

Networks of Complements^{*†}

Moshe Babaioff¹, Liad Blumrosen², and Noam Nisan³

- 1 Microsoft Research, Herzelia, Israel
moshe@microsoft.com
- 2 Hebrew University, Jerusalem, Israel
blumrosen@gmail.com
- 3 Hebrew University, Jerusalem, Israel, and
Microsoft Research, Herzelia, Israel
noam@cs.huji.ac.il

Abstract

We consider a network of sellers, each selling a single product, where the graph structure represents pair-wise complementarities between products. We study how the network structure affects revenue and social welfare of equilibria of the pricing game between the sellers. We prove positive and negative results, both of “Price of Anarchy” and of “Price of Stability” type, for special families of graphs (paths, cycles) as well as more general ones (trees, graphs). We describe best-reply dynamics that converge to non-trivial equilibrium in several families of graphs, and we use these dynamics to prove the existence of approximately-efficient equilibria.

1998 ACM Subject Classification J.4. Economics, G.2.2 Network Problems

Keywords and phrases Complements, Pricing, Networks, Game Theory, Price of Stability

Digital Object Identifier 10.4230/LIPIcs.ICALP.2016.140

1 Introduction

Sellers typically do not operate in isolated markets but in conjunction with other sellers and buyers. In particular, sellers need to take into account the fact that buyers may value the goods they sell as substitutes or complements to goods sold by others. This has a tremendous impact on how sellers compete with each other. Indeed, in Cournot’s [8] famous paper from 1838 about sellers who compete through quantities, Cournot also describes a model of a duopoly selling perfect complements, zinc and copper. In Cournot’s example, a manufacturer of zinc may observe that some of her major customers produce brass (made of zinc and copper); Therefore, zinc manufacturers compete not only with other zinc sellers, but they also indirectly compete with manufacturers of copper, as both target the money of brass producers. We are interested in a more complicated competition structure, where as zinc can be also used for Galvanization of iron, zinc sellers compete at the same time with sellers of iron. In a similar way, iron is also demanded by car manufacturers that need to purchase glass from other sellers, and so on.

Another classic example is by Ellet [10], who studied how owners of two consecutive segments of a canal determine the tolls for shippers; Clearly, every shipper must purchase a permit from both owners for being granted the right to cross the canal. Another, more contemporary, example might be a high-tech or pharmaceutical firm that must use two

* The full version of this extended abstract is available at <https://arxiv.org/abs/1605.00136>.

† Liad Blumrosen and Noam Nisan were supported by the Israeli Science Foundation grant number 230/10.



© Moshe Babaioff, Liad Blumrosen, and Noam Nisan;
licensed under Creative Commons License CC-BY

43rd International Colloquium on Automata, Languages, and Programming (ICALP 2016).

Editors: Ioannis Chatzigiannakis, Michael Mitzenmacher, Yuval Rabani, and Davide Sangiorgi;

Article No. 140; pp. 140:1–140:14



Leibniz International Proceedings in Informatics

Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany



registered patents to manufacture its product; The owners of the two patents quote prices for the usage rights, and these patents can be viewed as perfect complements for the firm. As a final example, consider an international trader who wishes to export goods from country X to country Y, and needs to pay license fees to both countries. In the last two examples, one can see how licenses have a network structure, as each patent may be needed for the production of several different products, and trade may occur between country X and country Y, but also between country X and another country Z, etc.

In this paper, we study markets where goods are complementary to several other goods, but not substitutes. Price discrimination is impossible, and sellers need to offer the same price in all markets. This situation creates a global price competition between sellers, which raises interesting questions we aim to explore: What kinds of equilibria exist in these games? How efficient are the equilibria in this game? Will natural dynamics reach highly efficient equilibria?


The structure of the market plays a central role in our analysis. We model the market using a weighted undirected graph, where each vertex represents a seller of a certain good. An edge with weight v between vertices i and j indicates that there is a buyer that is willing to pay an amount v for the bundle of goods $\{i, j\}$.¹ In the above example, a market with sellers of copper, zinc, iron and glass can be represented as a path-graph with 4 vertices (See Figure 1), where edges connect copper and zinc, zinc and iron, and iron and glass.

We study the following simultaneous full-information pricing game between sellers on a graph: the sellers observe the values of the buyers, which are common knowledge; each seller then posts a single price to all buyers; a buyer on an edge buys the two goods on this edge if the total price (the sum of the prices of the two goods) is no larger than her value. Sellers have zero manufacturing costs and unlimited supply, and the profit of each seller is the price she posted times the number of buyers that accepted this price. In a (pure) Nash equilibrium, no seller would benefit from changing its price given the prices offered by the other sellers.

There are two natural benchmarks for measuring the quality of equilibria in such games. The first is the *maximum welfare*, which is the sum of values of all buyers.² This is the welfare that would be achieved with zero prices for all goods. The second benchmark is the *optimal monopolist revenue*, which is the optimal total revenue achievable by a monopolist that owns all the goods in the network. (Clearly, the optimal monopolist revenue is always at most the maximum welfare.) Several papers studied the problem of computing the optimal pricing for a monopolist in our setting. The problem was proposed by [12] which showed that it is APX-hard and presented an approximation algorithm that is logarithmic in the number of buyers. A $1/4$ -approximation algorithm was later presented by [3] when buyers are interested in bundles of size 2, and recently [19] showed that this bound is tight under some computational assumptions. A similar algorithm for this problem was suggested by [20] in the context of setting up peering connections in networks. [18] provided an improved bound for monopolist revenue maximization with buyers having the same value but interested in one or two items. Unlike these papers which focus on the monopolist's algorithmic problem, we focus on analyzing equilibria in the game between competing sellers and on how the welfare and revenue of these equilibria approximate the above benchmarks.

¹ Throughout this paper, for the simplicity of presentation, we consider buyers that demand bundles of size 2 which we can model using graphs; Buyers that demand bundles of arbitrary sizes can be modeled as hyper-edges on a hyper-graph. We show how some of our results extend to hyper-graphs in the full paper.

² Assuming 0 production costs for sellers, and payments cancel out.



■ **Figure 1** A market with 4 sellers and 3 buyers. The weight on each edge is the amount this buyer is willing to pay for the bundle of the two adjacent goods. For example, the middle edge represents a buyer that is willing to pay \$6 for a bundle of zinc and iron, but has value 0 for each good alone.

It is straightforward to see that some equilibria in these games demonstrate a complete market failure. For example, a vector of prices of ∞ posted by all sellers forms a Nash equilibrium that yields zero welfare and revenue. Yet, there might be multiple equilibria in these games, and one may hope that other equilibria perform better. Indeed, we prove that for some families of graphs the best equilibria (“price of stability”) have high revenue and welfare. Observe that we cannot hope for full efficiency as even for path graphs, the most efficient equilibria are sometimes not fully efficient: For example, in Figure 1, if either the seller of zinc or of iron offer a price above 1, then one of its edges will not be sold. However, there is no equilibrium in which these two sellers offer prices at most 1: if this occurs, each seller gains at most 2 (from selling to its two edges) but they can get revenue of at least 5 by offering 5 and selling to the middle edge only.

Another way to alleviate the problem of low welfare and revenue of some equilibria is by restricting attention to equilibria with natural properties. For instance, it may be unreasonable for a seller that does not sell any good to insist on a very high price despite having zero costs; Such behavior might be considered as a malicious behavior towards his neighboring sellers. We therefore sometimes restrict attention to *non-malicious* equilibria, where all sellers that are not selling at all offer a price of 0. Non-malicious equilibria serve as a main tool for proving the existence of approximately-optimal equilibria: we show that the revenue of every such equilibrium approximates the maximum welfare, and we show cases where such equilibria exist.

Our Results

Our main goal is to understand how the network structure affects various properties of equilibria in markets. We aim to understand how well can the welfare and revenue in equilibrium approximate the maximum welfare and the optimal revenue. We explore whether best-response dynamics can lead to approximately-optimal equilibria, and we prove the existence of such equilibria in several settings. We study graphs of different complexities, like paths, cycles, trees and general graphs.

We start by considering some special families of graphs (paths and cycles) and show that there is a natural best response dynamics that reaches equilibria with a constant fraction of the optimal welfare as revenue. For path graphs, we consider a dynamics starting with all sellers pricing at infinity. Then, we change the price of a seller on one side of the path to zero, and let each seller in turn best reply to current prices. We show that when this best-reply process reaches the other end of the path, it ends up in an equilibrium. Moreover, for at least one of the end points, the revenue when starting from that point is at least half the maximum welfare. An extension of this dynamics for cycles converges to an equilibrium with revenue of at least one quarter of the optimal welfare.

After we prove that for some families of graphs there exist Nash equilibria with good revenue guarantees, the immediate question is whether this generalizes to more complex graphs. It turns out that non-maliciousness of equilibria can be used to prove bounds on the revenue. In particular, we show that for any tree, every non-malicious equilibrium has revenue that is at most $O(\log d)$ factor smaller than the optimal welfare, where d is the maximum degree. One might hope that such a result can be achieved for general graphs as well, but unfortunately this is not the case. For a clique of degree d , the loss can be linear in d ,³ and therefore we cannot hope for a logarithmic fraction for general graphs. Our main positive result presents a refined bound that depends not only on the maximum degree, but also on the *arboricity* of the graph (see, e.g., [22]). This additional parameter nicely captures the difference between trees and cliques (and other graphs). The arboricity of an undirected graph is the minimum number of forests into which its edges can be partitioned. (For example, a tree has arboricity 1, a cycle has arboricity 2, and a clique of n nodes has arboricity $n/2$.) We prove the following bound on the revenue of any non-malicious equilibrium:⁴

► **Theorem.** *In every graph with maximum degree d and arboricity w and every non-malicious equilibrium in it, the total revenue of all sellers is at least $\Omega(\frac{1}{w+\log d})$ of the maximum welfare.*

The above theorem does not claim anything about the existence of non-malicious equilibria, but only bounds the revenue obtained by such equilibria if they exist. For this result to imply a bound on the “price of stability” in such games,⁵ one needs to show that non-malicious equilibria actually exist. We prove the existence of non-malicious equilibria via the natural heuristic of repeated best-reply dynamics. Such a dynamics starts with arbitrary prices; At each step, a seller who is not best replying is chosen, and he updates his price to a best reply (breaking ties non-maliciously - towards 0 price). Our main result in this context shows that in tree graphs, a specific sequence of best replies does stop at a non-malicious equilibrium.

► **Theorem.** *In every tree, for every initial profile of prices, there exists a sequence of best replies that terminates in a non-malicious Nash equilibrium.*

In particular, this result implies the existence of non-malicious equilibria in trees, and together with the approximation theorem above, we conclude that the price of stability in trees (i.e., the approximation achieved by the best equilibrium in every tree) is at least $\Omega(\frac{1}{\log d})$. Of course, it would be interesting to strengthen this result and show that for any graph such dynamics terminate in a (non-malicious) equilibrium. Based on simulations we have executed, we conjecture that this is indeed the case.

► **Conjecture.** *In every graph any non-malicious best response dynamics starting from any price vector converges to an equilibrium in polynomial number of steps.*

This is the strongest version of our conjecture. One can also try to prove weaker versions of the conjecture by restricting the graph structure, order of plays, and relaxing the required time to convergence.

³ Consider a clique of degree d with all edges with the same value of 1. Such a clique has a non-malicious equilibrium in which one seller prices at 0 while all others price at 1. The welfare in this non-malicious equilibrium is only linear in d , while the optimal welfare is quadratic in d , and thus there is a linear loss in d .

⁴ We note that the theorem holds also for non-simple graphs that have parallel edges.

⁵ Our result above also has a “price-of-anarchy” flavor, as we prove that all non-malicious equilibria (when exist) exhibit the approximation guarantee.

We also prove several impossibility results. We first show that the $\Omega(\frac{1}{\log d})$ approximation for trees is tight, by presenting trees in which the welfare of every non-malicious equilibrium is at most an $\Omega(\frac{1}{\log d})$ fraction of the monopolist revenue (and thus of the optimal welfare). However, this bound is proved using constructions for which there are other (“malicious”) equilibria that achieve a constant approximation. We therefore strengthen this impossibility result and show instances where *all* equilibria, malicious or not, achieve bad results. In particular, we would like to find out whether the revenue of the best equilibrium is always a constant fraction of the monopolist’s revenue. Our main lower bounds show that the answer is negative, even for trees.

► **Theorem.** *For graphs with maximum degree d :*

- *There are graphs for which the welfare in every Nash equilibrium is at most $O(\frac{1}{\sqrt{\log d}})$ fraction of the monopolist’s revenue.*
- *There are trees for which the welfare in every Nash equilibrium is at most $O(\frac{1}{\log \log d})$ fraction of the monopolist’s revenue.*

Note that the fact that the revenue in every equilibrium is low compared to the monopolist revenue, implies that it is also low compared to the maximum welfare. Moreover, our bounds are actually stronger, showing that not only the revenue in equilibrium is low, but also the welfare.

Related Literature

The analysis of price competition goes back to Cournot [8] and Bertrand [4]. The famous competition model of [8] describes producers who compete through quantities, but in the same work he also described a model of a duopoly selling perfect complements.

[21] consider buyers interested in perfect complements, where each bidder is interested in one specific bundle (single-minded bidders); They study mechanism design for a single seller and buyers with private information, where our focus is on competition between multiple sellers with complete information. Auctions for networks of buyers and sellers with private information were studied in [17].

There is a line of literature studying interactions of sellers and buyers over networks. Among them, [16] considered bargaining in networks, where agents can choose whom they want to negotiate with and the solution implies matching of buyers and sellers. In our model, neighbouring sellers also “negotiate” on dividing the value of the buyer on the edge between them, but unlike [16], a seller can serve several neighbouring buyers simultaneously (with the constraint of non-discriminatory pricing). Sub-game perfect equilibria in bargaining games were studied in [7]. [13] studied general equilibrium models on networks with linear utilities of buyers.

[2] considered a similar model of sellers that compete in prices, and a trade network that is modeled as graph, with one main difference from this paper: items are substitutes for the buyers. Namely, a buyer is interested in buying *either* from one seller or from the other seller on the edge, while here we consider buyers that are interested in *both*. This difference implies significant differences in results. For example, there, unlike our paper, pure Nash equilibria hardly ever exist. For trees, equilibrium utilities of the sellers are uniquely defined, while here they are not.⁶

⁶ Our framework can be also viewed as a variant of graphical games [15]. There are several known algorithms for computing equilibria in graphical games (see survey [14]). Another related family of

2 Model

We study trade networks which are modeled as weighted undirected graphs, where sellers are represented by nodes and buyers by weighted edges. The set of edges is denoted by E , let N be the set of nodes and let $i, j \in N$ denote generic sellers. Sellers have no costs and can supply any number of items. The sellers are the players in the game, interested in maximizing revenue. Each buyer is single minded and is interested in the bundle of items sold by both sellers that lie on that edge. The weight $v_{i,j}$ on the edge represents the value of this buyer for buying the bundle $\{i, j\}$. In particular, buyers gain zero value from buying each item alone. A seller i posts a single price $p_i \geq 0$ that will be available to all incident edges (cannot price discriminate between buyers). A buyer of bundle $\{i, j\}$ buys if and only if $v_{i,j} \geq p_i + p_j$.

For a given price vector p we denote the set of sold edges by $S(p)$. The edge (i, j) is *tight* if $p_i + p_j = v_{i,j}$ (the sum of prices of the sellers on the edge equals the value of the edge). We say that *an edge (i, j) has slack* if $p_i + p_j < v_{i,j}$. For a given seller i and an edge (i, j) , *the slack of seller i on that edge* is $v_{i,j} - p_j$ (that is, it equals to the difference between the value of the edge and the price of the *other* seller on that edge).

For a given price vector p , the *revenue* r_i of seller i with price p_i is $p_i \cdot |\{j | (i, j) \in S(p)\}|$, that is, p_i for every adjacent edge that is sold. The *total revenue* is $\sum_i r_i$ and the *welfare* is $\sum_{(i,j) \in S(p)} v_{i,j}$.

The sellers compete in a game in which they simultaneously post prices. These prices form a (*pure*) *Nash equilibrium (NE)* if each seller maximizes his revenue given the prices of all other sellers. In this paper we only consider pure Nash equilibria.

For a given network, the *revenue of the monopolist* is the supremum of the total revenue over all price vectors for the sellers. The *maximum welfare* is simply the sum of values of all edges (this can be obtained when every seller prices at 0). We study the ratio between the revenue in equilibrium and either the revenue of the monopolist or the maximum welfare (whichever bound is harder to prove): For our positive results we show that the equilibrium revenue is some fraction of the maximum welfare (which clearly implies approximation to the monopolist revenue), while for our negative results we prove inapproximability of the monopolist revenue (which implies the same result with respect to the maximum welfare).

Non Malicious Equilibria. We have already observed that the welfare of the worst equilibrium is arbitrarily low. Thus, we sometimes study an equilibrium refinement in which sellers with 0 utility price at 0. We say that a losing seller (seller with 0 utility) is *non malicious* if she prices at 0. A price vector for the sellers is called *non malicious* if every losing seller is non malicious. This, in particular, implies that every seller's revenue is at least as high as the price he sets. We use the concept of non-malicious NE when proving our main result: we first show that such equilibria always approximate the optimal welfare, and we then show that they always exist in trees. This implies a positive result on the price-of-stability in trees.

games is Polymatrix games (see [5, 9, 6]), where an action of a player is played in several simultaneous bimatrix games (the games in our paper have continuous action spaces); This class of games played an important role in showing hardness of equilibrium computation results. Our goal in this paper is to find equilibria with good economic properties, and compare their properties to the optimal non-strategic solution. Our paper shows that allowing losing sellers to price "maliciously" might be beneficial to society as a whole, this phenomenon was termed the "Windfall of Malice" by [1].

3 Special Families of Graphs: Paths and Cycles

In this section, we present positive results for special cases of graphs and algorithms that compute Nash equilibria with high revenue for those cases. These algorithms consist of a particular sequence of best responses of the sellers. We present a linear time algorithm for path graphs that finds an equilibrium that obtains revenue that is at least half of the maximum welfare. We use it to present a linear time algorithm for cycles that computes an equilibrium with revenue at least a quarter of the maximum welfare.

3.1 Path Graphs

We first consider the simplest non-trivial graphs: path graphs. Our algorithm (Algorithm 1) assumes that sellers are indexed from left to right, with the leftmost seller indexed by 1. The algorithm starts with all sellers pricing at ∞ and then changes the price of the leftmost seller to 0. Then, it goes over the sellers from left to right, letting each seller to choose a best response. If the seller cannot get positive utility, he prices at the value of the previous edge (not at 0 - and thus the equilibrium might be malicious). The way we break ties for sellers with 0 utility ensures that after seller i best responds, all previous sellers are still best responding (as seller i leaves no slack to the previous seller). Our algorithm may achieve terrible revenue when executed in a certain direction on the path (for example, when the weights of the edges are monotonically decreasing), but in such cases we show that running it from for the opposite direction will perform well. This is essentially since every time the value of the edge is larger than the previous edge, it is sold tightly. The algorithm thus picks the better equilibrium out of executing the above procedure starting from one end, or from the other.⁷

Algorithm 1 A $\frac{1}{2}$ -Approximation for path graphs.

1. Initialize all prices to infinity, and the price of the first (leftmost) seller to 0.
 2. Starting from the second seller, go over the sellers from left to right and let each seller best respond to the current prices. If every price of the current seller i gets him 0 utility, set his price to the value of edge $(i - 1, i)$ (the value of the previous edge).
 3. Mirror the path left to right (mapping node i to become node $n + 1 - i$) and repeat the above algorithm. Output one of these two price vectors with the higher total revenue.
-

This algorithm terminates in a NE with revenue of at least half the maximum welfare (proof appears in the full paper):

► **Proposition 1.** *For any path graph, Algorithm 1 terminates in a NE after a linear number of steps. The total revenue in this equilibrium is at least half of the maximum welfare.*

3.2 Cycles

Our next result shows that in any cycle graph there is always a NE with high revenue. Moreover, we show how to compute such a NE in linear time; The algorithm runs the above

⁷ We illustrate the algorithm for the path graph of Figure 1: The algorithm chooses a price of 0 for the copper seller, then a price of 1 for the zinc seller, a price of 5 for iron and a price of 1 for glass. This is a NE with revenue 7 (out of maximum welfare of 8). By symmetry, the revenue starting from the other end would be the same.

path-graph algorithm for one full round on all sellers, and then initiates a best-response dynamic that will end up in such an equilibrium.

Fix an arbitrary seller and mark him as seller 1. Going clockwise along the cycle starting from seller 1, number the other sellers 2 to n . We present the algorithm as Algorithm 2.

Algorithm 2 A $\frac{1}{4}$ -Approximation for Cycles

1. Initialize all prices to be infinity, initialize the price of seller 1 to 0.
 2. Going clockwise: Starting from the second seller till the first seller (inclusive), go over the sellers clockwise and let each seller best response to the current prices. If every price of the considered seller i gets him 0 utility, set his price to the value of edge $(i - 1, i)$ (the value of the previous edge in the cycle, with edge $(n, 1)$ being the edge prior to 1).
 3. Run a best response dynamics (by iteratively letting a seller that is not best responding to update his price) till it stops.
 4. Mirror the cycle and repeat the above algorithm (i.e. the new order after seller 1 is $n, n - 1, \dots, 2, 1$). Output one of these two price vectors with the higher total revenue.
-

► **Theorem 2.** *For any cycle, Algorithm 2 terminates in a Nash equilibrium after linear number of steps. The total revenue in this equilibrium is at least one quarter of the maximum welfare.*

The details of the proof can be found in the full paper. We next present an informal description of the algorithm together with an outline of the proof. First, the algorithm essentially runs Algorithm 1 in one direction, starting with a price of 0 for seller 1 and going clockwise, letting each seller best respond (pricing at the value of the prior edge if his utility is 0) including the first seller as the last seller to change his price. We note that the first seller is now best responding (as the edge to the second seller had no slack, so he could not get any utility out of it). After that cycle completes, we argue that there is at most one seller (the second seller) that is not best responding. From this point on, the algorithm runs a best response dynamics till it stops. We argue that at each point, the only seller that might not be best responding is the next seller, and once he best responds, either the sellers are in a NE or the edge with the previous seller is tight. As we show, it follows that this dynamics will stop after a linear number of steps. We also show that the algorithm running either clockwise or counter-clockwise gets revenue that is at least one quarter of the maximum welfare: the path-graph algorithm gets at least half, and from then edges are only added (and sold tightly), except the edge following the last seller to update his price that might be dropped. We show that the revenue of the dropped edge is no larger than the remaining revenue, thus we lose at most half the revenue obtained by the path algorithm.

4 A Positive Result

We start by presenting our main positive result, showing that the revenue in some NE is at least some fraction of the maximum welfare (and thus also some fraction of the monopolist revenue). We prove this claim by showing that this is true for any non-malicious NE. For any network for which a non-malicious equilibrium exists (like in trees, see Theorem 9), this implies that at least the same fraction can be obtained in the highest revenue equilibrium.⁸

⁸ This result actually holds for a more general graph model, where some buyers are interested in a single good, not a pair, and also for graphs that contain parallel edges (that is, there might be multiple buyers

The bound we present is in terms of the maximum degree and the arboricity of the graph. The *maximum degree* is the largest degree of any node: the maximal number of buyers that demand an item from the same seller. The *arboricity* of an undirected graph is the minimum number of forests into which its edges can be partitioned. Equivalently, it is the minimum number of spanning forests needed to cover all the edges of the graph.

To gain some intuition for the notion of arboricity, here are some simple examples. A forest (and a tree) has arboricity 1. A cycle has arboricity 2 as it clearly cannot be spanned by a single forest, but it can be partitioned into two trees, a spanning tree and the missing edge. A clique of size n has arboricity $\lceil n/2 \rceil$: the arboricity is at least $n/2$ as any tree has at most $n - 1$ edges and the clique has $n(n - 1)/2$ edges, and it is at most $\lceil n/2 \rceil$ as it can be covered by $\lfloor n/2 \rfloor$ Hamiltonian paths, plus a star for odd n .

► **Theorem 3.** *In every graph with maximum degree d and arboricity w and every non-malicious NE in it, the total revenue of all sellers is at least $\frac{1}{2(w+1+\ln d)}$ fraction of the maximum welfare.*

Proof. Fix a graph and a non malicious Nash equilibrium in it. We will say that a vertex u is *low for the edge* (v, u) if u 's price in the equilibrium is at most half of the value of the edge (v, u) . We will say that an edge is *all-high* if both of the vertices on it are not low for the edge (in this case clearly the edge does not buy). Let E^H denote the set of all-high edges. For a vertex v let E^v denote the set of edges (v, u) (edges adjacent to v) such that u is low for (v, u) . Observe that since every edge is either all-high or low for some node, the set of all edges E is covered as follows: $E = E^H \cup (\cup_v E^v)$. Thus

$$\sum_{(v,u) \in E} v_{(v,u)} \leq \sum_{(v,u) \in E^H} v_{(v,u)} + \sum_v \sum_{(v,u) \in E^v} v_{(v,u)}.$$

Denote the total revenue by r . To complete the proof we bound each of the two terms separately, the first by $2w \cdot r$ and the second by $2(\ln d + 1) \cdot r$, together proving the theorem.

We start by bounding the edges that are not all-high. Let r_v denote the revenue of v . This claim is well known (e.g., [11]) and we present the proof for completeness in the full paper.⁹

► **Claim 4.** *For every node v it holds that $\sum_{(v,u) \in E^v} v_{(v,u)} \leq r_v \cdot 2(\ln d + 1)$*

By the claim we derive that

$$\sum_v \sum_{(v,u) \in E^v} v_{(v,u)} \leq \sum_v r_v \cdot 2(\ln d + 1) = 2(\ln d + 1) \cdot r.$$

We next bound the total value of the all-high edges. To take care of the total weight of the all-high edges we will use the fact that in a graph of arboricity w there exists a mapping from edges to vertices such that every edge is mapped to one of its two vertices and no vertex has more than w edges mapped to it (this is true as we can just root each tree and map every edge to its child node). Since the edge is all-high and the equilibrium is non malicious, the price – and thus also the revenue – of each of the two vertices on the edge is at least half the value of the edge. Summing, again, over all vertices, we get that the total weight of all

interested in each pair of goods).

⁹ The claim essentially says that a single price can gain revenue of at least a logarithmic fraction of the total demand (in our case, it is the residual demand given the prices of the neighbours- which is at least half the value of each edge as these edges are not all-high.).

all-high edges is at most $2w$ the total revenue of all vertices. Stated formally, consider the mapping M from E^H to the set of nodes N that maps each edge to an adjacent node and never maps more than w edges to the same node. For every node v that incident at least one high edge, define $u^*(v)$ to be a vertex such that for all u such that $(v, u) \in E^H$ we have $v_{(v, u^*(v))} \geq v_{(v, u)}$. Then, it holds that

$$\sum_{(v, u) \in E^H} v_{(v, u)} \leq \sum_v \sum_{u | M(v, u) = v} v_{(v, u)} \leq \sum_{v | \exists u M(v, u) = v} w \cdot v_{(v, u^*(v))} \quad (1)$$

$$\leq w \cdot \sum_v 2 \cdot r_v \leq 2w \cdot r \quad (2)$$

◀

For any network for which a non-malicious equilibrium exists, the theorem ensures that some Nash equilibrium has high revenue. Thus, it can be viewed as a “Price of Stability” result for such networks. In particular, it bounds the price-of-stability for trees, as for any tree a non-malicious equilibrium exists by Theorem 9. It also ensures that every non-malicious equilibrium has high revenue, thus can be viewed as a “Price of Anarchy” result for non-malicious equilibria.

The theorem implies a bound on trees that is only logarithmic in the maximum degree.

► **Corollary 5.** *In any tree with maximum degree d and every non-malicious equilibrium in it, the total revenue of all sellers is $\Omega(1/\ln d)$ fraction of the maximum welfare.*

5 Impossibility Results


Theorem 3 gives a positive result, ensuring that the revenue in equilibrium is some fraction of the welfare. Yet, it might be possible that an improved bound can be shown. Specifically, we would like to answer the following question: Is the revenue of best equilibrium always a constant fraction of the revenue of the monopolist? Unfortunately, the answer is no, even for trees. We first present a lower bound for general graphs. This bound is slightly weaker than the logarithmic bound (in the maximum degree) of Theorem 3. We then present a lower bound for trees, showing that the best equilibrium revenue is not necessarily a constant fraction of the monopolist revenue. For missing proofs see the full paper.

5.1 General Graphs

We start with the lower bound for general graphs. For proving the theorem, we construct graphs with arboricity that is much smaller than their maximum degree.

► **Theorem 6.** *There exists a family of graphs with maximum degree d and arboricity w for which $w^2 = O(\ln d)$ and the revenue that a monopolist seller can get is factor $\Omega(w)$ larger than the revenue (and welfare) in any Nash equilibrium. In terms of d , the factor can be as large as $\Omega(\sqrt{\ln d})$ when $w^2 = \Theta(\ln d)$.*

The bound is proven using the following construction (see Figure 2). There is a clique of size $w + 1$ and any node in the clique is connected to an “harmonic gadget”: d edges with values $1, 1/2, 1/3, \dots, 1/d$. The value of an edge connecting two nodes in the clique is 4. We first claim that for these parameters, a monopolist would price all the clique nodes at 0 and get full revenue from all the harmonic gadgets. In any NE, however, at most one seller prices below $1/w$, thus not much revenue is gained from the harmonic gadgets.



■ **Figure 2** The construction of Theorem 6. A clique of $w + 1$ sellers, each edge in it is of value 4. Each node in the clique is connected to an “harmonic gadget”.

5.2 Trees

For trees, the above bound (when $w = 1$) is a constant. We next present a different construction that shows that for some trees, the revenue of a monopolist can be factor $\Omega(\ln \ln d)$ larger than the revenue of the best equilibrium. In particular, it shows that a constant upper bound is impossible.

► **Theorem 7.** *There exists a family of trees with maximum degree d for which the revenue that a monopolist seller gets is factor $\Omega(\ln \ln d)$ larger than the revenue (and welfare) in any NE.*

For proving this impossibility result we construct the following graph, fixing an integer m (to be determined later). Consider a path with $2m + 1$ edges. The first edge has value 5. For any $j = 1, 2, \dots, m$, given that edge $2j - 1$ has value v , set the value of the edge $2j$ to be $2v + 2$ and the value of edge $2j + 1$ to be $2v + 6$. Any node of even index is additionally connected to an “harmonic gadget” with $d - 2$ spikes: $d - 2$ edges with values $1, 1/2, 1/3, 1/4, \dots, 1/(d - 2)$.

We first argue that in any Nash equilibrium there is at most a single seller of even index on the path with price that is at most 1. We then argue that a monopolist seller can get revenue of at least $m \ln(d - 2)$ by pricing every seller on the path at 0 and gaining all the revenue from the m harmonic gadgets. We conclude that the ratio of the monopolist revenue to equilibrium revenue (and welfare) is at least $\frac{\max\{m \ln(d - 2), 2^m\}}{3 + \ln d + 2^{m+4}}$. For $m = \Theta(\ln \ln d)$ this ratio tends to $\Omega(m)$ as we aimed to prove. Full details appear in the full paper.

For trees, there is still a gap between this $\Omega(\ln \ln d)$ bound and the upper bound of $O(\ln d)$ for non-malicious NE as implied by Corollary 5. Closing this gap is left as an open problem.

5.3 Non malicious Nash equilibrium

We have observed that the worst Nash equilibrium might have zero revenue, even for a single edge. In this section, we consider non-malicious NE and present two simple examples that show that the upper bound of Theorem 3 is tight in both parameters for the non-malicious NE with the worst revenue (“Price of Anarchy”).

► **Proposition 8.** *The following holds:*

- *For any $w \geq 2$ there exists a graph (symmetric clique) with $2w$ nodes and arboricity w for which the revenue in some non-malicious equilibrium is smaller than the monopolist revenue by a factor of at least w . Moreover, for that graph, every best-response dynamics starting at zero prices converges to such an equilibrium.*

- For any $d \geq 2$ there exists a tree (star) with maximum degree d for which the revenue in some non-malicious NE is smaller than the monopolist revenue by a factor of at least $\ln d$.

The two claims directly follow from the following simple examples. Consider first the clique graph of $n = 2w$ nodes with all weights being 1. Note that such a clique has arboricity w . Here is a non-malicious Nash equilibrium: one of the sellers prices at 0 and all others price at 1. The total revenue (and total welfare in this NE) is exactly $(n - 1)$. Note that there exists a fully-efficient fully-revenue-extracting non-malicious equilibrium where each seller prices at $1/2$. In this NE the revenue (and welfare) is exactly $n(n - 1)/2$. Thus, the ratio between the best non-malicious equilibrium revenue and the worst non-malicious equilibrium revenue is $n/2 = w$. Finally, observe that in any best response dynamics starting at 0, sellers keep increasing prices from 0 to 1, except when there is only one seller pricing at 0, and the others at 1. Thus, any best response dynamics starting at 0 prices will end at a non-malicious NE in which one seller prices at 0 and all others price at 1.

Consider now a star with d spikes, with edge i of value $1/i$. There is a non-malicious NE with revenue of 1: the center prices at 1 and all other price at 0. Revenue (and welfare) of $\sum_{i=1}^d 1/i \geq \ln d$ can be achieved by pricing the center at 0 and any other seller at the price of its edge. The second claim follows.

6 Best Reply Dynamics

Our main positive result (Theorem 3) ensures that for any graph for which a non-malicious equilibrium exists, there is an equilibrium with high revenue. Thus, one naturally wonders if all graphs admit a non-malicious equilibrium. If so, can one find such an equilibrium in polynomial time? Is there a natural dynamics that ends in such an equilibrium?

A natural procedure for converging to an equilibrium is repeated best-reply dynamics, and one might hope that such dynamics will indeed always converge to a non-malicious equilibrium in polynomial time. Such a dynamics starts with arbitrary prices, and at each step a seller that is not best replying is chosen, and he updates his price to a best reply. If the goal is finding a non-malicious equilibrium, then one needs to consider a seller with 0 utility that is not pricing at 0 as a seller that is not best replying. We conjecture that for any graph, such a process terminates in an equilibrium (which is clearly non-malicious). A stronger conjecture is that such an equilibrium will be reached in polynomial time.¹⁰

In this section, we prove the existence of non-malicious equilibria in *trees* by presenting some best-reply dynamic that converges to equilibrium in a finite number of iterations.¹¹

► **Theorem 9.** *In every tree, for every initial profile of prices, there exists a sequence of player best replies that terminates in a non-malicious Nash equilibrium.*

The full proof appears in the full paper. We now present a sketch of the proof, and recursively define the sequence of best responses. We pick a leaf u of the tree that is connected to the rest of the tree via vertex v . Let x_0 be the initial price of u and y_0 be the initial price of v . Our best reply dynamics will proceed by repeatedly, for $i = 1 \dots$, let x_i be u 's reply to y_{i-1} , and then recursively use a best reply sequence that updates the rest of the tree, assuming that u 's price is set to x_i . Note that this recursive best-reply sequence starts with

¹⁰We remark that we have simulated best response dynamics on general graphs with arbitrary order, and they always terminated in equilibrium very fast.

¹¹Our proof only ensures that the dynamics will stop, but does not ensure terminating in polynomial time.

v updating his price (since all the other vertices are already best-replying from the previous recursive call, but then other vertices may update their price and v may update again, and so on). The v 's price at the end of the recursive call is called y_i . The recursive call on the sub-tree terminates due to an inductive use of the theorem (i.e., the theorem is proved by induction on the number of vertices in the tree). To ensure that the induction hypothesis applies to the recursive call which is applied not just to a subtree, but rather to a subtree to which an extra leaf u with a *fixed* value x_i is attached, we prove the theorem (inductively) also for trees in which each vertex may have an arbitrary number of leaves with a fixed value attached to them.

References

- 1 Moshe Babaioff, Robert Kleinberg, and Christos H. Papadimitriou. Congestion games with malicious players. *Games and Economic Behavior*, 67(1):22–35, 2009.
- 2 Moshe Babaioff, Brendan Lucier, and Noam Nisan. Bertrand networks. In *ACM Conference on Electronic Commerce (ACM-EC)*, 2013.
- 3 Maria-Florina Balcan and Avrim Blum. Approximation algorithms and online mechanisms for item pricing. In *Proceedings of the 7th ACM Conference on Electronic Commerce, EC'06*, pages 29–35, 2006.
- 4 Joseph Louis François Bertrand. theorie mathematique de la richesse sociale. *Journal de Savants*, 67:499–508, 1883.
- 5 Yang Cai and Constantinos Daskalakis. On minmax theorems for multiplayer games. In *Proceedings of the Twenty-second Annual ACM-SIAM Symposium on Discrete Algorithms, SODA'11*, pages 217–234, 2011.
- 6 Xi Chen, Xiaotie Deng, and Shang-Hua Teng. Settling the complexity of computing two-player nash equilibria. *J. ACM*, 56(3):14:1–14:57, 2009.
- 7 Margarida Corominas-Bosch. Bargaining in a network of buyers and sellers. *Journal of Economic Theory*, 115(1):35–77, 2004.
- 8 Antoine Augustin Cournot. *Recherches sur les principes mathematiques de la theori des Richesses*. Trans. N.T. Bacon, New York: Macmillan, 1929, 1838.
- 9 Argyrios Deligkas, John Fearnley, Rahul Savani, and Paul Spirakis. Computing approximate nash equilibria in polymatrix games. In *Web and Internet Economics: 10th International Conference, WINE 2014*, pages 58–71, 2014.
- 10 C. Ellet. *An essay on the laws of trade in reference to the works of internal improvement in the United States*. Reprints of economic classics. A.M. Kelley, 1966.
- 11 Andrew V. Goldberg, Jason D. Hartline, Anna R. Karlin, Michael Saks, and Andrew Wright. Competitive auctions. *Games and Economic Behavior*, 55(2):242–269, 2006.
- 12 Venkatesan Guruswami, Jason D. Hartline, Anna R. Karlin, David Kempe, Claire Kenyon, and Frank McSherry. On profit-maximizing envy-free pricing. In *Proceedings of the Sixteenth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA'05*, pages 1164–1173, 2005.
- 13 Sham M. Kakade, Michael Kearns, Luis E. Ortiz, Robin Pemantle, and Siddharth Suri. Economic properties of social networks. In *Advances in Neural Information Processing Systems 17*, pages 633–640. MIT Press, 2005.
- 14 Michael Kearns. In *Noam Nisan, Tim Roughgarden, Eva Tardos and Vijay Vazirani (Editors), Algorithmic Game Theory. Chapter 7. Graphical Games*. Cambridge University Press, 2007.
- 15 Michael J. Kearns, Michael L. Littman, and Satinder P. Singh. Graphical models for game theory. In *Proceedings of the 17th Conference in Uncertainty in Artificial Intelligence, UAI'01*, pages 253–260, 2001.

- 16 Jon Kleinberg and Éva Tardos. Balanced outcomes in social exchange networks. In *Proceedings of the Fortieth Annual ACM Symposium on Theory of Computing*, STOC'08, pages 295–304, 2008.
- 17 E. Kranton and Deborah F. Minehart. A theory of buyer-seller networks. *American Economic Review*, 91:485–508, 2001.
- 18 Robert Krauthgamer, Aranyak Mehta, and Atri Rudra. Pricing commodities. *Theor. Comput. Sci.*, 412(7):602–613, 2011.
- 19 Euiwoong Lee. Hardness of graph pricing through generalized max-dicut. In *Proceedings of the Forty-Seventh Annual ACM Symposium on Theory of Computing*, STOC'15, pages 391–399, 2015.
- 20 E.w Lee, D. Buchfuhrer, L. Andrew, A. Tang, and S. Low. Progress on pricing with peering. In *Proceedings of the 45th Annual Allerton Conference on Computing, Communications and Control*, Allerton'07, pages 286–291, 2007.
- 21 Daniel Lehmann, Liadan Ita O'Callaghan, and Yoav Shoham. Truth revelation in approximately efficient combinatorial auctions. In *JACM 49(5)*, pages 577–602, Sept. 2002.
- 22 C. St.J. A. Nash-Williams. Edge-disjoint spanning trees of finite graphs. *Journal of the London Mathematical Society*, s1-36(1):445–450, 1961.