

The Price of Anarchy in Bertrand Games

Shuchi Chawla^{*}
Computer Sciences Dept.
University of Wisconsin - Madison
Madison, WI 53705
shuchi@cs.wisc.edu

Feng Niu[†]
Computer Sciences Dept.
University of Wisconsin - Madison
Madison, WI 53705
leonn@cs.wisc.edu

ABSTRACT

The Internet is composed of multiple economically-independent service providers that sell bandwidth in their networks so as to maximize their own revenue. Users, on the other hand, route their traffic selfishly to maximize their own utility. How does this selfishness impact the efficiency of operation of the network? To answer this question we consider a two-stage network pricing game where service providers first select prices to charge on their links, and users then pick paths to route their traffic. We give tight bounds on the price of anarchy of the game with respect to social value—the total value obtained by all the traffic routed. Unlike recent work on network pricing, in our pricing game users do not face congestion costs; instead service providers must ensure that capacity constraints on their links are satisfied. Our model extends the classic Bertrand game in economics to network settings.

Categories and Subject Descriptors

F.2 [Analysis of Algorithms and Problem Complexity]: Miscellaneous

General Terms

Theory

Keywords

Price of anarchy, bandwidth pricing, two-sided markets, Bertrand competition

1. INTRODUCTION

We study the following pricing game in two-sided combinatorial markets, with the goal of understanding how efficiently such markets operate. Each seller in the market owns a fixed inventory of a unique item and wants to obtain as much revenue as possible from

^{*}Supported in part by NSF awards CCF-0643763 and CCF-0830494.

[†]Supported in part by NSF award CCF-0643763.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

EC'09, July 6–10, 2009, Stanford, California, USA.

Copyright 2009 ACM 978-1-60558-458-4/09/07 ...\$5.00.

selling the item. Each consumer is interested in buying some bundle of items with the goal of maximizing its utility—the difference between its value for the bundle and the bundle price. The game proceeds with the sellers first setting prices on their items and the consumers then buying their favorite bundles. We call this game a *Bertrand game* following economics literature on price competition in homogeneous product markets.

The Bertrand game is a natural model for markets where prices are relatively static and non-discriminatory, and left-over inventory can be disposed off freely. A prime example of such a market is the Internet. The Internet is composed of a number of subnetworks, each owned by an economically independent service provider. Service providers sell bandwidth on their links with the goal of maximizing their own profit. They typically have fixed capacity limits on their links, and negligible marginal cost to routing traffic as long as it doesn't approach capacity. Likewise, users that are interested in routing traffic over the network themselves behave selfishly and optimize over the cost of routing as well as quality of service. The operation of the network therefore rests in the hands of multiple selfish entities. We aim to understand whether the resulting bandwidth allocation is (close to) optimal in terms of the total value obtained by all the traffic routed in the network.

We study price of anarchy type questions in Bertrand games in networks. Specifically we define the efficiency of allocation, or social value, of an outcome of the game to be the total value obtained by all the consumers from the bundles bought by them. The price of anarchy is then defined to be the ratio of the optimal efficiency achievable while observing capacity (supply) constraints to that of the worst Nash equilibrium in the game.

As mentioned earlier, the Bertrand game has been studied extensively in economics literature on homogeneous product markets, where all items are perfect substitutes (see, for example, [14, 11, 6] and references therein). It is well known [14], for example, that in markets with two or more sellers selling homogeneous products, the equilibrium price charged by each seller is equal to the marginal cost for the product. That is, the sellers make no profit and the entire social surplus is redistributed back to the consumers. This is known as the Bertrand paradox. Such markets always obtain optimal efficiency. On the other hand, in a market with a single seller (a monopoly), the seller can artificially restrict supply and raise prices so as to obtain a large profit, thereby hurting the efficiency of the market. Nevertheless, it is known [7] that the price of anarchy in such a market is no more than logarithmic in the disparity of consumer valuations.

In markets with multiple different items, the efficiency of allocation can depend on the specific substitute-complement relationships among different items and the distribution of demand for each collection of items. A primary goal of this paper is to understand

how this market and demand structure affects the price of anarchy. We focus on the network setting where desired bundles of items are paths in a directed graph. This setting already captures a rich class of substitute-complement relationships.

Chawla and Roughgarden [7] were the first to study the Bertrand game in combinatorial markets. They also focused on the network setting as in our work, and considered a simple model of consumer preferences where each consumer obtains a single fixed value for routing its traffic over any source-sink path. We call this the single-parameter setting. They showed that the price of anarchy of the game depends on the presence, and number, of monopolies in the network as well as the distribution of demand in the network.

Our work extends the work of Chawla and Roughgarden in several ways. First, we give a better and more general upper bound on the price of anarchy of the game in networks in the single-parameter setting. Second, we show that under a mild condition on the value distributions of consumers the price of anarchy improves considerably. Third, we consider a more general model of consumer preferences in two-node parallel-link networks, where each consumer can potentially have different values for routing its traffic over different source-sink paths. This “multi-parameter” setting allows us to incorporate the qualities of different paths into consumer preferences. We provide tight bounds on the price of anarchy of the game in each of these settings.

We remark that some instances of the Bertrand game do not admit (pure) Nash equilibria (see Section 2). Our results on the price of anarchy of the game do not apply to these instances. In effect our results show that under certain conditions on the demand and network structure the system will never get stuck in a very bad stable state. Chawla and Roughgarden previously addressed the issue of characterizing instances where equilibria exist and presented partial results. Obtaining a more complete understanding of the existence of Nash equilibria and exploring relaxed notions of stability for Bertrand games are important open questions.

Related work

There is a large body of work on analyzing equilibria in network games that involve only service providers (network formation games) [3, 4, 9] or only users (selfish routing games) [17, 18, 8, 10]. More recently, a sequence of works [13, 1, 2, 16, 15, 7] has considered network pricing games where edges selfishly select prices to maximize their profit and consumers pick paths for routing their demands in a utility-maximizing fashion. Most of these works assume that the consumer faces two kinds of costs for routing its traffic—the prices charged by the edges, and a latency cost owing to other traffic on the path. Our model differs from these other works in two key aspects. First, we have hard capacity constraints on edges in our model as opposed to elastic congestion-based cost functions used in prior work. Second, users do not face “congestion costs” in our model directly. Instead, service providers are incentivized to keep usage on links below their capacities, and pass on the “costs” associated with oversaturated edges to consumers in the form of higher usage prices. Although capacity constraints in our model mimic some congestion effects, the game admits non-trivial equilibria even in the absence of capacity constraints when the market contains monopolies. In effect, our work isolates the impact of selfish pricing on the efficiency of the network in the absence of congestion effects.

Our work is closely related to and extends the work of Chawla and Roughgarden [7]. Chawla et al. give a complete characterization of the price of anarchy in networks where all consumers have the same source and destination. In these instances, the worst-case price of anarchy depends on the number of monopolies in the net-

work. In more general networks, they show that a pure-strategy Nash equilibrium for the sellers does not always exist; when it does, the price of anarchy can be unbounded even in the absence of monopolies. However, in multiple-source single-sink instances under certain technical conditions on the structure of the network and the distribution of demands, the price of stability of the game (i.e., the performance of the best Nash equilibrium in comparison to the optimal solution) is at most the inverse of the sparsity of the network (a measure of congestion in the network).

Our results

We extend the work of Chawla et al. [7] in several ways. First we show (Section 3) that the bound from [7] on the price of stability in multiple-source single-sink networks applies also to the price of anarchy under mild conditions on the distribution of values in the network. Specifically, in the absence of monopolies, if every source in the network has the same distribution of values (but potentially different total amount of demand), the efficiency gap between *any* arbitrary equilibrium and the optimal solution is bounded by the inverse of the sparsity of the network. This result is tight; furthermore we show in Section 3.3 that no such bound holds in more general, multiple-source multiple-sink, settings—for every non-trivial value distribution, there exists a network where the price of anarchy is unbounded.

Next we examine the effect of the distribution of consumers’ values on the price of anarchy (Section 3.2). We show that if the values are drawn from a distribution with a monotone hazard rate (see Definition 1), the price of anarchy is no more than exponential in the length of the longest source-sink path, *even in the presence of monopolies in the network*. Thus low-diameter networks admit good equilibria. Once again this result holds in single-sink instances but (provably) cannot be extended to multiple-source multiple-sink instances.

Most interestingly, in Section 4 we consider a more general “multi-parameter” model of consumer preferences where consumers value different paths in the network at different amounts. We investigate this model in two-node parallel-link networks. When consumer values are related in that the difference between values for two paths is the same for any two consumers, the behavior of the game is similar to that in the single-parameter case—an equilibrium always exists and the price of anarchy of the game is 1. When values can be completely arbitrary, the game exhibits some of the properties of the multiple-source single-sink setting with single-parameter consumers, and consequently we show that pure strategy equilibria don’t always exist. Nevertheless, we can still bound the price of anarchy by a factor that is logarithmic in the ratio of the largest consumer value to the smallest value.

2. MODEL

A Bertrand game contains two kinds of players – sellers and consumers. Each seller sells a unique item. We use E to denote the set of sellers as well as the items they sell, and S to denote the set of consumers. Each seller $e \in E$ owns a fixed supply of n_e units of item e . The goal of the seller is to sell her items at an appropriate per-unit price so as to maximize her revenue. Each consumer s is interested in buying a bundle of items. The value of consumer s for bundle $B \subseteq E$ is denoted $v_{s,B}$. We denote by \mathcal{B}_s the collection of bundles for which consumer s has non-zero value; this is termed the consumer’s *type*¹. The goal of each consumer is to buy the bundle of items that maximizes her utility.

¹Our definition of type is different from what is commonly used in mechanism design settings; In particular, in our setting consumers

The game proceeds as follows. In the first stage sellers pick prices π_e ; in the second stage consumers select and buy bundles. The payoff to a consumer s buying a bundle $B \subseteq E$ is her utility— $v_{s,B} - \sum_{e \in B} \pi_e$. Consumers may randomize and/or buy fractional units of bundles (summing up to no more than 1). Let $f_{s,e}$ be the quantity of item e demanded at the current prices by consumer s and $f_e = \sum_s f_{s,e}$ the total quantity of e demanded. Then the payoff to seller e is $\pi_e \min(n_e, f_e)$.

Nash equilibrium

An outcome of the Bertrand game is a pair of vectors (π, f) , where π_e denotes the price for each seller and $f_{s,e}$ denotes the quantity of item e bought by consumer s . We study Nash equilibria in Bertrand games. An outcome is called a *subgame-perfect Nash equilibrium* if given the prices for each seller, no consumer can improve her payoff by changing her strategy unilaterally, and given the prices of other sellers, no seller can improve her payoff by changing her price unilaterally.

Consumers are price-takers and are at equilibrium whenever they follow best response to the sellers' prices. Specifically, at equilibrium, consumer s buys 1 unit of the bundle that maximizes her utility $v_{s,B} - \sum_{e \in B} \pi_e$, as long as this utility is positive. When the utility is zero, the consumer may buy up to 1 unit, and if there are multiple bundles that offer the maximum utility to the consumer, the consumer may split its demand arbitrarily over these bundles. However the consumer never buys bundles that bring it less than maximum utility.

Seller behavior at equilibrium is more interesting and complex. A seller is at equilibrium if it cannot improve its revenue by changing its price unilaterally. Note that by changing its price, a seller may affect the outcome of the second stage of the game, which in turn may cause the demand for other items to exceed their supply. We assume that the seller has no information about supplies of other items and does not utilize this information in computing its equilibrium strategy. Specifically, fix the prices of all items other than e , and suppose that when e is priced at p , $U_{e@p}$ denotes the set of consumers for whom *every*² maximum utility bundle contains the item e . The seller assumes that its revenue from changing its price to p is given by $p \min\{n_e, |U_{e@p}|\}$, and it is at equilibrium if for every price p this quantity is no larger than its current revenue $\pi_e \min\{n_e, f_e\}$.

Note that for instances of the game in which all the bundles in \mathcal{B}_s for all consumers contain at least two items, there always exists a trivial equilibrium in which all sellers charge extremely high prices and demand for all items is zero—in this case no seller can unilaterally lower its price and obtain non-zero demand. To avoid such trivial and uninteresting equilibria, we assume that at equilibrium a seller with $f_e = 0$ must charge a price of 0.

Finally, we note that sellers follow deterministic strategies; therefore Nash equilibria don't always exist. (Chawla and Roughgarden give a specific example of a Bertrand game in networks where an equilibrium does not exist.)

Prices of anarchy and stability

Our goal in this paper is to understand how efficiently two-sided markets with selfish buyers and sellers operate. We define the efficiency of allocation of an outcome of the Bertrand game, denoted

with the same type may have different values for the same bundles.

²The seller is being somewhat pessimistic by not including consumers for whom the item e belongs to *some* maximum utility bundle. However this is not an important issue from the point of view of stable outcomes—the seller can lower its price infinitesimally and break any ties in utility.

$\text{Val}(\pi, f)$, to be equal to the total value obtained by all the consumers. The *price of anarchy (POA)* of the game measures the largest gap between the efficiency of a Nash equilibrium of the game to the optimal achievable while satisfying capacity (supply) constraints. We provide upper bounds on the POA in terms of parameters such as the congestion in the network, the diameter of the network, the value-disparity among consumers, etc. Several of our bounds are tight—we give examples where *all* Nash equilibria achieve poor efficiency. Formally, we provide lower bounds on the *price of stability (POS)* of the game. The POS measures the gap between the efficiency of a Nash equilibrium of the game to the optimal achievable, for the *best* Nash equilibrium.

$$\text{POA} = \max_{\text{Nash equilibria}(\pi, f)} \frac{\text{Val}(\text{OPT})}{\text{Val}(\pi, f)}$$

$$\text{POS} = \min_{\text{Nash equilibria}(\pi, f)} \frac{\text{Val}(\text{OPT})}{\text{Val}(\pi, f)}$$

As mentioned earlier, some instances of the Bertrand game do not admit Nash equilibria. Usually in literature the POA of such instances is defined to be infinite. We focus on instances of the game that do admit equilibria and give bounds on the POA of such instances.

Single-parameter & uniform demand instances

In Section 3 we study a special case of the Bertrand game in which consumers are single-valued, that is, every consumer values all the bundles in its type equally. Specifically, consumer s has a value v_s with $v_{s,B} = v_s$ for $B \in \mathcal{B}_s$ and 0 otherwise. For ease of exposition, in these instances, we group the demand of consumers with identical type. We assume that values lie in the interval $[1, \mathcal{L}]$. We use the function $F_{\mathcal{B}} : [1, \mathcal{L}] \rightarrow [0, 1]$ to denote the value distribution function associated with type \mathcal{B} and $D_{\mathcal{B}}$ to denote the total demand for this type— $F_{\mathcal{B}}(v)$ is the fraction of demand of consumers of type \mathcal{B} that have value at most v , while $F_{\mathcal{B}}(v)D_{\mathcal{B}}$ gives the total demand of those consumers.

Some of our results apply to instances where the distribution function $F_{\mathcal{B}}$ is identical for all types \mathcal{B} (however, they may have different total demands $D_{\mathcal{B}}$ associated with them). We call such instances *uniform-demand instances*.

For a value distribution function F , we use $\bar{F}(v)$ as a shorthand for $1 - F(v)$. F^{-1} denotes the inverse function— $F^{-1}(\alpha) = v$ iff $F(v) = \alpha$.

Bertrand game in networks

As in the work of Chawla and Roughgarden, most of our results apply to instances of the Bertrand game that arise in networks. In these instances items are *edges* in a directed network; each consumer is associated with a *source-sink pair* in the network, with the type of the consumer being the collection of all directed paths from its source to its sink.

Monopolies and virtual monopolies

A key parameter controlling the price of anarchy in Bertrand games is the existence and number of monopolies in the market. A seller e is a monopoly if there exists a consumer s such that e belongs to all the bundles in \mathcal{B}_s . We note that monopolistic behavior can exist in a game even in the absence of actual monopolies. Specifically, in a given state of the system, let $\mathcal{A}_s \subseteq \mathcal{B}_s$ denote the collection of bundles bought by consumer s ; we call a seller e a *virtual monopoly* for a consumer s if e belongs to all the bundles in \mathcal{A}_s . Note that if

e is a monopoly and is purchased by s , then e is a virtual monopoly for s .

Let U_e denote the set of consumers that buy e , that is, $U_e = \{s : \exists B \in \mathcal{A}_s \text{ with } e \in B\}$.

PROPOSITION 1. *At equilibrium, every seller e either charges a 0 per-unit price, or sells its entire inventory, or is a virtual monopoly for every consumer $s \in U_e$ that buys from it.*

PROOF. Suppose that a seller e charges a positive price, which implies that $U_e \neq \emptyset$. Let s be any consumer in U_e , that is e belongs to some bundle $B \in \mathcal{A}_s$. If e is not a virtual monopoly for s , then there exists another bundle $B' \in \mathcal{A}_s$ s.t. $e \notin B'$. Now further suppose e has some residual stock. By lowering its price by a negligible amount, e would make B strictly preferable to B' for consumer s , who would therefore abandon B' and switch to buying more (a non-negligible amount) of B . That is, the payoff to e could have been improved, contradicting the assumption of equilibrium. \square

3. BERTRAND GAMES WITH SINGLE PARAMETER CONSUMERS

In this section we study single-parameter instances of the Bertrand game in networks. Previous work shows that when the network contains monopolies the price of anarchy can be unbounded; this holds even in uniform-demand instances. We therefore first focus on instances that do not contain monopolies. We show that in uniform-demand instances when all the consumers have a common sink the price of anarchy can be bounded in terms of the congestion in the graph. In Section 3.3 we show that such a bound cannot be generalized to multiple-source multiple-sink instances—for any “non-trivial” value distribution, it is possible to construct a uniform-demand instance with that distribution that has an unbounded price of anarchy.

When the Bertrand game contains monopolies it is still possible to bound the price of anarchy under additional (mild) assumptions on the value distributions. Specifically, in uniform-demand multiple-source single-sink instances, if the value distribution satisfies the monotone hazard rate (MHR) condition (Definition 1), the price of anarchy is at most e^k , where k is the length of the longest source-sink path. We note that without the MHR condition even the best equilibrium can be a factor of \mathcal{L}^k worse than the optimal in terms of efficiency (recall that \mathcal{L} is the ratio of the largest to the smallest consumer value); furthermore this holds even in single-source single-sink networks.

3.1 General value distributions

As mentioned earlier, in multiple-source single-sink uniform-demand instances, the price of anarchy depends on the congestion in the network. Specifically, given a directed graph $G = (V, E)$ with capacities c_e on edges and demands $d_{s,t}$ between pairs of nodes $(s, t) \in V \times V$, the *sparsity* of a cut $(S, V \setminus S)$, $S \subsetneq V$, is defined as:

$$\alpha_{G,c,d}(S, V \setminus S) = \frac{\sum_{e \in \delta(S)} c_e}{\sum_{(s,t) \in S \times (V \setminus S)} d_{s,t}}$$

Here $\delta(S)$ is the set of edges exiting a set S : $\delta(S) = E \cap (S \times (V \setminus S))$. The sparsity of the network is defined as the minimum such quantity over all cuts in the graph.

$$\alpha(G, c, d) = \min_{S \subsetneq V} \alpha_{G,c,d}(S, V \setminus S)$$

The sparsity of a network gives a bound on the maximum fraction of demand that can be routed simultaneously from each source to its

corresponding destination while satisfying the capacity constraints on the edges. In single-sink networks, this bound is tight—the maximum fraction of demand that can be routed concurrently is exactly equal to the sparsity. It is therefore a measure of congestion in the network—when sparsity is small, the network is more congested.

We show below that in multiple-source single-sink uniform demand instances, in every equilibrium in the game, every source routes at least an $\alpha(G, c, d)$ fraction of its demand to the sink. Since the optimal solution can route no more than the entire demand from every source, the price of anarchy is at most $1/\alpha(G, c, d)$. On the other hand, while the optimal solution cannot route more than an $\alpha(G, c, d)$ fraction of every demand simultaneously, it may be able to route the entire demand from some of the sources, so the price of anarchy may be arbitrarily close to $1/\alpha(G, c, d)$ and the bound in the theorem is tight in general. (Better bounds can be obtained, of course, for specific distributions.) In fact Chawla and Roughgarden give an example of an instance where the bound is tight. We reproduce this example below for completeness and to give some intuition for how sparsity is related to the POA.




Figure 1: A Bertrand game instance with POA equal to $1/\text{sparsity}$

We now give an instance where the price of anarchy is arbitrarily close to the inverse of the sparsity of the network. As shown in Fig. 1, the network consists of two sources s_1, s_2 and a sink t , as well as four edges with capacities labeled therein. Total demands from s_1 and s_2 are β and 1, respectively. The value distribution satisfies the condition that $\bar{F}(p_0) = 2\epsilon$, where $p_0 = \operatorname{argmax}_p \{p\bar{F}(p)\}$ is the price a monopoly would have charged to maximize its revenue. Clearly, in any equilibrium, π_2 , i.e. the price at s_2 , is at least p_0 . It follows that $\pi_1 \leq \pi_2$, because otherwise $(s_1 \rightarrow t)$ would have been better off to charge a price of p_0 . On the other hand, from $(s_1 \rightarrow t)$'s perspective, the optimal price in the range $[0, \pi_2]$ is p_0 . Notice that we may set β arbitrarily large; in the extreme case, only the upper 2ϵ fraction of the total demand is admitted, while obviously the optimal allocation can accommodate virtually its entirety.

THEOREM 2. *Consider a single-sink uniform-demand instance of the Bertrand game over a directed network $G = (V, E)$ with sink t , capacities c_e on edges, demands d_s at sources s , and distribution function F . If the graph contains no monopolies and the instance admits a Nash equilibrium, its price of anarchy is no more than $1/\alpha(G, d, c)$, where $\alpha(G, d, c)$ is the sparsity of the graph with demands d and capacities c .*

PROOF. Let's begin with some notations. Let S be the set of all sources (consumers). Recall that for any $s \in S$, we use \mathcal{B}_s to denote the set of all $s-t$ paths, and $\mathcal{A}_s \subseteq \mathcal{B}_s$ to denote the set of paths carrying flow from s in the equilibrium under consideration. Let π_s be the price charged to consumer s , i.e. $\pi_s = \min_{P \in \mathcal{A}_s} \sum_{e \in P} \pi_e$.

We say that a path is saturated if at least one edge in the path is saturated.

Given an equilibrium, we bound its performance by showing that the price charged to each source is not too high. By Proposition 1, a positive-priced edge is either saturated or is a virtual monopoly. As far as the performance of the equilibrium is concerned, the former type of edges do no harm; hence we focus on virtual monopolies, which *might* be unsaturated. Our proof hinges on the observation that, since there are no monopolies in the network, we can always bound the price of an unsaturated path by that of another path that the consumer could have used. We analyze the graph constructed according to this path-price ordering, arguing for the existence of a subset of sources separated from the sink by saturated edges only, and thereby relate path prices to the sparsity $\alpha(G, d, c)$.

Specifically, let S_0 be the set of all sources s with $\pi_s = 0$, S_{sat} the set of all other consumers s with all paths in \mathcal{A}_s saturated, and S_{vm} the set of all the remaining consumers. Proposition 1 implies that for all $s \in S_{\text{vm}}$ every path in \mathcal{A}_s contains a virtual monopoly; in fact all positive-priced edges on these paths are virtual monopolies because they are unsaturated.

We construct a (directed) *price order graph* $H(S_H, R)$ over all the sources in $S_H = S_{\text{sat}} \cup S_{\text{vm}}$; for any $s_1, s_2 \in S_H$, $(s_1 \rightarrow s_2) \in R$ implies $\pi_{s_1} \leq \pi_{s_2}$. Our goal is to find a collection of sources in this digraph that each pay the maximum total price, and together saturate a cut in the graph G . Then by the definition of sparsity, we can deduce that these sources route at least an $\alpha(G, d, c)$ fraction of their demand.

We add two kinds of arcs to H . First, each source has an arc in H to any other source that shares a positive-priced edge in G with it. Second, for sources that do not use all available paths (because some of them are too expensive), we add arcs to sources that use those other paths (and therefore must pay more). In particular if $s \rightarrow t$ is an arc of the second type, then the (high) price paid by s can be “blamed” on the (even higher) price paid by t .

Formally, for every source $s \in S_H$ we add outgoing edges to R as follows. Let $E_s = \{e \in E : \exists P \in \mathcal{B}_s \text{ such that } e \text{ is the first edge on } P \text{ with } \pi_e > 0\}$. For every edge $e \in E_s$, consider the set U_e of sources that route flow over e in this equilibrium. It is easy to see that $\forall s' \in U_e \setminus \{s\}$ we have $\pi_s \leq \pi_{s'}$. So for all $s' \in U_e \setminus \{s\}$ we add $r = (s \rightarrow s')$ into H and label it with $l(r) = e$. If $s \in U_e$, we say the newly added edges are formed by *sharing*; otherwise, we say they are formed by *yielding*.

We will now show that for every source $s \in S_H$ the price π_s is bounded above by $\bar{F}^{-1}(\alpha(G, d, c))$. Note that H is composed of strongly connected components with the prices on sources within the same strongly connected component being equal. By our construction of H , we only need to prove the price bound on sources in “sink” components (i.e. components with no outgoing edges).

Consider a sink component C . We first claim that either $C \cap S_{\text{sat}}$ or $C \cap S_{\text{vm}}$ is empty. If not, then there exist sources $s_1 \in C \cap S_{\text{sat}}$ and $s_2 \in C \cap S_{\text{vm}}$ with $r = (s_1 \rightarrow s_2) \in R$ and $\pi_{s_1} = \pi_{s_2}$. We first claim that r cannot be formed by sharing. In particular, both s_1 and s_2 belong to $U_{l(r)}$. Moreover, $s_2 \in S_{\text{vm}}$ implies that all positive-priced edges carrying the flow of s_2 must be unsaturated, and so by Proposition 1, $l(r)$ is a virtual monopoly for both s_1 and s_2 . But then it is easy to see that $s_1 \in S_{\text{sat}}$ implies $s_2 \in S_{\text{sat}}$ which contradicts our assumption that $s_2 \in S_{\text{vm}}$. Therefore r is formed by yielding. Once again $l(r)$ must be unsaturated but this time $s_1 \notin U_{l(r)}$. But then $\pi_{s_1} = \pi_{s_2}$ implies that $l(r)$ has an incentive to lower its price negligibly in order to steal flows from paths in \mathcal{A}_{s_1} , contradicting the assumption of equilibrium.

Next, suppose that $C \subseteq S_{\text{vm}}$. Note that since the graph contains no monopolies, every source in S_{vm} must have an outgoing edge

that is formed by yielding. Since C has no outgoing edges, there must be two sources $s_1, s_2 \in C$ with $r = (s_1 \rightarrow s_2) \in R$ being a yielding edge. Once again we can argue as before that $\pi_{s_1} = \pi_{s_2}$ implies that $l(r)$ has an incentive to lower its price negligibly in order to steal flows from $P \in \mathcal{A}_{s_1}$, contradicting the assumption of equilibrium.

Therefore we must have $C \subseteq S_{\text{sat}}$.

Consider the cut $E_C = \cup_{s \in C} E_s$. It is not hard to see that $\forall r = (x \rightarrow y) \in C$, $l(r)$ must be saturated (otherwise, r would be by yielding implying, as we argued earlier, $\pi_x < \pi_y$). So by definition, E_C is a saturated cut separating sources in C from t . Since there are no edges outgoing from C in H , no edge in E_C carries the flow of a source not in C . Also note that because all sources in C have the same price associated with them, no two edges in E_C belong to the same flow path. Therefore, the amount of total flow routed from sources in C equals the capacity of the cut E_C , i.e. sources in C have achieved maximum concurrent flow. By the definition of sparsity, the capacity of E_C , and so also this concurrent flow, must be at least a $\alpha(G, d, c)$ fraction of the total demand at sources in C . Since consumers are at equilibrium, we must have $\pi_s \leq \bar{F}^{-1}(\alpha(G, d, c))$ for $s \in C$. Therefore we are done. \square

3.2 Value distributions with monotone hazard rate

The above theorem holds only in networks that do not contain any monopolies. Chawla and Roughgarden show that in networks with monopolies the price of anarchy can be unbounded. We now show that when value-demand curves satisfy the “monotone hazard rate (MHR)” condition, the price of anarchy of a network improves considerably and can be bounded in terms of the longest source-sink paths in the graph. The MHR condition, defined below, is widely used in game theory and economics to characterize commonly occurring distributions, and is satisfied, for example, by the uniform, normal, exponential, power-law (for exponents greater than one), Laplace, and chi-square distributions [5].

DEFINITION 1. (THE MONOTONE HAZARD RATE (MHR) CONDITION). *Let F be a distribution defined on $[1, \mathcal{L}]$, with density function f . The MHR condition states that the “hazard rate” of F , $h(v) = \frac{f(v)}{1-F(v)} = \frac{f(v)}{F(v)}$, is a monotonically non-decreasing function of v . A useful equation is $\bar{F}(v) = \exp\{-\int_1^v h(x)dx\}$.*

It has been noted previously [12] that a monopolist selling a single item to a consumer whose value is drawn from a distribution satisfying the MHR condition sells with probability at least $1/e$. Our theorem below essentially shows that if a consumer faces a path with k monopolists, it is able to route at least an e^{-k} fraction of its flow at equilibrium. This bound on the amount of demand routed holds more generally in any single-sink uniform-demand instance of the Bertrand game. We show below that this bound is tight, and further that it fails to extend to multiple-source multiple-sink instances.

THEOREM 3. *Consider a single-sink uniform-demand instance of the Bertrand game over a directed network $G = (V, E)$ with sink t , capacities c_e on edges, demands d_s at sources s , and distribution function F . If F has a monotonically non-decreasing hazard rate h , and the instance admits a Nash equilibrium, its price of anarchy is no more than $\min(e^k, 1/\alpha(G, d, c))$, where k is the maximum hop length of a source-sink path, and $\alpha(G, d, c)$ is the sparsity of the graph with demands d and capacities c .*

PROOF. We again consider the price order graph H constructed in the proof of Theorem 2. If a sink component C in H contains no

sources routing over monopolies, the previous analysis still holds. On the other hand, sources that route over saturated monopolies route at least an $\alpha(G, d, c)$ fraction of their flow by definition. We therefore focus on sources that route over unsaturated monopolies and bound the price paid by them.

Let s be an arbitrary source routing over an unsaturated monopoly. Let m be an arbitrary edge in \mathcal{A}_s that charges a positive price p . Let U_m be the set of sources that route flow over m , including s . According to Proposition 1, m must be a virtual monopoly to each $u \in U_m$. Let $\{q_u + p\}$ be the total price at each $u \in U_m$. The revenue function for m is $p \sum_{u \in U_m} d_u \bar{F}(q_u + p)$. Setting its derivative to zero, we have

$$p = \frac{\sum_{u \in U_m} d_u \bar{F}_u}{\sum_{u \in U_m} d_u h_u \bar{F}_u}$$

where $\bar{F}_u = \bar{F}(q_u + p)$ and $h_u = h(q_u + p)$ is the corresponding hazard rate. Supposing $m = (x \rightarrow y)$, let $Q = \pi_x$ be the total price on a flow carrying path from x to t ; then Q serves as a lower bound to all $q_u + p$. It follows $\forall u \in U_m, h_u \geq h(Q)$, which implies $ph(Q) \leq 1$.

Now let m_1, m_2, \dots, m_k be the sequence of edges on an arbitrary path in \mathcal{A}_s , in the order of from t to s . Let x_i be the price charged by m_i , and $Q_i = \sum_{j=1}^i x_j$; in particular, $Q_0 = 0$. From the previous analysis, we have $x_i h(Q_i) \leq 1$; it follows that $\forall t \in [Q_{i-1}, Q_i] h(t) \leq 1/x_i$. Therefore,

$$\begin{aligned} \bar{F}(Q_k) &= \exp \left\{ - \int_0^{Q_k} h(t) dt \right\} \\ &= \exp \left\{ - \sum_{j=1}^k \int_{Q_{j-1}}^{Q_j} h(t) dt \right\} \\ &\geq \exp \left\{ - \sum_{j=1}^k \int_{Q_{j-1}}^{Q_{j-1} + x_j} \frac{1}{x_j} dt \right\} \\ &= e^{-k}. \end{aligned}$$

That is, at least the most valuable e^{-k} fraction of the demand from s is routed in the equilibrium, and $\pi_s \leq \bar{F}^{-1}(e^{-k})$. Therefore once again we have

$$\text{POA} \leq \min(e^k, 1/\alpha(G, d, c)).$$

□

We now show that the above theorem is nearly tight even in single-source single-sink instances.

THEOREM 4. *There exists a family of single-source single-sink instances of the Bertrand game over a directed network $G = (V, E)$ with unbounded capacities and distribution function F with a monotonically non-decreasing hazard rate such that $\text{POS} = e^k/(k+1)$, where k is the number of hops between source s and sink t .*

PROOF. The network is a chain of k edges from s to t . The demand function is defined as $\bar{F}(x) = e^{-x}$ for $x \geq 0$. The hazard rate for this distribution is constant, thereby satisfying the monotone property. At equilibrium, each edge tries to optimize its own revenue by maximizing its revenue function $R(x) = x\bar{F}(x + P)$, where x is its own price and the constant P is the sum of the prices charged by other edges. Setting $R'(x) = (1-x)e^{-(x+P)} = 0$, we see that each edge will charge a price of 1 in any equilibrium.

The total social value is $\int_0^\infty xe^{-x} dx = 1$, while the total value achieved by an equilibrium is $\int_k^\infty xe^{-x} dx = (k+1)/e^k$. Thus, $\text{POS} = \frac{e^k}{k+1}$. □

3.3 Multiple-source multiple-sink instances

Finally we show that Theorems 2 and 3 cannot be extended to multiple-source multiple-sink instances. In particular, for any non-trivial value distribution, there exists a uniform demand instance of the Bertrand game with unbounded price of anarchy.

For the theorem below, given a value distribution F , we define the ‘‘monopoly flow’’ for the distribution to be the fraction of demand admitted in an equilibrium by a single monopoly for a consumer with a value drawn from F , that is, $\text{argmax}_x x\bar{F}^{-1}(x)$.




Figure 2: (a) Multi-sink instance with unbounded POA; (b) Augmented version without monopolies

THEOREM 5. *Let F be any distribution with monopoly flow strictly less than 1. Then for every $\beta > 1$, there exists a uniform-demand instance of the Bertrand game with distribution F containing no monopolies, with price of anarchy β .*

PROOF. We first present an example containing monopolies, after which we show how to convert it into one without monopolies.

Let $[1, \mathcal{L}]$ be the support of the value distribution F ; that is, $F(1) = 0$ and $\mathcal{L} = \min_x F^{-1}(1)$. Consider the network of a single chain with nodes $\{0, 1, \dots, k\}$ and edges $(j \rightarrow j-1)$ for $1 \leq j \leq k$, as shown in Fig. 2(a). The capacity of each edge is unbounded. At each node $j > 0$ resides a source with the corresponding sink being $j-1$ and total demand α . In addition, there is a consumer s_0 desiring the whole chain, whose total demand is β .

Due to symmetry among the edges, without consumer s_0 , in an equilibrium each edge would behave exactly in the same way. Suppose each of them would set down to a price of $p_1 > 0$ and that $(k-1)p_1 > \mathcal{L}$ for some suitably chosen α , then the existence of s_0 would not violate the equilibrium, as no edge can lower its price unilaterally and admit the flow of s_0 . Since this is regardless of the value β , the price of anarchy ($> \beta/(k\alpha)$) is unbounded.

Fig. 2(b) illustrates how to augment the previous instance to one containing no monopolies. Specifically, for each link $(i \rightarrow i-1)$, we parallel it with a double-stranded chain (C2 hereafter) consisting of low-capacity edges. Each segment of C2 is demanded by an additional consumer with total demand 1. We can make the length of this side chain (denoted by k') large enough s.t. $k'\bar{F}^{-1}(2\epsilon) > \mathcal{L}$. The consequence is that $(i \rightarrow i-1)$ may happily ‘‘ignore’’ C2 and be a virtual monopoly, preserving the analysis on the previous example. □

4. BERTRAND GAMES WITH MULTI PARAMETER CONSUMERS

We now investigate a more general version of the Bertrand game in which consumers value different bundles at different levels. We study this model in two-node parallel-link networks.

We consider two different versions of this model. In the related-values version, every edge in the network has a cost q_e associated with it, and every consumer s has an intrinsic value ℓ_s . The value of consumer s for edge e is given by $v_{s,e} = \max(\ell_s - q_e, 0)$. (Note that in parallel-link networks each edge is a source-sink path.) This model captures markets where consumers experience different fixed “latency” costs on different links. We remark that this is different from congestion costs considered previously [13, 1, 15] that are functions of the traffic carried by a link. In Section 5 we discuss a few different extensions of this model to more general networks.

In the unrelated-values version the values $v_{s,P}$ are completely arbitrary, although we assume that they are non-negative for all P , and zero for collections of edges that are not source-sink paths.

We first show that for the related-values model in a parallel link network, the behavior of the game is identical to that of the single-parameter version. In particular, in the absence of monopolies, the game always admits an equilibrium, and all equilibria are optimal.

THEOREM 6. *Consider an instance of the Bertrand game in a two-node parallel link network containing multi-parameter consumers with related values. Then the instance admits a Nash equilibrium and has price of anarchy 1.*

PROOF. We first prove that $\mathbf{POA} = 1$. Let k be the number of parallel links and, without loss of generality, suppose that the costs associated with the links have the order $q_1 \leq q_2 \leq \dots \leq q_k$. Obviously, when all links are saturated, or when the entire demand is admitted and routed along the links with the smallest q 's, the network achieves optimal value. An equilibrium is also optimal if some link charges a price of zero, and therefore there is no residual demand. Now suppose that, in an equilibrium, every link carries some flow and charges a positive price. Note that, since each consumer s only considers buying from $\operatorname{argmin}_{e \in E} \{q_e + \pi_e\}$, which is independent of ℓ_s , it must be that $q_x + \pi_x = q_y + \pi_y$ for all $x, y \in E$. Accordingly, an unsaturated link would have an incentive to steal flows by lowering its price negligibly, which implies that the network must be saturated. In either case, an equilibrium achieves the optimal social value.

For the first part, we show how to construct an equilibrium that supports the optimal efficiency. Starting with an idle network, we fill the links one by one from e_1 through e_k ; the priorities of commodities are in accordance with their intrinsic values ℓ_s , i.e. the most valuable commodity is always satisfied first. This process continues until ℓ_s becomes smaller than q_j , or commodities are exhausted, or the entire network is saturated. In the first two cases, let j be the first *unsaturated* link; for $i \geq j$, set $\pi_i = 0$; for $i < j$, set $\pi_i = q_j - q_i$. In the last case, let $\ell_0 = \min\{\ell_s | s \text{ has been admitted}\}$; set $\pi_i = \ell_0 - q_i$. It's easy to verify that the resulting configuration is an equilibrium.

□

In the unrelated-values version, the behavior of the game is very different. In particular, equilibria don't always exist (Fact 7). Furthermore, even when there are no real monopolies in the network, some edges may behave essentially as monopolies because some consumers may prefer one edge much more than any other edges. Therefore, the price of anarchy may be large. Theorem 8 shows that the price of anarchy is nevertheless bounded by a factor logarithmic in \mathcal{L} .




Figure 3: Instance admitting no equilibria (see Fact 7)

FACT 7. *There exist instances of the Bertrand game with unrelated values in parallel-link networks that admit no equilibria.*

PROOF. We exploit an extreme case of valuation arbitrariness where some consumers at a source entirely ignore certain links. The network as depicted in Fig. 3 has two links: e_1 has unbounded capacity while the capacity of e_2 is 0.1. Consider the value distribution F with $\bar{F}(1) = 1$, $\bar{F}(2) = 0.1$, and $\bar{F}(3) = 0$. Consumers at s can be partitioned into s_1 and s_2 . For the group s_1 , the distribution of v_{i,e_1} follows F while $v_{i,e_2} \equiv 0$. For the group s_2 , $v_{j,e_1} = v_{j,e_2}$ also follows F .

In excluding all possible combinations of π_1 and π_2 , we conclude that this example admits no equilibria:

- If $\pi_1 > \pi_2$, e_2 would be the only choice of consumers in s_2 . Subject to the capacity constraint, e_2 has to charge a price no less than 2. It follows that the revenue of e_1 is no more than $3\bar{F}(2) = 0.3$, whereas the revenue with $\pi_1 = 1$ would be 2.
- If $\pi_1 < \pi_2$, then e_2 would be idle and motivated to lower its price.
- If $\pi_1 = \pi_2 > 0$, e_1 would be incentivized to lower its price negligibly and grab the flow on e_2 . Finally, as a monopoly to s_1 , e_1 would not charge a price of zero in an equilibrium.

□

THEOREM 8. *Consider an instance of the Bertrand game in a two-node parallel link network containing multi-parameter consumers with unrelated values. If the instance admits a Nash equilibrium, its price of anarchy is bounded by $3 + \log \mathcal{L}$.*

PROOF. Let $U^* = \{U_e^*\}$ be the set of commodities routed in an optimal allocation OPT, where U_e^* is the set of flow routed by link e . Given an equilibrium EQ routing commodities $U = \{U_e\}$, for each link e , define $R_e = U_e^* \setminus U$ and $A_e = (U_e^* \setminus U_e) \cap U$. In words, commodities routed by e in OPT but not in EQ are partitioned into two parts: R_e , those left out from EQ; and A_e , those routed by some other edge in EQ.

In order to bound the loss of social value $\text{LOSS} = \mathbf{Val}(U^*) - \mathbf{Val}(U)$, we look at that on each link, i.e. $\text{LOSS}_e = \mathbf{Val}(U_e^*) - \mathbf{Val}(U_e)$. Clearly, LOSS_e is no more than the loss incurred by not routing R_e , plus the loss incurred by routing commodities in A_e via potentially less valuable links. For each $i \in R_e$, the corresponding loss is exactly $\text{loss}_i = v_{i,e} \leq \pi_e$, since i has opted not to route. For each $i \in A_e$, $\text{loss}_i = v_{i,e} - v_{i,e'}$, assuming it's routed by some link e' in EQ. Observe that, as i has chosen e' over e , we have $v_{i,e} - \pi_e \leq v_{i,e'} - \pi_{e'}$; it follows that $\text{loss}_i \leq \pi_e - \pi_{e'} \leq \pi_e$.

For every $i \in R_e \cup A_e$, let $\theta_{i,e}$ denote the maximum price that e can charge in order to win over i ; we call this price the “pseudo-value” of i . For $i \in R_e$, this is just $v_{i,e} = \text{loss}_i$. For $i \in A_e$, this is $(v_{i,e} - v_{i,e'}) + \pi_e \geq \text{loss}_i$. In either case, loss_i is bounded by the marginal revenue gained by e on the part of i when e lowers its price to merely admitting i . For $i \in U_e$, we define $\theta_{i,e}$ to be π_e . Let Θ_e denote the value distribution defined by these pseudo-values for consumers in $R_e \cup A_e \cup U_e$, that is, $\Theta_e(v)$ is the total demand of all consumers with $v \leq \theta_{i,e}$.

Note that if e is saturated, then we can bound the total loss for e easily: $\sum_{i \in R_e \cup A_e} \text{loss}_i \leq c_e \pi_e \leq \mathbf{Val}(U_e)$ because $|R_e \cup A_e| \leq |U_e^*| \leq c_e$.

Suppose that e is unsaturated, and that it would be saturated at the threshold price π_e^0 , provided that other links would stick to their current prices in EQ. In particular, $\pi_e^0 = 0$ if e would not be saturated at any price. (E.g. e might have unbounded capacity.) Let $\pi_0 = \max\{\pi_e^0, 1\}$. Define $X_e = \{i \in R_e \cup A_e \mid \text{loss}_i \geq \pi_0\}$, and $Y_e = \{i \in R_e \cup A_e \mid \text{loss}_i < \pi_0\}$. We first bound the loss of consumers in X_e .

Let $\mathbf{Pro}_e \leq \mathbf{Val}(U_e)$ be the revenue of e and f_e be the amount of flow through e in EQ; then since π_e is an optimal price for e , for $\pi_0 \leq x \leq \pi_e$, we have $\Theta_e(x) \leq \mathbf{Pro}_e/x$. Therefore,

$$\begin{aligned} \sum_{i \in X_e} \text{loss}_i &\leq \sum_{i \in X_e} \theta_{i,e} \leq \int_{\pi_0}^{\pi_e} (\Theta_e(\theta) - f_e) d\theta \\ &\leq \int_{\pi_0}^{\pi_e} \frac{\mathbf{Pro}_e}{\theta} d\theta \\ &\leq \mathbf{Pro}_e \log \pi_e \leq \mathbf{Val}(U_e) \log \mathcal{L}. \end{aligned}$$

It remains to bound the loss of consumers in Y_e . If $\pi_e^0 \geq 1$, which implies that e would be saturated at price π_e^0 and that $\pi_0 = \pi_e^0$, then

$$\sum_{i \in Y_e} \text{loss}_i \leq c_e \pi_e^0 \leq \mathbf{Pro}_e \leq \mathbf{Val}(U_e).$$

Otherwise, $\pi_e^0 < 1$ and $\pi_0 = 1$. Note that in this case, $Y_e \subseteq A_e$. Accordingly, $\forall i \in Y_e$, $\text{loss}_i < 1 \leq v_{i,e'}$, i.e. the loss on the part of Y_e can be covered by the corresponding value realized in EQ.

$$\sum_{i \in Y_e} \text{loss}_i \leq \mathbf{Val}(A_e).$$

Consequently,

$$\begin{aligned} \text{LOSS} &= \mathbf{Val}(U^*) - \mathbf{Val}(U) \\ &\leq \sum_e (1 + \log \mathcal{L}) \mathbf{Val}(U_e) + \sum_e \mathbf{Val}(A_e) \\ &\leq (2 + \log \mathcal{L}) \mathbf{Val}(U). \end{aligned}$$

□

We point out that this bound is nearly tight: the source may have very low valuation for all links but one—the effective monopoly—imitating monopoly pricing from the single-parameter setting. Specifically for a lower bound, we can use the single-source single-sink example from Theorem 2 of [7] with $\mathbf{POA} = \Theta(\log \mathcal{L})$, where $\bar{F}(x) = \epsilon + 1/x$, and construct a network with two links containing two kinds of consumers; the first type has value $1/2$ for the first link, and distributed according to F for the second link; the second type has value $1/2$ for the second link, and distributed according to F for the first link; then each link behaves as a monopoly for one of the types, and the price of anarchy is $\Theta(\log \mathcal{L})$.

5. DISCUSSION AND OPEN QUESTIONS

We have provided a nearly complete characterization of the price of anarchy of the Bertrand game in networks with single-parameter consumers. While in multiple-source multiple-sink networks the price of anarchy can be large regardless of the underlying distribution, it may be possible to bound the price of stability under certain assumptions on the value distribution. It would also be interesting to extend our results to markets that are more general than multi-source single-sink networks but less general than arbitrary multi-source multi-sink networks.

The most interesting open question related to our work is to study equilibria in the related values model in more general networks. Specifically, we can incorporate (traffic independent) latency costs q_e for each edge e by setting $\ell_{s,P} = \ell_s - \sum_{e \in P} q_e$ for a source-sink path P . Alternately, if q_e denotes the “quality of service” on edge e , then the value from a path P can be defined to be $\ell_{s,P} = \ell_s + \min_e q_e$. In the parallel-links case both of these models are identical to the one we considered in Section 4. In more general networks (even single-source single-sink networks), however, the two models may behave quite differently. We conjecture that the latency-costs model will continue to display the same equilibrium behavior as the single-parameter case in single-source single-sink networks. Another interesting model for path values is to set $\ell_{s,P}$ to be equal to $\ell_s \prod_e (1 - q_e)$ where $q_e \in [0, 1]$ are packet-loss probabilities on edges e .

Finally, an interesting open problem related to our work is finding Nash equilibria in Bertrand games when they exist.

6. REFERENCES

- [1] D. Acemoglu and A. Ozdaglar. Competition and efficiency in congested markets. *Mathematics of Operations Research*, 32(1):1–31, 2007.
- [2] D. Acemoglu and A. Ozdaglar. Competition in parallel-serial networks. *IEEE Journal on Selected Areas in Communications, special issue on Non-cooperative Behavior in Networking*, 25(6):1180–1192, 2007.
- [3] E. Anshelevich, A. Dasgupta, J. Kleinberg, É. Tardos, T. Wexler, and T. Roughgarden. The price of stability for network design with fair cost allocation. In *Proc. 45th IEEE Symp. Foundations of Computer Science*, pages 295–304, 2004.
- [4] E. Anshelevich, A. Dasgupta, É. Tardos, and T. Wexler. Near-optimal network design with selfish agents. In *Proc. ACM Symp. Theory of Computing*, pages 511–520, 2003.
- [5] Mark Bagnoli and Ted Bergstrom. Log-concave probability and its applications. *Economic Theory*, 26(2):445–469, 08 2005.
- [6] Michael R. Baye and John Morgan. A folk theorem for one-shot bertrand games. *Economics Letters*, 65(1):59–65, 1999.
- [7] Shuchi Chawla and Tim Roughgarden. Bertrand competition in networks. In *Symposium on Algorithmic Game Theory*, pages 70–82, 2008.
- [8] R. Cole, Y. Dodis, and T. Roughgarden. Pricing network edges for heterogeneous selfish users. In *Proc. 34th ACM Symp. Theory of Computing*, 2003.
- [9] A. Fabrikant, A. Luthra, E. Maneva, C. Papadimitriou, and S. Shenker. On a network creation game. In *Proc. ACM Symp. Principles of Distributed Computing*, pages 247–251, 2003.
- [10] L. Fleischer, K. Jain, and M. Mahdian. Tolls for heterogeneous selfish users in multicommodity networks and

- generalized congestion games. In *Proc. 45th IEEE Symp. Foundations of Computer Science*, pages 277–285, 2004.
- [11] Joseph Jr. Harrington. A re-evaluation of perfect competition as the solution to the bertrand price game. *Mathematical Social Sciences*, 17(3):315–328, 1989.
- [12] Jason Hartline, Vahab Mirrokni, and Mukund Sundararajan. Optimal marketing strategies over social networks. In *WWW '08: Proceeding of the 17th international conference on World Wide Web*, pages 189–198, 2008.
- [13] A. Hayrapetyan, É. Tardos, and T. Wexler. A network pricing game for selfish traffic. In *Proc. Symp. Principles of distributed Computing*, pages 284–291, 2005.
- [14] A. Mas-Colell, M. D. Whinston, and J. R. Green. *Microeconomic Theory*. Oxford, 1995.
- [15] J. Musacchio and S. Wu. The price of anarchy in a network pricing game. In *45th Annual Allerton Conference on Communication and Control*, 2007.
- [16] A. Ozdaglar and R. Srikant. Incentives and pricing in communication networks. In *Algorithmic Game Theory*. Cambridge Press, 2007.
- [17] T. Roughgarden. *Selfish Routing and the Price of Anarchy*. MIT Press, 2005.
- [18] T. Roughgarden and É. Tardos. How bad is selfish routing? *Journal of the ACM*, 49(2):236–259, 2002.