

Bargaining dynamics in exchange networks

Mohsen Bayati¹ Christian Borgs² Jennifer Chayes² Yashodhan Kanoria³
Andrea Montanari⁴

November 1, 2011

¹Operations and Information Technology, Graduate School of Business, Stanford University

²Microsoft Research New England

³Department of Electrical Engineering, Stanford University, ykanoria@stanford.edu, (650) 353-0476

⁴Department of Electrical Engineering and Department of Statistics, Stanford University

Abstract

We consider a one-sided assignment market or exchange network with transferable utility and propose a model for the dynamics of bargaining in such a market. Our dynamical model is local, involving iterative updates of ‘offers’ based on estimated best alternative matches, in the spirit of pairwise Nash bargaining. We establish that when a balanced outcome (a generalization of the pairwise Nash bargaining solution to networks) exists, our dynamics converges rapidly to such an outcome. We extend our results to the cases of (i) general agent ‘capacity constraints’, i.e., an agent may be allowed to participate in multiple matches, and (ii) ‘unequal bargaining powers’ (where we also find a surprising change in rate of convergence).

Keywords: Nash bargaining, network, dynamics, convergence, matching, assignment, exchange, balanced outcomes.

JEL classification: C78

1 Introduction

Bargaining has been heavily studied in the economics and sociology literature, e.g., [32, 36, 26, 35]. While the case of bargaining between two agents is fairly well understood [32, 36, 26], less is known about the results of bargaining on networks. We consider exchange networks [13, 27], also called assignment markets [39, 33], where agents occupy the nodes of a network, and edges represent potential partnerships between pairs of agents, which can generate some additional value for these agents. To form a partnership, the pair of agents must reach an agreement on how to split the value of the partnership. Agents are constrained on the number of partnerships they can participate in, for instance, under a matching constraint, each agent can participate in at most one partnership. The fundamental question of interest is: Who will partner with whom, and on what terms? Such a model is relevant to the study of the housing market, the labor market, the assignment of interns to hospitals, the marriage market and so on. An assignment model is suitable for markets with heterogeneous indivisible goods, where trades are constrained by a network structure.

Balanced outcomes¹ generalize the pairwise Nash bargaining solution to this setting. The key issue here is the definition of the threats or best alternatives of participants in a match — these are defined by assuming the incomes of other potential partners to be fixed. In a balanced outcome on a network, each pair plays according to the local Nash bargaining solution thus defined. Balanced outcomes have been found to possess some favorable properties:

Predictive power. The set of balanced outcomes refines the set of stable outcomes (the core), where players have no incentive to deviate from their current partners. For instance, in the case of a two player network, all possible deals are stable, but there is a unique balanced outcome. Balanced outcomes have been found to capture various experimentally observed effects on small networks [13, 27].

Computability. Kleinberg and Tardos [27] provide an efficient centralized algorithm to compute balanced outcomes. They also show that balanced outcomes exist if and only if stable outcomes exist.

Connection to cooperative game theory. The set of balanced outcomes is identical to the core intersection prekernel of the corresponding cooperative game [33, 4].

However, this leaves unanswered the question of whether balanced outcomes can be predictive on large networks, since there was previously no dynamical description of how players can find such an outcome via a bargaining process.

In this work, we consider an assignment market with arbitrary network structure and edge weights that correspond to the value of the potential partnership on that edge. We present a model for the bargaining process in such a network setting, under a capacity constraint on the number of exchanges that each player can establish. We define a model that satisfies two important requirements: locality and convergence.

Locality. It is reasonable to assume that agents have accurate information about their neighbors in the network (their negotiation partners). Specifically, an agent is expected to know the weights of the edges with each of her negotiation partners. Further, for $(ij) \in E$, agent i may estimate the ‘best alternative’ of agent j in the current ‘state’ of negotiations.² On the other hand, it is unlikely that the offer made to an agent during a negotiation depends on arbitrary other agents in the network. We thus require that agent make ‘myopic’ choices on the basis of their neighborhood in the exchange networks. This is consistent with the bulk of the game theory literature on learning

¹Also called *symmetrically pairwise-bargained allocations* in [33] or *Nash bargaining solutions* [27].

²For instance, i may form this estimate based on her conversation with j . (We do not present a game theoretic treatment with fully rational/strategic agents in this work, cf. Section 6.)

[18, 10].

Convergence. The duration of a negotiation is unlikely to depend strongly on the overall network size. For instance, the negotiation on the price of a house, should not depend too much on the size of the town in which it takes place, all other things being equal. Thus a realistic model for negotiation should converge to a fixed point (hence to a set of exchange agreements) in a time roughly independent of the network size.

Our dynamical model is fairly simple. Players compute the current best alternative to each exchange, both for them, and for their partner. On the basis of that, they make a new offer to their partner according to the pairwise Nash bargaining solution. This of course leads to a change in the set of best alternatives at the next time step. We make the assumption that ‘pairing’ occurs at the end, or after several iterative updates, thus suppressing the effect of agents pairing up and leaving.

This dynamics is of course local. Each agent only needs to know the offers she is receiving, as well as the offers that her potential partner is receiving. The technical part of this paper is therefore devoted to the study of the convergence properties of this dynamics.

Remarkably, we find that it converges rapidly, in time roughly independent of the network size. Further its fixed points are not arbitrary, but rather a suitable generalization of Nash bargaining solutions [13, 33] introduced in the context of assignment markets and exchange networks (also referred to as ‘balanced outcomes’ [27] or ‘symmetrically pairwise-bargained allocations’ [33]). Thus, our work provides a dynamical justification of Nash bargaining solutions in assignment markets.

We now present the mathematical definitions of bargaining networks and balanced outcomes.

The network consists of a graph $G = (V, E)$, with positive weights $w_{ij} > 0$ associated to the edges $(i, j) \in E$. A player sits at each node of this network, and two players connected by edge (i, j) can share a profit of w_{ij} dollars if they agree to trade with each other. Each player can trade with at most one of her neighbors (this is called the *1-exchange rule*), so that a set of valid trading pairs forms a matching M in the graph G .

We define an *outcome* or *trade outcome* as a pair $(M, \underline{\gamma})$ where $M \subseteq E$ is a matching of G , and $\underline{\gamma} = \{\gamma_i : i \in V\}$ is the vector of players’ profits. This means, $\gamma_i \geq 0$, and $(i, j) \in M$ implies $\gamma_i + \gamma_j = w_{ij}$, whereas for every unmatched node $i \notin M$ we have $\gamma_i = 0$.

A balanced outcome, or Nash bargaining (NB) solution, is a trade outcome that satisfies the additional requirements of *stability* and *balance*. Denote by ∂i the set of neighbors of node i in G . *Stability.* If player i is trading with j , then she cannot earn more by simply changing her trading partner. Formally $\gamma_i + \gamma_j \geq w_{ij}$ for all $(i, j) \in E \setminus M$.

Balance. If player i is trading with j , then the surplus of i over her best alternative must be equal to the surplus of j over his best alternative. Mathematically,

$$\gamma_i - \max_{k \in \partial i \setminus j} (w_{ik} - \gamma_k)_+ = \gamma_j - \max_{l \in \partial j \setminus i} (w_{jl} - \gamma_l)_+ \quad (1)$$

for all $(i, j) \in M$. Here $(x)_+$ refers to the non-negative part of x , i.e. $(x)_+ \equiv \max(0, x)$.

It turns out that the interplay between the 1-exchange rule and the stability and balance conditions results in highly non-trivial predictions regarding the influence of network structure on individual earnings.

We conclude with some examples of networks and corresponding balanced outcomes (see Figure 1).

The network G_1 has a unique balanced outcome with the nodes a and c forming a partnership with a split of $\gamma_a = 0.5, \gamma_c = 1.5$. Node d remains isolated with $\gamma_d = 0$. The best alternative of




Figure 1: Examples of networks and corresponding balanced outcomes. The network G_1 admits a unique balanced outcome, G_2 admits multiple balanced outcomes, and G_3 admits no balanced outcome. For G_2 one solution is shown inside the square and the other solution is outside.

node c is $(w_{cd} - \gamma_d)_+ = 1$, whereas it is 0 for node a , and the excess of $2 - 1 = 1$ is split equally between a and c , so that each earns a surplus of 0.5 over their outside alternatives.

The network G_2 admits multiple balanced outcomes. Each balanced outcome involves the pairing $M = \{(e, f), (h, i)\}$. The earnings $\gamma_e = 0.5, \gamma_f = 1.5, \gamma_h = 2, \gamma_i = 1$ are balanced, and so is the symmetric counterpart of this earnings vector $\gamma_e = 1.5, \gamma_f = 0.5, \gamma_h = 1, \gamma_i = 2$. In fact, every convex combination of these two earnings vectors is also balanced.

The network G_3 does not admit any stable outcome, and hence does not admit any balanced outcomes. To see this, observe that for any outcome, there is always a pair of agents who can benefit by deviating.

1.1 Related work

Recall the LP relaxation to the maximum weight matching problem

$$\begin{aligned}
 & \text{maximize} && \sum_{(i,j) \in E} w_{ij} x_{ij}, && (2) \\
 & \text{subject to} && \sum_{j \in \partial i} x_{ij} \leq 1 \quad \forall i \in V, \\
 & && x_{ij} \geq 0 \quad \forall (i, j) \in E.
 \end{aligned}$$

The dual problem to (2) is

$$\begin{aligned}
 & \text{minimize} && \sum_{i \in V} y_i, && (3) \\
 & \text{subject to} && y_i + y_j \geq w_{ij} \quad \forall (i, j) \in E, \\
 & && y_i \geq 0 \quad \forall i \in V.
 \end{aligned}$$

Stable outcomes were studied by Sotomayor [40].

Proposition 1. [40] *Stable outcomes exist if and only if the linear programming relaxation (2) of the maximum weight matching problem on G admits an integral optimum. Further, if $(M, \underline{\gamma})$ is a stable solution then M is a maximum weight matching and $\underline{\gamma}$ is an optimum solution to the dual LP (3).*

Following [33, 13], Kleinberg and Tardos [27] first considered balanced outcomes on general exchange networks and proved that: *a network G admits a balanced outcome if and only if it admits a stable outcome.*

The same paper describes a polynomial algorithm for constructing balanced outcomes. This is in turn based on the dynamic programming algorithm of Aspvall and Shiloach [1] for solving systems of linear inequalities. However, [27] left open the question of whether the actual bargaining process converges to balanced outcomes.

Rochford [33], and recent work by Bateni et al [4], relate the assignment market problem to the extensive literature on cooperative game theory. They find that balanced outcomes correspond to the core intersect pre-kernel of the corresponding cooperative game. A consequence of the connection established is that the results of Kleinberg and Tardos [27] are implied by previous work in the economics literature. The existence result follows from Proposition 1, and the fact that if the core of a cooperative game is non-empty then the core intersect prekernel is non-empty. Efficient computability follows from work by Faigle et al [17], who provide a polynomial time algorithm for finding balanced outcomes ³

However, [33, 4] also leave open the twin questions of finding (i) a natural model for bargaining, and (ii) convergence (or not) to NB solutions.

Azar and co-authors [2] studied the question as to whether a balanced outcome can be produced by a local dynamics, and were able to answer it positively⁴. Their results left, however, two outstanding challenges: (I) The algorithm analyzed by these authors first selects a matching M in G using the message passing algorithm studied in [5, 22, 6, 37], corresponding to the pairing of players that trade. In a second phase the algorithm determines the profit of each player. While such an algorithm can be implemented in a distributed way, Azar et al. point out that it is not entirely realistic. Indeed the rules of the dynamics change abruptly after the matching is found. Further, if the pairing is established at the outset, the players lose their bargaining power; (II) The bound on the convergence time proved in [2] is exponential in the network size, and therefore does not provide a solid justification for convergence to NB solutions in large networks.

1.2 Our contribution

The present paper (based partly on our recent conference papers [25, 23]) aims at tackling these challenges. First we introduce a natural dynamics that is interpretable as a realistic negotiation process. We show that the fixed points of the dynamics are in one to one correspondence with NB solutions, and prove that it converges to such solutions. Moreover, we show that the convergence to approximate NB solutions is fast. Furthermore we are able to treat the more general case of nodes with unsymmetrical bargaining powers and generalize the result of [27] on existence of NB solutions to this context. These results are obtained through a new and seemingly general analysis method, that builds on powerful quantitative estimates on mappings in the Banach spaces [3]. For instance, our approach allows us to prove that a simple variant of the edge balancing dynamics of [2] converges in polynomial time (see Section 7).

We consider various modifications to the model and analyze the results. One direction is to allow arbitrary integer ‘capacity constraints’ that capture the maximum number of deals that a particular node is able to simultaneously participate in (the model defined above corresponds to

³In fact, Faigle et al [17] work in the more general setting of cooperative games. The algorithm involves local ‘transfers’, alternating with a *non-local* LP based step after every $O(n^2)$ transfers.

⁴Stearns [43] defined a very similar dynamics and proved convergence for general cooperative games. The dynamics can be interpreted in terms of the present model using the correspondence with cooperative games discussed in [4].

a capacity of one for each node). Such a model would be relevant, for example, in the context of a job market, where a single employer may have more than one opening available. We show that many of our results generalize to this model in Section 5.

Another well motivated modification is to depart from the assumption of symmetry/balance and allow nodes to have different ‘bargaining powers’. Rochford and Crawford [34] mention this modification in passing, with the remark that it “...seems to yield no new insights”. Indeed, we show here (see Section 4) that our asymptotic convergence results generalize to the unsymmetrical case. However, surprisingly, we find that the natural dynamics may now take exponentially long to converge. We find that this can occur even in a two-sided network with the ‘sellers’ having slightly more bargaining power than the ‘buyers’. Thus, a seemingly minor change in the model appears to drastically change the convergence properties of our dynamics. Other algorithms like that of Kleinberg and Tardos [27] and Faigle et al [17] also fail to generalize, suggesting that, in fact, we may lose computability of solutions in allowing asymmetry. However, we show that a suitable modification to the bargaining process yields a fully polynomial time approximation scheme (FPTAS) for the unequal bargaining powers. The caveat is that this algorithm, though local, is not a good model for bargaining because it fixes the matching at the outset (cf. comment (I) above).

Our dynamics and its analysis have similarities with a series of papers on using max-product belief propagation for the weighted matching problems [5, 22, 6, 37]. We discuss that connection and extensions of our results to those settings in one of our conference papers [25, Appendix F]. We obtain a class of new message passing algorithms to compute the maximum weight matching, with belief propagation and our dynamics being special cases.

Related work in sociology

Besides economists, sociologists have been interested in such markets, called *exchange networks* in that literature. The key question addressed by *network exchange theory* is that of how network structure influences the power balance between agents. Numerous predictive frameworks have been suggested in this context including generalized Nash bargaining solutions [13]. Moreover, controlled experiments [45, 29, 41] have been carried out by sociologists. The typical experimental set-up studies exactly the model of assignment markets proposed by economists [39, 33]. It is often the case that players are provided information only about their immediate neighbors. Typically, a number of ‘rounds’ of negotiations are run, with no change in the network, so as to allow the system to reach an ‘equilibrium’. Further, players are usually not provided much information beyond who their immediate neighbors are, and the value of the corresponding possible deals.

In addition to balanced outcomes [13], other frameworks have been suggested to predict/explain the outcomes of these experiments [12, 9, 41].

1.3 Outline of the paper

Our dynamical model of bargaining in a network is described in Section 2. We state our main results characterizing fixed points and convergence of the dynamics in Section 3.

We present two extensions of our model. In Section 4 we investigate the unsymmetrical case with nodes having different bargaining powers. We find that the (generalized) dynamics may take exponentially long to converge in this case, but we provide a modified local algorithm that provides an FPTAS. We consider the case of general capacity constraints in Section 5, and show that our main results generalize to this case.

In Section 7, we prove Theorems 2 and 3 on convergence of our dynamics. We characterize fixed points in Section 8 with a proof of Theorem 1 (the proof of Theorem 4 is deferred to Appendix D). Section 9 shows polynomial time convergence on bipartite graphs (proof of Theorem 5).

We present a discussion of our results in Section 6.

Appendix A contains a discussion on variations of the natural dynamics including time and node varying damping factors and asynchronous updates.

2 Dynamical model

Consider a bargaining network $G = (V, E)$, where the vertices represent agents, and the edges represent potential partnerships between them. There is a positive weight $w_{ij} > 0$ on each edge $(i, j) \in E$, representing the fact that players connected by edge (i, j) can share a profit of w_{ij} dollars if they agree to trade with each other. Each player can trade with at most one of her neighbors (this is called the *1-exchange rule*), so that a set of valid trading pairs forms a matching M in the graph G . We define a *trade outcome* as in Section 1, in accordance with the above constraints.

We expect natural dynamical description of a bargaining network to have the following properties: It should be *local*, i.e. involve limited information exchange along edges and processing at nodes; It should be *time invariant*, i.e. the players' behavior should be the same/similar on identical local information at different times; It should be *interpretable*, i.e. the information exchanged along the edges should have a meaning for the players involved, and should be consistent with reasonable behavior for players.

In the model we propose, at each time t , each player sends a message to each of her neighbors. The message has the meaning of 'best current alternative'. We denote the message from player i to player j by $\alpha_{i \setminus j}^t$. Player i is telling player j that she (player i) currently estimates earnings of $\alpha_{i \setminus j}^t$ elsewhere, if she chooses not to trade with j .

The vector of all such messages is denoted by $\underline{\alpha}^t \in \mathbb{R}_+^{2|E|}$. Each agent i makes an 'offer' to each of her neighbors, based on her own 'best alternative' and that of her neighbor. The offer from node i to j is denoted by $m_{i \rightarrow j}^t$ and is computed according to

$$m_{i \rightarrow j}^t = (w_{ij} - \alpha_{i \setminus j}^t)_+ - \frac{1}{2}(w_{ij} - \alpha_{i \setminus j}^t - \alpha_{j \setminus i}^t)_+. \quad (4)$$

It is easy to deduce that this definition corresponds to the following policy: (i) An offer is always non-negative, and a positive offer is never larger than $w_{ij} - \alpha_{i \setminus j}^t$ (no player is interested in earning less than her current best alternative); (ii) Subject to the above constraints, the surplus $(w_{ij} - \alpha_{i \setminus j}^t - \alpha_{j \setminus i}^t)$ (if non-negative) is shared equally. We denote by $\underline{m}^t \in \mathbb{R}_+^{2|E|}$ the vector of offers.

Notice that \underline{m}^t is just a deterministic function of $\underline{\alpha}^t$. In the rest of the paper we shall describe the network status uniquely through the latter vector, and use $\underline{m}|_{\underline{\alpha}^t}$ to denote \underline{m}^t defined by (4) when required so as to avoid ambiguity.

Each node can estimate its potential earning based on the network status, using

$$\gamma_i^t \equiv \max_{k \in \partial i} m_{k \rightarrow i}^t, \quad (5)$$

the corresponding vector being denoted by $\underline{\gamma}^t \in \mathbb{R}_+^{|V|}$. Notice that $\underline{\gamma}^t$ is also a function of $\underline{\alpha}^t$.

Messages are updated synchronously through the network, according to the rule

$$\alpha_{i \setminus j}^{t+1} = (1 - \kappa) \alpha_{i \setminus j}^t + \kappa \max_{k \in \partial i \setminus j} m_{k \rightarrow i}^t. \quad (6)$$

Here $\kappa \in (0, 1]$ is a ‘damping’ factor: $(1 - \kappa)$ can be thought of as the inertia on the part of the nodes to update their current estimates (represented by outgoing messages). The use of $\kappa < 1$ eliminates pathological behaviors related to synchronous updates. In particular, we observe oscillations on even-length cycles in the undamped synchronous version. We mention here that in Appendix A we present extensions of our results to various update schemes (e.g., asynchronous updates, time-varying damping factor).

Remark 2. *An update under the natural dynamics requires agent i to perform $O(|\partial i|)$ arithmetic operations, and $O(|E|)$ operations in total.*

Let $W_{\max} \equiv \max_{(ij) \in E} w_{ij}$. Often in the paper we take $W_{\max} = 1$, since this can always be achieved by rescaling the problem, which is the same as changing units. It is easy to see that $\underline{\alpha}^t \in [0, W_{\max}]^{2|E|}$, $\underline{m}^t \in [0, W_{\max}]^{2|E|}$ and $\underline{\gamma}^t \in [0, W_{\max}]^{|V|}$ at all times (unless the initial condition violates this bounds). Thus we call $\underline{\alpha}$ a ‘valid’ message vector if $\underline{\alpha} \in [0, W_{\max}]^{2|E|}$.

2.1 An Example

We consider a simple graph G with $V = \{A, B, C, D\}$, $E = \{(A, B), (B, C), (C, D)\}$, $w_{AB} = 8$, $w_{BC} = 6$ and $w_{CD} = 2$. The unique maximum weight matching on this graph is $M = \{(A, B), (C, D)\}$. By Proposition 1, stable outcomes correspond to matching M and can be parameterized as

$$\underline{\gamma} = (8 - \gamma_B, \gamma_B, \gamma_C, 2 - \gamma_C)$$

where (γ_B, γ_C) are constrained as

$$\begin{aligned} \gamma_B &\in [0, 8] \\ \gamma_C &\in [0, 2] \\ \gamma_B + \gamma_C &\geq 6 \end{aligned}$$

For instance, the set of stable outcomes (all on matching M) includes $(0, 8, 2, 0)$, $(4, 4, 2, 0)$, $(3, 5, 1, 1)$ and so on. Now suppose we impose the balance condition Eq. (1) in addition, i.e., we look for balanced outcomes. Using the algorithm of Kleinberg and Tardos [27], we find that the network admits a unique balanced outcome $\underline{\gamma} = (1.5, 6.5, 1, 1)$.

Now we consider the evolution of the natural dynamics proposed above on the graph G . We arbitrarily choose to study the initialization $\underline{\alpha}^0 = \underline{0}$, i.e., each node initially estimates its best alternatives to be 0 with respect to each neighbor. We set $\kappa = 1$ for simplicity⁵. The evolution of the estimates and offers under the dynamics is shown in Figure 2.1. We now comment on a few noteworthy features demonstrated by this example. In the first step, nodes A and B receive their best offers from each other, node C receives its best offer from B and node D receives its best offer from C . Thus, we might expect nodes A and B to be considering the formation of a partnership already (though the terms are not yet clear), but this is not the case for C and D . After one iteration, at $t = 1$, both pairs (A, B) and (C, D) receive their best offers from each other. In fact, this property remains true at all future times (the case $t = 2$ is shown). However, the vectors $\underline{\alpha}$ and \underline{m} continue to evolve from one iteration to the next. At iteration $t = 6$, a fixed point is reached, i.e., $\underline{\alpha}$ and \underline{m} remain unchanged for $t \geq 6$. Moreover, we notice that the fixed point captures the

⁵Our results assume $\kappa < 1$ to avoid oscillatory behavior. However, it turns out that on graphs with no even cycles, for instance the graph G under consideration, oscillations do not occur. We choose to consider $\kappa = 1$ for simplicity of presentation.




Figure 2: Progress of the natural dynamics on a graph with four nodes and three edges. We (arbitrarily) choose the initialization $\underline{\alpha}^0 = \underline{0}$. A fixed point is reached at $t = 6$.

unique balanced outcome on this graph, with the matching M and the splits ($\gamma_A = 1.5, \gamma_B = 6.5$) and ($\gamma_C = 1, \gamma_D = 1$) emerging from the fixed point \underline{m}^* .

We remark here that convergence to a fixed point in finite number of iterations is not a general phenomenon. This occurs as a consequence of the simple example considered and the choice $\kappa = 1$. However, as we prove below, we always obtain rapid convergence of the dynamics, and fixed points always correspond to balanced outcomes, on *any graph* possessing balanced outcomes, and for *any initialization*.

3 Main results: Fixed point properties and convergence

Our first result is that fixed points of the update equations (4), (6) (hereafter referred to as ‘natural dynamics’) are indeed in correspondence with Nash bargaining solutions when such solutions exist. *Note that the fixed points are independent of the damping factor κ .* The correspondence with NB solutions includes pairing between nodes, according to the following notion of induced matching.

Definition 3. *We say that a state $(\underline{\alpha}, \underline{m}, \underline{\gamma})$ (or just $\underline{\alpha}$) induces a matching M if the following happens. For each node $i \in V$ receiving non-zero offers ($m_{\rightarrow i} > 0$), i is matched under M and gets its unique best offer from node j such that $(i, j) \in M$. Further, if $\gamma_i = 0$ then i is not matched in M . In other words, pairs in M receive unique best offers that are positive from their respective matched neighbors whereas unmatched nodes receive no non-zero offers.*

Consider the LP relaxation to the maximum weight matching problem (2). A feasible point \underline{x} for LP (2) is called *half-integral* if for all $e \in E$, $x_e \in \{0, 1, \frac{1}{2}\}$. It is well known that problem (2) always has an optimum \underline{x}^* that is half-integral [38]. An LP with a fully integer \underline{x}^* ($x_e^* \in \{0, 1\}$) is called *tight*.

Theorem 1. *Let G be an instance admitting one or more Nash bargaining solutions, i.e. the LP (2) admits an integral optimum.*

(a) *Unique LP optimum (generic case): Suppose the optimum is unique corresponding to matching M^* . Let $(\underline{\alpha}, \underline{m}, \underline{\gamma})$ be a fixed point of the natural dynamics. Then $\underline{\alpha}$ induces matching M^* and $(M^*, \underline{\gamma})$ is a Nash bargaining solution. Conversely, every Nash bargaining solution $(M, \underline{\gamma}_{\text{NB}})$ has $M = M^*$ and corresponds to a unique fixed point of the natural dynamics with $\underline{\gamma} = \underline{\gamma}_{\text{NB}}$.*

(b) *Let $(\underline{\alpha}, \underline{m}, \underline{\gamma})$ be a fixed point of the natural dynamics. Then $(M^*, \underline{\gamma})$ is a Nash bargaining solution for any integral maximum weight matching M^* . Conversely, if $(M, \underline{\gamma}_{\text{NB}})$ is a Nash bargaining solution, M is a maximum weight matching and there is a unique fixed point of the natural dynamics with $\underline{\gamma} = \underline{\gamma}_{\text{NB}}$.*

We prove Theorem 1 in Section 8. Theorem 12 in Appendix C extends this characterization of fixed points of the natural dynamics to cases where Nash bargaining solutions do not exist.

Remark 4. *The condition that a tight LP (2) has a unique optimum is generic (see Appendix C, Remark 18). Hence, fixed points induce a matching for almost all instances (cf. Theorem 1(a)). Further, in the non-unique optimum case, we cannot expect an induced matching, since there is always some node with two equally good alternatives.*

The existence of a fixed point of the natural dynamics is immediate from Brouwer’s fixed point theorem. Our next result says that the natural dynamics always converges to a fixed point. The proof is in Section 7.

Theorem 2. *The natural dynamics has at least one fixed point. Moreover, for any initial condition with $\underline{\alpha}^0 \in [0, W]^{2|E|}$, $\underline{\alpha}^t$ converges to a fixed point.*

Note that Theorem 2 does not require any condition on LP (2). It also does not require uniqueness of the fixed point.

With Theorems 1 and 2, we know that in the limit of a large number of iterations, the natural dynamics yields a Nash bargaining solution. However, this still leaves unanswered the question of the rate of convergence of the natural dynamics. Our next theorem addresses this question, establishing fast convergence to an approximate fixed point.

However, before stating the theorem we define the notion of approximate fixed point.

Definition 5. *We say that $\underline{\alpha}$ is an ϵ -fixed point, or ϵ -FP in short, if, for all $(i, j) \in E$ we have*

$$\left| \alpha_{i \setminus j} - \max_{k \in \partial i \setminus j} m_{k \rightarrow i} \right| \leq \epsilon, \quad (7)$$

and similarly for $\alpha_{j \setminus i}$. Here, \underline{m} is obtained from $\underline{\alpha}$ through Eq. (4) (i.e., $\underline{m} = \underline{m} |_{\underline{\alpha}}$).

Note that ϵ -fixed points are also defined independently of the damping κ .

Theorem 3. *Let $G = (V, E)$ be an instance with weights $(w_e, e \in E) \in [0, 1]^{|E|}$. Take any initial condition $\underline{\alpha}^0 \in [0, 1]^{2|E|}$. Take any $\epsilon > 0$. Define*

$$T^*(\epsilon) = \frac{1}{\pi \kappa (1 - \kappa) \epsilon^2}. \quad (8)$$

Then for all $t \geq T^*(\epsilon)$, $\underline{\alpha}^t$ is an ϵ -fixed point. (Here $\pi = 3.14159\dots$)

Thus, if we wait until time t , we are guaranteed to obtain an $(1/\sqrt{\pi \kappa (1 - \kappa)t})$ -FP. Theorem 3 is proved in Section 7.

Remark 6. *For any $\epsilon > 0$, it is possible to construct an example such that it takes $\Omega(1/\epsilon)$ iterations to reach an ϵ -fixed point. This lower bound can be improved to $\Omega(1/\epsilon^2)$ in the unequal bargaining powers case (cf. Section 4). However, in our constructions, the size of the example graph grows with decreasing ϵ in each case.*

We are left with the problem of relating approximate fixed points to approximate Nash bargaining solutions. We use the following definition of ϵ -Nash bargaining solution, that is analogous to the standard definition of ϵ -Nash equilibrium (e.g., see [14]).

Definition 7. *We say that $(M, \underline{\gamma})$ is an ϵ -Nash bargaining solution if it is a valid trade outcome that is stable and satisfies ϵ -balance. ϵ -balance means that for every $(i, j) \in M$ we have*

$$\left| [\gamma_i - \max_{k \in \partial i \setminus j} (w_{ik} - \gamma_k)_+] - [\gamma_j - \max_{l \in \partial j \setminus i} (w_{jl} - \gamma_l)_+] \right| \leq \epsilon. \quad (9)$$

A subtle issue needs to be addressed. For an approximate fixed point to yield an approximate Nash bargaining solution, a suitable pairing between nodes is needed. Note that our dynamics does not force a pairing between the nodes. Instead, a pairing should emerge quickly from the dynamics. In other words, nodes on the graph should be able to identify their trading partners from the messages being exchanged. As before, we use the notion of an induced matching (see Definition 3).

Definition 8. Consider LP (2). Let \mathcal{H} be the set of half integral points in the primal polytope. Let $\underline{x}^* \in \mathcal{H}$ be an optimum. Then the LP gap g is defined as $g = \min_{\underline{x} \in \mathcal{H} \setminus \{\underline{x}^*\}} \sum_{e \in E} w_e x_e^* - \sum_{e \in E} w_e x_e$.

Theorem 4. Let G be an instance for which the LP (2) admits a unique optimum, and this is integral, corresponding to matching M^* . Let the gap be $g > 0$. Let $\underline{\alpha}$ be an ϵ -fixed point of the natural dynamics, for some $\epsilon < g/(6n^2)$. Let $\underline{\gamma}$ be the corresponding earnings estimates. Then $\underline{\alpha}$ induces the matching M^* and $(\underline{\gamma}, M^*)$ is an (6ϵ) -Nash bargaining solution. Conversely, every ϵ -Nash bargaining solution $(M, \underline{\gamma}_{\text{NB}})$ has $M = M^*$ for any $\epsilon > 0$.

Note that $g > 0$ is equivalent to the unique optimum condition (cf. Remarks 2, 5). The proof of this theorem requires generalization of the analysis used to prove Theorem 1 to the case of approximate fixed points. Since its proof is similar to the proof of Theorem 1, we defer it to Appendix D. We stress, however, that Theorem 4 is not, in any sense, an obvious strengthening of Theorem 1. In fact, this is a delicate property of approximate fixed points that holds *only in the case of balanced outcomes*. This characterization breaks down in the face of a seemingly benign generalization to unequal bargaining powers (cf. Section 4 and [23, Section 4]).

Theorem 4 holds for all graphs, and is, in a sense, the best result we can hope for. To see this, consider the following immediate corollary of Theorems 3 and 4.

Corollary 9. Let $G = (V, E)$ be an instance with weights $(w_e, e \in E) \in [0, 1]^{|E|}$. Suppose LP (2) admits a unique optimum, and this is integral, corresponding to matching M^* . Let the gap be $g > 0$. Then for any $\underline{\alpha}^0 \in [0, 1]^{2|E|}$, there exists $T^* = O(n^4/g^2)$ such that for any $t \geq T^*$, $\underline{\alpha}^t$ induces the matching M^* and $(\underline{\gamma}^t, M^*)$ is an $(6/\sqrt{\pi\kappa(1-\kappa)t})$ -NB solution.

Proof. Choose T^* as $T^*(g/(10n^2))$ as defined in (8). Clearly, $T^* = O(n^4/g^2)$. From Theorem 3, $\underline{\alpha}^t$ is an $\epsilon(t)$ -FP for $\epsilon(t) = 1/\sqrt{\pi\kappa(1-\kappa)t}$. Moreover, for all $t \geq T^*$, $\epsilon(t) \leq g/(10n^2)$. Hence, by Theorem 4, $\underline{\alpha}^t$ induces the matching M^* and $(\underline{\gamma}^t, M^*)$ is a $(6\epsilon(t))$ -NB solution for all $t \geq T^*$. \square

Corollary 9 implies that for any $\epsilon > 0$, the natural dynamics finds an ϵ -NB solution in time $O(\max(n^4/g^2, 1/\epsilon^2))$.

This result is the essentially the strongest bound we can hope for in the following sense. First, note that we need to find M^* (see converse in Theorem 4) *and* balance the allocations. Max product belief propagation, a standard local algorithm for computing the maximum weight matching, requires $O(n/g)$ iterations to converge, and *this bound is tight* [6]. Similar results hold for the Auction algorithm [7] which also locally computes M^* . Moreover, max product BP and the natural dynamics are intimately related (see [25]), with the exception that *max product is designed to find M^* , but this is not true for the natural dynamics*. Corollary 9 shows that natural dynamics only requires a time that is polynomial in the same parameters n and $1/g$ to find M^* , while it simultaneously takes rapid care of balancing the outcome.

3.1 Example: Polynomial convergence to ϵ -NB solution on bipartite graphs.

The next result further shows a concrete setting in which Corollary 9 leads to a strong guarantee on quickly reaching an approximate NB solution.

Theorem 5. Let $G = (V, E)$ be a bipartite graph with weights $(w_e, e \in E) \in [0, 1]^{|E|}$. Take any $\xi \in (0, 1), \eta \in (0, 1)$. Construct a perturbed problem instance with weights $\bar{w}_e = w_e + \eta U_e$, where U_e are independent identically distributed random variables uniform in $[0, 1]$. Then there exists

$C = C(\kappa) < \infty$, such that for

$$T^* = C \left(\frac{n^2 |E|}{\eta \xi} \right)^2, \quad (10)$$

the following happens for all $t \geq T^*$ with probability at least $1 - \xi$. State $\underline{\alpha}^t$ induces a matching M that is independent of t . Further, $(\underline{\gamma}^t, M)$ is a $\epsilon(t)$ -NB solution for the perturbed problem, with $\epsilon(t) = 12/\sqrt{\pi\kappa(1-\kappa)t}$.

ξ represents our target in the probability that a pairing does not emerge, while η represents the size of perturbation of the problem instance.

Theorem 5 implies that for any fixed η and ξ , and any $\epsilon > 0$, we find an ϵ -NB solution in time $\tau(\epsilon) = K \max(n^4 |E|^2, 1/\epsilon^2)$ with probability at least $1 - \xi$, where $K = K(\eta, \xi, \kappa) < \infty$. Theorem 5 is proved in Section 9.

3.2 Other results

A different analysis allows us to prove *exponentially fast convergence* to a unique Nash bargaining solution. We describe this briefly in Section 7.1, referring to an unpublished manuscript [24] for the proof, in the interest of space.

We investigate the case of *nodes with unsymmetrical bargaining powers* in Section 4. We show that generalizations of the Theorems 1, 2 and 3 hold for a suitably modified dynamics. Somewhat surprisingly, we find that the modified dynamics may take exponentially long to converge (so Theorem 3 does *not* generalize). However, we find a different local procedure that efficiently finds approximate solutions.

In Section 5 we consider the case where agents have arbitrary integer capacity constraints on the number of partnerships they can participate in, instead of the one-matching constraint. We generalize our dynamics and the notion of balanced outcomes to this case. We show that Theorems 1, 2 and 3 generalize. As a corollary, we establish the existence of balanced outcomes whenever stable outcomes exist (Corollary 14) in this general setting⁶.

Appendix A presents extensions of our convergence results to cases where the damping factor varies in time or from node to node, and when updates are asynchronous. This shows that our insights are robust to variations in the natural of damping and the timing of iterative updates.

4 Unequal bargaining powers

It is reasonable to expect that not all edge surpluses on matching edges are divided equally in an exchange network setting. Some nodes are likely to have more ‘bargaining power’ than others. This bargaining power can arise, for example, from ‘patience’; a patient agent is expected to get more than half the surplus when trading with an impatient partner. This phenomenon is well known in the Rubinstein game [36] where nodes alternately make offers to each other until an offer is accepted – the node with a smaller discount factor earns more in the subgame perfect Nash equilibrium. Moreover, a recent experimental study of bargaining in exchange networks [11] found that patience correlated positively with earnings.

A reasonable approach to model this effect would be to assign a positive ‘bargaining power’ to each node, and postulate that if a pair of nodes trade with each other, then the edge surplus is

⁶The caveat here is that Corollary 14 does not say anything about the corner case of non-unique maximum weight **b**-matching.

divided in the ratio of their bargaining powers. We choose instead, a more general setting where on each edge (ij) there is an expected surplus split fraction quantified by $r_{ij} \in (0, 1)$. Namely, r_{ij} is the fraction of surplus that goes to i if i and j trade with each other, and similarly for r_{ji} . Note that we have $r_{ij} + r_{ji} = 1$. We call a weighted graph G along with the postulated split fraction vector \underline{r} an *unequal division (UD)* instance.

The balance condition is replaced by *correct division* condition

$$\begin{aligned} [r_{ij}]^{-1} [\gamma_i - \max_{k \in \partial i \setminus j} (w_{ik} - \gamma_k)_+] \\ \stackrel{\text{UD}}{=} [r_{ji}]^{-1} [\gamma_j - \max_{l \in \partial j \setminus i} (w_{jl} - \gamma_l)_+], \end{aligned} \quad (11)$$

on matched edges (ij) . We retain the stability condition. We call trade outcomes satisfying (11) and stability *unequal division (UD) solutions*. A natural modification to our dynamics in this situation consists of the following redefinition of offers.

$$m_{i \rightarrow j}^t \stackrel{\text{UD}}{=} (w_{ij} - \alpha_{i \setminus j}^t)_+ - r_{ij} (w_{ij} - \alpha_{i \setminus j}^t - \alpha_{j \setminus i}^t)_+. \quad (12)$$

We call the dynamics resulting from (12) and the update rule (6) the *unsymmetrical natural dynamics*. One can check that \mathbb{T} defined in (22) is nonexpansive for offers defined as in (12). *It follows that Theorems 2 and 3 hold for the UD-natural dynamics with damping.* (We retain Definition 5 of an ϵ -FP). Further, Theorem 1 can also be extended to this case. The proof involves exactly the same steps as for the natural dynamics (cf. Section 8). Properties 1-6 in the direct part all hold (proofs nearly verbatim) and an identical construction works for the converse.

Theorem 6. *Let G be an instance for which the LP (2) admits an integral optimum. Let $(\underline{\alpha}, \underline{m}, \underline{\gamma})$ be a fixed point of the UD-natural dynamics. Then $(M^*, \underline{\gamma})$ is a UD solution for any maximum weight matching M^* . Conversely, for any UD solution $(M, \underline{\gamma}_{\text{UD}})$, matching M is a maximum weight matching and there is a unique fixed point of the UD-natural dynamics with $\underline{\gamma} = \underline{\gamma}_{\text{UD}}$.*

Further, if the LP (2) has a unique integral optimum, corresponding to matching M^ , then any fixed point $\underline{\alpha}$ induces matching M^* .*

We note that the following generalization of the result on existence of Nash bargaining solutions [27] follows from Theorem 6 and the existence of fixed points.

Lemma 1. *UD solutions exist if and only if a stable outcome exists (i.e. LP (2) has an integral optimum.)*

Proof. The direct part of Theorem 6, along with the existence of fixed points of the UD natural dynamics (from Brouwer's fixed point theorem, also first part of Theorem 2 for UD) shows that UD solutions exist if LP (2) has an integral optimum. The converse is trivial since if LP (2) has no integral optimum, then there are no stable solutions (see Proposition 1) and hence no UD solutions. \square

4.1 Exponential convergence time in the UD case

It is possible to derive a characterization similar to Theorem 4 also for the UD case. However, the bound on ϵ needed to ensure that the right pairing emerges in an ϵ -FP turns out to be exponentially small in n . As such, we are only able to show that a pairing emerges in time $2^{O(n)}/g^2$. In fact, as the example below shows, it *does* take exponentially long for a pairing to emerge in worst case.

Let $n = |V|$. In Appendix G.1, we construct a sequence of instances $(I_n, n \geq 16)$, such that for each instance in the sequence the following holds.

- (a) The instance admits a UD solution, and the gap $g \geq 1$.
- (b) There is a message vector $\underline{\alpha}$ that is an ϵ -FP for $\epsilon = 2^{-cn}$ such that $\underline{\alpha}$ does *not* induce any matching. In fact, for any matching in our example, there exists j such that j gets an offer from outside the matching that exceeds the offer of its partner under the matching (if any) by at least 1. Thus, $\underline{\alpha}$ is ‘far’ from inducing a matching.

Further, split fractions are bounded within $[r, 1 - r]$ for arbitrary desired $r \in (0, 1/2)$ (c depends on r). Also, the weights are uniformly bounded by a constant $W(r)$.

Now, if we start at an ϵ -FP, the successive iterates of our dynamics are all ϵ -FPs (this follows from the nonexpansivity of the update operator, cf. Section 8). Hence, no offer can change by more than ϵ in each iteration. Thus, in our constructed instances I_n , it requires at least $2^{\Omega(n)}$ iterations starting with $\underline{\alpha}^0 = \underline{\alpha}$, before $\underline{\alpha}^t$ induced a matching. Thus, our construction with the above properties implies that the unsymmetrical natural dynamics can take exponentially long to induce a matching, even on well behaved instances ($g = \Omega(1)$). Moreover, as discussed in Appendix G.1, our construction corresponds to a plausible two-sided market, so this is not to be dismissed as a unrealistic special case that can be ignored.

4.2 A fully polynomial time approximation scheme

We show that though the natural dynamics may take exponentially long to converge, there is a polynomial time iterative ‘re-balancing’ algorithm that enables us to compute an approximate UD solution.

First we define an approximate version of correct division, asking that Eq. (11) be satisfied to within an additive ϵ , for all matched edges. For each edge $(ij) \in M$, we define the ‘edge surplus’ as the excess of w_{ij} over the sum of best alternatives, i.e.,

$$\text{Surp}_{ij}(\underline{\gamma}) = w_{ij} - \max_{k \in \partial i \setminus j} (w_{ik} - \gamma_k)_+ - \max_{l \in \partial j \setminus i} (w_{jl} - \gamma_l)_+. \quad (13)$$

Definition 10 (ϵ -Correct division). *An outcome $(\underline{\gamma}, M)$ is said to satisfy ϵ -correct division if, for all $(ij) \in M$,*

$$|\gamma_i - \max_{k \in \partial i \setminus j} (w_{ik} - \gamma_k)_+ - r_{ij} \text{Surp}_{ij}(\underline{\gamma})| \leq \epsilon, \quad (14)$$

where $\text{Surp}_{ij}(\cdot)$ is defined by Eq. (13).

We define approximate UD solutions as follows:

Definition 11 (ϵ -UD solution). *An outcome $(\underline{\gamma}, M)$ is an ϵ -UD solution for $\epsilon \geq 0$ if it is stable and it satisfies ϵ -correct division (cf. Definition 10).*

It follows from Lemma 1 that ϵ -UD solutions exist if and only if the LP (2) admits an integral optimum, which is the same as the requirement for existence of UD solutions. We prove the following:

Theorem 7. *There is an algorithm that is polynomial in the input and $1/\epsilon$, such that for any problem instance with weights uniformly bounded by 1, i.e., $(w_e, e \in E) \in (0, 1]^{|E|}$:*

- *If the instance admits a UD solution, the algorithm finds an ϵ -UD solution.*
- *If the instance does not admit a UD solution, the algorithm returns the message UNSTABLE.*

Our approach to finding an ϵ -UD solution consists of two main steps:

1. Find a maximum weight matching M^* and a dual optimum $\underline{\gamma}$ (solution to the dual LP (3)). Thus, form a stable outcome $(\underline{\gamma}, M^*)$. Else certify that the instance has no UD solution.
2. Iteratively update the allocation $\underline{\gamma}$ without changing the matching. Updates are local, and are designed to converge fast to an allocation satisfying the ϵ -correct division solution *while maintaining stability*. Thus, we arrive at an ϵ -UD solution.

As mentioned earlier, this is similar to the approach of [2]. The crucial differences (enabling our results) are: (i) we stay within the space of stable outcomes, (ii) our analysis of convergence.

The algorithm leading to Theorem 7 and the corresponding analysis are presented in Appendix 4.2. We remark that our algorithm can easily be made *local*, without sacrificing efficiency (see [23] for details).

5 General capacity constraints

In several situations, agents may be less restricted: Instead of an agent being allowed to enter at most one agreement, for each agent i , there may be an integer *capacity* constraint b_i specifying the maximum number of partnerships that i can enter into. For instance, in a labor market for full time jobs, an employer j may have 4 openings for a particular role ($b_j = 4$), another employer may have 6 openings for a different role, and so on, but the job seekers can each accept at most one job. In this section, we describe a generalization of our dynamical model to the case of general capacity constraints, in an attempt to model behavior in such settings. We find that most of our results from the one-matching case, cf. Section 3, generalize.

5.1 Preliminaries

Now a bargaining network is specified by an undirected graph $G = (V, E)$ with positive weights on the edges $(w_{ij})_{(ij) \in E}$, and integer capacity constraints associated to the nodes $(b_i)_{i \in V}$. We generalize the notion of ‘matching’ to sets of edges that satisfy the given capacity constraints: Given capacity constraints $\mathbf{b} = (b_i)$, we call a set of edges $M \subseteq E$ a *\mathbf{b} -matching* if the degree $d_i(M)$ of i in the graph (V, M) is at most b_i , for every $i \in V$. We say that i is *saturated* under M if $d_i(M) = b_i$.

We assume that there are no double edges between nodes⁷. Thus, an agent can use at most one unit of capacity with any one of her neighbors in the model we consider.

A *trade outcome* is now a pair (M, Γ) , where M is a \mathbf{b} -matching and $\Gamma \in [0, 1]^{2|E|}$ is a splitting of profits $\Gamma = (\gamma_{i \rightarrow j}, \gamma_{j \rightarrow i})_{(ij) \in E}$, with $\gamma_{i \rightarrow j} = 0$ if $(ij) \notin M$, and $\gamma_{i \rightarrow j} + \gamma_{j \rightarrow i} = w_{ij}$ if $(ij) \in M$.

Define $\gamma_i = \min_{j: (ij) \in M} \gamma_{j \rightarrow i}$ if i is saturated (i.e. $d_i(M) = b_i$) and $\gamma_i = 0$ if i is not saturated. Note that this definition is equivalent to $\gamma_i = (b_i^{\text{th-max}})_{j \in \partial_i} \gamma_{j \rightarrow i}$. Here $(b^{\text{th-max}}) : \mathbb{R}_+^* \rightarrow \mathbb{R}_+$ denotes the b -th largest of a set of non-negative reals, being defined as 0 if there are less than b numbers in the set. It is easy to see that our definition of γ_i here is consistent with the definition for the one-exchange case. (But Γ is not consistent with $\underline{\gamma}$, which is why we use different notation.)

We say that a trading outcome is *stable* if $\gamma_i + \gamma_j \geq w_{ij}$ for all $ij \notin M$. This definition is natural; a selfish agent would want to switch partners if and only if he can gain more utility elsewhere.

⁷This assumption was not needed in the one-exchange case since, in that case, utility maximizing agents i and j will automatically discard all but the heaviest edge between them. This is no longer true in the case of general capacity constraints.

An outcome (M, γ) is said to be *balanced* if

$$\gamma_{j \rightarrow i} - (b_i^{\text{th-max}})_{k \in \partial i \setminus j} (w_{ik} - \gamma_k)_+ = \gamma_{i \rightarrow j} - (b_j^{\text{th-max}})_{l \in \partial j \setminus i} (w_{jl} - \gamma_l)_+ \quad (15)$$

for all $(ij) \in M$.

Note that the definitions of stability and balance generalize those for the one-exchange case.

An outcome (M, Γ) is a Nash bargaining solution if it is both stable and balanced.

Consider the problem of finding the maximum weight (not necessarily perfect) \mathbf{b} -matching on a weighted graph $G = (V, E)$. The LP-relaxation of this problem and its dual are given by

$$\begin{array}{l|l} \max & \sum_{(ij) \in E} x_{ij} w_{ij} \\ \text{subject to} & \sum_{j \in N(i)} x_{ij} \leq b_i \quad \forall i \\ & 0 \leq x_{ij} \leq 1 \quad \forall (ij) \in E \\ & \text{Primal LP} \end{array} \quad \left| \quad \begin{array}{l|l} \min & \sum_{i \in V} b_i y_i + \sum_{(ij) \in E} y_{ij} \\ \text{subject to} & y_{ij} + y_i + y_j - w_{ij} \geq 0 \quad \forall (ij) \in E \\ & y_{ij} \geq 0 \quad \forall (ij) \in E \\ & y_i \geq 0 \quad \forall i \in V \\ & \text{Dual LP.} \end{array} \right. \quad (16)$$

Complementary slackness says that a pair of feasible solutions is optimal if and only if

- For all $ij \in E$; $x_{ij}^* (-w_{ij} + y_{ij}^* + y_i^* + y_j^*) = 0$.
- For all $ij \in E$; $(x_{ij}^* - 1) y_{ij}^* = 0$.
- For all $i \in V$; $(\sum_{j \in N(i)} x_{ij}^* - b_i) y_i^* = 0$.

Lemma 2. *Consider a network $G = (V, E)$ with edge weights $(w_{ij})_{(ij) \in E}$ and capacity constraints $\mathbf{b} = (b_i)$. There exists a stable solution if and only if the primal LP (16) admits an integer optimum. Further, if (M, Γ) is a stable outcome, then M is a maximum weight \mathbf{b} -matching, and $y_i = \gamma_i$ for all $i \in V$ and $y_{ij} = (w_{ij} - y_i - y_j)_+$ for all $(ij) \in E$ is an optimum solution to the dual LP.*

Proof. If x^* is an integer optimum for the primal LP, and $H^* \subset G$ is the corresponding \mathbf{b} -matching), the complementary slackness conditions read

- (i) For all $ij \in E(H^*)$; $w_{ij} = y_{ij}^* + y_i^* + y_j^*$.
- (ii) For all $ij \notin E(H^*)$; $y_{ij}^* = 0$.
- (iii) For all i with $d_i(H^*) < b_i$; $y_i^* = 0$.

We can construct a stable outcome (H^*, γ) by setting $\gamma_{i \rightarrow j} = y_j^* + y_{ij}^*/2$ for $(ij) \in H^*$, and $\gamma_{i \rightarrow j} = 0$ otherwise: Using (iii) above, $\gamma_i \geq y_i^*$ (cf. definition of γ_i above), so for any $(ij) \notin H^*$, we have $\gamma_i + \gamma_j \geq y_i^* + y_j^* \geq w_{ij}$, using (ii) above. It is easy to check that $\gamma_{i \rightarrow j} + \gamma_{j \rightarrow i} = w_{ij}$ for any $(ij) \in H^*$ using (i) above. Thus, (H^*, γ) is a stable outcome.

For the converse, consider a stable allocation (M, γ) . We claim that M forms a (integer) primal optimum. For this we simply demonstrate that there is a feasible point in the dual with the same value as the primal value at M : Take $y_i = \gamma_i$, and $y_{ij} = w_{ij} - y_i - y_j$ for edges in M , and 0 otherwise. The dual objective is then exactly equal to the weight of M . This also proves the second part of the lemma. \square

5.2 Dynamical model

We retain the notation $\alpha_{i \setminus j}^t$ for the ‘best alternative’ estimated in iteration t . As before, ‘offers’ are determined as

$$m_{i \rightarrow j}^t = (w_{ij} - \alpha_{i \setminus j}^t)_+ - \frac{1}{2}(w_{ij} - \alpha_{i \setminus j}^t - \alpha_{j \setminus i}^t)_+,$$

in the spirit of the pairwise Nash bargaining solution.

Now the best alternative $\alpha_{i \setminus j}^t$ should be the estimated income from the ‘replacement’ partnership, if i and j do not reach an agreement with each other. This ‘replacement’ should be the one corresponding to the b_i^{th} largest offer received by i from neighbors other than j . Hence, the update rule is modified to

$$\alpha_{i \setminus j}^{t+1} = (1 - \kappa) \alpha_{i \setminus j}^t + \kappa (b_i^{\text{th-max}})_{k \in \partial i \setminus j} m_{k \rightarrow i}^t, \quad (17)$$

where $\kappa \in (0, 1)$ is the damping factor.

Further, we define $\Gamma = (\gamma_{i \rightarrow j}, \gamma_{j \rightarrow i})_{(ij) \in E}$ by

$$\gamma_{j \rightarrow i}^t \equiv \begin{cases} m_{j \rightarrow i}^t & \text{if } m_{j \rightarrow i}^t \text{ is among top } b_i \text{ incoming offers to } i \\ 0 & \text{otherwise.} \end{cases} \quad (18)$$

Here ties are broken arbitrarily in ordering incoming offers. Finally, we define

$$\gamma_i^t \equiv (b_i^{\text{th-max}})_{k \in \partial i} \gamma_{k \rightarrow i}^t = (b_i^{\text{th-max}})_{k \in \partial i} m_{k \rightarrow i}^t \quad (19)$$

5.3 Results

Our first result is that fixed points of the new update equations (4), (17) are again in correspondence with Nash bargaining solutions when such solutions exist (analogous to Theorem 1). *Note that the fixed points are independent of the damping factor κ .* First, we generalize the notion of an induced matching.

Definition 12. *We say that a state $(\underline{\alpha}, \underline{m}, \Gamma)$ (or just $\underline{\alpha}$) induces a \mathbf{b} -matching M if the following happens. For each node $i \in V$ receiving at least b_i non-zero offers ($m_{\cdot \rightarrow i} > 0$): there is no tie for the $(b_i^{\text{th-max}})$ incoming offer to i , and node i is matched under M to the b_i neighbors from whom it is receiving its b_i highest offers. For each node $i \in V$ receiving less than b_i non-zero offers: node i is matched under M to all its neighbors from whom it is receiving positive offers.*

Consider the LP relaxation to the maximum weight matching problem (2). A feasible point \underline{x} for LP (2) is called *half-integral* if for all $e \in E$, $x_e \in \{0, 1, \frac{1}{2}\}$. Again, it can be easily shown that the primal LP (16) always has an optimum \underline{x}^* that is half-integral [38, Chapter 31]. As before, an LP with a fully integer \underline{x}^* (i.e., $x_e^* \in \{0, 1\}$ for all $e \in E$) is called *tight*.

Theorem 8. *Let $G = (V, E)$ with edge weights $(w_{ij})_{(ij) \in E}$ and capacity constraints $\mathbf{b} = (b_i)$ be an instance such that the primal LP (16) has a unique optimum that is integral, corresponding to matching M^* . Let $(\underline{\alpha}, \underline{m}, \Gamma)$ be a fixed point of the natural dynamics. Then $\underline{\alpha}$ induces matching M^* and (M^*, Γ) is a Nash bargaining solution. Conversely, every Nash bargaining solution (M, Γ_{NB}) has $M = M^*$ and corresponds to a unique fixed point of the natural dynamics with $\Gamma = \Gamma_{\text{NB}}$.*

We prove Theorem 8 in Appendix E.

Remark 13. *The condition that a tight primal LP (16) has a unique optimum is generic (analogous to Appendix C, Remark 18). Hence, Theorem 8 applies to ‘almost all’ problems with for which there exists a stable solution (cf. Lemma 2).*

Corollary 14. *Let $G = (V, E)$ with edge weights $(w_{ij})_{(ij) \in E}$ and capacity constraints $\mathbf{b} = (b_i)$ be an instance such that the primal LP (16) has a unique optimum that is integral. Then the instance possesses a Nash bargaining solution.*

Thus, we obtain a (almost) tight characterization of the when NB solutions exist in the case of general capacity constraints⁸.

Our convergence results, Theorems 2 and 3, generalize immediately, with the proofs (cf. Section 7) going through nearly verbatim:

Theorem 9. *Let $G = (V, E)$ with edge weights $(w_{ij})_{(ij) \in E}$ and capacity constraints $\mathbf{b} = (b_i)$ be any instance. The natural dynamics has at least one fixed point. Moreover, for any initial condition with $\underline{\alpha}^0 \in [0, W]^{2|E|}$, $\underline{\alpha}^t$ converges to a fixed point.*

We retain the Definition 5 for an ϵ -fixed point.

Theorem 10. *Let $G = (V, E)$ with weights $(w_e, e \in E) \in [0, 1]^{|E|}$ and capacity constraints $\mathbf{b} = (b_i)$ be any instance. Take any initial condition $\underline{\alpha}^0 \in [0, 1]^{2|E|}$. Take any $\epsilon > 0$. Define*

$$T^*(\epsilon) = \frac{1}{\pi \kappa (1 - \kappa) \epsilon^2}. \quad (20)$$

Then for all $t \geq T^(\epsilon)$, $\underline{\alpha}^t$ is an ϵ -fixed point. (Again $\pi = 3.14159\dots$)*

6 Discussion

Our results provide a dynamical justification for balanced outcomes, showing that agents bargaining with each other in a realistic, local manner can find such outcomes quickly. Refer to Section 1.2 for a summary of our contributions.

Some caution is needed in the interpretation of our results. Our dynamics avoids the question of how and when a pair of agents will cease to make iterative updates, and commit to each other. We showed that the right pairing will be found in time polynomial in the network size n and the parameter LP parameter g . But how will agents find out when this convergence has occurred? After all, agents are not likely to know n , and even less likely to know g . Further, why should agents wait for the right pairing to be found? It may be better for them to strike a deal after a few iterative updates because (i) they may estimate that they are unlikely to get a better deal later, (ii) they may be impatient, (iii) the convergence time may be very large on large networks. If a pair of agents *do* pair up and leave, then this changes the situation for the remaining agents, some of whom may have lost possible partners ([30] studies a model with this flavor). Our dynamics does not deal with this. A possible approach to circumventing some of these problems is to interpret our model in the context of a repeated game, where agents can pair up, but still continue to renegotiate their partnerships. Formalizing this is an open problem.

Related to the above discussion is the fact that our agents are not strategic. Though our dynamics admits interpretation as a bargaining process, it is unclear how, for instance, agent j becomes

⁸For simplicity, we have stated and proved, in Theorem 8, a generalization of only part (a) of Theorem 1. However, we expect that part (b) also generalizes, which would then lead to an exact characterization of when NB solutions exist in this case.

aware of the best alternative $\alpha_{i \setminus j}$ of a neighbor i . In the case of a fixed best alternative, the work of Rubinstein [36] justifies the pairwise Nash bargaining solution, but in our case the best alternative estimates evolve in time. Thus, it is unclear how to explain our dynamics game theoretically. However, we do not consider this to be a major drawback of our approach. Non-strategic agent behavior is commonly assumed in the literature on learning in games [18], even in games of only two players. Alternative recent approaches to bargaining in networks assume strategic agents, but struggle to incorporate reasonable informational assumptions (e.g. [30] assumes common knowledge of the network and perfect information of all prior events). Prima facie, it appears that bounded rationality models like ours may be more realistic.

Several examples admit multiple balanced outcomes (see Example 3 in Section 1). In fact, this is a common feature of two-sided assignment markets, which typically contain multiple even cycles. It would be very interesting to investigate whether our dynamics favors some balanced outcomes over others. If this is the case, it may improve our ability to predict outcomes in such markets.

Our model assumes the network to be exogenous, which does not capture the fact that agents may strategically form links. It would be interesting (and very challenging) to endogenize the network. A perhaps less daunting proposition is to characterize bargaining on networks that experience shocks, like the arrival of new agents, the departure of agents or the addition/deletion of links. Our result showing convergence to an approximate fixed point in time independent of the network size provides hope of progress on this front.

Finally, we remark on the computational problem of computing exact UD solutions in the unsymmetric case (recall that we give an FPTAS). We conjecture that the problem is computationally hard (cf. Section 1.2), with the recently introduced complexity class *continuous local search* [15] providing a possible way forward. We leave it as a challenging open problem to prove or refute this conjecture.

7 Convergence to fixed points: Proofs of Theorems 2 and 3

Theorems 2 and 3 admit a surprisingly simple proofs, that build on powerful results in the theory of nonexpansive mappings in Banach spaces.

Definition 15. *Given a normed linear space L , and a bounded domain $D \subseteq L$, a nonexpansive mapping $\mathbb{T} : D \rightarrow L$ is a mapping satisfying $\|\mathbb{T}x - \mathbb{T}y\| \leq \|x - y\|$ for all $x, y \in D$.*

Mann [31] first considered the iteration $x^{t+1} = (1-\kappa)x^t + \kappa\mathbb{T}x^t$ for $\kappa \in (0, 1)$, which is equivalent to iterating $\mathbb{T}_\kappa = (1-\kappa)I + \kappa\mathbb{T}$. Ishikawa [20] and Edelstein-O'Brien [16] proved the surprising result that, if the sequence $\{x^t\}_{t \geq 0}$ is bounded, then $\|\mathbb{T}x^t - x^t\| \rightarrow 0$ (the sequence is *asymptotically regular*) and indeed $x^t \rightarrow x^*$ with x^* a fixed point of \mathbb{T} .

Baillon and Bruck [3] recently proved a powerful quantitative version of Ishikawa's theorem: If $\|x^0 - x^t\| \leq 1$ for all t , then

$$\|\mathbb{T}x^t - x^t\| < \frac{1}{\sqrt{\pi\kappa(1-\kappa)t}}. \quad (21)$$

The surprise is that such a result holds irrespective of the mapping \mathbb{T} and of the normed space (in particular, of its dimensions). Theorems 2 and 3 immediately follow from this theory once we recognize that the natural dynamics can be cast into the form of a Mann iteration for a mapping which is nonexpansive with respect to a suitably defined norm.

Let us stress that the nonexpansivity property does not appear to be a lucky mathematical accident, but rather an intrinsic property of bargaining models under the one-exchange constraint.

It loosely corresponds to the basic observation that if earnings in the neighborhood of a pair of trade partners change by amounts N_1, N_2, \dots, N_k , then the balanced split for the partners changes at most by $\max(N_1, N_2, \dots, N_k)$, i.e., the largest of the neighboring changes.

Our technique seems therefore applicable in a broader context. (For instance, it can be applied successfully to prove fast convergence of a synchronous and damped version of the edge-balancing dynamics of [2].)

Proof (Theorem 2). We consider the linear space $L = \mathbb{R}^{2|E|}$ indexed by directed edges in G . On the bounded domain $D = [0, W]^{2|E|}$ we define the mapping $\mathsf{T} : \underline{\alpha} \mapsto \mathsf{T}\underline{\alpha}$ by letting, for $(i, j) \in E$,

$$(\mathsf{T}\underline{\alpha})_{i \setminus j} \equiv \max_{k \in \partial i \setminus j} m_{k \rightarrow i} |_{\underline{\alpha}}, \quad (22)$$

where $m_{k \rightarrow i} |_{\underline{\alpha}}$ is defined by Eq. (4). It is easy to check that the sequence of best alternatives produced by the natural dynamics corresponds to the Mann iteration $\underline{\alpha}^t = \mathsf{T}_\kappa^t \underline{\alpha}^0$. Also, T is nonexpansive for the ℓ_∞ norm

$$\|\underline{\alpha} - \underline{\beta}\|_\infty = \max_{(i,j) \in E} |\alpha_{i \setminus j} - \beta_{i \setminus j}|. \quad (23)$$

Non-expansivity follows from:

(i) The ‘max’ in Eq. (22) is non expansive.

(ii) An offer $m_{i \rightarrow j}$ as defined by Eq. (4) is nonexpansive. To see this, note that $m_{i \rightarrow j} = f(\alpha_{i \setminus j}, \alpha_{j \setminus i})$, where $f(x, y) : \mathbb{R}_+^2 \rightarrow \mathbb{R}_+$ is given by

$$f(x, y) = \begin{cases} \frac{w_{ij} - x + y}{2} & x + y \leq w_{ij} \\ (w_{ij} - x)_+ & \text{otherwise.} \end{cases} \quad (24)$$

It is easy to check that f is continuous everywhere 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 $\|\nabla f\|_1 = |\frac{\partial f}{\partial x}| + |\frac{\partial f}{\partial y}| \leq 1$. Hence, f is Lipschitz continuous in the L^∞ norm, with Lipschitz constant 1, i.e., it is nonexpansive in sup norm.

Notice that T_κ maps $D \equiv [0, W]^{2|E|}$ into itself. The thesis follows from [20, Corollary 1]. \square

Proof (Theorem 3). With the definitions given above, consider $W = 1$ (whence $\|\mathsf{T}\underline{\alpha}^t - \underline{\alpha}^0\|_\infty \leq 1$ for all t) and apply [3, Theorem 1]. \square

7.1 Exponentially fast convergence to unique Nash bargaining solution

Convergence of the natural dynamics was studied in an earlier version of this paper using a different (and much more laborious) technique [24]. While the results in Section 3 constitute a large improvement in elegance and generality over those of [24], the latter retain an independent interest. Indeed the analysis of [24] shows that convergence is exponentially fast in a well defined class of instances. We decided therefore to retain the main result of that analysis (recast from [24]).

Theorem 11. *Assume $W = 1$. Let G be an instance having unique Nash bargaining solution $(M, \underline{\gamma})$ with KT gap $\sigma > 0$, and let $\underline{\gamma}$ denote the corresponding allocation. Then, for any $\epsilon \in (0, \sigma/4)$, there exists $T_*(n, \sigma, \epsilon) = C n^7 [1/\sigma + \log(\sigma/\epsilon)]$, such that, for any initial condition with $\underline{\alpha}^0 \in [0, 1]^{2|E|}$, and any $t \geq T_*$ the natural dynamics yields earning estimates $\underline{\gamma}^t$, with $|\gamma_i^t - \gamma_i| \leq \epsilon$ for all $i \in V$. Moreover, $\underline{\alpha}^t$ induces the matching M and $(M, \underline{\gamma}^t)$ is a (4ϵ) -NB solution for any $t \geq T_*$.*

We refer to Appendix F for a definition of the KT gap σ (here KT stands for Kleinberg-Tardos). Suffice it to say that it is related to the Kleinberg-Tardos decomposition of G and that it is polynomially computable [27].

As mentioned above, the proof is based on a very different technique, namely on ‘approximate decoupling’ of the natural dynamics on different KT structures under the assumptions $\sigma > 0$ (which is generic) and that there is a unique NB solution. See preprint [24] for a complete proof.

Let us stress here that, for fixed σ , $T_*(n, \sigma, \epsilon)$ is only logarithmic in $(1/\epsilon)$ while it is proportional to $1/\epsilon^2$ in Theorem 3. In other words, for instances with KT gap bounded away from 0, the natural dynamics converges exponentially fast, while Theorem 3 guarantees inverse polynomial convergence in the general case.

8 Fixed point properties: Proof of Theorem 1

Let \mathcal{S} be the set of optimum solutions of LP (2). We call $e \in E$ a *strong-solid edge* if $x_e^* = 1$ for all $x^* \in \mathcal{S}$ and a *non-solid edge* if $x_e^* = 0$ for all $x^* \in \mathcal{S}$. We call $e \in E$ a *weak-solid edge* if it is neither strong-solid nor non-solid.

Proof of Theorem 1: From fixed points to NB solutions. The direct part follows from the following set of fixed point properties. The proofs of these properties are given in Appendix C. Throughout $(\underline{\alpha}, \underline{m}, \underline{\gamma})$ is a fixed point of the dynamics (4), (6) (with $\underline{\gamma}$ given by (5)).

(1) Two players $(i, j) \in E$ are called *partners* if $\gamma_i + \gamma_j = w_{ij}$. Then the following are equivalent: (a) i and j are partners, (b) $w_{ij} - \alpha_{i \setminus j} - \alpha_{j \setminus i} \geq 0$, (c) $\gamma_i = m_{j \rightarrow i}$ and $\gamma_j = m_{i \rightarrow j}$.

(2) Let $P(i)$ be the set of all partners of i . Then the following are equivalent: (a) $P(i) = \{j\}$ and $\gamma_i > 0$, (b) $P(j) = \{i\}$ and $\gamma_j > 0$, (c) $w_{ij} - \alpha_{i \setminus j} - \alpha_{j \setminus i} > 0$, (d) i and j receive unique best positive offers from each other.

(3) We say that (i, j) is a *weak-dotted edge* if $w_{ij} - \alpha_{i \setminus j} - \alpha_{j \setminus i} = 0$, a *strong-dotted edge* if $w_{ij} - \alpha_{i \setminus j} - \alpha_{j \setminus i} > 0$, and a *non-dotted edge* otherwise. If i has no adjacent dotted edges, then $\gamma_i = 0$.

(4) An edge is strong-solid (weak-solid) if and only if it is strongly (weakly) dotted.

(5) The balance property (1), holds at every edge $(i, j) \in E$ (with both sides being non-negative).

(6) $\underline{\gamma}$ is an optimum solution for the dual LP (3) to LP (2) and $m_{i \rightarrow j} = (w_{ij} - \gamma_i)_+$ holds for all $(i, j) \in E$.

Proof of Theorem 1 (a), direct implication. Assume that the LP (2) has a unique optimum that is integral. Then, by property 4, the set of strong-dotted edges form the unique maximum weight matching M^* and all other edges are non-dotted. By property 3 for i that is unmatched under M^* , $\gamma_i = 0$. Hence by property 2, $\underline{\alpha}$ induces the matching M^* . Finally, by properties 6 and 5, the pair $(M^*, \underline{\gamma})$ is stable and balanced respectively, and thus forms a NB solution. \square

The corresponding result for the non-unique optimum case (part (b)) can be proved similarly: it follows immediately Theorem 12, Appendix C.

Remark 16. *Properties 1-6 hold for any instance. This leads to the general result Theorem 12 in Appendix C shows that in general, fixed points correspond to dual optima satisfying the unmatched balance property (1).*

Proof of Theorem 1: From NB solutions to fixed points.

Proof. Consider any NB solution $(M, \underline{\gamma}_{\text{NB}})$. Using Proposition 1, M is a maximum weight matching. Construct a corresponding FP as follows. Set $m_{i \rightarrow j} = (w_{ij} - \gamma_{\text{NB},i})_+$ for all $(i, j) \in E$. Compute $\underline{\alpha}$ using $\alpha_{i \setminus j} = \max_{k \in \partial i \setminus j} m_{k \rightarrow i}$. We claim that this is a FP and that the corresponding $\underline{\gamma}$ is $\underline{\gamma}_{\text{NB}}$. To prove that we are at a fixed point, we imagine updated offers $\underline{m}^{\text{upd}}$ based on $\underline{\alpha}$, and show $\underline{m}^{\text{upd}} = \underline{m}$.

Consider a matching edge $(i, j) \in M$. We know that $\gamma_{\text{NB},i} + \gamma_{\text{NB},j} = w_{ij}$. Also stability and balance tell us $\gamma_{\text{NB},i} - \max_{k \in \partial i \setminus j} (w_{ik} - \gamma_{\text{NB},k})_+ = \gamma_{\text{NB},j} - \max_{l \in \partial j \setminus i} (w_{jl} - \gamma_{\text{NB},l})_+$ and both sides are non-negative. Hence, $\gamma_{\text{NB},i} - \alpha_{i \setminus j} = \gamma_{\text{NB},j} - \alpha_{j \setminus i} \geq 0$. Therefore $\alpha_{i \setminus j} + \alpha_{j \setminus i} \leq w_{ij}$,

$$\begin{aligned} m_{i \rightarrow j}^{\text{upd}} &= \frac{w_{ij} - \alpha_{i \setminus j} + \alpha_{j \setminus i}}{2} = \frac{w_{ij} - \gamma_{\text{NB},i} + \gamma_{\text{NB},j}}{2} \\ &= \gamma_{\text{NB},j} = w_{ij} - \gamma_{\text{NB},i} = m_{i \rightarrow j}. \end{aligned}$$

By symmetry, we also have $m_{j \rightarrow i}^{\text{upd}} = \gamma_{\text{NB},i} = m_{j \rightarrow i}$. Hence, the offers remain unchanged. Now consider $(i, j) \notin M$. We have $\gamma_{\text{NB},i} + \gamma_{\text{NB},j} \geq w_{ij}$ and, $\gamma_{\text{NB},i} = \max_{k \in \partial i \setminus j} (w_{ik} - \gamma_{\text{NB},k})_+ = \alpha_{i \setminus j}$. Similar equation holds for $\gamma_{\text{NB},j}$. The validity of this identity can be checked individually in the cases when $i \in M$ and $i \notin M$. Hence, $\alpha_{i \setminus j} + \alpha_{j \setminus i} \geq w_{ij}$. This leads to $m_{i \rightarrow j}^{\text{upd}} = (w_{ij} - \alpha_{i \setminus j})_+ = (w_{ij} - \gamma_{\text{NB},i})_+ = m_{i \rightarrow j}$. By symmetry, we know also that $m_{j \rightarrow i}^{\text{upd}} = m_{j \rightarrow i}$.

Finally, we show $\underline{\gamma} = \underline{\gamma}_{\text{NB}}$. For all $(i, j) \in M$, we already found that $m_{i \rightarrow j} = \gamma_j$ and vice versa. For any edge $(ij) \notin M$, we know $m_{i \rightarrow j} = (w_{ij} - \gamma_{\text{NB},i})_+ \leq \gamma_{\text{NB},j}$. This immediately leads to $\underline{\gamma} = \underline{\gamma}_{\text{NB}}$. It is worth noting that making use of the uniqueness of LP optimum we know that $M = \bar{M}^*$, and we can further show that $\gamma_i = m_{j \rightarrow i} > \alpha_{i \setminus j}$ if and only if $(ij) \in M$, i.e., the fixed point reconstructs the pairing $M = M^*$. \square

9 Polynomial convergence on bipartite graphs: Proof of Theorem 5

Theorem 5 says that on a bipartite graph, under a small random perturbation on any problem instance, the natural dynamics is likely to quickly find the maximum weight matching. Now, in light of Corollary 9, this simply involves showing that the gap g of the perturbed problem instance is likely to be sufficiently large. We use a version of the well known Isolation Lemma to for this. Note that on bipartite graphs, there is always an integral optimum to the LP (2).

Next, is our Isolation lemma (recast from [21]). For the proof, see Appendix B.

Lemma 3 (Isolation Lemma). *Consider a bipartite graph $G = (V, E)$. Choose $\eta > 0$, $\xi > 0$. Edge weights are generated as follows: for each $e \in E$, \bar{w}_e is chosen uniformly in $[w_e, w_e + \eta]$. Denote by \mathcal{M} the set of matchings in G . Let M^* be a maximum weight matching. Let M^{**} be a matching having the maximum weight in $\mathcal{M} \setminus M^*$. Denote by $\bar{w}(M)$ the weight of a matching M . Then*

$$\Pr[\bar{w}(M^*) - \bar{w}(M^{**}) \geq \eta\xi/(2|E|)] \geq 1 - \xi \quad (25)$$

Proof of Theorem 5. Using Lemma 3, we know that the gap of the perturbed problem satisfies $\bar{g} \geq \eta\xi/(2|E|)$ with probability at least $1 - \xi$. Now, the weights in the perturbed instance are bounded by $\bar{W} = 2$. Rescale by dividing all weights and messages by 2, and use Corollary 9. The theorem follows from the following two elementary observations. First, an $(\epsilon/2)$ -NB solution for the rescaled problem corresponds to an ϵ -NB solution for the original problem. Second, induced matchings are unaffected by scaling. \square

We remark that Theorem 5 does not generalize to any (non-bipartite) graph with edge weights such that the LP (2) has an integral optimum, for the following reason. We can easily generalize the Isolation Lemma to show that the gap g of the perturbed problem is likely to be large also in this case. However, there is a probability arbitrarily close to 1 (depending on the instance) that a random perturbation will result in an instance for which LP (2) does *not* have an integral optimum, i.e. the perturbed instance does not have any Nash bargaining solutions!

Acknowledgements. We thank Eva Tardos for introducing us to network exchange theory and Daron Acemoglu for insightful discussions. We also thank the anonymous referees for their comments.

A large part of this work was done while Y. Kanoria, M. Bayati and A. Montanari were at Microsoft Research New England. This research was partially supported by NSF, grants CCF-0743978 and CCF-0915145, and by a Terman fellowship. Y. Kanoria is supported by a 3Com Corporation Stanford Graduate Fellowship.

References

- [1] B. Aspvall and Y. Shiloach, *A polynomial time algorithm for solving systems of linear inequalities with two variables pre inequality*, Proc. 20th IEEE Symp. Found. of Comp. Sc., 1979.
- [2] Y. Azar, B. Birnbaum, L. Elisa Celis, N. R. Devanur and Y. Peres, *Convergence of Local Dynamics to Balanced Outcomes in Exchange Networks*, Proc. 50th IEEE Symp. Found. of Comp. Sc., 2009.
- [3] J. Baillon and R. E. Bruck, *The rate of asymptotic regularity is $O(1/\sqrt{n})$* , in: A.G. Kartsatos (ed.), *Theory and applications of nonlinear operators of accretive and monotone type*, Lecture Notes in Pure and Appl. Math. 178, Marcel Dekker, Inc., New York, 1996, pp. 51–81.
- [4] M. Bateni, M. Hajiaghayi, N. Immorlica and H. Mahini, *The cooperative game theory foundations of network bargaining games*, Proc. Intl. Colloq. Automata, Languages and Programming, 2010.
- [5] M. Bayati, D. Shah and M. Sharma, *Max-Product for Maximum Weight Matching: Convergence, Correctness, and LP Duality*, IEEE Trans. Inform. Theory, 54 (2008), pp. 1241–1251.
- [6] M. Bayati, C. Borgs, J. Chayes, R. Zecchina, *On the exactness of the cavity method for Weighted b -Matchings on Arbitrary Graphs and its Relation to Linear Programs*, arXiv:0807.3159, (2007).
- [7] D. P. Bertsekas, *The Auction Algorithm: A Distributed Relaxation Method for the Assignment Problem*, Annals of Operations Research, 14 (1988), pp. 105–123.
- [8] K. Binmore, A. Rubinstein and A. Wolinsky, *The Nash Bargaining Solution in Economic Modeling*, Rand Journal of Economics, 17(2) (1986), pp. 176–188.
- [9] N. Braun, T. Gautschi, *A Nash bargaining model for simple exchange networks*, Social Networks, 28(1) (2006), pp. 1–23.
- [10] M. Kandori, G. J. Mailath, and R. Rob, "Learning, Mutation, and Long Run Equilibria in Games," *Econometrica* 61(1) (1993).

- [11] T. Chakraborty, S. Judd, M. Kearns, J. Tan, *A Behavioral Study of Bargaining in Social Networks*, Proc. 10th ACM Conf. Electronic Commerce, 2010.
- [12] T. Chakraborty, M. Kearns and S. Khanna, *Network Bargaining: Algorithms and Structural Results*, Proc. 10th ACM Conf. Electronic Commerce, July 2009.
- [13] K. S. Cook and T. Yamagishi, *Power exchange in networks: A power-dependence formulation*, Social Networks, 14 (1992), pp. 245–265.
- [14] C. Daskalakis, C. Papadimitriou, *On oblivious PTAS's for nash equilibrium*, Proc. ACM Symp. Theory of Computing, 2009.
- [15] C. Daskalakis, C. Papadimitriou, *Geometrically Embedded Local Search*, Proc. ACM-SIAM Symp. Discrete Algorithms, 2011.
- [16] M. Edelstein and R.C. O'Brien, *Nonexpansive mappings, asymptotic regularity, and successive approximations*, J. London Math. Soc. 1 (1978), pp. 547–554.
- [17] U. Faigle, W. Kern, and J. Kuipers, *On the computation of the nucleolus of a cooperative game*, International Journal of Game Theory, 30 (2001), pp. 79–98.
- [18] D. Fudenberg and D. K. Levine, *The Theory of Learning in Games*, The MIT Press, 1998.
- [19] Gabow, H. N., Tarjan, R. E.: Faster scaling algorithms for general graph-matching problems. J. ACM, **38**(4):815-853, 1991.
- [20] S. Ishikawa, *Fixed points and iteration of a nonexpansive mapping in a Banach space*, Proc. American Mathematical Society, 59(1) (1976).
- [21] D. Gamarnik, D. Shah and Y. Wei, *Belief Propagation for Min-cost Network Flow: Convergence and Correctness*, Proc. ACM-SIAM Symp. Disc. Alg., 2010.
- [22] B. Huang, T. Jebara, *Loopy belief propagation for bipartite maximum weight b-matching*, Artificial Intelligence and Statistics (AISTATS), March, 2007.
- [23] Y. Kanoria, *An FPTAS for Bargaining Networks with Unequal Bargaining Powers*, Proc. Workshop on Internet and Network Economics, 2010, (full version arXiv:1008.0212).
- [24] Y. Kanoria, M. Bayati, C. Borgs, J. Chayes, and A. Montanari, *A Natural Dynamics for Bargaining on Exchange Networks*, arXiv:0912.5176 (2009).
- [25] Y. Kanoria, M. Bayati, C. Borgs, J. Chayes, and A. Montanari, *Fast Convergence of Natural Bargaining Dynamics in Exchange Networks*, Proc. ACM-SIAM Symp. on Discrete Algorithms (2011), pp. 1518–1537.
- [26] E. Kalai, M. Smorodinsky, *Other Solutions to Nash's Bargaining Problem*, Econometrica, 43(3) (1975).
- [27] J. Kleinberg and E. Tardos, *Balanced outcomes in social exchange networks*, in Proc. ACM Symp. Theory of Computing, 2008.
- [28] U. Kohlenback, *A Quantitative Version Of A Theorem Due To Borwein-Reich-Shafir*, Numerical Functional Analysis and Optimization, 22(5-6), August 2001.

- [29] J.W. Lucas, C.W. Younts, M.J. Lovaglia, and B. Markovsky, *Lines of power in exchange networks*, Social Forces, 80 (2001), pp. 185–214.
- [30] M. Mihai, D. Abreu, “Markov Equilibria in a Model of Bargaining in Networks,” to appear in Games and Economic Behavior <http://economics.mit.edu/files/6985>.
- [31] W. R. Mann, *Mean value methods in iteration*, Proc. Amer. Math Soc., 4 (1953), pp. 506–510.
- [32] J. Nash, *The bargaining problem*, Econometrica, 18 (1950), pp. 155–162.
- [33] S.C. Rochford, *Symmetric pairwise-bargained allocations in an assignment market*, in J. Economic Theory, 34 (1984), pp. 262–281.
- [34] V. P. Crawford, S. C. Rochford, *Bargaining and Competition in Matching Markets*, Intl. Economic Review 27(2) (1986), pp. 329–348.
- [35] J. Z. Rubin, B. R. Brown, *The Social Psychology of Bargaining and Negotiation*, Academic Press 1975.
- [36] A. Rubinstein: *Perfect equilibrium in a bargaining model*, Econometrica, 50 (1982), pp. 97–109.
- [37] S. Sanghavi, D. Malioutov, A. Willsky, *Linear Programming Analysis of Loopy Belief Propagation for Weighted Matching*, Proc. Neural Inform. Processing Systems, 2007.
- [38] A. Schrijver, *Combinatorial Optimization*, Springer-Verlag, Vol. A (2003).
- [39] L. Shapley and M. Shubik, “The Assignment Game I: The Core,” Intl. J. Game Theory, 1 (1972), pp. 111-130.
- [40] M. Sotomayor, *On The Core Of The One-Sided Assignment Game*, 2005, http://www.usp.br/feaecon/media/fck/File/one_sided_assignment_game.pdf.
- [41] J. Skvoretz and D. Willer, *Exclusion and power: A test of four theories of power in exchange networks*, American Sociological Review, 58 (1993), pp. 801–818.
- [42] F. Spitzer, *Principles of Random Walk*, Springer, 2005.
- [43] R. E. Stearns, *Convergent transfer schemes for n-person games*, Transactions of American Mathematical Society, 134 (1968), pp. 449-459.
- [44] M. J. Wainwright, T. S. Jaakkola, and A. S. Willsky, *Exact MAP estimates via agreement on (hyper)trees: Linear programming and message-passing*, IEEE Trans. Inform. Theory, 51 (2005), pp. 3697–3717.
- [45] D. Willer (ed.) *Network Exchange Theory*, Praeger, 1999.

A Variations of the natural dynamics

Now that we have a reasonable dynamics that converges fast to balanced outcomes, it is natural ask whether variations of the natural dynamics also yield balanced outcomes. What happens in the case of asynchronous updates, different nodes updating at different rates, damping factors that vary across nodes and in time, and so on? We discuss some of these questions in this section, focussing

on some situations in which we *can* prove convergence with minimal additional work. Note that we are only concerned with extending our convergence results since the fixed point properties remain unchanged.

A.1 Node dependent damping

Consider that the damping factor may be different for different nodes, but unchanging over time. Denote by $\kappa(v)$, the damping factor for node v . Assume that $\kappa(v) \in [1 - \kappa^*, \kappa^*] \forall v \in V$ for some $\kappa^* \in [0.5, 1)$, i.e. damping factors are uniformly bounded away from 0 and 1. Define operator $\mathbb{T} : [0, 1]^{|E|} \rightarrow [0, 1]^{|E|}$ by

$$(\mathbb{T}\underline{\alpha})_{i \setminus j} = \left(\frac{\kappa(i)}{\kappa^*} \right) \max_{k \in \partial i \setminus j} m_{k \rightarrow i} + \left(\frac{\kappa^* - \kappa(i)}{\kappa^*} \right) \alpha_{i \setminus j} \quad (26)$$

\mathbb{T} is nonexpansive. Now, the dynamics can be written as $\underline{\alpha}^{t+1} = \kappa^* \mathbb{T} \underline{\alpha}^t + (1 - \kappa^*) \underline{\alpha}^t$. Clearly, convergence to fixed points (Theorem 2) holds in this situation. Note that fixed points of \mathbb{T} are the same as fixed points of the natural dynamics. Moreover, we can use [3] to assert that $\|\underline{\alpha}^t - \mathbb{T} \underline{\alpha}^t\|_\infty = O(1/\sqrt{t})$ and hence $\underline{\alpha}^t$ is an $O(1/\sqrt{t})$ -FP. In short, we don't lose anything with this generalization!

A.2 Time varying damping

Now consider instead that the damping may change over time, but is the same for all nodes. Denote by κ_t the damping factor at time t , i.e.

$$\underline{\alpha}^{t+1} = \kappa_t \max_{k \in \partial i \setminus j} m_{k \rightarrow i}^t + (1 - \kappa_t) \underline{\alpha}^t \quad (27)$$

The result of [20] implies that, as long as $\sum_{t=0}^{\infty} \kappa_t = \infty$ and $\lim_{t \rightarrow \infty} \sup \kappa_t < 1$, the dynamics is guaranteed to converge to a fixed point. Note that again the fixed points are unchanged. [28] provides a quantitative estimate of the rate of convergence in this case, guaranteeing in particular that an ϵ -FP is reached in time $\exp(O(1/\epsilon))$ if κ_t is uniformly bounded away from 0 and 1. Note that this estimate is much weaker than the one provided by [3], leading to Theorem 3. It seems intuitive that the stronger $O(1/\sqrt{k})$ bound holds also for the time varying damping case in the general nonexpansive operator setting, but a proof has remained elusive thus far.

A.3 Asynchronous updates

Finally, we look at the case of asynchronous updates, i.e., one message $\alpha_{i \setminus j}$ is updated in any given step while the others remain unchanged. Define $\mathbb{T}_{i \setminus j} : [0, 1]^{|E|} \rightarrow [0, 1]^{|E|}$ by

$$(\mathbb{T}_{i \setminus j} \underline{\alpha})_{i' \setminus j'} = \begin{cases} \max_{k \in \partial i' \setminus j'} m_{k \rightarrow i'} & \text{if } (i, j) = (i', j') \\ \alpha_{i' \setminus j'} & \text{otherwise} \end{cases} \quad (28)$$

Let $m \equiv |E|$. There are $2m$ such operators, two for each edge. Clearly, each $\mathbb{T}_{i \setminus j}$ is nonexpansive in sup-norm. Now, consider an arbitrary permutation of the $2m$ directed edges $((i_1, j_1), (i_2, j_2), \dots)$. Consider the updates induced by $\mathbb{T}_{i_1 \setminus j_1}, \mathbb{T}_{i_2 \setminus j_2}, \dots$ in order, each with a damping factor of $1/(2m)$.

Consider the resulting product

$$\begin{aligned}
& \left((1/2m)\mathbb{T}_{i_1 \setminus j_1} + (1 - (1/2m))I \right) \cdot \\
& \cdot \left((1/2m)\mathbb{T}_{i_2 \setminus j_2} + (1 - (1/2m))I \right) \cdot \dots \\
& = \left(1 - (1 - (1/2m))^{2m} \right) \mathbb{T} + (1 - (1/2m))^{2m} I
\end{aligned} \tag{29}$$

Here (29) defines \mathbb{T} , and I is the identity operator. It is easy to deduce that \mathbb{T} above is nonexpansive from the following elementary facts – the product of nonexpansive operators is nonexpansive, and the convex combination of nonexpansive operators is nonexpansive. Also, $(1 - (1/2m))^{2m} \in [1/4, 1/e] \forall m$. Thus, if we repeat these asynchronous updates periodically in a series of ‘update cycles’, we are guaranteed to quickly converge to an ϵ -FP of \mathbb{T} ([3]).

Proposition 17. *An ϵ -FP of \mathbb{T} is an $O(m\epsilon)$ -FP of the natural dynamics.*

Proof. Suppose we start an update cycle at $\underline{\alpha}$, an ϵ -FP of \mathbb{T} . Then we know that at the end of the update cycle, no coordinate changes by more than $(1 - 1/4)\epsilon \leq \epsilon$. Note that among the $2m$ steps in a cycle, any particular $i \setminus j$ ‘coordinate’ only changes in one step. Thus, each such coordinate change is bounded by ϵ . Consider the s -th step in the update cycle. The state before the s -th step, call it $\underline{\alpha}(s - 1)$, is ϵ -close to $\underline{\alpha}$. Also, we know that the $(i_s \setminus j_s)$ coordinate changes by at most ϵ in this step. Hence,

$$\begin{aligned}
& \|\mathbb{T}_{i_s \setminus j_s} \underline{\alpha}(s - 1) - \underline{\alpha}(s - 1)\|_\infty \leq (2m)\epsilon \\
\Rightarrow & \|\mathbb{T}_{i_s \setminus j_s} \underline{\alpha} - \underline{\alpha}\|_\infty \leq (2m + 2)\epsilon
\end{aligned}$$

This holds for $s = 1, 2, \dots, 2m$. Hence the result. \square

Note that with $\epsilon = 0$, Proposition 17 tells us that fixed points of \mathbb{T} are fixed points of the natural dynamics. Thus, we are immediately guaranteed convergence to fixed points of the natural dynamics. Moreover, the quantitative estimate in Proposition 17 guarantees that in a small number of update cycles we reach approximate fixed points of the natural dynamics.

Finally, we comment that instead of ordering updates by a permutation of directed edges, we could have an arbitrary periodic sequence of updates satisfying non-starvation and obtain similar results. For example, this would include cases where some nodes update more frequently than others. Also, note that the damping factors of $(1/2m)$ were chosen for simplicity and to ensure *fast* convergence. Any non-trivial damping would suffice to guarantee convergence.

It remains an open question to show convergence for non-periodic asynchronous updates.

B Proof of Isolation lemma

Our proof of the isolation lemma is adapted from [21].

Proof of Lemma 3. Fix $e \in E$ and fix $\bar{w}_{e'}$ for all $e' \in E \setminus e$. Let M_e be a maximum weight matching among matchings that strictly include edge e , and let $M_{\sim e}$ be a maximum weight matching among matchings that exclude edge e . Clearly, M_e and $M_{\sim e}$ are independent of \bar{w}_e . Define

$$\begin{aligned}
f_e(\bar{w}_e) &\equiv \bar{w}(M_e) = f_e(0) + \bar{w}_e \\
f_{\sim e} &\equiv \bar{w}(M_{\sim e}) = \text{const} < \infty
\end{aligned}$$

Clearly, $f_e(0) \leq f_{\sim e}$, since we cannot do worse by forcing exclusion of a zero weight edge. Thus, there is some unique $\theta \geq 0$ such that $f_e(\theta) = f_{\sim e}$. Define $\delta = \eta\xi/2|E|$. Let $D(e)$ be the event that $|\bar{w}(M_e) - \bar{w}(M_{\sim e})| < \delta$. It is easy to see that $D(e)$ occurs if and only if $\bar{w}_e \in (\theta - \delta, \theta + \delta)$. Thus, $\Pr[D(e)] \leq 2\delta/\eta = \xi/|E|$. Now,

$$\left\{ \bar{w}(M^*) - \bar{w}(M^{**}) < \delta \right\} = \bigcup_{e \in E} D(e) \quad (30)$$

and the lemma follows by union bound. □

C Proofs of fixed point properties

In this section we state and prove the fixed point properties that were used for the proof of Theorem 1 in Section 8. Before that, however, we remark that the condition: “LP (2) has a unique optimum” in Theorem 1(a) is almost always valid.

Remark 18. *We argue that the condition “LP (2) has a unique optimum” is generic in instances with integral optimum:*

Let $\mathbf{G}_I \subset [0, W]^{|E|}$ be the set of instances having an integral optimum. Let $\mathbf{G}_{UI} \subset \mathbf{G}_I$ be the set of instances having a unique integral optimum. It turns out that \mathbf{G}_I has dimension $|E|$ (i.e. the class of instances having an integral optimum is large) and that \mathbf{G}_{UI} is both open and dense in \mathbf{G}_I .

Notation. In proofs of this section and Section D we denote surplus $w_{ij} - \alpha_{i \setminus j} - \alpha_{j \setminus i}$ of edge (ij) by Surp_{ij} .

Lemma 4. $\underline{\gamma}$ satisfies the constraints of the dual problem (3).

Proof. Since offers $m_{i \rightarrow j}$ are by definition non-negative therefore for all $v \in V$ we have $\gamma_v \geq 0$. So we only need to show $\gamma_i + \gamma_j \geq w_{ij}$ for any edge $(ij) \in E$. It is easy to see that $\gamma_i \geq \alpha_{i \setminus j}$ and $\gamma_j \geq \alpha_{j \setminus i}$. Therefore, if $\alpha_{i \setminus j} + \alpha_{j \setminus i} \geq w_{ij}$ then $\gamma_i + \gamma_j \geq w_{ij}$ holds and we are done. Otherwise, for $\alpha_{i \setminus j} + \alpha_{j \setminus i} < w_{ij}$ we have $m_{i \rightarrow j} = \frac{w_{ij} - \alpha_{i \setminus j} + \alpha_{j \setminus i}}{2}$ and $m_{j \rightarrow i} = \frac{w_{ij} - \alpha_{j \setminus i} + \alpha_{i \setminus j}}{2}$ which gives $\gamma_i + \gamma_j \geq m_{i \rightarrow j} + m_{j \rightarrow i} = w_{ij}$. □

Recall that for any $(ij) \in E$, we say that i and j are ‘partners’ if $\gamma_i + \gamma_j = w_{ij}$ and $P(i)$ denotes the partners of node i . In other words $P(i) = \{j : j \in \partial i, \gamma_i + \gamma_j = w_{ij}\}$.

Lemma 5. *The following are equivalent:*

- (a) i and j are partners,
 - (b) $\text{Surp}_{ij} \geq 0$.
 - (c) $\gamma_i = m_{j \rightarrow i}$ and $\gamma_j = m_{i \rightarrow j}$.
- Moreover, if $\gamma_i = m_{j \rightarrow i}$ and $\gamma_j > m_{i \rightarrow j}$ then $\gamma_i = 0$.

Proof. We will prove (a) \Rightarrow (b) \Rightarrow (c) \Rightarrow (a).

(a) \Rightarrow (b): Since $\gamma_i \geq \alpha_{i \setminus j}$ and $\gamma_j \geq \alpha_{j \setminus i}$ always holds then $w_{ij} = \gamma_i + \gamma_j \geq \alpha_{i \setminus j} + \alpha_{j \setminus i}$.

(b) \Rightarrow (c): If $\text{Surp}_{ij} \geq 0$ then $(w_{ij} - \alpha_{i \setminus j} + \alpha_{j \setminus i})/2 \geq \alpha_{j \setminus i}$. But $m_{i \rightarrow j} = (w_{ij} - \alpha_{i \setminus j} + \alpha_{j \setminus i})/2$ therefore $\gamma_j = m_{i \rightarrow j}$. The argument for $\gamma_i = m_{j \rightarrow i}$ is similar.

(c) \Rightarrow (a): If $\text{Surp}_{ij} \geq 0$ then $m_{i \rightarrow j} = (w_{ij} - \alpha_{i \setminus j} + \alpha_{j \setminus i})/2$ and $m_{j \rightarrow i} = (w_{ij} - \alpha_{j \setminus i} + \alpha_{i \setminus j})/2$ which gives $\gamma_i + \gamma_j = m_{i \rightarrow j} + m_{j \rightarrow i} = w_{ij}$ and we are done. Otherwise, we have $\gamma_i + \gamma_j = m_{i \rightarrow j} + m_{j \rightarrow i} \leq$

$(w_{ij} - \alpha_{i \setminus j})_+ + (w_{ij} - \alpha_{j \setminus i})_+ < \max \left[(w_{ij} - \alpha_{i \setminus j})_+, (w_{ij} - \alpha_{j \setminus i})_+, 2w_{ij} - \alpha_{i \setminus j} - \alpha_{j \setminus i} \right] \leq w_{ij}$ which contradicts Lemma 4 that $\underline{\gamma}$ satisfies the constraints of the dual problem (3).

Finally, we need to show that $\gamma_i = m_{j \rightarrow i}$ and $\gamma_j > m_{i \rightarrow j}$ give $\gamma_i = 0$. First note that by equivalence of (b) and (c) we should have $w_{ij} < \alpha_{i \setminus j} + \alpha_{j \setminus i}$. On the other hand $\alpha_{i \setminus j} \leq \gamma_i = m_{j \rightarrow i} \leq (w_{ij} - \alpha_{j \setminus i})_+$. Now if $w_{ij} - \alpha_{j \setminus i} > 0$ we get $\alpha_{i \setminus j} \leq w_{ij} - \alpha_{j \setminus i}$ which is a contradiction. Therefore $\gamma_i = (w_{ij} - \alpha_{j \setminus i})_+ = 0$. \square

Lemma 6. *The following are equivalent:*

- (a) $P(i) = \{j\}$ and $\gamma_i > 0$,
- (b) $P(j) = \{i\}$ and $\gamma_j > 0$,
- (c) $w_{ij} - \alpha_{i \setminus j} - \alpha_{j \setminus i} > 0$.
- (d) i and j receive unique best positive offers from each other.

Proof. (a) \Rightarrow (c) \Rightarrow (b): (a) means that for all $k \in \partial i \setminus j$, $\text{Surp}_{ik} < 0$. This means $m_{k \rightarrow i} = (w_{ik} - \alpha_{k \setminus i})_+ < \alpha_{i \setminus k} = m_{j \rightarrow i}$ (using $\gamma_i > 0$). Hence, $\alpha_{i \setminus j} < m_{j \rightarrow i}$. From (a), it also follows that $m_{j \rightarrow i} > 0$ or $(w_{ij} - \alpha_{j \setminus i})_+ = w_{ij} - \alpha_{j \setminus i}$. Therefore, $m_{j \rightarrow i} \leq (w_{ij} - \alpha_{j \setminus i})_+ = w_{ij} - \alpha_{j \setminus i}$ which gives $w_{ij} - \alpha_{i \setminus j} - \alpha_{j \setminus i} > 0$ or (c). From this we can explicitly write $m_{i \rightarrow j} = (w_{ij} - \alpha_{i \setminus j} + \alpha_{j \setminus i})/2$ which is strictly bigger than $\alpha_{j \setminus i}$. Hence we obtain (b).

By symmetry (b) \Rightarrow (c) \Rightarrow (a). Thus, we have shown that (a), (b) and (c) are equivalent.

(c) \Rightarrow (d): (c) implies that $m_{i \rightarrow j} = (w_{ij} - \alpha_{i \setminus j} + \alpha_{j \setminus i})/2 > \alpha_{j \setminus i} = \max_{k \in \partial j \setminus i} m_{k \rightarrow j}$. Thus, j receives its unique best positive offer from i . Using symmetry, it follows that (d) holds.

(d) \Rightarrow (c): (d) implies $\gamma_i = m_{j \rightarrow i}$ and $\gamma_j = m_{i \rightarrow j}$. By Lemma 5, i and j are partners, i.e. $\gamma_i + \gamma_j = w_{ij}$. Hence, $m_{i \rightarrow j} + m_{j \rightarrow i} = w_{ij}$. But since (d) holds, $\alpha_{i \setminus j} < m_{j \rightarrow i}$ and $\alpha_{j \setminus i} < m_{i \rightarrow j}$. This leads to (c).

This finishes the proof. \square

Recall that (ij) is a weak-dotted edge if $w_{ij} - \alpha_{i \setminus j} - \alpha_{j \setminus i} = 0$, a strong-dotted edge if $w_{ij} - \alpha_{i \setminus j} - \alpha_{j \setminus i} > 0$, and a non-dotted edge otherwise. Basically, for any dotted edge (ij) we have $j \in P(i)$ and $i \in P(j)$.

Corollary 19. *A corollary of Lemmas 5-6 is that strong-dotted edges are only adjacent to non-dotted edges. Also each weak-dotted edge is adjacent to at least one weak-dotted edge at each end (assume that the earnings of the two endpoints are non-zero).*

Lemma 7. *If i has no adjacent dotted edges, then $\gamma_i = 0$*

Proof. Assume that the largest offer to i comes from j . Therefore, $\alpha_{i \setminus j} \leq m_{j \rightarrow i} \leq (w_{ij} - \alpha_{j \setminus i})_+$. Now if $w_{ij} - \alpha_{j \setminus i} > 0$ then $\alpha_{i \setminus j} \leq w_{ij} - \alpha_{j \setminus i}$ or (ij) is dotted edge which is impossible. Thus, $w_{ij} - \alpha_{j \setminus i} = 0$ and $\gamma_i = 0$. \square

Lemma 8. *The following are equivalent:*

- (a) $\alpha_{i \setminus j} = \gamma_i$,
- (b) $\text{Surp}_{ij} \leq 0$,
- (c) $m_{i \rightarrow j} = (w_{ij} - \alpha_{i \setminus j})_+$.

Proof. (a) \Rightarrow (b): “not (b)” $\Rightarrow m_{j \rightarrow i} = (w_{ij} - \alpha_{j \setminus i} + \alpha_{i \setminus j})/2 > \alpha_{i \setminus j} \Rightarrow$ “not (a)”.

(b) \Rightarrow (c): Follows from the definition of $m_{i \rightarrow j}$.

(c) \Rightarrow (a): From $m_{i \rightarrow j} = (w_{ij} - \alpha_{i \setminus j})_+$ we have $\text{Surp}_{ij} \leq 0$. Therefore, $m_{j \rightarrow i} = (w_{ij} - \alpha_{j \setminus i})_+ \leq \max [w_{ij} - \alpha_{j \setminus i}, 0] \leq \alpha_{i \setminus j}$. \square

Note that (b) is symmetric in i and j , so (a) and (c) can be transformed by interchanging i and j .

Corollary 20. $\alpha_{i \setminus j} = \gamma_i$ if and only if $\alpha_{j \setminus i} = \gamma_j$

Lemma 9. $m_{i \rightarrow j} = (w_{ij} - \gamma_i)_+$ holds $\forall (ij) \in E$

Proof. If $w_{ij} - \alpha_{i \setminus j} - \alpha_{j \setminus i} \leq 0$ then the result follows from Lemma 8. Otherwise, (ij) is strongly dotted and $\gamma_i = m_{j \rightarrow i} = (w_{ij} - \alpha_{j \setminus i} + \alpha_{i \setminus j})/2$, $\gamma_j = m_{i \rightarrow j} = (w_{ij} - \alpha_{i \setminus j} + \alpha_{j \setminus i})/2$. From here we can explicitly calculate $w_{ij} - \gamma_i = (w_{ij} - \alpha_{i \setminus j} + \alpha_{j \setminus i})/2 = m_{i \rightarrow j}$. \square

Lemma 10. *The unmatched balance property, equation (1), holds at every edge $(ij) \in E$, and both sides of the equation are non-negative.*

Proof. In light of lemma 9, (1) can be rewritten at a fixed point as

$$\gamma_i - \alpha_{i \setminus j} = \gamma_j - \alpha_{j \setminus i} \quad (31)$$

which is easy to verify. The case $\text{Surp}_{ij} \leq 0$ leads to both sides of Eq. (31) being 0 by Corollary 20. The other case $\text{Surp}_{ij} > 0$ leads to

$$m_{i \rightarrow j} - \alpha_{j \setminus i} = m_{j \rightarrow i} - \alpha_{i \setminus j} = \frac{\text{Surp}_{ij}}{2} \quad (32)$$

Clearly, we have $\gamma_i = m_{j \rightarrow i}$ and $\gamma_j = m_{i \rightarrow j}$. So Eq. (31) holds. \square

Next lemmas show that dotted edges are in correspondence with the solid edges that were defined in Section 8.

Lemma 11. *A non-solid edge cannot be a dotted edge, weak or strong.*

Before proving the lemma let us define alternating paths. A path $P = (i_1, i_2, \dots, i_k)$ in G is called *alternating path* if: (a) There exist a partition of edges of P into two sets A, B such that either $A \subset M^*$ or $B \subset M^*$. Moreover A (B) consists of all *odd* (*even*) edges; i.e. $A = \{(i_1, i_2), (i_3, i_4), \dots\}$ ($B = \{(i_2, i_3), (i_4, i_5), \dots\}$). (b) The path P might intersect itself or even repeat its own edges but no edge is repeated immediately. That is, for any $1 \leq r \leq k - 2$: $i_r \neq i_{r+1}$ and $i_r \neq i_{r+2}$. P is called an *alternating cycle* if $i_1 = i_k$.

Also, consider \underline{x}^* and \underline{y}^* that are optimum solutions for the LP and its dual, (2) and (3). The complementary slackness conditions (see [38]) for more details) state that for all $v \in V$, $y_v^*(\sum_{e \in \partial v} x_e^* - 1) = 0$ and for all $e = (ij) \in E$, $x_e^*(y_i^* + y_j^* - w_{ij}) = 0$. Therefore, for all solid edges the equality $y_i^* + y_j^* = w_{ij}$ holds. Moreover, any node $v \in V$ is adjacent to a solid edge if and only if $y_v^* > 0$.

Proof of Lemma 11. First, we refine the notion of solid edges by calling an edge e , $1-\underline{x}^*$ -solid ($\frac{1}{2}$ - \underline{x}^* -solid) whenever $x_e^* = 1$ ($x_e^* = \frac{1}{2}$).

We need to consider two cases:

Case (I). Assume that LP has an optimum solution \underline{x}^* that is integral as well (having a tight LP).

The idea of the proof is that if there exists a non-solid edge e which is dotted, we use a similar analysis to [6] to construct an alternating path consisting of dotted and \underline{x}^* -solid edges that leads to creation of at an optimal solution to LP (2) that assigns a positive value to e . This contradicts the non-solid assumption on e .

Now assume the contrary: take (i_1, i_2) that is a non-solid edge but it is dotted. Consider an endpoint of (i_1, i_2) . For example take i_2 . Either there is a \underline{x}^* -solid edge attached to i_2 or not. If there is not, we stop. Otherwise, assume (i_2, i_3) is a \underline{x}^* -solid edge. Using Lemma 7, either $\gamma_{i_3} = 0$ or there is a dotted edge connected to i_3 . But if this dotted edge is (i_2, i_3) then $P(i_2) \supseteq \{i_1, i_3\}$. Therefore, by Lemma 6 there has to be another dotted edge (i_3, i_4) connected to i_3 . Now, depending on whether i_4 has (has not) an adjacent \underline{x}^* -solid edge we continue (stop) the construction. A similar procedure could be done by starting at i_1 instead of i_2 . Therefore, we obtain an alternating path $P = (i_{-k}, \dots, i_{-1}, i_0, i_1, i_2, \dots, i_\ell)$ with all odd edges being dotted and all even edges being \underline{x}^* -solid. Using the same argument as in [6] one can show that one of the following four scenarios occur.

Path: Before P intersects itself, both end-points of the path stop. Either the last edge is \underline{x}^* -solid (then $\gamma_v = 0$ for the last node) or the last edge is a dotted edge. Now consider a new solution \underline{x}' to LP (2) by $x'_e = x_e^*$ if $e \notin P$ and $x'_e = 1 - x_e^*$ if $e \in P$. It is easy to see that \underline{x}' is a feasible LP solution at all points $v \notin P$ and also for internal vertices of P . The only nontrivial case is when $v = i_{-k}$ (or $v = i_\ell$) and the edge (i_{-k}, i_{-k+1}) (or $(i_{\ell-1}, i_\ell)$) is dotted. In both of these cases, by construction v is not connected to an \underline{x}^* -solid edge outside of P . Hence, making any change inside of P is safe. Now denote the weight of all solid (dotted) edges of P by $w(P_{\text{solid}})$ ($w(P_{\text{dotted}})$). Here, we only include edges outside P_{solid} in P_{dotted} . Clearly,

$$\sum_{e \in E} w_e x_e^* - \sum_{e \in E} w_e x'_e = w(P_{\text{solid}}) - w(P_{\text{dotted}}). \quad (33)$$

But $w(P_{\text{dotted}}) = \sum_{v \in P} \gamma_v$. Moreover, from Lemma 4, $\underline{\gamma}$ is dual feasible which gives $w(P_{\text{solid}}) \leq \sum_{v \in P} \gamma_v$. We are using the fact that if there is a \underline{x}^* -solid edge at an endpoint of P the γ of the endpoint should be 0. Now Eq. (33) reduces to $\sum_{e \in E} w_e x_e^* - \sum_{e \in E} w_e x'_e \leq 0$. This contradicts that $e = (i_1, i_2)$ is non-solid since $x'_e > 0$.

Cycle: P intersects itself and will contain an even cycle C_{2s} . This case can be handled very similar to the path by defining $x'_e = x_e^*$ if $e \notin C_{2s}$ and $x'_e = 1 - x_e^*$ if $e \in C_{2s}$. The proof is even simpler since the extra check for the boundary condition is not necessary.

Blossom: P intersects itself and will contain an odd cycle C_{2s+1} with a path (stem) P' attached to the cycle at point u . In this case let $x'_e = x_e^*$ if $e \notin P' \cup C_{2s+1}$, and $x'_e = 1 - x_e^*$ if $e \in P'$, and $x'_e = \frac{1}{2}$ if $e \in C_{2s+1}$. From here, we drop the subindex $2s+1$ to simplify the notation. Since the cycle has odd length, both neighbors of u in C have to be dotted. Therefore,

$$\begin{aligned} & \sum_{e \in E} w_e x_e^* - \sum_{e \in E} w_e x'_e \\ &= w(P'_{\text{solid}}) + w(C_{\text{solid}}) - w(P'_{\text{dotted}}) \\ & \quad - \frac{w(C_{\text{dotted}}) + w(C_{\text{solid}})}{2} \\ &= w(P'_{\text{solid}}) + \frac{w(C_{\text{solid}})}{2} - w(P'_{\text{dotted}}) - \frac{w(C_{\text{dotted}})}{2}. \end{aligned}$$

Plugging $w(P'_{\text{solid}}) \leq \sum_{v \in P'} \gamma_v$, $w(C_{\text{solid}}) \leq \sum_{v \in C} \gamma_v - \gamma_u$, $w(P'_{\text{dotted}}) = \sum_{v \in P'} \gamma_v - \gamma_u$ and $w(C_{\text{dotted}}) = \sum_{v \in C} \gamma_v + \gamma_u$, we obtain

$$\sum_{e \in E} w_e x_e^* - \sum_{e \in E} w_e x'_e \leq 0,$$

which is again a contradiction.

Bicycle: P intersects itself at least twice and will contain two odd cycles C_{2s+1} and $C'_{2s'+1}$ with a path (stem) P' that is connecting them. Very similar to Blossom, let $x'_e = x_e^*$ if $e \notin P' \cup C \cup C'$, $x'_e = 1 - x_e^*$ if $e \in P'$, and $x'_e = \frac{1}{2}$ if $e \in C \cup C'$. The proof follows similar to the case of blossom.

Case (II). Assume that there is an optimum solution \underline{x}^* of LP that is not necessarily integral.

Everything is similar to Case (I) but the algebraic treatments are slightly different. Some edges e in P can be $\frac{1}{2}$ - \underline{x}^* -solid ($x_e^* = \frac{1}{2}$). In particular some of the odd edges (dotted edges) of P can now be $\frac{1}{2}$ - \underline{x}^* -solid. But the subset of $\frac{1}{2}$ - \underline{x}^* -solid edges of P can be only sub-paths of odd length in P . On each such sub-path defining $\underline{x}' = 1 - \underline{x}^*$ means we are not affecting \underline{x}^* . Therefore, all of the algebraic calculations should be considered on those sub-paths of P that have no $\frac{1}{2}$ - \underline{x}^* -solid edge which means both of their boundary edges are dotted.

Path: Define \underline{x}' as in Case (I). Using the discussion above, let $P^{(1)}, \dots, P^{(r)}$ be disjoint sub-paths of P that have no $\frac{1}{2}$ - \underline{x}^* -solid edge. Thus, $\sum_{e \in E} w_e x_e^* - \sum_{e \in E} w_e x'_e = \sum_{i=1}^r [w(P_{\text{solid}}^{(i)}) - w(P_{\text{dotted}}^{(i)})]$. Since in each $P^{(i)}$ the two boundary edges are dotted, $w(P_{\text{solid}}^{(i)}) \leq \sum_{v \in P^{(i)}} \gamma_v$ and $\sum_{v \in P^{(i)}} \gamma_v = w(P_{\text{dotted}}^{(i)})$. The rest can be done as in Case (I).

Cycle, Blossom, Bicycle: These cases can be done using the same method of breaking the path and cycles into sub-paths $P^{(i)}$ and following the case of path. □

Lemma 12. *Every strong-solid edge is a strong-dotted edge. Also, every weak-solid edge is a weak-dotted edge.*

Proof. We rule out all alternative cases one by one. In particular we prove:

(i) *A strong-solid edge cannot be weak-dotted.* If an edge (i, j) is strong-solid then it cannot be adjacent to another solid edge (weak or strong). Therefore, using Lemma 11 none of adjacent edges to (i, j) are dotted. However, if (i, j) is weak-dotted by Lemma 6 it is adjacent to at least one other weak-dotted edge (since at least one of γ_i and γ_j is positive) which is a contradiction. Thus (i, j) cannot be weak-dotted.

(ii) *A strong-solid edge cannot be non-dotted.* Similar to (i), if an edge (i, j) is strong-solid it cannot be adjacent to dotted edges. Now, if (i, j) is non-dotted then $\gamma_i = \gamma_j = 0$ using Lemma 7. Hence $w_{ij} < \gamma_i + \gamma_j = 0$ which is contradiction since we assumed all weights are positive.

(iii) *A weak-solid edge cannot be strong-dotted.* Assume, (i_1, i_2) is weak-solid and strong-dotted. Then we can show an optimum to LP (2) can be improved which is a contradiction. The proof is very similar to proof of Lemma 11. Since (i_1, i_2) is weak-solid, there is a half-integral matching \underline{x}^* that is optimum to LP and puts a mass $1/2$ or 0 on (i_1, i_2) . Then either there is an adjacent \underline{x}^* -solid edge (i_2, i_3) or an adjacent \underline{x}^* -solid edge (i_0, i_1) with mass at least $1/2$ or we stop. In the latter case, increasing the value of $x_{i_1 i_2}^*$ increases $\sum_{e \in E} w_e x_e^*$ while keeping it LP feasible which is a contradiction. Otherwise, by strong-dotted assumption on (i_1, i_2) ((i_0, i_1)), the new edge (i_2, i_3) is not dotted. Now we select a dotted edge (i_3, i_4) if it exists (otherwise we stop and in that case $\gamma_{i_3} = 0$). This process is repeated as in proof of Lemma 11 in both directions to obtain an alternating path $P = (i_{-k}, \dots, i_{-1}, i_0, i_1, i_2, \dots, i_\ell)$ with all odd edges being dotted with \underline{x}^* value at most $1/2$ and all even edges being \underline{x}^* -solid with mass at least $1/2$. We discuss the case of P being a simple path (not intersecting itself) here, and other cases: cycle, bicycle and blossom can be treated similar to path as in proof of Lemma 11.

Construct LP solution \underline{x}' that is equal to \underline{x}^* outside of P and inside it satisfies $x'_e = x_e^* + 1/2$ if e is an odd edge that is $e = (i_{2k-1}, i_{2k})$, and $x'_e = x_e^* - 1/2$ when e is an even edge that is $e = (i_{2k}, i_{2k+1})$. It is easy to see that \underline{x}' is a feasible LP solution. And since for all edges (i_j, i_{j+1}) we have $\gamma_{i_j} + \gamma_{i_{j+1}} \geq w_{i_j i_{j+1}}$ and on dotted edges we have equality $\gamma_{i_j} + \gamma_{i_{j+1}} = w_{i_j i_{j+1}}$ then

$\sum_{e \in E} w_e x_e^* - \sum_{e \in E} w_e x'_e = \frac{w(P_{\text{dotted}}) - w(P_{\text{solid}})}{2} \geq \frac{\gamma_{i_2} + \gamma_{i_3} - w_{i_2 i_3}}{2} > 0$ where the last inequality follows from the fact that (i_2, i_3) is not-dotted. Hence we reach a contradiction.

(iv) *A weak-solid edge cannot be non-dotted.* Assume, (i_1, i_2) is weak-solid and non-dotted. Similar to (iii) we can show the best solution to LP (2) can be improved which is a contradiction. Since (i_1, i_2) is weak-solid we can choose a half-integral \underline{x}^* that puts a mass at least 1/2 on (i_1, i_2) . Also, this time the alternation in P is the opposite of (iii). That is we choose (i_2, i_3) to be dotted (if it does not exist $\gamma_{i_2} = 0$ and we stop.) The solution \underline{x}' is constructed as before: equal to \underline{x}^* outside of P , $x'_e = x_e^* + 1/2$ if e is odd and $x'_e = x_e^* - 1/2$ if it is even. Hence, $\sum_{e \in E} w_e x_e^* - \sum_{e \in E} w_e x'_e \geq \frac{\gamma_{i_1} + \gamma_{i_2} - w_{i_1 i_2}}{2} > 0$, using the non-dotted assumption on (i_1, i_2) . Hence, we obtain another contradiction. \square

Lemma 13. $\underline{\gamma}$ is an optimum for the dual problem (3)

Proof. Lemma 4 guarantees feasibility. Optimality follows from lemmas 7, 11 and 12 as follows. Take any optimum half integral matching \underline{x}^* to LP. Now using Lemma 12: $\sum_v \gamma_v = \sum_{e \in E} w_e x_e^*$ which finishes the proof. \square

Theorem 12. Let \mathcal{BALOPT} be the set of optima of the dual problem (3) satisfying the unmatched balance property, Eq. (1), at every edge. If $(\underline{\alpha}, \underline{m}, \underline{\gamma})$ is a fixed point of the natural dynamics then $\underline{\gamma} \in \mathcal{BALOPT}$. Conversely, for every $\underline{\gamma}_{\text{BO}} \in \mathcal{BALOPT}$, there is a unique fixed point of the natural dynamics with $\underline{\gamma} = \underline{\gamma}_{\text{BO}}$.

Proof. The direct implication is immediate from Lemmas 10 and 13. The converse proof here follows the same steps as for Theorem 1, proved in Section 8. Instead of separately analyzing the cases $(ij) \in M$ and $(ij) \notin M$, we study the cases $\gamma_i + \gamma_j = w_{ij}$ and $\gamma_i + \gamma_j > w_{ij}$. \square

D ϵ -fixed point properties: Proof of Theorem 4

In this section we prove Theorem 4, stated in Section 3. In this section we assume that $\underline{\alpha}$ is an ϵ -fixed point with corresponding offers \underline{m} and earnings $\underline{\gamma}$. That is, for all i, j

$$\begin{aligned} \epsilon &\geq |\alpha_{i \setminus j} - \max_{k \in \partial i \setminus j} m_{k \rightarrow i}|, \\ m_{i \rightarrow j} &= (w_{ij} - \alpha_{i \setminus j})_+ - \frac{(w_{ij} - \alpha_{i \setminus j} - \alpha_{j \setminus i})_+}{2}, \\ \gamma_i &= \max_{k \in \partial i} m_{k \rightarrow i}. \end{aligned}$$

Definition 21. An edge (ij) is called δ -dotted ($\delta \geq 0$) if $\gamma_i + \gamma_j \leq w_{ij} + \delta$.

Lemma 14. For all edge $(ij) \in E$ and all $\delta, \delta_1, \delta_2 \in \mathbb{R}$ the following hold:

- (a) If (ij) is δ -dotted then $\text{Surp}_{ij} \geq -(2\epsilon + \delta)$.
- (b) If $\text{Surp}_{ij} \geq -\delta$ then $m_{i \rightarrow j} \geq \gamma_j - (\epsilon + \delta)$ and $m_{j \rightarrow i} \geq \gamma_i - (\epsilon + \delta)$.
- (c) If $m_{i \rightarrow j} \geq \gamma_j - \delta_1$ and $m_{j \rightarrow i} \geq \gamma_i - \delta_2$ then (ij) is $(\delta_1 + \delta_2)$ -dotted.
- (d) If $\gamma_i - \delta \leq m_{j \rightarrow i}$ and $\gamma_j > m_{i \rightarrow j} + 2\epsilon + \delta$ then $\gamma_i = 0$.
- (e) If $\gamma_i > 0$ and $m_{j \rightarrow i} \geq \gamma_i - \delta$ then (ij) is $(2\delta + 2\epsilon)$ -dotted.
- (f) For $\gamma_i, \gamma_j > 0$, $m_{j \rightarrow i} \leq \alpha_{i \setminus j} + \delta$ if and only if $m_{i \rightarrow j} \leq \alpha_{j \setminus i} + \delta$.
- (h) For all (ij) , $|m_{i \rightarrow j} - (w_{ij} - \gamma_i)_+| \leq \epsilon$.
- (i) For all (ij) , $\gamma_i - (w_{ij} - \gamma_j)_+ \geq -\epsilon$ and $\gamma_i + \gamma_j \geq w_{ij} - \epsilon$.
- (j) For all i , if $\gamma_i > 0$ then there is at least a 2ϵ -dotted edge attached to i .

Proof. (a) Since $\underline{\alpha}$ is ϵ -fixed point, $\gamma_i \geq m_{i \rightarrow j} - \epsilon$ and $\gamma_j \geq m_{j \rightarrow i} - \epsilon$. Therefore, $\text{Surp}_{ij} = w_{ij} - m_{i \rightarrow j} - m_{j \rightarrow i} \geq w_{ij} - \gamma_i - \gamma_j - (2\epsilon) \geq -(2\epsilon + \delta)$.

(b) First consider the case $\text{Surp}_{ij} \leq 0$. Then, $m_{i \rightarrow j} = (w_{ij} - \alpha_{i \setminus j})_+ \geq w_{ij} - \alpha_{i \setminus j} \geq \alpha_{j \setminus i} - \delta \geq \max_{\ell \in \partial j \setminus i}(m_{\ell \rightarrow j}) - \delta - \epsilon$, which yields $m_{i \rightarrow j} \geq \gamma_j - (\epsilon + \delta)$. The proof of $m_{j \rightarrow i} \geq \gamma_i - (\epsilon + \delta)$ is similar.

For the case $\text{Surp}_{ij} > 0$, $m_{i \rightarrow j} = \frac{w_{ij} - \alpha_{i \setminus j} + \alpha_{j \setminus i}}{2} = \frac{\text{Surp}_{ij}}{2} + \alpha_{i \setminus j} \geq \max(\frac{-\delta}{2}, 0) + \max_{\ell \in \partial j \setminus i}(m_{\ell \rightarrow j}) - \epsilon$, and the rest follows as above.

(c) Note that $\gamma_i + \gamma_j \leq m_{i \rightarrow j} + m_{j \rightarrow i} + \delta_1 + \delta_2$. If $\text{Surp}_{ij} \geq 0$ then the result follows from $m_{i \rightarrow j} + m_{j \rightarrow i} = w_{ij}$. For $\text{Surp}_{ij} < 0$ the result follows from $m_{i \rightarrow j} + m_{j \rightarrow i} \leq \max[(w_{ij} - \alpha_{i \setminus j})_+, (w_{ij} - \alpha_{j \setminus i})_+, 2w_{ij} - \alpha_{i \setminus j} - \alpha_{j \setminus i}] \leq w_{ij}$.

(d) We need to show that when $\gamma_i \leq m_{j \rightarrow i} + \delta$ and $\gamma_j > m_{i \rightarrow j} + 2\epsilon + \delta$ then $\gamma_i = 0$. From part (b) that was just shown, the surplus should satisfy $\text{Surp}_{ij} < -(\epsilon + \delta)$. On the other hand $\alpha_{i \setminus j} - \epsilon \leq \max_{k \in \partial i \setminus j}(m_{k \rightarrow i}) \leq \gamma_i \leq m_{j \rightarrow i} + \delta \leq (w_{ij} - \alpha_{i \setminus j})_+ + \delta$. Now, if $\gamma_i > 0$ then $w_{ij} - \alpha_{i \setminus j} > 0$ which gives $\alpha_{i \setminus j} - \epsilon \leq w_{ij} - \alpha_{i \setminus j} + \delta$. This is equivalent to $\text{Surp}_{ij} \geq -(\epsilon + \delta)$ which is a contradiction. Hence $\gamma_i = 0$.

(e) Using part (d) we should have $m_{i \rightarrow j} \geq \gamma_j - (2\epsilon + \delta)$. Now applying part (c) the result follows.

(f) If $\text{Surp}_{ij} \geq 0$ then $\frac{w_{ij} - \alpha_{j \setminus i} + \alpha_{i \setminus j}}{2} = m_{j \rightarrow i} \leq \alpha_{i \setminus j} + \delta$. This inequality is equivalent to $m_{i \rightarrow j} = \frac{w_{ij} - \alpha_{i \setminus j} + \alpha_{j \setminus i}}{2} \leq \alpha_{j \setminus i} + \delta$, which proves the result. If $\text{Surp}_{ij} < 0$ then $w_{ij} - \alpha_{j \setminus i} \leq (w_{ij} - \alpha_{j \setminus i})_+ \leq \alpha_{i \setminus j} + \delta$. This is equivalent to $w_{ij} - \alpha_{i \setminus j} \leq \alpha_{j \setminus i} + \delta$ which yields the result.

(h) If $\text{Surp}_{ij} \geq 0$ then by part (b), $m_{i \rightarrow j} + \epsilon \geq \gamma_j$ and $m_{j \rightarrow i} + \epsilon \geq \gamma_i$. Therefore, using $\gamma_j \geq m_{i \rightarrow j}$, $\gamma_i \geq m_{j \rightarrow i}$ and $m_{j \rightarrow i} + m_{i \rightarrow j} = w_{ij}$ we have, $m_{i \rightarrow j} \geq w_{ij} - m_{j \rightarrow i} \geq w_{ij} - \gamma_i \geq w_{ij} - m_{j \rightarrow i} - \epsilon \geq m_{i \rightarrow j} - \epsilon$, which gives the result.

If $\text{Surp}_{ij} < 0$ then $m_{i \rightarrow j} = (w_{ij} - \alpha_{i \setminus j})_+ < \alpha_{j \setminus i}$ this gives $\gamma_j - \epsilon < \alpha_{j \setminus i}$. On the other hand $\alpha_{j \setminus i} \leq \gamma_j + \epsilon$ holds. Similarly, $\gamma_i + \epsilon \geq \alpha_{i \setminus j} \geq \gamma_i - \epsilon$ that leads to $|(w_{ij} - \alpha_{i \setminus j})_+ - (w_{ij} - \gamma_i)_+| \leq \epsilon$. Hence, the result follows from $m_{i \rightarrow j} = (w_{ij} - \alpha_{i \setminus j})_+$.

(i) Using part (h), $m_{j \rightarrow i} + \epsilon \geq (w_{ij} - \gamma_j)_+$. Now result follows using $\gamma_i \geq m_{j \rightarrow i}$.

(j) There is at least one neighbor $j \in \partial i$ that sends the maximum offer $m_{j \rightarrow i} = \gamma_i$. Using part (d) we should have $m_{i \rightarrow j} \geq \gamma_j - 2\epsilon$ and now the result follows from part (c). \square

Lemma 15. *For any edge $(ij) \in E$ the earnings estimate $\underline{\gamma}$ satisfies 6ϵ -balanced property (i.e., Eq. (9) holds for 6ϵ instead of ϵ).*

Proof. Using Lemma 14(h), $\alpha_{i \setminus j} - 2\epsilon \leq \max_{k \in \partial i \setminus j}(m_{k \rightarrow i}) - \epsilon \leq \max_{k \in \partial i \setminus j}[(w_{ik} - \gamma_k)_+] \leq \max_{k \in \partial i \setminus j}(m_{k \rightarrow i}) + \epsilon \leq \alpha_{i \setminus j} + 2\epsilon$, or

$$\left| \max_{k \in \partial i \setminus j} [(w_{ik} - \gamma_k)_+] - \alpha_{i \setminus j} \right| \leq 2\epsilon \quad (34)$$

Now, if $\text{Surp}_{ij} \leq 0$ then $m_{j \rightarrow i} = (w_{ij} - \alpha_{j \setminus i})_+ \leq \alpha_{i \setminus j}$ which gives $|\gamma_i - \alpha_{i \setminus j}| \leq \epsilon$ or, $|\gamma_i - \max_{k \in \partial i \setminus j} [(w_{ik} - \gamma_k)_+]| \leq 3\epsilon$. Therefore, 6ϵ -balance property holds.

And if $\text{Surp}_{ij} > 0$, by Lemma 14(b) we have $m_{j \rightarrow i} + \epsilon \geq \gamma_i$. Hence, $\frac{\text{Surp}_{ij}}{2} + \epsilon = m_{j \rightarrow i} - \alpha_{i \setminus j} + \epsilon \geq \gamma_i - \alpha_{i \setminus j} \geq m_{j \rightarrow i} - \alpha_{i \setminus j} = \frac{\text{Surp}_{ij}}{2}$. Same bound holds for $\gamma_j - \alpha_{j \setminus i}$ by symmetry. Therefore, using Eq. (34), $|\gamma_i - \max_{k \in \partial i \setminus j} [(w_{ik} - \gamma_k)_+]|$ and $|\gamma_j - \max_{\ell \in \partial j \setminus i} [(w_{j\ell} - \gamma_\ell)_+]|$ are within $3\epsilon \leq 6\epsilon$ of each other. \square

Lemma 16. *If (ij) is δ -dotted for $k \in \partial i \setminus j$ and if $\gamma_k > \max(\delta, \epsilon) + 6\epsilon$, then there exists $r \in \partial k \setminus i$ such that (rk) is $(\max(\delta, \epsilon) + 6\epsilon)$ -dotted.*

Proof. Using, $\gamma_i + \gamma_j \leq w_{ij} + \delta$ and Lemma 14(i),

$$-\epsilon \leq \gamma_i - \max_{s \in \partial i \setminus k} [(w_{is} - \gamma_s)_+] \leq \gamma_i - (w_{ij} - \gamma_j)_+ \leq \delta.$$

Therefore, $|\gamma_i - \max_{s \in \partial i \setminus k} [(w_{is} - \gamma_s)_+]| \leq \max(\delta, \epsilon)$ which combined with Lemma 15 gives

$$|\gamma_k - \max_{r \in \partial k \setminus i} [(w_{rk} - \gamma_r)_+]| \leq \max(\delta, \epsilon) + 6\epsilon.$$

This fact and $\gamma_k > \max(\delta, \epsilon) + 6\epsilon$, show that there exists an edge $r \in \partial k \setminus i$ with $|\gamma_k - (w_{rk} - \gamma_r)_+| \leq \max(\delta, \epsilon) + 6\epsilon$ and the result follows. \square

Lemma 17. *A non-solid edge cannot be a δ -dotted edge for $\delta \leq 4\epsilon$.*

Note that this Lemma holds even for the more general case of M^* being non-integral.

The proof is a more complex version of proof of Lemma 11. Recall the notion of alternating path from that proof.

Also, consider \underline{x}^* and \underline{y}^* that are optimum solutions for the LP and its dual, (2) and (3). Also recall that by complementary slackness conditions, for all solid edges the equality $y_i^* + y_j^* = w_{ij}$ holds. Moreover, any node $v \in V$ is adjacent to a solid edge if and only if $y_v^* > 0$.

Proof of Lemma 17. We need to consider two cases:

Case (I). Assume that the optimum LP solution \underline{x}^* is integral (having a tight LP). Now assume the contrary: take (i_1, i_2) that is a non-solid edge but it is δ -dotted. Consider an endpoint of (i_1, i_2) . For example take i_2 . Either there is a solid edge attached to i_2 or not. If there is not, we stop. Otherwise, assume (i_2, i_3) is a solid edge. Using Lemma 16, either $\gamma_{i_3} > 10\epsilon$ or there is a 10ϵ -dotted edge (i_3, i_4) connected to i_3 . Now, depending on whether i_4 has (has not) an adjacent solid edge we continue (stop) the construction. Similar procedure could be done by starting at i_1 instead of i_2 . Therefore, we obtain an alternating path $P = (i_{-k}, \dots, i_{-1}, i_0, i_1, i_2, \dots, i_\ell)$ with each (i_{2k}, i_{2k+1}) being $(6k + 4)\epsilon$ -dotted and all (i_{2k-1}, i_{2k}) being solid. Using the same argument as in [6] one can show that one of the following four scenarios occur.

Path: Before P intersects itself, both end-points of the path stop. At each end of the path, either the last edge is solid (then $\gamma_v < (3n + 4)\epsilon$ for the last node v) or the last edge is a $(3n + 4)$ -dotted edge with no solid edge attached to v . Now consider a new solution \underline{x}' to LP (2) by $x'_e = x_e^*$ if $e \notin P$ and $x'_e = 1 - x_e^*$ if $e \in P$. It is easy to see that \underline{x}' is a feasible LP solution at all points $v \notin P$ and also for internal vertices of P . The only nontrivial case is when $v = i_{-k}$ (or $v = i_\ell$) and the edge (i_{-k}, i_{-k+1}) (or $(i_{\ell-1}, i_\ell)$) is $(3n + 4)\epsilon$ -dotted. In both of these cases, by construction no solid edge is attached to v outside of P so making any change inside of P is safe. Now denote the weight of all solid (remaining) edges of P by $w(P_{\text{solid}})$ ($w(P_{\text{dotted}})$). Hence, $\sum_{e \in E} w_e x_e^* - \sum_{e \in E} w_e x'_e = w(P_{\text{solid}}) - w(P_{\text{dotted}})$.

But $w(P_{\text{dotted}}) + (3n^2 + 16n)\epsilon/4 \geq \sum_{v \in P} \gamma_v$. Moreover, from Lemma 14(i), $\gamma_i + \gamma_j \geq w_{ij} - \epsilon$ for all $(ij) \in P$ which gives $w(P_{\text{solid}}) \leq \sum_{v \in P} \gamma_v + n\epsilon/2$. Now $\sum_{e \in E} w_e x_e^* - \sum_{e \in E} w_e x'_e = w(P_{\text{solid}}) - w(P_{\text{dotted}})$ yields $w_e x_e^* - \sum_{e \in E} w_e x'_e \leq (3n^2 + 18n)\epsilon/4 \leq n(n + 5)\epsilon$. For $\epsilon < g/(6n^2)$ This contradicts the tightness of LP relaxation (2) since $x'_e \neq x_e^*$ holds at least for $e = (i_1, i_2)$.

Cycle: P intersects itself and will contain an even cycle C_{2s} . This case can be handled very similar to the path by defining $x'_e = x_e^*$ if $e \notin C_{2s}$ and $x'_e = 1 - x_e^*$ if $e \in C_{2s}$. The proof is even simpler since the extra check for the boundary condition is not necessary.

Blossom: P intersects itself and will contain an odd cycle C_{2s+1} with a path (stem) P' attached to the cycle at point u . In this case let $x'_e = x_e^*$ if $e \notin P' \cup C_{2s+1}$, and $x'_e = 1 - x_e^*$ if $e \in P'$, and

$x'_e = \frac{1}{2}$ if $e \in C_{2s+1}$. From here, we drop the subindex $2s + 1$ to simplify the notation. Since the cycle has odd length, both neighbors of u in C have to be dotted. Therefore,

$$\begin{aligned} \sum_{e \in E} w_e x_e^* - \sum_{e \in E} w_e x'_e &= w(P'_{\text{solid}}) + w(C_{\text{solid}}) - w(P'_{\text{dotted}}) - \frac{w(C_{\text{dotted}}) + w(C_{\text{solid}})}{2}, \\ &= w(P'_{\text{solid}}) + \frac{w(C_{\text{solid}})}{2} - w(P'_{\text{dotted}}) - \frac{w(C_{\text{dotted}})}{2} \\ &< \sum_{v \in P'} \gamma_v + \left\lceil \frac{|P|}{2} \right\rceil \epsilon + \frac{\sum_{v \in C} \gamma_v - \gamma_u}{2} + s\epsilon - \sum_{v \in P'} \gamma_v + \gamma_u \\ &\quad + \left(\frac{3|P|^2 + 16|P|}{4} \right) \epsilon - \frac{\sum_{v \in C} \gamma_v + \gamma_u}{2} + \left(\frac{3s^2 + 16s}{4} \right) \epsilon. \end{aligned}$$

But the last term is at most $n(n + 5)\epsilon$ which is again a contradiction.

Bicycle: P intersects itself at least twice and will contain two odd cycles C_{2s+1} and $C'_{2s'+1}$ with a path (stem) P' that is connecting them. Very similar to Blossom, let $x'_e = x_e^*$ if $e \notin P' \cup C \cup C'$, $x'_e = 1 - x_e^*$ if $e \in P'$, and $x'_e = \frac{1}{2}$ if $e \in C \cup C'$. The proof follows similar to the case of blossom.

Case (II). Assume that the optimum LP solution \underline{x}^* is not necessarily integral.

Everything is similar to Case (I) but the algebraic treatments are slightly different. Some edges e in P can be $\frac{1}{2}$ -solid ($x_e^* = \frac{1}{2}$). In particular some of the odd edges (dotted edges) of P can now be $\frac{1}{2}$ -solid. But the subset of $\frac{1}{2}$ -solid edges of P can be only sub-paths of odd length in P . On each such sub-path defining $\underline{x}' = 1 - \underline{x}^*$ means we are not affecting \underline{x}^* . Therefore, all of the algebraic calculations should be considered on those sub-paths of P that have no $\frac{1}{2}$ -solid edge which means both of their boundary edges are dotted.

Path: Define \underline{x}' as in Case (I). Using the discussion above, let $P^{(1)}, \dots, P^{(r)}$ be disjoint sub-paths of P that have no $\frac{1}{2}$ -solid edge. Thus, $\sum_{e \in E} w_e x_e^* - \sum_{e \in E} w_e x'_e = \sum_{i=1}^r [w(P^{(i)}_{\text{solid}}) - w(P^{(i)}_{\text{dotted}})]$. Since in each $P^{(i)}$ the two boundary edges are dotted, $w(P^{(i)}_{\text{solid}}) \leq \sum_{v \in P^{(i)}} \gamma_v + |P^{(i)}|\epsilon/2$ and $\sum_{v \in P^{(i)}} \gamma_v \leq w(P^{(i)}_{\text{dotted}}) + (3|P^{(i)}|^2 + 16|P^{(i)}|)\epsilon/4$. The rest can be done as in Case (I).

Cycle, Blossom, Bicycle: These cases can be done using the same method of breaking the path and cycles into sub-paths $P^{(i)}$ and following the case of path. □

The direct part of Theorem 4 follows from the next lemma.

Lemma 18. $\underline{\alpha}$ induces the matching M^* .

Proof. From Lemma 17 it follows that the set of 2ϵ -dotted edges is a subset of the solid edges. In particular, when the optimum matching M^* is integral, no node can be adjacent to more than one 2ϵ -dotted edges. If we define a \underline{x}' to be zero on all edges and $x'_e = 1$ for all 2ϵ -dotted edges (ij) with $\gamma_i + \gamma_j > 0$. Then clearly \underline{x}' is feasible to (2). On the other hand, using the definition of 2ϵ -dotted for all e' with $x_{e'} = 1$, and Lemma 14(j) that each node with $\gamma_i > 0$ is adjacent to at least one 2ϵ -dotted edge we can write $\sum_{e \in E} w_e x'_e \geq \sum_{v \in V} \gamma_v - n\epsilon$. Separately, from Lemma 14(i) we have $\sum_{v \in V} \gamma_v \geq \sum_{e \in E} w_e x_e^* - \frac{n\epsilon}{2}$, which shows that \underline{x}' is also an optimum solution to (2) (when $\epsilon < g/(6n^2)$). From the uniqueness assumption on \underline{x}^* we obtain that M^* is equal to the set of all 2ϵ -dotted edges with at least one endpoint having a positive earning estimate. We would like to show that for any such edge (ij) , both earning estimates γ_i and γ_j are positive.

Assume the contrary, i.e., without loss of generality $\gamma_i = 0$. Then, $\text{Surp}_{ij} \leq 0$ and $0 = m_{j \rightarrow i} = (w_{ij} - \alpha_{j \setminus i})_+$ that gives $\alpha_{j \setminus i} \geq w_{ij}$ or

$$m_{\ell \rightarrow j} \geq \alpha_{j \setminus i} - \epsilon \geq w_{ij} - \epsilon \geq (w_{ij} - \alpha_{i \setminus j})_+ - \epsilon = \gamma_j - \epsilon.$$

for some $\ell \in \partial j \setminus i$. Now using Lemma 14(e) the edge $(j\ell)$ is 4ϵ -dotted which contradicts Lemma 17.

Finally, the endpoints of the matched edges provide each other their unique best offers. This latter follows from the fact that each node with $\gamma_i > 0$ receives an offer equal to γ_i and the edge corresponding to that offer has to be 2ϵ -dotted using Lemma 14(d). The nodes with no positive offer $\gamma_i = 0$ are unmatched in M^* as well. \square

Proof of Theorem 4.

Proof. For any $\epsilon < g/(6n^2)$, an ϵ -fixed point induces the matching M^* using Lemma 18. Additionally, the earning vector $\underline{\gamma}$ is (6ϵ) -balanced using Lemma 15. Next we show that $(\underline{\gamma}, M^*)$ is a stable trade outcome.

Lemma 19. *The earnings estimates $\underline{\gamma}$ is an optimum solution to the dual (3). In particular the pair $(\underline{\gamma}, M^*)$ is a stable trade outcome.*

Proof. Using Lemma 17, we can show that for any non-solid edge (ij) , stability holds, i.e. $\gamma_i + \gamma_j \geq w_{ij}$.

Now let (i, j) be a solid edge. Then i and j are sending each other their best offers. If $\text{Surp}_{ij} \geq 0$ we are done using $\gamma_i + \gamma_j = m_{j \rightarrow i} + m_{i \rightarrow j} = \frac{w_{ij} - \alpha_{i \setminus j} + \alpha_{j \setminus i}}{2} + \frac{w_{ij} - \alpha_{j \setminus i} + \alpha_{i \setminus j}}{2} = w_{ij}$. And if $\text{Surp}_{ij} < 0$ then $\gamma_i = m_{j \rightarrow i} = (w_{ij} - \alpha_{j \setminus i})_+ \leq \alpha_{i \setminus j}$. Similarly, $\gamma_j \leq \alpha_{j \setminus i}$. This means there exist $k \in \partial i \setminus j$ with $m_{k \rightarrow i} \geq \alpha_{i \setminus j} - \epsilon \geq \gamma_i - \epsilon$. But, from Lemma 14(e) the edge (ik) would become 4ϵ -dotted which is a contradiction. \square

The converse of Theorem 4 is trivial since any ϵ -NB solution $(M, \underline{\gamma}_{\text{NB}})$ is stable and produces a trade outcome by definition, hence it is a dual optimal solution which means $M = M^*$. \square

E Proof of Theorem 8

Theorem 13. *Let $G = (V, E)$ with edge weights $(w_{ij})_{(ij) \in E}$ and capacity constraints $\mathbf{b} = (b_i)$ be an instance such that the primal LP (16) has a unique optimum that is integral, corresponding to matching M^* . Let $(\underline{\alpha}, \underline{m}, \Gamma)$ be a fixed point of the natural dynamics. Then $\underline{\alpha}$ induces matching M^* and (M^*, Γ) is a Nash bargaining solution. Conversely, every Nash bargaining solution (M, Γ_{NB}) has $M = M^*$ and corresponds to a unique fixed point of the natural dynamics with $\Gamma = \Gamma_{\text{NB}}$.*

Let \mathcal{S} be the set of optimum solutions of LP (16). As in the one-matching case, we call $e \in E$ a *strong-solid edge* if $x_e^* = 1$ for all $x^* \in \mathcal{S}$ and a *non-solid edge* if $x_e^* = 0$ for all $x^* \in \mathcal{S}$. We call $e \in E$ a *weak-solid edge* if it is neither strong-solid nor non-solid.

Proof of Theorem 8: From fixed points to NB solutions. The direct part follows from the following set of fixed point properties, similar to those for the one-matching case. Throughout $(\underline{\alpha}, \underline{m}, \Gamma)$ is a fixed point of the dynamics (17) (with Γ given by (18), and \underline{m} given by (4)). The properties are proved for the case when the primal LP in (16) has a unique integral optimum (which implies that there are no weak-solid edges).

(1) Two players $(i, j) \in E$ are called *partners* if $\gamma_i + \gamma_j \leq w_{ij}$. Then the following are equivalent: (a) i and j are partners, (b) $w_{ij} - \alpha_{i \setminus j} - \alpha_{j \setminus i} \geq 0$, (c) $\gamma_i \leq m_{j \rightarrow i}$ and $\gamma_j \leq m_{i \rightarrow j}$.

(2) The following are equivalent: (a) $w_{ij} - \alpha_{i \setminus j} - \alpha_{j \setminus i} > 0$, (b) $\gamma_{j \rightarrow i} > \alpha_{i \setminus j} = (b_i^{\text{th-max}})_{k \in \partial i \setminus j} m_{k \rightarrow i}$. Denote this set of edges by M .

(3) We say that (i, j) is a *weak-dotted edge* if $w_{ij} - \alpha_{i \setminus j} - \alpha_{j \setminus i} = 0$, a *strong-dotted edge* if $w_{ij} - \alpha_{i \setminus j} - \alpha_{j \setminus i} > 0$, and a *non-dotted edge* otherwise. If i has less than b_i adjacent dotted edges, then $\gamma_i = 0$.

(4) Each strong solid edge is strong dotted, and each non-solid edge is non-dotted.

(5) The balance property (15), holds at every edge $(i, j) \in M$.

(6) We have

$$m_{i \rightarrow j} = \begin{cases} \gamma_{i \rightarrow j} & \text{for } (ij) \in M, \\ (w_{ij} - \gamma_i)_+ & \text{for } (ij) \notin M. \end{cases}$$

(7) An optimum solution for the dual LP in (16) can be constructed as $y_i = \gamma_i$ for all $i \in V$ and:

$$y_{ij} = \begin{cases} w_{ij} - \gamma_i - \gamma_j & \text{for } (ij) \in M, \\ 0 & \text{for } (ij) \notin M. \end{cases}$$

Proof of Theorem 8, direct implication. Assume that the primal LP in (16) has a unique optimum that is integral. Then, by property 4, the set of strong-dotted edges M is the unique maximum weight matching M^* , i.e. $M = M^*$, and all other edges are non-dotted. By property 3, for i that has less than b_i partners under M^* , we have $\gamma_i = 0$. Hence by property 2, we know that for i not saturated under M^* , for every $(ij) \notin M^*$ since (ij) is a non-dotted edge $\gamma_{j \rightarrow i} = \alpha_{i \setminus j} = 0$, and for every $(ij) \in M^*$ node i gets a positive incoming offer $\gamma_{j \rightarrow i} = m_{j \rightarrow i}$. For i saturated under M^* , property 2 yields that the $(b_i^{\text{th}}\text{-max})$ highest incoming offers to i come from neighbors in M^* (without ties). It follows that $\underline{\alpha}$ induces the matching M^* . Also, we deduce that $\gamma_{i \rightarrow j} = \gamma_{j \rightarrow i} = 0$ for $(ij) \notin M^*$.

From property 6 we deduce that $\gamma_{i \rightarrow j} + \gamma_{j \rightarrow i} = w_{ij}$ for $(ij) \in M$ and from property 1, we deduce that $m_{j \rightarrow i} < \gamma_i$ for $(ij) \notin M$. It follows that (M, Γ) is a trade outcome. Finally, by properties 7 and 5, the pair $(M^*, \underline{\gamma})$ is stable and balanced respectively, and thus forms a NB solution. \square

Proof of Theorem 8: From NB solutions to fixed points.

Proof. Consider any NB solution (M, Γ_{NB}) . Using Proposition 1, $M = M_*$, the unique maximum weight matching. Construct a corresponding FP as follows. Set

$$m_{i \rightarrow j} = \begin{cases} \gamma_{\text{NB}, i \rightarrow j} & \text{for } (ij) \in M, \\ (w_{ij} - \gamma_{\text{NB}, i})_+ & \text{for } (ij) \notin M. \end{cases}$$

Compute $\underline{\alpha}$ using $\alpha_{i \setminus j} = (b_i^{\text{th}}\text{-max})_{k \in \partial i \setminus j} m_{k \rightarrow i}$. We claim that this is a FP and that the corresponding Γ is Γ_{NB} .

To prove that we are at a fixed point, we imagine updated offers $\underline{m}^{\text{upd}}$ based on $\underline{\alpha}$, and show $\underline{m}^{\text{upd}} = \underline{m}$.

Consider a matching edge $(i, j) \in M$. We know that $\gamma_{\text{NB}, i \rightarrow j} + \gamma_{\text{NB}, j \rightarrow i} = w_{ij}$. Also, stability and balance tell us

$$\gamma_{\text{NB}, j \rightarrow i} - (b_i^{\text{th}}\text{-max})_{k \in \partial i \setminus j} (w_{ik} - \gamma_{\text{NB}, k})_+ = \gamma_{\text{NB}, i \rightarrow j} - (b_j^{\text{th}}\text{-max})_{l \in \partial j \setminus i} (w_{jl} - \gamma_{\text{NB}, l})_+$$

and both sides are non-negative. For $(i, k) \in M$, we know that $(w_{ik} - \gamma_{\text{NB}, k})_+ \geq (w_{ik} - \gamma_{\text{NB}, i \rightarrow k})_+ = \gamma_{\text{NB}, k \rightarrow i} = m_{k \rightarrow i}$. It follows that

$$(b_i^{\text{th}}\text{-max})_{k \in \partial i \setminus j} (w_{ik} - \gamma_{\text{NB}, k})_+ = (b_i^{\text{th}}\text{-max})_{k \in \partial i \setminus j} m_{k \rightarrow i} = \alpha_{i \setminus j}$$

Hence, $\gamma_{\text{NB},j \rightarrow i} - \alpha_{i \setminus j} = \gamma_{\text{NB},i \rightarrow j} - \alpha_{j \setminus i} \geq 0$. Therefore $\alpha_{i \setminus j} + \alpha_{j \setminus i} \leq w_{ij}$,

$$\begin{aligned} m_{i \rightarrow j}^{\text{upd}} &= \frac{w_{ij} - \alpha_{i \setminus j} + \alpha_{j \setminus i}}{2} = \frac{w_{ij} - \gamma_{\text{NB},j \rightarrow i} + \gamma_{\text{NB},i \rightarrow j}}{2} \\ &= \gamma_{\text{NB},i \rightarrow j} = m_{i \rightarrow j}. \end{aligned}$$

By symmetry, we also have $m_{j \rightarrow i}^{\text{upd}} = \gamma_{\text{NB},j \rightarrow i} = m_{j \rightarrow i}$. Hence, the offers remain unchanged. Now consider $(i, j) \notin M$. We have $\gamma_{\text{NB},i} + \gamma_{\text{NB},j} \geq w_{ij}$ and, $\gamma_{\text{NB},i} = (b_i^{\text{th-max}})_{k \in \partial i \setminus j} \gamma_{\text{NB},k \rightarrow i} = \alpha_{i \setminus j}$. A similar equation holds for $\gamma_{\text{NB},j}$. The validity of this identity can be checked individually in the cases when i is saturated under M and i is not saturated under M . Hence, $\alpha_{i \setminus j} + \alpha_{j \setminus i} \geq w_{ij}$. This leads to $m_{i \rightarrow j}^{\text{upd}} = (w_{ij} - \alpha_{i \setminus j})_+ = (w_{ij} - \gamma_{\text{NB},i})_+ = m_{i \rightarrow j}$. By symmetry, we know also that $m_{j \rightarrow i}^{\text{upd}} = m_{j \rightarrow i}$.

Finally, we show $\Gamma = \Gamma_{\text{NB}}$. Note that since we have already established $\underline{\alpha}$ is a fixed point, we know from the direct part that $\underline{\alpha}$ induces the matching M , so there is no tie breaking required to determine the b_i highest incoming offers to node $i \in V$. For all $(i, j) \in M$, we already found that $m_{i \rightarrow j} = \gamma_{\text{NB},i \rightarrow j}$ and vice versa. For any edge $(ij) \notin M$, we know $m_{i \rightarrow j} = (w_{ij} - \gamma_{\text{NB},i})_+ \leq \gamma_{\text{NB},j}$. This immediately leads to $\Gamma = \Gamma_{\text{NB}}$. \square

E.1 Proof of properties used in direct part

Now we prove the fixed point properties that were used in the direct part of the proof of Theorem 8. Before that, however, we remark that the condition: “the primal LP in (16) has a unique optimum” in Theorem 8 is almost always valid.

Remark 22. *We argue that the condition “the primal LP in (16) has a unique optimum” is generic in instances with integral optimum:*

Let $\mathbf{G}_I \subset [0, 1]^{|E|}$ be the set of instances having an integral optimum, given a graph G with capacity constraints \mathbf{b} . Let $\mathbf{G}_{UI} \subset \mathbf{G}_I$ be the set of instances having a unique integral optimum. It turns out that \mathbf{G}_I has dimension $|E|$ (i.e. the class of instances having an integral optimum is large) and that \mathbf{G}_{UI} is both open and dense in \mathbf{G}_I .

Again, we denote surplus $w_{ij} - \alpha_{i \setminus j} - \alpha_{j \setminus i}$ of edge (ij) by Surp_{ij} .

Lemma 20. *The following are equivalent:*

- (a) $\gamma_i + \gamma_j \leq w_{ij}$,
- (b) $\text{Surp}_{ij} \geq 0$.
- (c) $\gamma_i \leq m_{j \rightarrow i}$ and $\gamma_j \leq m_{i \rightarrow j}$.

Moreover, if $\gamma_i \leq m_{j \rightarrow i}$ and $\gamma_j > m_{i \rightarrow j}$ then $\gamma_i = 0$.

Proof. We will prove $(a) \Rightarrow (b) \Rightarrow (c) \Rightarrow (a)$.

(a) \Rightarrow (b): Since $\gamma_i \geq \alpha_{i \setminus j}$ and $\gamma_j \geq \alpha_{j \setminus i}$ always holds then $w_{ij} \geq \gamma_i + \gamma_j \geq \alpha_{i \setminus j} + \alpha_{j \setminus i}$.

(b) \Rightarrow (c): If $\text{Surp}_{ij} \geq 0$ then $m_{i \rightarrow j} = (w_{ij} - \alpha_{i \setminus j} + \alpha_{j \setminus i})/2 \geq \alpha_{j \setminus i}$. So $m_{i \rightarrow j}$ is among the b_j best offers received by node j , implying $\gamma_j \leq m_{i \rightarrow j}$. The argument for $\gamma_i \leq m_{j \rightarrow i}$ is similar.

(c) \Rightarrow (a): If $\text{Surp}_{ij} \geq 0$ then $m_{i \rightarrow j} = (w_{ij} - \alpha_{i \setminus j} + \alpha_{j \setminus i})/2$ and $m_{j \rightarrow i} = (w_{ij} - \alpha_{j \setminus i} + \alpha_{i \setminus j})/2$ which gives $\gamma_i + \gamma_j \leq m_{i \rightarrow j} + m_{j \rightarrow i} = w_{ij}$ and we are done. Otherwise, we have $\gamma_i + \gamma_j \leq m_{i \rightarrow j} + m_{j \rightarrow i} = (w_{ij} - \alpha_{i \setminus j})_+ + (w_{ij} - \alpha_{j \setminus i})_+ < \max \left[(w_{ij} - \alpha_{i \setminus j})_+, (w_{ij} - \alpha_{j \setminus i})_+, 2w_{ij} - \alpha_{i \setminus j} - \alpha_{j \setminus i} \right] \leq w_{ij}$.

Finally, we suppose $\gamma_i = m_{j \rightarrow i}$ and $\gamma_j > m_{i \rightarrow j}$. First note that by equivalence of (b) and (c) we should have $w_{ij} < \alpha_{i \setminus j} + \alpha_{j \setminus i}$. On the other hand $\alpha_{i \setminus j} \leq \gamma_i \leq m_{j \rightarrow i} \leq (w_{ij} - \alpha_{j \setminus i})_+$. Now if

$w_{ij} - \alpha_{j \setminus i} > 0$ we get $\alpha_{i \setminus j} \leq w_{ij} - \alpha_{j \setminus i}$ which is a contradiction. Therefore $\gamma_i \leq (w_{ij} - \alpha_{j \setminus i})_+ = 0$, implying $\gamma_i = 0$. \square

Lemma 21. *The following are equivalent:*

- (a) $m_{j \rightarrow i} > \alpha_{i \setminus j}$,
- (b) $m_{i \rightarrow j} > \alpha_{j \setminus i}$,
- (c) $w_{ij} - \alpha_{i \setminus j} - \alpha_{j \setminus i} > 0$.
- (d) i and j receive positive offers from each other, with $m_{j \rightarrow i} > ((b_i + 1)^{\text{th-max}})_{k \in \partial i} m_{k \rightarrow i}$ and similarly for i .

These conditions imply that $m_{j \rightarrow i} = \gamma_{j \rightarrow i}$ and $m_{i \rightarrow j} = \gamma_{i \rightarrow j}$.

Proof. (a) \Rightarrow (c) \Rightarrow (b): (a) implies that $(w_{ij} - \alpha_{j \setminus i})_+ \geq m_{j \rightarrow i} > \alpha_{i \setminus j}$, which yields (c). From this we can explicitly write $m_{i \rightarrow j} = (w_{ij} - \alpha_{i \setminus j} + \alpha_{j \setminus i})/2$ which is strictly bigger than $\alpha_{j \setminus i}$. Hence we obtain (b).

By symmetry (b) \Rightarrow (c) \Rightarrow (a). Thus, we have shown that (a), (b) and (c) are equivalent.

(c) \Rightarrow (d): (c) implies that $m_{i \rightarrow j} = (w_{ij} - \alpha_{i \setminus j} + \alpha_{j \setminus i})/2 > \alpha_{j \setminus i} = (b_j^{\text{th-max}})_{k \in \partial j \setminus i} m_{k \rightarrow j}$. Using symmetry, it follows that (d) holds.

(d) \Rightarrow (a) is easy to check.

This finishes the proof of equivalence. The implication follows from the definition of $\gamma_{i \rightarrow j}$. \square

Recall that (ij) is a weak-dotted edge if $w_{ij} - \alpha_{i \setminus j} - \alpha_{j \setminus i} = 0$, a strong-dotted edge if $w_{ij} - \alpha_{i \setminus j} - \alpha_{j \setminus i} > 0$, and a non-dotted edge otherwise.

Lemma 22. *If $\gamma_i > 0$, then i has b_i adjacent strong dotted edges, or at least $b_i + 1$ adjacent dotted edges.*

Proof. Suppose $\gamma_i > 0$. Then the b_i largest incoming offers to i are all strictly positive. Suppose one of these offers comes from j . Then, $\alpha_{i \setminus j} \leq m_{j \rightarrow i} \leq (w_{ij} - \alpha_{j \setminus i})_+$. Now $m_{i \rightarrow j} > 0$, implying that $\alpha_{i \setminus j} \leq w_{ij} - \alpha_{j \setminus i}$ or (ij) is dotted edge. If there is strict inequality for all j , this means that we have at least b_i strong dotted edges adjacent to i . If we have equality for some j , that means there is a tie for the b_i highest offer incoming to i . We deduce that at least $b_i + 1$ dotted edges adjacent to i . \square

Lemma 23. *The following are equivalent:*

- (a) $\text{Surp}_{ij} \leq 0$,
 - (b) $m_{i \rightarrow j} = (w_{ij} - \alpha_{i \setminus j})_+$.
- Moreover, these conditions imply $\alpha_{i \setminus j} = \gamma_i$ and $\alpha_{j \setminus i} = \gamma_j$.

Proof. Equivalence of (a) and (b) follows from the definitions. (a) implies $m_{i \rightarrow j} \leq \alpha_{j \setminus i}$ which yields $\alpha_{j \setminus i} = \gamma_j$. By symmetry, we can also deduce $\alpha_{i \setminus j} = \gamma_i$. \square

Note that (a) is symmetric in i and j , so (b) can be transformed by interchanging i and j .

Lemma 24. *A non-solid edge cannot be a dotted edge, weak or strong.*

The proof of this lemma is very similar to that of Lemma 11: we consider optimal solutions of the primal and dual LPs (16), and construct an alternating path consisting of alternate (i) non-solid dotted edges, and (ii) strong solid, non-strong dotted edges. We omit the proof.

Lemma 25. *Every strong-solid edge is a strong-dotted edge.*

Again, the proof is very similar to the proof of Lemma 12, and we omit it.



Figure 3: Examples of basic structures: path, blossom, bicycle, and cycle (matched edges in bold).

Lemma 26. Consider the set of edges $M \equiv \{(ij) : \text{Surp}_{ij} > 0\}$. The balance property (15) is satisfied for all $(ij) \in M$.

Proof. Consider any $(ij) \in M$. From Lemma 21, we know that

$$\gamma_{j \rightarrow i} = m_{j \rightarrow i} = \alpha_{i \setminus j} + \text{Surp}_{ij}/2$$

To prove balance, it then suffices to establish

$$(b_i^{\text{th-max}})_{k \in \partial i \setminus j} m_{k \rightarrow i} = \alpha_{i \setminus j}. \quad (35)$$

Now node i can have at most b_i adjacent strong dotted edges, from Lemma 6 (d). One of these is (ij) . (35) follows from the property $m_{i \rightarrow j} = (w_{ij} - \gamma_i)_+$ on non-strong dotted edges (from Lemma 23). \square

Lemma 27. An optimum solution for the dual LP in (16) can be constructed as $y_i = \gamma_i$ for all $i \in V$ and:

$$y_{ij} = \begin{cases} w_{ij} - \gamma_i - \gamma_j & \text{for } (ij) \in M, \\ 0 & \text{for } (ij) \notin M. \end{cases}$$

Proof. We first show that this construction satisfies the dual constraints. $y_{ij} \geq 0$ follows from Lemma 20 (a) and (b). We have

$$y_i + y_j + y_{ij} = w_{ij} \quad \text{for } (ij) \in M \quad (36)$$

by construction. For $(ij) \notin M$, we have $y_i + y_j = \gamma_i + \gamma_j \geq w_{ij}$ from Lemma 23. This completes our proof of feasibility.

To show optimality, we establish that the weight of matching M is the same as the dual objective value $\sum_{i \in V} b_i y_i + \sum_{e \in E} y_{ij}$ at the chosen \underline{y} . Lemma 22 guarantees $\gamma_i = 0$ for i having less than b_i adjacent dotted edges. Using Lemmas 24 and 25, we know that M is a valid \mathbf{b} matching consisting of strong dotted edges, and all other edges are non-dotted. We deduce

$$w(M) = \sum_{(ij) \in M} w_{ij} = \sum_{(ij) \in M} y_i + y_j + y_{ij} = \sum_{i \in V} b_i y_i + \sum_{e \in E} y_{ij},$$

using (36). This completes our proof of optimality. \square

F The Kleinberg-Tardos construction and the KT gap

Let G be an instance which admits at least one stable outcome, M^* be the corresponding matching (recall that this is a maximum weight matching), and consider the Kleinberg-Tardos (KT) procedure for finding a NB solution [27]. Any NB solution $\underline{\gamma}^*$ can be constructed by this procedure with appropriate choices at successive stages. At each stage, a linear program is solved with variables

γ_i attached to node i . The linear program maximizes the minimum ‘slack’ of all unmatched edges and nodes, whose values have not yet been set (the slack of edge $(i, j) \notin M$ is $\gamma_i + \gamma_j - w_{ij}$).

At the first stage, the set of nodes that remain unmatched (i.e. are not part of M^*) is found, if such nodes exist. Call the set of unmatched nodes \mathcal{C}_0 .

After this, at successive stages of the KT procedure, a sequence of structures $\mathcal{C}_1, \mathcal{C}_2, \dots, \mathcal{C}_k$ characterizing the LP optimum are found. We call this the *KT sequence*. Each such structure is a pair $\mathcal{C}_q = (V(\mathcal{C}_q), E(\mathcal{C}_q))$ with $V(\mathcal{C}_q) \subseteq V$, $E(\mathcal{C}_q) \subseteq E$. According to [27] \mathcal{C}_q belongs to one of four topologies: alternating path, blossom, bicycle, alternating cycle (Figure 3). The q -th linear program determines the value of γ_i^* for $i \in V(\mathcal{C}_q)$. Further, one has the partition $E(\mathcal{C}_q) = E_1(\mathcal{C}_q) \cup E_2(\mathcal{C}_q)$ with $E_1(\mathcal{C}_q)$ consisting of all matching edges along which nodes in $V(\mathcal{C}_q)$ trade, and $E_2(\mathcal{C}_q)$ consists of edges (i, j) such that some $i \in V(\mathcal{C}_q)$ receives its second-best, positive offer from j .

The γ values for nodes on the limiting structure are uniquely determined if the structure is an alternating path, blossom or bicycle⁹. In case of an alternating cycle there is one degree of freedom – setting a value γ_i^* for one node $i \in \mathcal{C}_q$ fully determines the values at the other nodes.

We emphasize that, within the present definition, \mathcal{C}_q is not necessarily a subgraph of G , in that it might contain an edge (i, j) but not both its endpoints. On the other hand, $V(\mathcal{C}_q)$ is always subset of the endpoints of $E(\mathcal{C}_q)$. We denote by $V_{\text{ext}}(\mathcal{C}_q) \supseteq V(\mathcal{C}_q)$ the set of nodes formed by all the endpoints of edges in $E(\mathcal{C}_q)$.

For all nodes $i \in V(\mathcal{C}_q)$ the second best offer is equal to $\gamma_i^* - \sigma_q$, where σ_q is the slack of the q -th structure. Therefore

$$\gamma_i^* + \gamma_j^* - w_{ij} = \begin{cases} 0 & \text{if } (i, j) \in E_1(\mathcal{C}_q), \\ \sigma_q & \text{if } (i, j) \in E_2(\mathcal{C}_q). \end{cases}$$

The slacks form an increasing sequence $(\sigma_1 \leq \sigma_2 \leq \dots \leq \sigma_k)$.

Definition 23. We say that a unique Nash bargaining solution $\underline{\alpha}^*$ has a *KT gap* σ if

$$\sigma \leq \min \{ \sigma_1; \sigma_2 - \sigma_1; \dots; \sigma_k - \sigma_{k-1} \},$$

and if for each edge (i, j) such that $i, j \in V_{\text{ext}}(\mathcal{C}_q)$ and $(i, j) \notin E(\mathcal{C}_q)$,

$$\gamma_i^* + \gamma_j^* - w_{ij} \geq \sigma_q + \sigma.$$

It is possible to prove that the positive gap condition is generic in the following sense. The set of all instances such that the NB solution is unique can be regarded as a subset $\mathbf{G} \subseteq [0, W]^{|E|}$ (W being the maximum edge weight). It turns out that \mathbf{G} has dimension $|E|$ (i.e. the class of instances having unique NB solution is large) and that the subset of instances with gap $\sigma > 0$ is both *open* and *dense* in \mathbf{G} .

G Appendix to Section 4

G.1 UD example showing exponentially slow convergence

We construct a sequence of instances I_n along with particular initializations such that convergence to an approximate fixed point is exponentially slow, cf. Section 4.1. Let us first consider $n = 8N$, where $N \geq 2$. The construction can be easily extended to arbitrary $n \geq 16$. The graph $G_n = (V_n, E_n)$ we

⁹In [27] it is claimed that the γ values ‘may not be fully determined’ also in the case of bicycles. However it is not hard to prove that γ values are, in fact, uniquely determined in bicycles.

will consider is a simple ‘ring’. More precisely, $V_n = \{1, 2, \dots, n\}$ and $E_n = \{(1, 2), (2, 3), \dots, (n-1, n), (n, 1)\}$. All edges have the same weight W , with the exception $w_{4N, 4N+1} = W - 1$. (We will choose W later.) Thus, G_n has a unique maximum weight matching $M_* = \{(1, 2), (3, 4), \dots, (n-1, n)\}$. Moreover, M_* solves LP (2), implying that G_n possesses a UD solution (using Lemma 1), and the LP gap $g = 1$ provided $W \geq 1$. Given any $r \in (0, 1/2)$, we define the split fractions as follows¹⁰: We first define the split fractions on edges $E_{n, \text{left}} = \{(1, 2), (2, 3), \dots, (4N-1, 4N)\}$. We set

$$\begin{aligned} r_{3,2} = r_{5,4} = \dots = r_{2N-1, 2N-2} &= r \\ r_{2N+1, 2N+2} = r_{2N+3, 2N+4} = \dots = r_{4N-3, 4N-2} &= r \end{aligned}$$

and the split fractions on all other edges in $E_{n, \text{left}}$ to $1/2$. As before, $r_{i, i+1} = 1 - r_{i+1, i}$ is implicit.

For edges on the ‘right’ side we define split fractions in a symmetrical way. Define ‘reflection’ $\mathcal{R} : \{4N+1, 4N+2, \dots, 8N\} \rightarrow \{1, 2, \dots, 4N\}$ as

$$\mathcal{R}(l) = n - l + 1 \tag{37}$$

We set $r_{i, i+1} = r_{\mathcal{R}(i), \mathcal{R}(i+1)}$ for all $i \in \{4N+1, 4N+2, \dots, 8N-1\}$.

Finally, we set $r_{n, 1} = r_{4N, 4N+1} = 1/2$.

Now we show how to construct $\underline{\alpha}$ satisfying properties (a) and (b) above. Let $\beta \equiv r/(1-r) < 1$.

Set $\alpha_{n \setminus 1} = \alpha_{1 \setminus n} = W/2 - 1$.

For $0 \leq i \leq N-1$, define¹¹

$$\alpha_{2i+1 \setminus 2i+2} = W/2 - \sum_{j=0}^{i-1} \beta^j, \quad \alpha_{2i+2 \setminus 2i+1} = W/2 + \sum_{j=0}^i \beta^j.$$

For $1 \leq i \leq N-1$, define

$$\alpha_{2i \setminus 2i+1} = W/2 + \sum_{j=0}^{i-2} \beta^j, \quad \alpha_{2i+1 \setminus 2i} = W/2 - \sum_{j=0}^i \beta^j.$$

Let $\alpha_{2N \setminus 2N+1} = W/2 + \sum_{j=0}^{N-1} \beta^j$, $\alpha_{2N+1 \setminus 2N} = W/2 - \sum_{j=0}^{N-1} \beta^j$. For $0 \leq i \leq N-2$, define

$$\alpha_{2N+2i+1 \setminus 2N+2i+2} = W/2 - \sum_{j=0}^{N-i-1} \beta^j, \quad \alpha_{2N+2i+2 \setminus 2N+2i+1} = W/2 + \sum_{j=0}^{N-i-3} \beta^j.$$

For $1 \leq i \leq N-1$, define

$$\alpha_{2N+2i \setminus 2N+2i+1} = W/2 + \sum_{j=0}^{N-i-1} \beta^j, \quad \alpha_{2N+2i+1 \setminus 2N+2i} = W/2 - \sum_{j=0}^{N-i-2} \beta^j.$$

Next, let $\alpha_{4N-1 \setminus 4N} = \alpha_{4N \setminus 4N-1} = W/2 - 1$. $\alpha_{4N \setminus 4N+1} = \alpha_{4N+1 \setminus 4N} = W/2$. Again, the messages on the ‘right’ side are defined in terms of the reflection Eq. (37). We define $\alpha_{i \setminus i+1} = \alpha_{\mathcal{R}(i) \setminus \mathcal{R}(i+1)}$, and $\alpha_{i+1 \setminus i} = \alpha_{\mathcal{R}(i+1) \setminus \mathcal{R}(i)}$, for all $i \in \{4N+1, 4N+2, \dots, 8N-1\}$.

¹⁰It turns out the values of the split fraction on approximately half the edges is irrelevant for our choice of $\underline{\alpha}$.

¹¹The summations can be written in closed form if desired, but it turns out to be easier to work with the summations directly.

The corresponding offers $\underline{m}(\underline{\alpha})$ can be easily computed. We have $m_{1 \rightarrow n} = m_{n \rightarrow 1} = W/2$. For $0 \leq i \leq N - 1$, we have

$$m_{2i+1 \rightarrow 2i+2} = W/2 + \sum_{j=0}^{i-1} \beta^j, \quad m_{2i+2 \rightarrow 2i+1} = W/2 + \sum_{j=0}^i \beta^j.$$

For $1 \leq i \leq N - 1$, we have

$$m_{2i \rightarrow 2i+1} = W/2 - \sum_{j=0}^{i-1} \beta^j, \quad m_{2i+1 \rightarrow 2i} = W/2 + \sum_{j=0}^{i-1} \beta^j.$$

We have $m_{2N \rightarrow 2N+1} = W/2 - \sum_{j=0}^{N-1} \beta^j$, $m_{2N+1 \rightarrow 2N} = W/2 + \sum_{j=0}^{N-1} \beta^j$. For $0 \leq i \leq N - 2$, we have

$$m_{2N+2i+1 \rightarrow 2N+2i+2} = W/2 + \sum_{j=0}^{N-i-2} \beta^j, \quad m_{2N+2i+2 \rightarrow 2N+2i+1} = W/2 - \sum_{j=0}^{N-i-2} \beta^j.$$

For $1 \leq i \leq N - 1$, we have

$$m_{2N+2i \rightarrow 2N+2i+1} = W/2 - \sum_{j=0}^{N-i-1} \beta^j, \quad m_{2N+2i+1 \rightarrow 2N+2i} = W/2 + \sum_{j=0}^{N-i-2} \beta^j.$$

Finally, we have $m_{4N-1 \rightarrow 4N} = m_{4N \rightarrow 4N-1} = W/2$, $m_{4N \rightarrow 4N+1} = m_{4N+1 \rightarrow 4N} = W/2 - 1$. Offers on the ‘right’ side are clearly the same as the corresponding offer on the left.

Importantly, note that it suffices to have $W \geq 2 + 2/(1 - \beta)$ to ensure that all messages and offers are bounded below by 1 (in particular, none of them is negative). Choose $W = 2 + 2/(1 - \beta)$ (for example).

Next, note that all the fixed point conditions are satisfied, except the update rules for $\alpha_{2N \setminus 2N+1}$, $\alpha_{2N+1 \setminus 2N}$ and the corresponding messages in the reflection. However, it is easy to check that these update rules are each violated by only β^{N-1} . Thus, $\underline{\alpha}$ is an ϵ_n -FP for $\epsilon_n = \beta^{N-1} \leq 2^{-cn}$ for appropriate $c = c(r) > 0$.

G.2 Proof of Theorem 7

We show how to obtain an FPTAS using the two steps stated in Section 4.2.

Step 1 can be carried out by finding a maximum weight matching M^* (see, e.g., [19]) and also solving the the dual linear program (3). For the dual LP, let v be the optimum value and let $\underline{\gamma}$ be an optimum solution. We now use Proposition 1. If the weight of M^* is smaller than v , we return UNSTABLE, since we know that no stable outcome exists, hence no UD solution (or ϵ -UD solution) exists. Else, $(\underline{\gamma}, M^*)$ is a stable outcome. This completes step 1! The computational effort involved is polynomial in the input size. All unmatched nodes have earnings of 0.

In step 2, we fix the matching M^* , and rebalance the matched edges iteratively. It turns out to be crucial that our iterative updates preserve stability. An example similar to the one in Section G.1 demonstrates that the rebalancing procedure can take an exponentially large time to reach an approximate UD solution if stability is not preserved [23].

We now motivate the rebalancing procedure briefly, before we give a detailed description and state results. Imagine an edge $(i, j) \in M^*$. Since we start with a stable outcome, the edge weight

Table 1: Local algorithm that converts stable outcome to ϵ -UD solution

EDGE REBALANCING (Instance I , Stable outcome $(\underline{\gamma}, M)$, Damping factor κ , Error target ϵ)	
1:	Check $\kappa \in (0, 1/2]$, $\epsilon > 0$, $(\underline{\gamma}, M)$ is stable outcome \top
2:	If (Check fails) Return ERROR
3:	$\underline{\gamma}^0 \leftarrow \underline{\gamma}$
4:	$t \leftarrow 0$
5:	Do
6:	ForEach $(i, j) \in M$
7:	$\gamma_i^{\text{reb}} \leftarrow \max_{k \in \partial i \setminus j} (w_{ik} - \gamma_k^t)_+ + r_{ij} \text{Surp}_{ij}(\underline{\gamma}^t)$
8:	$\gamma_j^{\text{reb}} \leftarrow \max_{l \in \partial j \setminus i} (w_{jl} - \gamma_l^t)_+ + r_{ji} \text{Surp}_{ij}(\underline{\gamma}^t)$
9:	End ForEach
10:	ForEach $i \in V$ that is unmatched under M
11:	$\gamma_i^{\text{reb}} \leftarrow 0$
12:	End ForEach
13:	If $(\ \underline{\gamma}^{\text{reb}} - \underline{\gamma}^t\ _\infty \leq \epsilon)$ Break Do
14:	$\underline{\gamma}^{t+1} = \kappa \underline{\gamma}^{\text{reb}} + (1 - \kappa) \underline{\gamma}^t$
15:	$t \leftarrow t + 1$
16:	End Do
17:	Return $(\underline{\gamma}^t, M)$

w_{ij} is at least the sum of the best alternatives, i.e. $\text{Surp}_{ij} \geq 0$. Suppose we change the division of w_{ij} into γ'_i, γ'_j so that the Surp_{ij} is divided as per the prescribed split fraction r_{ij} . Earnings of all other nodes are left unchanged. Since $r_{ij} \in (0, 1)$, γ'_i is at least as large as the best alternative of i , as was the case for γ_i . This leads to $\gamma'_i + \gamma_k \geq w_{ik}$ for all $k \in \partial i \setminus j$. A similar argument holds for node j . In short, *stability is preserved!*

It turns out that the analysis of convergence is simpler if we analyze synchronous updates, as opposed to asynchronous updates as described above. Moreover, we find that simply choosing an appropriate ‘damping factor’ allows us to ensure that stability is preserved even with synchronous updates. We use a powerful technique introduced in our recent work [27] to prove convergence.

Table 1 shows the algorithm EDGE REBALANCING we use to complete step 2. Note that each iteration of the loop can requires $O(|E|)$ simple operations.

Correctness of EDGE REBALANCING:

It is easy to check that $(\underline{\gamma}^{\text{reb}}, M)$ and $(\underline{\gamma}^t, M)$ are valid outcomes for all t . We show that $\underline{\gamma}^t$ computed by EDGE REBALANCING is, in fact, a stable allocation:

Lemma 28. *If EDGE REBALANCING is given a valid input satisfying the ‘Check’ on line 1, then $(\underline{\gamma}^t, M)$ is a stable outcome for all $t \geq 0$.*

Proof. We prove this lemma by induction on time t . Clearly $(\underline{\gamma}^0, M)$ is a stable outcome, since the input is valid. Suppose $(\underline{\gamma}^t, M)$ is a stable outcome.

Consider any $(i, j) \in M$. It is easy to verify that $\gamma_i^{\text{reb}} + \gamma_j^{\text{reb}} = w_{ij}$, for $\underline{\gamma}^{\text{reb}}$ computed from $\underline{\gamma}^t$ in Lines 6-9 of EDGE REBALANCING. Also, we know that $\gamma_i^t + \gamma_j^t = w_{ij}$. It follows that $\gamma_i^{t+1} + \gamma_j^{t+1} = w_{ij}$ as needed. For $i \in V$ unmatched under M , $\gamma_i^t = 0$ by hypothesis and as per Lines 10-12, $\gamma_i^{\text{reb}} = 0 \Rightarrow \gamma_i^{t+1} = 0$ as needed.

Consider any $(i, k) \in E \setminus M$. We know that $\gamma_i^t + \gamma_k^t \geq w_{ik}$. We want to show the corresponding inequality at time $t + 1$. Define $\sigma_{ik}^t \equiv \gamma_i^t + \gamma_k^t - w_{ik} \geq 0$.

Claim 24. $\gamma_i^{\text{reb}} \geq \gamma_i^t - \sigma_{ik}^t$.

If we prove the claim, it follows that a similar inequality holds for γ_k^{reb} , and hence $\gamma_i^{\text{reb}} + \gamma_k^{\text{reb}} \geq \gamma_i^t + \gamma_k^t - 2\sigma_{ik}^t = w_{ik} - \sigma_{ik}^t$. It then follows from the definition in Line 14 that $\gamma_i^{t+1} + \gamma_k^{t+1} \geq w_{ik}$, for any $\kappa \in (0, 1/2]$. This will complete our proof that $(\underline{\gamma}^{t+1}, M)$ is a stable outcome.

Let us now prove the claim. Suppose i is matched under M . Using the definition in Line 7 (Line 8 contains a symmetrical definition), $\gamma_i^{\text{reb}} \geq \max_{k' \in \partial i \setminus j} (w_{ik'} - \gamma_{k'}^t)_+$ since $\text{Surp}_{ik}(\underline{\gamma}^t) \geq 0$. Hence,

$$\gamma_i^{\text{reb}} \geq (w_{ik} - \gamma_k^t)_+ \geq (w_{ik} - \gamma_k^t) = \gamma_i^t - \sigma_{ik}^t,$$

as needed. If i is not matched under M , then $\gamma_i^t = \gamma_i^{\text{reb}} = 0$, so the claim follows from $\sigma_{ik}^t \geq 0$. \square

Convergence of EDGE REBALANCING:

Note that the termination condition $\|\underline{\gamma}^{\text{reb}} - \underline{\gamma}^t\|_\infty \leq \epsilon$ on Line 13 is equivalent to ϵ -correct division. We show that the rebalancing algorithm terminates fast:

Lemma 29. *For any instance with weights bounded by 1, i.e. $(w_e, e \in E) \in (0, 1]^{|E|}$, if EDGE REBALANCING is given a valid input, it terminates in T iterations, where*

$$T \leq \left\lceil \frac{1}{\pi\kappa(1-\kappa)\epsilon^2} \right\rceil. \quad (38)$$

and returns an outcome satisfying ϵ -correct division (cf. Definition 10). Here $\pi = 3.14159\dots$

sketch only. This result is proved using the same technique as in the proof of Theorem 3. The iterative updates of EDGE REBALANCING can be written as

$$\underline{\gamma}^{t+1} = \kappa \mathbb{T} \underline{\gamma}^t + (1 - \kappa) \underline{\gamma}^t, \quad (39)$$

where \mathbb{T} is a self mapping of the (convex) set of allocations corresponding to matching M , $\mathcal{A}_M \subseteq [0, 1]^{|E|}$. The ‘edge balancing’ operator \mathbb{T} essentially corresponds to Lines 6-9 in Table 1. It is fairly straightforward to show that \mathbb{T} is nonexpansive with respect to sup norm. The main theorem in [3] then tells us that

$$\|\mathbb{T} \underline{\gamma}^t - \underline{\gamma}^t\|_\infty \leq \frac{1}{\sqrt{\pi\kappa(1-\kappa)t}}. \quad (40)$$

The result follows.

For the full proof, see [23]. \square

Using Lemmas 28 and Lemmas 29, we immediately obtain our main result, Theorem 7.

of *Theorem 7*. We showed that step 1 can be completed in polynomial time. If the instance has no UD solutions then the algorithm returns UNSTABLE. Else we obtain a stable outcome and proceed to step 2.

Step 2 is performed using EDGE REBALANCING. The input is the instance, the stable outcome obtained from step 1, $\kappa = 1/2$ (for example) and the target error value $\epsilon > 0$. Lemmas 28 and 29 show that EDGE REBALANCING terminates after at most $\lceil 1/(\pi\kappa(1 - \kappa)\epsilon^2) \rceil$ iterations, returning a outcome that is stable and satisfies ϵ -correct division, i.e. an ϵ -UD solution. Moreover, each iteration requires $O(|E|)$ simple operations. Hence, step 2 is completed in $O(|E|/\epsilon^2)$ simple operations.

The total number of operations required by the entire algorithm is thus polynomial in the input and in $(1/\epsilon)$. \square