Price Competition in Networked Markets: How Do Monopolies Impact Social Welfare?

Elliot Anshelevich and Shreyas Sekar(✉)

Rensselaer Polytechnic Institute, Troy, NY, USA
eanshel@cs.rpi.edu, sekars@rpi.edu

Abstract. We study the efficiency of allocations in large markets with a network structure where every seller owns an edge in a graph and every buyer desires a path connecting some nodes. While it is known that stable allocations can be very inefficient, the exact properties of equilibria in markets with multiple sellers are not fully understood, even in single-source single-sink networks. In this work, we show that for a large class of buyer demand functions, equilibrium always exists and allocations can often be close to optimal. In the process, we characterize the structure and properties of equilibria using techniques from min-cost flows, and obtain tight bounds on efficiency in terms of the various parameters governing the market, especially the number of monopolies M.

Although monopolies can cause large inefficiencies in general, our main results for single-source single-sink networks indicate that for several natural demand functions the efficiency only drops linearly with M. For example, for concave demand we prove that the efficiency loss is at most a factor $1+\frac{M}{2}$ from the optimum, for demand with monotone hazard rate it is at most $1+M$, and for polynomial demand the efficiency decreases logarithmically with M. In contrast to previous work that showed that monopolies may adversely affect welfare, our main contribution is showing that monopolies may not be as 'evil' as they are made out to be. Finally, we consider more general, multiple-source networks and show that in the absence of monopolies, mild assumptions on the network topology guarantee an equilibrium that maximizes social welfare.

1 Introduction

The mechanism governing large decentralized markets is often straightforward: sellers post prices for their goods and buyers buy bundles that meet their requirements. Given this framework, the challenge faced by researchers has been to characterize the equilibrium states at which these markets operate. More concretely, consider a market with multiple sellers that can be represented by a directed graph G as follows:

- Every seller owns an item, which is a link in the network.
- Every infinitesimal buyer seeks to purchase a path in the network (set of items) connecting some pair of nodes.

This work was partially supported by NSF awards CCF-1101495 and CNS-1218374.

E. Markakis and G. Schäfer (Eds.): WINE 2015, LNCS 9470, pp. 16–30, 2015.
DOI: 10.1007/978-3-662-48995-6_2

In addition to actual bandwidth markets where users purchase capacity on links for routing traffic, networks are commonly used in the literature to model combinatorial markets where the items are a mix of substitutes and complements [3,14,18]. For instance, in a computer market, each link could represent some component (e.g., a processor or video card) and buyers require a set of parts to assemble a complete computer system. In ad-markets, the buyers (advertisers) may want to purchase ads from a satisfactory combination of websites to reach a target audience. Our goal in this paper is to analyze the effects of *price competition* in such networked markets, i.e., the pricing strategies employed by competing sellers and their effect on equilibrium welfare.

An extensive body of work has culminated in the design of pricing mechanisms for a variety of markets with a single central seller (for example, see [9,15,17] and the references therein). In contrast, there has been very little focus on even simple decentralized markets where multiple price-setting sellers operate, and buyers require bundles of goods. With the exception of a few specific but incomparable settings (homogeneous goods [7,8], single buyer [10]), our understanding of how different parameters affect equilibrium in markets with price competition is quite limited. With this in mind, we seek to answer the following questions:

1. What conditions on the market structure guarantee equilibrium existence?
2. How efficient are the equilibrium allocations and how do they depend on *buyer demand* and *network topology*?

Model and Equilibrium Concept. We model the interaction between buyers and sellers as a two-stage pricing game. Each seller e controls a single good or link in a network G; he can produce any quantity x of this good incurring a production cost of $C_e(x)$. Every buyer i in the market wants to purchase an infinitesimal amount of some path connecting a source and a sink node for which she receives a value v_i. For the majority of this work, we will focus on *single-source, single-sink networks*, i.e., markets where every buyer wants to purchase a path between the same source node s and sink node t. Such networks capture combinatorial markets where buyers are interested in a single type of good; e.g., all buyers desire a computer but may have different valuations (v_i) for the same.

We consider a full information game where sellers can estimate the aggregate demand. In the first stage of the game, sellers set prices on the edges and in the second stage, buyers buy edges along a path. For any seller e, if at a price of p_e per unit amount of the good, a population x_e of buyers purchase the good, then the profit is $p_e x_e - C_e(x_e)$. The buyer's utility is v_i minus the total price paid. A solution is said to be a Nash Equilibrium if (i) Every buyer receives a utility maximizing bundle, i.e., the cheapest s-t path with price at most v_i, (ii) No seller can unilaterally change his price and improve his profit at the new allocation.

Bertrand Competition with Monopolies. Our work is most closely related to the model of Bertrand Competition in networks with supply limited sellers studied in [14] and later in [13]. Our model is more general as the production costs (that we consider) are a substantial generalization of limited supply. The behavior of Bertrand networks with seller costs was posed as an open question in [14]. We address this question by applying techniques from the theory of min-cost flows. The above papers also considered the efficiency of supply-limited markets and showed that in the worst case the equilibrium solution can be arbitrarily worse than the social optimum, and in some special cases it decreases exponentially with the length of longest s-t path in the network. Our paper provides a nuanced understanding of efficiency in terms of the buyer demand and the network topology. One of our high-level contributions in this paper is breaking down the dependence of efficiency on topology into a single parameter M: the number of monopoly edges in the graph $G = (V, E)$.

$$\text{(Set of monopolies in } G) \, \mathcal{M} := \{e \mid (s,t) \text{ are disconnected in } (V, E - \{e\})\}.$$

Monopolies offer a natural market interpretation: these are the items which are not substitutable. It is not surprising, although also not obvious, that monopolies are the main cause of inefficiencies in markets where the items are a mix of substitutes and complements. What may be extremely surprising, and what we view as one of the main contributions of our paper, is that in many reasonable settings the effect of the monopolies on equilibrium efficiency is very limited. Our results show that having a few monopolies is still not so bad: high inefficiency only occurs when the number of monopolies is large. This is in contrast to conventional wisdom that monopolies are 'evil', and even a single monopoly can cause a significant loss in social welfare [22]. More concretely, our main result is that for a large class of natural demand functions, equilibrium not only exists, but the loss in efficiency is at most a factor $(1 + M)$ from the optimum solution. We interpret this as a positive result for the following reason:

- Given previous results [13,14] that in the worst-case, social welfare can drop exponentially as M increases, a linear loss in welfare for many natural market types establishes a crucial separation between theoretical worst-case analysis and settings that are more likely to arise.

The Inverse Demand Function. In this work, our primary focus will be on single-source single-sink networks where every buyer has a different value v_i, although we do look at more general models in Sect. 5. In markets with many buyers, it is common to consider a 'full information in the large' game where the sellers know exactly how many buyers value the s-t path at v or more. This can be estimated, for instance, using prior data. Formally, we define an inverse demand function $\lambda(x)$ such that for any v, $\lambda(x) = v$ implies that exactly x amount of buyers value the path at v or larger. For example, suppose that $\lambda(x) = 1 - x$. Then, $\lambda(0.25) = 0.75$, i.e., one-fourth of the buyers have a value of 0.75 or more for the s-t paths.

1.1 Our Contributions

Our objective in this paper is to characterize the quality of equilibrium in terms of the inverse demand function, and specifically to show the effect of monopolies on efficiency. Therefore, our efficiency bounds depend only on the number of monopolies $M = |\mathcal{M}|$. Note that we define efficiency to be the ratio of the optimum social welfare of the market to that at equilibrium.

Single-Source Single-Sink Networked Markets
Our first results concern existence and uniqueness. We show that:

1. There exists a Nash Equilibrium Pricing in every market under a very mild assumption on the demand function. Moreover, there exists a Nash Equilibrium Pricing satisfying several desirable properties, including individual rationality, Pareto-optimality, and robustness to small perturbations. We call such a solution a *focal equilibrium.*
2. We further prove the uniqueness of focal equilibria. Our result is constructive: we explicitly characterize the prices and allocations in the focal equilibrium and provide an algorithm to efficiently compute them.

Since the focal equilibrium solution is the unique one satisfying many properties that one would expect from a market equilibrium, we believe that this is the correct equilibrium to study, and the one which is likely to arise in a real system. Because of this we mainly focus on analyzing the efficiency of focal equilibria.

Efficiency. We consider the following hierarchy of inverse demand functions

$$\text{Uniform} \subset \text{Polynomial} \subset \text{Concave} \subset \text{Log-Concave} = \text{MHR}.$$

Our main result is that for every function in this hierarchy, the efficiency of the focal equilibrium drops only linearly as the number of monopolies increases. Specifically, we show the following,

(Informal Theorem). *If the inverse demand function has a monotone hazard rate (MHR), the loss in efficiency at equilibrium is bounded by a factor of* $1+M$.

This result is quite general as the MHR class encapsulates all demand functions satisfying log-concavity. Moreover, some of the popular demand functions considered in the literature happen to be Concave or Polynomial (see Sect. 2 for examples). We show improved efficiency bounds for these classes, namely,

- (Uniform Demand) The Nash equilibrium maximizes welfare.
- (Polynomial Demand) Efficiency drops logarithmically as M increases.
- (Concave Demand) The efficiency loss is $1 + \frac{M}{2}$.

All of our efficiency bounds are tight. The main conclusion to draw from this is that monopolies do not completely destroy efficiency: it crucially depends on the nature of buyer demand and the *number* of these monopolies. We reiterate that since production costs strictly generalize limited supply, all of our efficiency bounds hold for the type of market considered in [13,14] as well. We make absolutely no assumption on the production cost function other than convexity, which is standard in the literature.

Multiple-Source Networks. We provide a first step towards understanding efficiency in multiple-source networked markets by tackling a question of special interest: what conditions cause equilibrium to be fully efficient in such markets? Our main result is the following: *even when buyers desire different paths, as long as the network has a series-parallel topology, the absence of monopolies guarantees an efficient equilibrium*. In contrast, without the series-parallel structure, even simple networks with no monopolies may have inefficient equilibria. We also show conditions on the buyer demand that lead to optimal equilibrium. We briefly discuss our novel contributions and the techniques that enable our results:

1. Production costs are a non-trivial addition to the Bertrand model. In particular, the pricing strategies used in [13,14] do not extend to our model as we cannot price all non-monopoly edges at zero and choose equal prices for the monopolies. Instead, we extensively apply techniques from the theory of min-cost flows to compute equilibrium prices. Specifically, the property that the flow is 'balanced' across paths is utilized to set prices on the edges.
2. In order to show efficiency bounds for MHR demand, we establish a new connection between the sellers' profit and the 'lost welfare' at equilibrium. This approach may be useful in other settings involving MHR functions.

Relation to Other Concepts. For markets with multiple buyers and sellers, the standard solution concept used in the literature is the *Walrasian Equilibrium*: a set of prices such that when both buyers and sellers act as 'price-takers', the market clears. Walrasian Equilibria are indeed attractive: they always exist in large markets [6] and are often guaranteed to be optimal. However, the idea that prices are just 'handed out' so that the market clears may not be applicable in a decentralized market. In contrast, the body of work on *price-setting* sellers (e.g., [2,7,21]) takes the view that the sellers control their own prices in order to maximize profit. Therefore, our motivation is to analyze the two-stage game where sellers set prices and buyers purchase bundles. Our work also differs from the papers in Mechanism Design that study settings with strategic buyers and a single seller [17]. Instead, we consider a market with many strategic sellers and a continuum of buyers. In such a model, it is reasonable to assume that buyers behave as price-taking agents since their individual demand is infinitesimal.

1.2 Related Work

As mentioned earlier, the study of Bertrand competition in networks was initiated in [13,14], which gave worst case bounds on efficiency over all demand functions. Despite our model being more general, we show that for many important classes of demand, the efficiency is much better than the bound shown in the above papers. Price competition between sellers was also studied in [10], where it was shown that in markets with a single buyer, equilibrium allocations are efficient. The Uniform demand case that we study is similar in spirit to what they consider, but our main results are for more complex demand functions. Finally, our work bears broad similarities to recent papers that also study existence or efficiency in somewhat specific settings with multiple sellers [7,8,19]. However, their models are not comparable to ours. In [8], all sellers possess a single homogeneous good but buyers may not have access to all of them; in [7,19], there is a single buyer but sellers may own more than one good. In contrast, we consider a market with multiple buyers where every seller controls one good but the goods are not homogeneous.

Some researchers have also considered more sophisticated pricing mechanisms like *non-linear pricing*(see [16,18]). While complex mechanisms do sometimes lead to an improvement in efficiency, they are not commonly used as they impose a large overhead on buyers who have to anticipate the change in price due to others' demands. In this work, we study the more natural fixed pricing mechanism and attempt to provide additional insight on the quality of equilibrium.

Finally, one line of research that has gained traction in recent years [2,3] is pricing in networked markets with congestion, i.e., buyers pay the price on each edge, but also incur a delay due to congestion. In contrast, we share the view taken by Shenker et al. [21] that 'congestion costs are inherently inaccessible to the network'. Due to the underlying complexities of this model, most of the results are only known for simple networks such as parallel paths. One exception is [20], which considers a unique one-sided model where the routing decisions are taken locally by sellers and not buyers as in our paper. They show that in the absence of monopolies, local decisions by sellers can result in efficient solutions.

2 Definitions and Preliminaries

An instance of our two-stage game is specified by a directed graph $G = (V, E)$, a source and a sink (s, t), an inverse demand function $\lambda(x)$ and a cost function $C_e(x)$ on each edge. There is a population T of infinitesimal buyers; every buyer wants to purchase edges on some s-t path and x amount of buyers hold a value of $\lambda(x)$ or more for these paths. A buyer is satisfied if she purchases all the edges on some path connecting s and t and is indifferent among the different paths.

We define M to be the number of monopolies in the market: an edge e is a monopoly if removing it disconnects the source and sink. We make the following standard assumptions on the demand and cost functions.

1. The inverse demand function $\lambda(x)$ is continuous on $[0, T]$ and non-increasing, implying that demand decreases as price increases.
2. $C_e(x)$ is non-decreasing and convex $\forall e$, which is the standard way to model production costs. Moreover, $C_e(x)$ is continuous, differentiable, and its derivative $c_e(x) = \frac{d}{dx}C_e(x)$ satisfies $c_e(0) = 0$.

Nash Equilibrium Pricing. A solution of our two-stage game is a vector of prices on each item $\boldsymbol{p}$ and an allocation or flow $\boldsymbol{x}$ of the amount of each s-t path purchased, representing the strategies of the sellers and buyers respectively. The total flow or market demand is equal to the number of buyers with non-zero allocation $x = \sum_{P \in \mathbb{P}} x_P$, where $\mathbb{P}$ is the set of s-t paths. We can also decompose this flow $\boldsymbol{x}$ into the amount of each edge purchased by the buyers $(x_e)_{e\in E}$. Given this solution, the total utility of the sellers is $\sum_{e\in E}(p_e x_e - C_e(x_e))$ and the aggregate utility of the buyers is $\int_{t=0}^{x}\lambda(t)dt - \sum_{e\in E} p_e x_e$. The social welfare is simply $\int_{t=0}^{x}\lambda(t)dt - \sum_e C_e(x_e)$, i.e., prices are intrinsic to the system and do not appear in the welfare.

We use the standard definition of Nash equilibrium for two-stage games to model the stable states of our market. Formally, an allocation $\boldsymbol{x}$ is said to be a **best-response** by the buyers to prices $\boldsymbol{p}$ if buyers only buy the cheapest paths and for any cheapest path P, $\lambda(x) = \sum_{e\in P} p_e$. That is, buyers act as price-takers and any buyer whose value is at least the price of the cheapest path will purchase some such path. A solution $(\boldsymbol{p}, \boldsymbol{x})$ is a Nash equilibrium if $\boldsymbol{x}$ is a best-response allocation to the prices and, $\forall e$ if the seller unilaterally changes his price from p_e to p'_e, then for *every* feasible best-response flow (x'_e) for the new prices, seller e's profit cannot increase, i.e., $p_e x_e - C_e(x_e) \geq p'_e x'_e - C_e(x'_e)$. Our notion of equilibrium is quite strong as the seller does not have to anticipate the resulting flow: for every best-response by the buyers, the seller's profit should not increase.

Classes of inverse demand functions that we are interested in For ease of exposition, we assume that both the inverse demand and the production costs are continuously differentiable. However, **all our results** hold exactly even without this assumption. Note that $\lambda'(x)$ cannot be positive since $\lambda(x)$ is non-increasing. The reader is asked to refer to the full version of this paper [5] for additional discussion on these classes of demand.

Uniform Demand: $\lambda(x) = \lambda_0 > 0$ for $x \leq T$. In other words, a population of T buyers all have the same value λ_0 for the bundles.

Polynomial Demand: $\lambda(x) = \lambda_0(a - x^\alpha)$ for $\alpha \geq 1$. Polynomial demand functions are quite popular [11], especially linear inverse demand ($\lambda(x) = a - x$).

Concave Demand: $\lambda'(x)$ is a non-increasing function of x.

Monotone Hazard Rate (MHR) Demand: $\frac{\lambda'(x)}{\lambda(x)}$ is non-increasing or $h(x) = \frac{|\lambda'(x)|}{\lambda(x)}$ is non-decreasing in x. This is equivalent to the class of *log-concave* functions [4] where $\log(\lambda(x))$ is concave. Example function: $\lambda(x) = e^{-x}$.

It is not hard to see that Uniform[1] $\subset$ Polynomial $\subset$ Concave $\subset$ MHR. We remark that the MHR and Concave classes are quite general whereas Uniform or Polynomial demand are more common due their tractability.

Min-Cost Flows and the Social Optimum: Since an allocation vector $\boldsymbol{x}$ is equivalent to a s-t flow, we briefly dwell upon minimum cost flows. Formally, we define $R(x)$ to be the cost $\sum_e C_e(x_e)$ of the min-cost flow of magnitude $x \geq 0$ and $r(x)$, its derivative, i.e., $r(x) = \frac{d}{dx}R(x)$. Both the flow and its cost can be computed via a simple convex program given the graph. The min-cost function $R(x)$ obeys several desirable properties that we use later including:

[1] Uniform $= \lim_{\alpha\to\infty} \lambda_0(1 - x^\alpha)$.

Proposition 1. *$R(x)$ is continuous, non-decreasing, differentiable, and convex.*

From the KKT conditions, we have that for a min-cost flow $\boldsymbol{x}$, $r(x) = \sum_{e \in P} c_e(x_e)$ for any path P with non-zero flow. Using this property, we obtain the following characterization of the welfare maximizing solution in terms of $R(x)$.

Proposition 2. *The solution maximizing social welfare is a min-cost flow of magnitude x^* satisfying $\lambda(x^*) \geq r(x^*)$. Moreover, $\lambda(x^*) = r(x^*)$ unless $x^* = T$.*

3 Existence, Uniqueness, and Computation

In this section, we prove that a Nash equilibrium is guaranteed to exist under the very mild assumption that the demand function has a monotone *price elasticity.* Moreover, we show that there always exists a unique 'focal equilibrium' that satisfies several desirable properties. We also provide an algorithm to compute this important equilibrium.

Before proving our general existence result, it is important to understand the different types of equilibria that may exist in networked markets. In markets such as ours, an existence result by itself is meaningless because a large sub-class of instances admit trivial and unrealistic equilibria.

Trivial Equilibrium: In a networked market where all paths have a length of at least 2, it is easy to see that every seller setting an unreasonably high price (say larger than $\lambda(0)$) would result in a Nash equilibrium with zero flow. The existence of such unrealistic equilibria was also observed in [13], where they were referred to as trivial equilibria.

Our goal in this paper is to analyze the equilibrium operating states of actual markets. Given that our model admits such uninteresting equilibria, it is important that any existence result be characterized by properties that one might come to expect from equilibria that are likely to arise in practice; for example, one might expect that a meaningful equilibrium has non-zero flow, is not dominated by other equilibria and most importantly from the perspective of a large market, is robust to small perturbations (we define these formally below). Our main existence result is that under a very mild condition on the demand, there exists a 'nice' equilibrium that satisfies many such desiderata.

We first formally define what it means for the price elasticity of a demand function to be monotone. This condition is quite minimal: it is obeyed by almost all of the demand functions in the literature (for example: [1,4,11,13]).

Definition 3. ***Monotone Price Elasticity*** *(MPE) An inverse demand function $\lambda(x)$ is said to have a monotone price elasticity if its price elasticity $\frac{x|\lambda'(x)|}{\lambda(x)}$ is a non-decreasing function of x which approaches zero as $x \to 0$.*

All the classes of demand functions listed in the previous section satisfy the MPE condition. At a high level, the MPE condition simply implies that a market's responsiveness at low prices cannot be too large compared to its responsiveness at a high price. Even more intuitively, MPE functions are concave if plotted on a log-log plot, and are essentially all functions which are "less convex" than x^{-r}.

Theorem 4. *For any given instance of a networked market where the inverse demand function $\lambda(x)$ has a monotone price elasticity, there exists a Nash Equilibrium $(p)_{e\in E}, (\tilde{x})_{e\in E}$ satisfying the following properties*

1. *Non-Trivial Pricing (Non-zero flow)*
2. *Recovery of Production Costs (Individual Rationality)*
3. *Pareto-Optimality*
4. *Local Dominance (Robustness to small perturbations)*

We now formally define these properties and argue why it is reasonable to expect an actual market equilibrium to satisfy them. For example, although they are not stable solutions for our price-setting sellers, it is not hard to see that Walrasian Equilibria satisfy all of these properties.

1. **(Non-Trivial Pricing):** Every edge that does not admit flow must be priced at 0. This guarantees that the equilibrium has non-zero flow.
2. **(Recovery of Production Costs):** Given an equilibrium $(\boldsymbol{p}, \boldsymbol{x})$, every item's price is at least $c_e(x_e)$. This property is similar in spirit to *individual rationality* and ensures that the prices are fair to the sellers. Suppose that $p_e < c_e(x_e)$, this means that the seller is selling at least some fraction of his items at price smaller than its cost of production, and therefore, would have no incentive to produce the given quantity of items.
3. **(Pareto-Optimality):** A Pareto-optimal solution over the space of equilibria is an equilibrium solution such that for any other equilibrium, at least one agent (buyer or seller) prefers the former solution to the latter. Pareto-Optimality is often an important criterion in games with multiple equilibria; research suggests that in Bertrand Markets, Pareto optimal equilibria are the solutions that arise in practice [12].
4. **(Local Dominance):** Given an equilibrium $(\boldsymbol{p}, \boldsymbol{x})$, consider a different flow assignment for the same prices $(\boldsymbol{p}, \boldsymbol{x'})$, differing only in which cheapest paths are taken by the buyers. Local Dominance means that the profit of each seller must be larger at the equilibrium solution than at any $(\boldsymbol{p}, \boldsymbol{x'})$. The essence of this property is that the solution is resilient against small buyer perturbations. In other words, if instead of changing his price (which we know no seller would do at equilibrium), a seller instead convinced some buyers to take different paths of the same total price, then this seller still could not benefit from the resulting new flow. If this were not the case, then a seller may be able to attract a small fraction of buyers towards his item and improve his profit, indicating that the original equilibrium is not robust.

(Proof Sketch of Theorem 4) The proof proceeds by analyzing the behavior of monopolies and non-monopolies at equilibrium: every monopoly behaves as if it is a part of a two-link serial network where the rest of the network can be composed into a single serial link. This allows us to derive a sufficient condition on a

monopoly's price at equilibrium that is independent of every other link (namely, $p_e = c_e(\tilde{x}) + \tilde{x}|\lambda'(\tilde{x})|$). In contrast, non-monopolies at equilibrium behave as if they are a part of a two-edge parallel network. A crucial ingredient of our result is the application of min-cost flows to link the behavior of monopolies and non-monopolies. Namely, the property that the (marginal) flow cost is balanced across all paths is used to choose the price of every edge. Once we have explicitly constructed the equilibrium prices, the rest of the theorem involves showing that these prices result in a non-trivial best-response flow. Note that standard techniques such as fixed point theorems cannot be used here since the solution space is not convex: small changes in price may result in large deviations. ■

The full proofs of all the theorems can be found in a full version of this paper [5]. The next corollary, which is the main ingredient in all of our efficiency bounds essentially characterizes the equilibrium structure by expressing the equilibrium flow $(\tilde{x})$ as a function of only the number of monopolies in the network.

Corollary 5. *For any demand λ satisfying the MPE condition, $\exists$ a Nash equilibrium with a min-cost flow $(\tilde{x}_e)$ of size $\tilde{x} \leq x^*$ such that,*

$$\textit{Either}\ \ \frac{\lambda(\tilde{x}) - r(\tilde{x})}{M} = \tilde{x}|\lambda'(\tilde{x})|\ \ \textit{or}\ \ \tilde{x} = x^*,\ \ \textit{the optimum solution.}$$

We now show that the equilibrium from Theorem 4 (which we will refer to as the *focal equilibrium*) is the unique solution that satisfies the useful desiderata defined above. In order to truly understand the equilibrium efficiency of our two-stage game, it does not make sense to show a blanket bound on all stable solutions since some of these are highly unrealistic (for example, Price of Anarchy is almost always unbounded due to the presence of trivial equilibria). However, since the focal equilibrium solution is the unique one satisfying many properties that one would expect from a market equilibrium, we focus on analyzing its efficiency in the rest of this paper.

Theorem 6. *For any given instance with strictly monotone MPE demand and non-zero costs, we are guaranteed that one of the following is always true:*

1. *There is a unique non-trivial equilibrium that satisfies Local Dominance.* (or)
2. *All non-trivial equilibria that satisfy Local Dominance maximize welfare.*

Moreover, we can compute this equilibrium efficiently.

For the purposes of studying efficiency, the above theorem provides a useful baseline: either all equilibria are fully efficient or it suffices to bound the efficiency of the unique equilibrium that satisfies Corollary 5 (which we do in Sect. 4). As always in the case of real-valued settings (e.g., convex programming, etc.), "computing" a solution means getting within arbitrary precision of the desired solution; the exact solution could be irrational.

4 Effect of Monopolies on the Efficiency of Equilibrium

In this paper, we are interested in settings where approximately efficient outcomes are reached despite the presence of self-interested sellers with monopolizing power. While for general functions $\lambda(x)$ obeying the MPE condition, the efficiency can be exponentially bad, we show that for many natural classes of functions it is much better, even in the presence of monopolies.

We begin with a more fundamental result that reinforces the fact that even in arbitrarily large networks (not necessarily parallel links), competition results in efficiency, i.e., when $M = 0$, the efficiency is 1. This result is only a starting point for us since it is the addition of monopolies that leads to interesting behavior.

Claim 7. *In any network with no monopolies (i.e., you cannot disconnect s, t by removing any one edge), there exists a focal Nash Equilibrium maximizing social welfare.*

We remark that our notion of a "no monopoly" graph is weaker than what has been considered in some other papers [16,20] and therefore, our result is stronger. We are now in a position to show our main theorem. The largest class of inverse demand functions that we consider are the MHR or Log-Concave functions. Note that all MHR functions satisfy the MPE condition and thus existence is guaranteed. Our main result is that for all demand functions in this class, the efficiency loss compared to the optimum solution is $1 + M$. We believe that this result has strong implications. First, log-concavity is a very natural assumption on the demand; these functions have received considerable attention in Economics literature(see [4] and follow-ups). Secondly, it is reasonable to assume that even in multi-item markets, the number of purely monopolizing goods is not too large: in such cases the equilibrium quality is high.

Theorem 8. *The social welfare of the Nash equilibrium from Sect. 3 is always within a factor of $1 + M$ of the optimum for MHR λ, and this bound is tight.*

(Proof Sketch). The proof relies crucially on our characterization of equilibria obtained in Corollary 5, and the following interesting claim for MHR functions linking the welfare loss at equilibrium to the profit made by all the sellers: 'the loss in welfare is at most a factor M times the total profit in the market at equilibrium'. In addition, it is also not hard to see that in any market, the profit cannot exceed the total social welfare of a solution.

Why is this claim useful? Using profit as an intermediary, we can now compare the welfare lost at equilibrium to the welfare retained. This implies that the welfare loss cannot be too high because that would mean that the profit and hence the welfare retained is also high. But then, the sum of welfare lost + retained is the optimum welfare and is bounded. Therefore, we can immediately bound the overall efficiency. Mathematically, our key claim is,

$$\text{Lost Welfare} = \int_{\tilde{x}}^{x^*} \lambda(x)dx - [R(x^*) - R(\tilde{x})] \leq M(p\tilde{x} - R(\tilde{x})),$$

where p is the payment made by every buyer, $\tilde{x}$ is the amount of buyers in the equilibrium solution and x^*, in the optimum. The integral in the LHS can be rewritten as $\int_{\tilde{x}}^{x^*}(\lambda(x)-r(x))dx$. Now, we apply some fundamental properties of MHR functions (λ) and show that for all $x \geq \tilde{x}$, the following is true, $\frac{\lambda(x)-r(x)}{|\lambda'(x)|} \leq \frac{\lambda(\tilde{x})-r(\tilde{x})}{|\lambda'(\tilde{x})|} = M\tilde{x}$. The final equality comes from our equilibrium characterization in Corollary 5. Therefore, we can prove the key claim as follows:

$$\begin{aligned}\int_{\tilde{x}}^{x^*}(\lambda(x)-r(x))dx &\leq M\tilde{x}\int_{\tilde{x}}^{x^*}|\lambda'(x)|dx \\ &\leq M\tilde{x}(\lambda(\tilde{x})-\lambda(x^*)) \quad (\ \lambda(x)\text{is non-increasing and } \tilde{x} \leq x^*) \\ &\leq M(\lambda(\tilde{x})\tilde{x} - R(\tilde{x})) \qquad (\lambda(x^*)\tilde{x} \geq r(x^*)\tilde{x} \geq r(\tilde{x})\tilde{x} \geq R(\tilde{x}))\end{aligned}$$

The total payment p on any path must exactly equal $\lambda(\tilde{x})$ ■

Tighter Bounds for Sub-classes. We now consider log-concave demand functions that satisfy additional requirements, namely Uniform, Polynomial, and Concave demand. For these classes, we show much stronger bounds on the efficiency loss at equilibrium.

Theorem 9. *The following bounds on the efficiency are tight*

- *Every instance with **Uniform** demand admits a fully efficient focal Nash equilibrium.*
- *For any instance with **Polynomial** demand, the inefficiency of focal equilibrium is at most $(1+M\alpha)^{\frac{1}{\alpha}}$, where $\alpha \geq 1$ is the degree of the polynomial. When $\alpha \geq M$, this quantity is approximately $1+\frac{\log(M\alpha)}{\alpha}$.*
- *When the demand is **Concave**, the inefficiency of focal equilibrium is $1+\frac{M}{2}$.*

The efficiency bound for Polynomial demand extends to more general polynomials of the form $\lambda(x) = a_0 - \sum_{i=1}^{k} a_i x^{\alpha_i}$ with α now defined as $\min_i \alpha_i$.

Concave-Log Demand. In this paper, we considered MPE functions (concave on a log-log plot) and log-concave functions (concave on a semi-log plot). For the sake of completeness, we also consider functions that are not log-concave but still obey the MPE condition. One such important class consists of Concave-Log demand functions, which are in some sense the opposite of log-concave functions; in other words $\lambda(x)$ is concave against a logarithmically varying buyer demand (i.e., $\lambda(log(x))$ is concave). This class of functions was considered in [13], where an efficiency bound of e^D was shown: D being the length of the longest s-t path in the network which could potentially be much larger than M. We generalize their results to markets with cost functions, and further are able to improve upon the bound in [13].

Claim 10. *For any instance with concave-log demand, the inefficiency of the focal equilibrium is at most $\frac{M}{M-1}e^{M-1}$ for $M \geq 2$.*

We reiterate here that all of our results require no assumption on the graph structure and only the ones mentioned in Sect. 2 for the cost functions.

5 Generalizations: Multiple-Source Networks

We now move on to more general networks where different buyers have different s_i-t_i paths that they wish to connect and the demand function can be different for different sources. Unfortunately, our intuition from the previous sections does not carry over. Even when buyers have Uniform demand, Nash equilibrium may not even exist whereas in the single-source case, equilibrium was efficient. Perhaps more surprisingly, we give relatively simple examples in which perfect efficiency is no longer achieved in the absence of monopolies. Nevertheless, we prove that for some interesting special cases, fully efficient Nash equilibrium still exists even when buyers desire different types of bundles. In particular, we believe that our result on series-parallel networks is an important starting point for truly understanding multiple-source networks.

Claim 11. *There exist simple instances with two sources and one sink such that*

1. *Nash equilibrium may not exist even when the buyers at each source have Uniform demand.*
2. *All Nash equilibria are inefficient even when no edge is a monopoly.*

Series-Parallel Networks: In some sense, Claim 7 embodies the very essence of the Bertrand paradox, the fact that competition leads to efficiency. So it is surprising that this does not hold in general networks. However, we now show that for a large class of markets which have the series-parallel structure, the absence of monopolies still gives us efficient equilibria. Series-Parallel networks have been commonly used [16,18] to model the substitute and complementary relationship that exists between various products in combinatorial markets.

We define a multiple-source single-sink graph to be a series-parallel graph if the super graph of the given network obtained by adding a super-source and connecting it to all the sources has the series-parallel structure. The notion of "no-monopolies" for a complex network has the same idea as a single-source network: there is no edge in the graph such that its removal would disconnect any source from the sink. We are now in a position to show our result.

Theorem 12. *A multiple-source single-sink series-parallel network with no monopolies admits a welfare-maximizing Nash Equilibrium for any given demand.*

Finally, we show additional conditions on both the network topology and demand that lead to efficient equilibria, even in the presence of monopolies.

Claim 13. *There exists a fully efficient equilibrium in multiple-source multiple-sink networks with Uniform demand buyers at each source if one of the following is true:* (*i*) *Buyers have a large demand and production costs are strictly convex,* (*ii*) *Every source node is a leaf in the network.*

The second case commonly arises in several telecommunication networks, where the *last mile* between a central hub and the final user is often controlled by a local monopoly and thus the source is a leaf.

6 Conclusions

In this work, we initiate the study of Bertrand price competition in networked markets with *production costs*. Our results provide an improved understanding of how monopolies affect welfare in large, decentralized markets. Our main contribution is that as long as the inverse demand obeys a natural condition (monotone hazard rate), the efficiency loss is at most $1 + M$ for single-source single-sink networks, with stronger results for other important classes. Cast in the light of previous work [13,14], our result establishes that the inefficiency for commonly used demand is much better than the worst-case exponential inefficiency. Finally, for markets where buyers desire different paths, we identify series-parallel networks topology as a condition for efficiency. We believe this result is a useful first step in understanding the impact of monopolies on multiple-source networks. In a full version of this paper [5], we extend all our results to markets without a network structure where all buyers desire the same bundles.

References

1. Abolhassani, M., Bateni, M.H., Hajiaghayi, M.T., Mahini, H., Sawant, A.: Network cournot competition. In: Liu, T.-Y., Qi, Q., Ye, Y. (eds.) WINE 2014. LNCS, vol. 8877, pp. 15–29. Springer, Heidelberg (2014)
2. Acemoglu, D., Ozdaglar, A.: Competition and efficiency in congested markets. Math. Oper. Res. **32**(1), 1–31 (2007)
3. Acemoglu, D., Ozdaglar, A.: Competition in parallel-serial networks. IEEE J. Sel. Areas Commun. **25**(6), 1180–1192 (2007)
4. Amir, R.: Cournot oligopoly and the theory of supermodular games. Games Econ. Behav. **15**(2), 132–148 (1996)
5. Anshelevich, E., Sekar, S.: Price competition in networked markets: how do monopolies impact social welfare? arXiv preprint arXiv:1410.1113 (2015)
6. Azevedo, E.M., Weyl, E.G., White, A.: Walrasian equilibrium in large, quasilinear markets. Theor. Econ. **8**(2), 281–290 (2013)
7. Babaioff, M., Leme, R.P., Sivan, B.: Price competition, fluctuations and welfare guarantees. In: Proceedings of EC (2015)
8. Babaioff, M., Lucier, B., Nisan, N.: Bertrand networks. In: Proceedings of EC (2013)
9. Babaioff, M., Lucier, B., Nisan, N., Leme, R.P: On the efficiency of the walrasian mechanism. In: Proceedings of EC (2014)
10. Babaioff, M., Nisan, N., Leme, R.P: Price competition in online combinatorial markets. In: Proceedings of WWW (2014)
11. Bulow, J.I., Pfleiderer, P.: A note on the effect of cost changes on prices. J. Polit. Econ. **91**, 182–185 (1983)
12. Cabon-Dhersin, M.-L., Drouhin, N.: Tacit collusion in a one-shot game of price competition with soft capacity constraints. J. Econ. Manage. Strategy **23**(2), 427–442 (2014)
13. Chawla, S., Niu, F.: The price of anarchy in bertrand games. In: Proceedings of EC (2009)
14. Chawla, S., Roughgarden, T.: Bertrand competition in networks. In: Monien, B., Schroeder, U.-P. (eds.) SAGT 2008. LNCS, vol. 4997, pp. 70–82. Springer, Heidelberg (2008)

15. Chawla, S., Sivan, B.: Bayesian algorithmic mechanism design. SIGecom Exch. **13**(1), 5–49 (2014)
16. Correa, J.R., Figueroa, N., Lederman, R., Stier-Moses, N.E.: Pricing with markups in industries with increasing marginal costs. Math. Program., 1–42 (2008)
17. Hassidim, A., Kaplan, H., Mansour, Y., Nisan, N.: Non-price equilibria in markets of discrete goods. In: Proceedings of EC (2011)
18. Kuleshov, V., Wilfong, G.: On the efficiency of the simplest pricing mechanisms in two-sided markets. In: Goldberg, P.W. (ed.) WINE 2012. LNCS, vol. 7695, pp. 284–297. Springer, Heidelberg (2012)
19. Lev, O., Oren, J., Boutilier, C., Rosenschein, J.S.: The pricing war continues: on competitive multi-item pricing. In: Proceedings of AAAI (2015)
20. Papadimitriou, C.H., Valiant, G: A new look at selfish routing. In: Proceedings of ICS (2010)
21. Shenker, S., Clark, D., Estrin, D., Herzog, S.: Pricing in computer networks: reshaping the research agenda. ACM SIGCOMM Comput. Commun. Rev. **26**(2), 19–43 (1996)
22. Tullock, G.: The welfare costs of tariffs, monopolies, and theft. Econ. Inq. **5**(3), 224–232 (1967)

Springer

http://www.springer.com/978-3-662-48994-9

Web and Internet Economics
11th International Conference, WINE 2015, Amsterdam,
The Netherlands, December 9-12, 2015, Proceedings
Markakis, E.; Schäfer, G. (Eds.)
2015, XIX, 442 p. 34 illus. in color., Softcover
ISBN: 978-3-662-48994-9