposition 6(h) we have $d_{e_1} < -(g - 3\delta) + \delta = -(g - 4\delta)$. Therefore, $d_{e_0} + d_{e_1} < -(\epsilon - g) - (g - 4\delta) = -(\epsilon - 4\delta)$.

In both cases (a) and (b), we have $w(\mathcal{Q}^+) - w(\mathcal{Q}^-) > 0$. □

References

1. Azar, Y., Birnbaum, B., Elisa Celis, L., Devanur, N.R., Peres, Y.: Convergence of local dynamics to balanced outcomes in exchange networks. FOCS 2009, pp. 293–302 (2009)
2. Azar, Y., Devanur, N.R., Jain, K., Rabani, Y.: Monotonicity in bargaining networks. SODA 2010, pp. 817–826 (2010)
3. Baillon, J., Bruck, R.E.: The rate of asymptotic regularity is O($1/\sqrt{n}$). 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.H., Hajiaghayi, M.T., Immorlica, N., Mahini, H.: The cooperative game theory foundations of network bargaining games. ICALP 2010, pp. 67–78 (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., Borgs, C., Chayes, J.T., Kanoria, Y.ash., Montanari, A.: Bargaining dynamics in exchange networks. J. Econ. Theory **156**, 417–454 (2015)

7. Bayati, M., Shah, D., Sharma, M.: Max-product for maximum weight matching Convergence, correctness, and lp duality. IEEE Trans. Inf. Theory **54**(3), 1241–1251 (2008)
8. Binmore, K.: Game Theory and the Social Contract, vol. 2, MIT Press (1998)
9. Elisa Celis, L., Devanur, N.R., Peres, Y.: Local dynamics in bargaining networks via random-turn games. WINE 2010, pp. 133–144 (2010)
10. Chakraborty, T., Kearns, M.: Bargaining solutions in a social network. WINE 2008, pp. 548–555 (2008)
11. Chakraborty, T., Kearns, M., Khanna, S.: Network bargaining: algorithms and structural results. EC 2009, pp. 159–168 (2009)
12. Cook, K.S., Yamagishi, T.: Power in exchange networks A power-dependence formulation. Soc. Networks **14**, 245–265 (1992)
13. Driessen T.S.H.: Cooperative Games: Solutions and Applications. Kluwer Academic Publishers (1988)
14. 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)
15. Kanoria, Y.: An fptas for bargaining networks with unequal bargaining powers. WINE 2010, pp. 282–293 (2010)
16. Kanoria, Y., Bayati, M., Borgs, C., Chayes, J.T., Montanari, A.: Fast convergence of natural bargaining dynamics in exchange networks. SODA 2011 (2011)
17. Kleinberg, J., Tardos, É.: Balanced outcomes in social exchange networks. STOC' 08, pp. 295–304
18. Koufogiannakis, C., Young, N.E.: Distributed fractional packing and maximum weighted b-matching via tail-recursive duality. DISC'09 (2009)
19. Lotker, Z., Patt-Shamir, B., Pettie, S.: Improved distributed approximate matching. SPAA 2008, pp. 129–136 (2008)
20. Nash, J.: The bargaining problem. Econometrica **18**, 155–162 (1950)
21. 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 '08, pp. 87–92. ACM, New York (2008)
22. Sanghavi, S., Malioutov, D.M., Willsky, A.S.: Linear programming analysis of loopy belief propagation for weighted matching. NIPS' 07 (2007)
23. Schrijver, A.: Combinatorial optimization: polyhedra and efficiency, vol. 24, Springer Science & Business Media (2003)

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Affiliations

T.-H. Hubert Chan[1] · Fei Chen[2] · Li Ning[3]

Fei Chen
chen.f@huawei.com

Li Ning
li.ning@siat.ac.cn

[1] Department of Computer Science, The University of Hong Kong, Pokfulam Road, Pok Fu Lam, Hong Kong

[2] Huawei Noah's Ark Lab, Pok Fu Lam, Hong Kong

[3] Shenzhen Institutes of Advanced Technology, Chinese Academy of Science, Shenzhen, China