

Bargaining for Revenue Shares on Tree Trading Networks

Arpita Ghosh*

Information Science Dept.
Cornell University, Ithaca, NY
arpitaghosh@cornell.edu

Kevin Lang

Yahoo! Research
Sunnyvale, CA
langk@yahoo-inc.com

Satyen Kale*

IBM Watson
Yorktown Heights, NY
sckale@us.ibm.com

Benjamin Moseley†

TTI-Chicago
Chicago, IL
moseley@ttic.edu

Abstract

We study trade networks with a *tree* structure, where a seller with a single indivisible good is connected to buyers, each with some value for the good, via a unique path of intermediaries. Agents in the tree make multiplicative revenue share offers to their parent nodes, who choose the best offer and offer part of it to their parent, and so on; the winning path is determined by who finally makes the highest offer to the seller. In this paper, we investigate how these revenue shares might be set via a natural bargaining process between agents on the tree, specifically, *egalitarian* bargaining between endpoints of each edge in the tree. We investigate the fixed point of this system of bargaining equations and prove various desirable properties for this solution concept, including (i) existence, (ii) uniqueness, (iii) efficiency, (iv) membership in the core, (v) strict monotonicity, (vi) polynomial-time computability to any given accuracy. Finally, we present numerical evidence that asynchronous dynamics with randomly ordered updates always converges to the fixed point, indicating that the fixed point shares might arise from decentralized bargaining amongst agents on the trade network.

1 Introduction

Motivated by applications to ad exchanges such as the Yahoo!’s Right Media Exchange [Yah, 2007], we consider a theoretical model of trade networks which take the form of a rooted tree. In this model, publishers selling impressions can be connected via a string of *intermediary* ad-networks [Feldman *et al.*, 2010] to advertisers interested in buying these impressions at the leaf nodes. These intermediaries want a cut

of the surplus generated when a trade facilitated by them occurs. Typically, these cuts are specified as multiplicative *revenue shares or cuts* on edges that link a pair of entities. In practice, the value of these revenue shares would be set by business negotiations between the entities. A natural theoretical question, is what constitutes a reasonable set of values for these revenue shares. Of course, a complete solution to this problem would require analyzing a very complex setting with advertisers and intermediaries optimizing over multiple heterogeneous impressions and publishers in a network setting; in this paper, we take the first steps towards understanding this problem by analyzing the sale of a single impression.

In our model, each buyer makes an offer to pay its parent intermediary in the tree a revenue share in the form of some fraction of its value for the item, *i.e.*, for being matched to the seller; the parent intermediary chooses the highest offer from all the buyers it is connected to. Each such intermediary then makes an offer to its parent, who selects the highest offer, and so on. Finally, the seller selects the highest offer it receives from its children in the tree, which determines the winning buyer. Given the tree structure and buyer values, the revenue shares completely specify the winning path and all winners’ payoffs.

The two-player bargaining problem widely studied in cooperative game theory, where the seller and a single buyer with value v must fairly divide the value v generated from their trade, is a special case of this setting: in the simplest version, the bargaining solution is to split the value v equally amongst the two agents. Now consider a tree network where the child and parent nodes on each edge bargain about the revenue share using two-player bargaining. Here, the child might want to offer a revenue share greater than $1/2$ for two reasons: first, the parent node might have other children to trade with that this node needs to beat out. Second, and more unique to our setting, even if the child does beat out its siblings, the *parent* may not be able to make an adequately large offer to beat out *its* siblings higher up in the tree, and so on— if this happens, neither the parent nor the child belongs to the winning path, and the value actually *realized* by the child is zero. So how much the child offers its parent, accounting for

*Work done while the author was at Yahoo! Research.

†Work done while the author was at Yahoo! Research and the University of Illinois. Partially supported by NSF grant CCF-1016684.

both of these effects, will depend on shares elsewhere in the tree, *i.e.*, the revenue share negotiated on an edge depends on shares elsewhere in the tree. The key question we consider in this paper is whether there is a set of *mutually consistent* shares on all edges, and if yes, what kinds of outcomes it generates.

While there has been plenty of work on network bargaining problems building on the seminal work of Kleinberg and Tardos [2008], the model considered in those papers is unsuitable to our problem since the values that are being bargained on the edges are *exogenous*. In our setting, the value being split on an edge is *endogenous*, depending on the splits elsewhere in the tree. In the language of bargaining games, in the Kleinberg-Tardos model for network bargaining the feasible set for the bargaining problem on each edge is independent of shares on other edges (although the disagreement point is not), whereas in our setting the feasible set for an edge in the tree bargaining problem changes with shares elsewhere in the tree.

Such endogenous values on edges arise naturally in bargaining networks arising from trading settings, where there is competition for goods being sold. We consider the simplest possible version of this new bargaining model, which is bargaining on a tree. There are certainly many possible generalizations, but the goal of this paper is to analyze the simplest setting fully. Thus, while many ad-hoc schemes can be proposed to compute these values, in this paper we investigate a natural bargaining game motivated by the fact that entities in the trade network negotiate revenue shares, and show that the outcome corresponding to the *unique fixed point* of this bargaining game on the trade network has many desirable properties.

Overview of Conceptual Contributions and Technical Results. Our key conceptual contribution in this paper is the formulation of a bargaining game on the trading tree and a new solution concept for the game based on fixed points of the bargaining game. In this bargaining game, the nodes at the endpoints of each edge in the tree negotiate about how to split the value arriving at the child node according to two-player *egalitarian*, or proportional, bargaining [Kalai, 1977], given the splits elsewhere in the tree (§3). Each such two-player bargaining game gives us a (non-linear) equation for the revenue share on that edge, and a fixed point of the game is a simultaneous solution to this system of equations.

We first show a reduction from tree bargaining to path bargaining (§4). We prove that for any tree and any set of buyer values, there is a reduced *path* such that every fixed point of the system of bargaining equations on the path corresponds to a fixed point of the system of bargaining equations on the tree, and vice versa. This reduction to path bargaining allows us to analyze a smaller set of simultaneous equations, with one variable for each edge in the *path*, which we use to prove the following set of results:

1. *Existence and uniqueness:* We show that a fixed point to the bargaining equations always exists; further, the fixed point is unique.
2. *Efficiency:* The winner is always a buyer with highest

value. (This is not true with Nash bargaining.)

3. *Core:* The payoffs given by the fixed point of the bargaining equations belong to the *core* of the natural cooperative game corresponding to our setting.
4. *Strict Monotonicity:* If the maximum value in the subtree rooted at a node in the winning path increases, the node’s final payoff strictly increases as well.
5. *Computability:* The fixed point of the system of bargaining equations can be efficiently computed to any accuracy in polynomial time by a centralized algorithm.

Finally, in §6.4, we present exhaustive numerical simulations indicating that *asynchronous dynamics*, where at each step a random edge in the tree renegotiates the share x_e given the current shares in the remaining edges, converges rapidly to the fixed point—this suggests that decentralized bargaining on edges should lead to the shares specified by this fixed point.

The outcome corresponding to the fixed point of the bargaining equations can be thought of as a solution concept for the corresponding cooperative game. A natural question is the suitability of other solution concepts such as the Shapley value or the nucleolus for our setting, or using Nash bargaining instead to define the solution concept: all these candidates seem to have some deficiency compared to our concept. We refer the interested reader to the full version [Ghosh *et al.*, 2013] for a discussion.

Techniques. Our results are based on several analytical and combinatorial techniques. First, we prove several structural properties that *any* fixed point solution, if one exists, must satisfy, which allows us to reduce the general tree bargaining problem to a structurally simpler path bargaining problem (see §4). Next, for the path bargaining problem, we use analytic techniques to deduce certain monotonicity properties of any fixed point solution. These properties directly give us uniqueness of the fixed point, assuming it exists. To show existence, we appeal to Brouwer’s fixed point theorem by constructing a continuous mapping that is closely related to the bargaining equations. Our proofs of the core and strict monotonicity properties of the fixed point are again based on analytic techniques, and the use of an optimal substructure result for the fixed point which follows from our uniqueness result. Finally, by refining our monotonicity arguments quantitatively, we give an algorithm based on binary search to compute the fixed point to any accuracy, with running time that is polynomial in the number of nodes and the *logarithms* of the accuracy parameter and the gap between the highest and second highest values.

Related work. The problem we study relates to many well-studied branches of the economics and computer science literature. The question of how agents on the winning path should split the generated value can be thought of as a *fair division* or revenue-sharing problem on which there is an extensive literature, albeit in settings different from ours; for an overview, see [Moulin, 2004]. The work of Blume *et al.* [2007] is perhaps the most similar in spirit to ours from this literature, though it looks at a different setting where traders set prices strategically and buyers and sellers react to these

offers in a general trade network, and investigates subgame perfect Nash equilibria.

There is much recent work on bargaining in social networks, starting with the work of Kleinberg and Tardos [2008]. This work extends the classic two-player bargaining problem to a network where pairs of agents, instead of bargaining in isolation, can choose which neighbor to bargain with¹. A number of papers since [Kleinberg and Tardos, 2008] have addressed computational and structural aspects of the network bargaining problem, as well as extensions to the model and dynamics; see [Chakraborty and Kearns, 2008; Chakraborty *et al.*, 2009; Azar *et al.*, 2009; Celis *et al.*, 2010; Azar *et al.*, 2010; Kanoria, 2010]. While there are similarities between the network bargaining and our model, there are also fundamental differences: there, an outcome is a *matching* on the network, whereas we seek a *path*. More importantly, the values that are being bargained over on the edges are *exogenous* in their model, while in ours the value being split on an edge itself *depends* on the splits elsewhere in the tree: in the language of bargaining games, the feasible set for the bargaining problem on each edge is *independent* of shares on other edges in the network bargaining problem (although the disagreement point is not), whereas the feasible set for an edge in the tree bargaining problem *changes* with shares elsewhere in the tree.

2 Model

There is a seller selling a single item, buyers, each of whom derives some value from the item, and a number of intermediaries who assist in connecting buyers to the seller. The trade network between these agents is given by a *rooted tree* T : the leaf nodes in T (denoted generically by l) are the buyers, the root r is the seller, and the internal nodes (denoted generically by i) are the intermediaries. We use v_l to denote leaf l 's value for the item. The tree structure of the trade network means that each buyer has a *unique* path to the seller. An *instance* (T, \vec{v}) of the tree bargaining problem is specified by the tree topology T , and the values v_l at the leaves of T .

We use e to denote edges and p to denote paths connecting the seller and a buyer in T . Given a path $p = \{r, i_1, \dots, i_k, l\}$, we define the value of the path $v(p) = v_l$. For any two nodes t_1 and t_2 let $p_{t_1 t_2}$ denote the unique path from t_1 to t_2 in the tree T . A child node in the tree makes an offer to its parent, who chooses the highest of these and offers part of it to *its* parent, and so on, as described next.

The endpoints of each edge $e = (t, s)$ in T *split* the value that arrives at the child node t , specifying what portion of this value t retains and what portion it is willing to pass up to s . We use x_e , where $x_e \in [0, 1]$, to denote the multiplicative split or ‘revenue share’ on edge e : if the potential value² arriving at t is w_t , t keeps $w_t(1 - x_e)$ and passes up $w_t x_e$ to s . We use the multiplicative split x_e rather than an additive split for convenience in correctly writing the bargaining equations.

¹For a nice survey of the literature on network exchange theory as well as two-player bargaining, see [Kleinberg and Tardos, 2008; Chakraborty and Kearns, 2008] respectively.

²We say potential value because this value is realized only if these nodes belong to the winning path.

Note that the value of x_e can, of course, depend on w_t , as well as the splits $x_{e'}$ on other edges $e' \in T$.

Given an instance (T, \vec{v}) , an outcome consists of a *winning buyer* l^* , which also specifies the *winning path* $p^* = p_{l^* r}$, and a split of the value v_{l^*} amongst the nodes on the winning path (including the leaf and the root).

The set of revenue shares x_e completely specifies the outcome for an instance (T, v) as follows. Every node in the tree, when presented with multiple children offering different payoffs, chooses to transact with the child that gives her the highest payoff. Define the value reaching a non-leaf node $s \in T$, w_s , recursively as follows. Set $w_l = v_l$ for all leaves l , and let C_s be the set of children of s in T . Then, we have $w_s = \max_{t \in C_s} x_{ts} w_t$.

Let $t^*(s) = \arg \max_{t \in C_s} x_{ts} w_t$, with ties broken arbitrarily, denote the ‘winning child’ of the parent node s . The path $p^* = (r, t^*(r), t^*(t^*(r)), \dots, l^*)$ from root r to leaf l^* is the *winning path*, and l^* is the *winning buyer*. The value v_{l^*} , generated by matching l^* to r is split among the nodes on p^* using the revenue shares on edges of p^* . For all other nodes in the tree, the payoff is zero.

This setting can also be modeled as a cooperative game; we do this in §6.

3 Bargaining on Trees

Given an instance (T, \vec{v}) , the splits x_e on the edges $e \in T$ completely specify the outcome, namely who the winning agents are, and what payoffs they receive. How might these splits x_e be determined?

We consider a bargaining-based determination of the shares x_e . We suppose that the agents corresponding to the endpoint of each edge negotiate according to two-player bargaining about how to split the value arriving at that edge. The trading tree structure affects the two-player bargaining that takes place on each edge in two ways: first, the disagreement point for the parent node is determined by the offers it negotiates with its other children, and second, the feasible set of splits depends on the revenue share on the edge connecting the parent node to *its* parent, because the parent node must pass up this fraction of the value that it receives from the split. Note that the revenue shares on these edges all influence each other, since the split of the value on one edge influences the bargaining power and therefore the split of the value on a different edge.

A natural choice for x_e , then, would be a *fixed point* to the system of bargaining equations, that is, a set of splits that are mutually consistent in the following sense: given the shares $x_{e'}$ on all remaining edges e' , the solution y_e to the two-player bargaining problem on any edge e with parameters specified by the remaining $x_{e'}$ is precisely x_e . It is not clear if such a fixed point exists, and even if it does, whether the final winner in a fixed point is the buyer with highest value.

Bargaining equations. The egalitarian, or proportional, bargaining solution [Kalai, 1977] for the two-player bargaining problem on the edge $e = (t, s)$, given the shares $x_{e'}$ on all other edges $e' \in T$, specifies that the parent and child node

each receive an equal incremental benefit from participating in the transaction.

Let C_s be the set of children of s in T . Let s' be the parent of s (if $s = r$, we consider a fictitious parent r' of r , with the revenue share on edge (r, r') always set to 0). Define $w_{s \setminus t} = \max_{t' \in C_s \setminus t} w_{t'} x_{t's}$. This is the maximum value that would reach s given a set of shares x if t did not exist as a child of s . Then, the two-player egalitarian bargaining solution on (t, s) specifies splitting w_t , the value reaching node t , according to x_{ts} where $x_{ts} \in [0, 1]$ satisfies

$$(1 - x_{ts})w_t = (1 - x_{ss'}) (\max(w_{s \setminus t}, w_t x_{ts}) - w_{s \setminus t}). \quad (1)$$

The left-hand side is the incremental benefit to node t from transacting with s : it receives a payoff of $(1 - x_{ts})w_t$ if it retains the edge with s , and nothing if it cuts off the edge. The right-hand side is the incremental benefit to the parent node s : if it retains the edge (t, s) , s can choose the highest payoff from C_s of which it will keep a $(1 - x_{ss'})$ share (since it needs to share this payoff with its parent); if it cuts off the edge (t, s) , it only gets $(1 - x_{ss'})$ times the highest payoff from the set $C_s \setminus t$.

The system of bargaining equations is given by writing (1) for all edges in the tree. A solution to this system is a fixed point of the bargaining game on the tree.

Note: It may seem that Equation (1) implicitly assumes that the parent node s indeed lies on the winning path because the payoff to s is $(1 - x_{ss'}) (\max_{t \in C_s} x_{ts} w_t)$ *only if* s lies on the winning path, and is 0 otherwise. However, we can show that (see full version [Ghosh *et al.*, 2013] for proof) when x is *fixed point* of these equations as opposed to an arbitrary set of shares, the right hand side is indeed s 's payoff irrespective of whether or not it lies in the winning path.

4 Reduction to path bargaining

The fixed point computation on the tree can be reduced to finding a fixed point of bargaining equations on a *single path*—the path from the least common ancestor of the highest value leaves to the root (if there is a unique leaf with highest value, this is the path from that leaf to the root). For want of space we omit this reduction.

We summarize the reduction as follows. Let $v^* = \max_l v_l$ be the maximum value in T . Find the least common ancestor s_0 of the leaves $\{l_1, \dots, l_k\}$ with $v_{l_i} = v^*$. Remove the entire subtree rooted at s_0 , and replace it with a fictitious buyer with value $d_0 = v^*$ at $l^* = s_0$.

Let the path from l^* to the root be of length n ; call this path P^* . We relabel nodes from l^* to the root $0, 1, \dots, n$ (so that l^* is 0 and the root is n). For $i \in [n]$, e_i is the edge connecting $i - 1$ to i . We can show (see full version [Ghosh *et al.*, 2013] for proof) that $x_e = 1$ for all other edges $e \in T$. So to each node $i = 1, \dots, n$, we can add a single edge with $x_e = 1$, to a fictitious buyer—this fictitious buyer's value is the largest value excluding v^* in the subtree rooted at i . Call this value d_i ; this is node i 's disagreement point, and we may also think of d_i as node i 's bid for the item being sold. We refer to this reduced instance as a path because the only edges with unknown revenue shares x_i lie on a path. Denote this new path bargaining instance by (P^*, \vec{d}) . Note that d_i is

strictly less than d_0 for $i = 1, \dots, n$. The following theorem summarizes this reduction:

Theorem 1. *Given an instance (T, \vec{v}) of the tree bargaining problem, construct the path bargaining instance (P^*, \vec{d}) as described above. Then, \vec{x} is a fixed point for T if and only if $x_e = 1$ for $e \notin P^*$, and the shares x_e , $e \in P^*$ constitute a fixed point to the path bargaining problem (P^*, \vec{d}) .*

5 Existence and Uniqueness of Fixed Point

We now investigate fixed points of the path bargaining problem, having shown that every tree bargaining instance can be reduced to a path bargaining instance.

Recall that the value at node 0 is d_0 and the remaining values at the leaves d_1, d_2, \dots, d_n are all strictly less than d_0 . The share on edge e_i is x_i . For notational convenience, we assume there is a fictitious edge e_{n+1} going up from the root to a fictitious node labeled $n + 1$ with share $x_{n+1} := 0$. For $i = 0, 1, \dots, n$, define $w_i = d_0 \prod_{j=1}^i x_j$, i.e. the value that reaches node i .

A fixed point solution $\vec{x} = \langle x_1, x_2, \dots, x_n \rangle$ satisfies the bargaining equations (2) for all edges i , with $x_i \in [0, 1]$: that is, it simultaneously solves the following system of equations, one for each edge e_i :

$$(1 - x_i)w_{i-1} = (1 - x_{i+1})(x_i w_{i-1} - d_i). \quad (2)$$

We note that in replacing the $\max\{x_i w_{i-1}, d_i\}$ term by $x_i w_{i-1}$ on the right-hand side of the bargaining equation, we have used the fact (see full version [Ghosh *et al.*, 2013] for proof) that we must have $w_{i-1} x_i \geq d_i$ in *any* fixed point x_i since $d_i < d_0$.

We can rewrite each bargaining equation in two ways: the “upward equation” gives x_{i+1} in terms of x_i :

$$x_{i+1} = 1 - \frac{(1 - x_i)w_{i-1}}{x_i w_{i-1} - d_i} = 1 - \frac{w_{i-1} - w_i}{w_i - d_i}. \quad (3)$$

The “downward equation” gives x_i in terms of x_{i+1} :

$$x_i = \frac{w_{i-1} + (1 - x_{i+1})d_i}{(2 - x_{i+1})w_{i-1}}. \quad (4)$$

Now we show that a fixed point to the path bargaining equations always exists, and is unique. The existence proof is via Brouwer's fixed point theorem. We show that the mapping f that is (essentially) obtained by simultaneous updates to the shares on all edges using the downward equations (4) is a continuous mapping from $[0, 1]^n$ to itself. The uniqueness proof requires more effort. We write two equations for x_n in terms of x_1 : one by using the upward equations (3) and one by using the downward equations (4). These equations can be represented by two curves, and any intersection point of the two curves leads to a fixed point. We next show that in the feasible range for the curves, one is strictly increasing, and the other strictly decreasing; thus there is a unique intersection point. We now formalize this.

First, we use the upward equations to write x_2, x_3, \dots, x_n in terms of x_1 and \vec{d} . However, not every value of $x_1 \in [0, 1]$ will give us to values of x_i in $[0, 1]$ and $w_i > d_i$. We will say

that x_1 is *feasible* if it does lead to $x_i \in [0, 1]$ and $w_i > d_i$. The following lemma (see full version [Ghosh *et al.*, 2013] for proof) characterizes some monotonicity properties of the x_i 's and w_i 's when written in terms of x_1 .

Lemma 1. *If $x_1' < 1$ is feasible, then for all $x_1 \in [x_1', 1]$, and for all $i = 1, 2, \dots, n$:*

1. $x_i \in [x_i', 1]$
2. $w_i > d_i$.
3. $\frac{dx_i}{dx_1} > 0$ (so x_i is strictly increasing as a function of x_1).
4. $\frac{dw_i}{dx_1} > 0$ (so w_i is strictly increasing as a function of x_1).

Here, x_i, w_i (resp. x_i', w_i') etc. are defined by x_1 (resp. x_1') using the upward equations (3).

Since $x_i = 1$ for all i is a feasible solution, in particular $x_1 = 1$ is feasible, and Lemma 1 immediately implies the following structure of the feasible region:

Lemma 2. *Let $x_1^\circ = \inf\{x_1 : x_1 \text{ is feasible}\}$. Then the feasible region for x_1 is either the interval $[x_1^\circ, 1]$ or $(x_1^\circ, 1]$, depending on whether x_1° is feasible or not.*

If x_1 is feasible, and x_2, \dots, x_n , are computed using the upward equations, then the balance conditions for edges e_1, e_2, \dots, e_{n-1} are automatically satisfied. The equation for e_n may not be satisfied, however. A fixed point is obtained precisely when x_n satisfies the balance condition for e_n . Geometrically, equations (3) and (4) for x_n define two curves, the upward curve, and the downward curve respectively. A fixed point is obtained at any intersection point of the two curves for x_n in the feasible region of x_1 . The following lemma (see full version [Ghosh *et al.*, 2013] for proof) gives monotonicity properties of the two curves:

Lemma 3. *In the feasible region for x_1 , the upward curve for x_n is strictly increasing, and the downward curve for x_n is strictly decreasing.*

We immediately get our uniqueness result:

Theorem 2 (Uniqueness). *If a fixed point to the equations (2) exists, then it is unique.*

Proof. This is immediate from Lemma 3: a strictly increasing and strictly decreasing curve can intersect in at most 1 point. \square

Finally, using Brouwer's fixed point theorem we can show (see full version [Ghosh *et al.*, 2013] for proof) that a fixed point always exists:

Theorem 3 (Existence). *A fixed point to the bargaining equations (2) exists.*

Briefly, we consider the following function $f : [0, 1]^n \rightarrow [0, 1]^n$, which represents a simultaneous update of the shares vector \vec{x} on all edges using the downward equations:

$$f_i(x) = \min \left\{ \frac{w_{i-1} + (1 - x_{i+1})d_i}{(2 - x_{i+1})w_{i-1}}, 1 \right\},$$

where $w_{i-1} = d_0 \prod_{j=1}^{i-1} x_j$ as usual, $x_{n+1} := 0$, and we make the convention that when $w_{i-1} = 0$, the first expression in the minimum above is $+\infty$, so that $f_i(x) = 1$. The above function is continuous, and its domain $[0, 1]^n$ is a convex,

compact set. By Brouwer's Fixed Point Theorem, f has a fixed point. The main work in the proof of Theorem 3 then consists in showing that any fixed point of f is a fixed point to the bargaining equations (2).

6 Properties of the Fixed Point

6.1 Core Property

The setting we study is naturally modeled as a cooperative game (T, V) , where the agents are the nodes in the trading tree T , and the coalition values V are defined as follows. The value of the coalition consisting of nodes on a path $p = (l, i_1, \dots, i_k, r)$ is $V(p) = v_l$. A coalition cannot generate value unless it contains a path from a leaf to the root; if it does contain such paths, its value is the maximum value amongst these paths: $V(S) = \max_{p \in S} V(p)$, and $V(S) = 0$ if S does not contain any such path p . Note specifically that $V(S) = 0$ for all sets that do not contain the seller r , and that $V(T) = V(p^*) = v^*$.

The *core* [Leyton-Brown and Shoham, 2008] of a cooperative game (N, V) is defined as a set of nonnegative payoff vectors (u_1, \dots, u_N) with $\sum u_i = V(N)$ such that every coalition's total payoff is at least as much as the value it generates: $\sum_{i \in N} u_i \geq V(S) \quad \forall S$. The core consists of the set of payoff vectors that are *not blocked* by any coalition which can increase its total payoff by splitting from the grand coalition and playing amongst themselves — an outcome not in the core is unlikely to occur in practice since there is a coalition that can benefit by deviating. In general, the core of a game can be empty, but our particular cooperative game does have a non-empty core, and in fact, our fixed point lies in the core. We can show the following theorem (see full version [Ghosh *et al.*, 2013] for proof):

Theorem 4. *The payoff vector u^* belongs to the core of (T, V) .*

6.2 Monotonicity

Monotonicity, which means that increasing the bargaining power of an agent increases his payoff, is a desirable property for a solution concept to our game. We establish a *strict* monotonicity property for the payoff to all nodes on the winning path in terms of their bargaining power. Since we are only interested in nodes on the winning path, we can restrict ourselves to discussing reduced path instance P^* .³ (We note that strict monotonicity cannot hold for nodes outside the winning path since the outcome itself must change for these nodes to receive a nonzero payoff; however, a weak monotonicity condition trivially holds.) We can prove the following strict monotonicity property (see full version [Ghosh *et al.*, 2013] for proof):

Theorem 5. *Consider the path bargaining problem. If any d_i is increased (but is still kept less than d_0) while the remaining $d_{i'}$ are unchanged, the payoff of i strictly increases.*

³When there is more than one leaf with value v^* , P^* does not contain all nodes on the winning path, but the strict monotonicity result extends easily to that case since an increase in bargaining power for a winning node not in P^* means that there is now a leaf with value greater than v^* .

6.3 Computability

We know that there exists a unique fixed point, x^* , of the bargaining equations. We now turn to computability of the fixed point. Note that since the shares affect the bids multiplicatively, the fixed point solution is *scale-free*: if we scale all bids by the same amount, the fixed point stays the same. So to simplify calculations, we assume that the maximum bid, d_0 , is normalized to 1, and all other bids d_i are less than 1. We can give a polynomial-time algorithm to compute an ε -fixed point: i.e., a set of shares such that all bargaining equations are satisfied within an additive ε error. For the original unscaled bids where the maximum bid may not be equal to 1, the additive error gets scaled by the maximum bid as well.

We now state our theorem (see full version [Ghosh *et al.*, 2013] for proof) regarding computability of an approximate fixed point. It is given in terms of a parameter $\gamma = \min\{1 - \max_{i>0}\{d_i\}, \frac{1}{n+2}\}$, which is essentially how close the second highest bid is to the maximum. Note that the dependence on the error parameter ε and γ is only poly-logarithmic. In practice, the algorithm converges extremely fast.

Theorem 6. *There is an algorithm that, for any given $\varepsilon > 0$, computes an ε -fixed point to the bargaining equations (2) in $\text{poly}(n, \log(1/\gamma\varepsilon))$ time.*

The algorithm essentially works by running a binary search to find the intersection point of the upward and downward curves for x_n . The parameter γ is important in giving bounds on the number of iterations needed in the binary search to obtain the desired accuracy, essentially by obtaining quantitative versions of the arguments of §5.

6.4 Dynamics

We have already seen in the previous sections that the solution prescribed by the fixed point of the bargaining equations has several desirable properties. A natural question is whether the agents on the tree would, without help from a centralized authority, be able to converge to this fixed point. We now present numerical evidence that this is indeed likely. Our experiments suggest that a natural dynamics consisting of *asynchronous updates*—where in each step a random edge e

updates x_e according to the two-player egalitarian bargaining equation (1), using the current values of $x_{e'}$ on other edges—indeed converges to the fixed point.

We run 10,000 tries of the following experiment: generate random bids at the leaves of a depth-8 balanced binary tree with 256 leaves and 510 edges; this is a convenient size that permits 10,000 tries to be run in a few hours. The bids are drawn from the lognormal distribution $e^{(1+N)}$ where N is the normal distribution with zero mean and unit variance. We initialize all 510 edge multipliers to the arbitrary value 0.99, and then repeatedly re-negotiate the edge multipliers one at a time in a random order: the negotiation for each edge consists of solving equation (2) for that edge (while freezing the values of all other multipliers). More specifically, binary search down to a tolerance of 1.0×10^{-15} is used to solve the equation. The edge updates are organized into “rounds” during each of which every edge is individually updated in a random order specified by a different random permutation for every round.

We continue iterating until the solution is close enough to the fixed point computed using the reduction to the path and the algorithm in §6.3. The efficient fixed point finding algorithm uses the reduction of §4 to convert the tree problem to a path problem. This path problem is then solved using a heuristic program (not described here) that uses the algorithm of §6.3 as a subroutine and computes the 8 multipliers to a nominal accuracy of 2.0×10^{-16} . The multipliers for the original tree are obtained by copying those 8 values onto the winning path and then setting the 502 multipliers lying on side branches to the value 1.0.

Every one of the 10,000 tries converged to the desired tolerance. The plot in Figure 1 shows the average convergence rate summarizing those 10,000 tries. It is clear that the shares always converge to within the desired accuracy at a reasonable rate. While we do not include the figures here, we also observed similar convergence behavior on trees with different structures and sizes, as well as for several different bid distributions.

7 Further Directions

In this paper, we defined a bargaining game on trees motivated by a fair division question in display ad exchanges, and investigated the properties of its fixed point. There are a number of interesting directions for further work. The most interesting open question is proving the convergence of dynamics, since numerical simulations strongly suggest that even asynchronous dynamics converge to the fixed point. Another interesting direction is that of a Bayesian model for values—suppose instead of values v_i at the leaves, we had *distributions* of values. The problem of solving the bargaining equations to set the shares x_e in this case is a very meaningful one, but also one that appears to be technically extremely challenging. Finally, there are questions related to extending the trade network model itself: for example, in this paper, we only consider a single seller and a tree topology. The question of how to model and solve for multiple sellers, and how the fixed point behaves if the underlying trade network is a directed acyclic graph instead of a tree, are also interesting

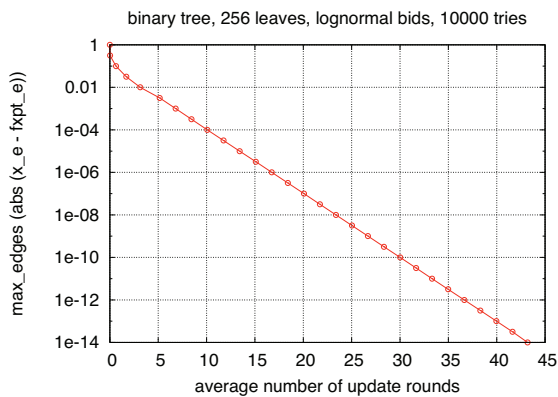


Figure 1: Asynchronous dynamics convergence: accuracy vs. average number of rounds to achieve accuracy.

directions for further work.

Acknowledgments. We are very grateful to Nikhil Devanur and Mohammad Mahdian for insightful discussions on network bargaining and cooperative games, and to Matt Jackson, Preston McAfee, Herve Moulin, Michael Schwarz, and Mukund Sundararajan for helpful comments and pointers to relevant literature.

References

- [Azar *et al.*, 2009] Y. Azar, B. E. Birnbaum, L. E. Celis, N. R. Devanur, and Yuval Peres. Convergence of local dynamics to balanced outcomes in exchange networks. In *FOCS*, pages 293–302, 2009.
- [Azar *et al.*, 2010] Y. Azar, N. R. Devanur, K. Jain, and Y. Rabani. Monotonicity in bargaining networks. In *SODA*, pages 817–826, 2010.
- [Blume *et al.*, 2007] L. Blume, D. A. Easley, J. M. Kleinberg, and É. Tardos. Trading networks with price-setting agents. In *EC*, pages 143–151, 2007.
- [Celis *et al.*, 2010] L. E. Celis, N. R. Devanur, and Y. Peres. Local dynamics in bargaining networks via random-turn games. In *WINE*, pages 133–144, 2010.
- [Chakraborty and Kearns, 2008] T. Chakraborty and M. Kearns. Bargaining solutions in a social network. In *WINE*, pages 548–555, 2008.
- [Chakraborty *et al.*, 2009] T. Chakraborty, M. Kearns, and S. Khanna. Network bargaining: algorithms and structural results. In *EC*, pages 159–168, 2009.
- [Feldman *et al.*, 2010] J. Feldman, V. S. Mirrokni, S. Muthukrishnan, and M. M. Pai. Auctions with intermediaries: extended abstract. In *EC*, pages 23–32, 2010.
- [Ghosh *et al.*, 2013] A. Ghosh, S. Kale, K. Lang, and B. Moseley. Bargaining for revenue shares on tree trading networks. *arXiv*, 2013.
- [Kalai, 1977] E. Kalai. Proportional solutions to bargaining situations: Interpersonal utility comparisons. *Econometrica*, 45(7):1623–30, October 1977.
- [Kanoria, 2010] Y. Kanoria. An FPTAS for bargaining networks with unequal bargaining powers. In *WINE*, pages 282–293, 2010.
- [Kleinberg and Tardos, 2008] J. M. Kleinberg and É. Tardos. Balanced outcomes in social exchange networks. In *STOC*, pages 295–304, 2008.
- [Leyton-Brown and Shoham, 2008] K. Leyton-Brown and Y. Shoham. *Essentials of Game Theory: A Concise, Multidisciplinary Introduction*. Morgan and Claypool Publishers, 2008.
- [Moulin, 2004] H. Moulin. *Fair Division and Collective Welfare*. The MIT Press, September 2004.
- [Yah, 2007] <http://advertising.yahoo.com/article/right-media.html>, 2007.