On the Efficiency of Markets with Two-sided Proportional Allocation Mechanisms

Volodymyr Kuleshov and Adrian Vetta*

Department of Mathematics and Statistics, and School of Computer Science, McGill University
volodymyr.kuleshov@mail.mcgill.ca, vetta@math.mcgill.ca

Abstract. We analyze the performance of resource allocation mechanisms for markets in which there is competition amongst both consumers and suppliers (namely, two-sided markets). Specifically, we examine a natural generalization of both Kelly's proportional allocation mechanism for demand-competitive markets [9] and Johari and Tsitsiklis' proportional allocation mechanism for supply-competitive markets [7].

We first consider the case of a market for one divisible resource. Assuming that marginal costs are convex, we derive a tight bound on the price of anarchy of about 0.5887. This worst case bound is achieved when the demand-side of the market is highly competitive and the supply-side consists of a duopoly. As more firms enter the market, the price of anarchy improves to 0.64. In contrast, on the demand side, the price of anarchy improves when the number of consumers decreases, reaching a maximum of 0.7321 in a monopsony setting. When the marginal cost functions are concave, the above bound smoothly degrades to zero as the marginal costs tend to constants. For monomial cost functions of the form $C(x) = cx^{1+\frac{1}{d}}$, we show that the price of anarchy is $\Omega(\frac{1}{d^2})$.

We complement these guarantees by identifying a large class of two-sided single-parameter market-clearing mechanisms among which the proportional allocation mechanism uniquely achieves the optimal price of anarchy. We also prove that our worst case bounds extend to general multi-resource markets, and in particular to bandwidth markets over arbitrary networks.

1 Introduction

How to produce and allocate scarce resources is the most fundamental question in economics.[1] The standard tool for guiding production and allocation is a pricing mechanism. However, different mechanisms will have different performance attributes: no two mechanisms are equal. Of particular interest to computer scientists is the fact that there will typically be an inherent trade-off between the economic efficiency of a mechanism (measured in terms of social welfare) and its computational efficiency (both time and communication complexity). Socially optimal allocations can be achieved using pricing mechanisms based on classical VCG results, but implementing such mechanisms generally induces excessively high informational and computational costs [13]. In this paper, we study this tradeoff from the opposite viewpoint: we examine the level of social welfare that can be achieved by mechanisms performing *minimal* amounts of computation. In particular, we restrict our attention to so-called *scalar-parametrized* pricing mechanisms. Each participant submits only a single scalar bid that is used to set a unique market-clearing price for each good. Evidently, such mechanisms are computationally trivial to handle; more surprisingly, they can produce high welfare.

The chief practical motivation for considering scalar-parametrized mechanisms (both in our work and in the existing literature) is the problem of bandwidth sharing. Namely, how should we allocate capacity amongst users that want to transmit data over a network link? The use of market mechanisms for this task has been studied in Asynchronous Transfer Mode (ATM) networks [16] and the Internet [15]. The Internet is made up of smaller interconnected networks that buy capacities from each other,

* The authors were supported in part by NSERC grant 28833

[1] In fact, economics is often defined as "the study of scarcity".

and the market mechanisms we consider are closely inspired by the structure of the Internet. Specifically, we are restricting our attention to mechanisms that are scalable to very large networks. This requirement for scalability forces us consider only simple mechanisms, such as those that set a unique market clearing price. The computational requirements of more complex systems, e.g. mechanisms that perform price discrimination, become impractical on large networks [1].

We remark that unique price mechanisms are also intuitively "fair", as every participant is treated equally. This fairness is appealing from a social and political perspective, and indeed these systems are used in many real-world settings, such as electricity markets [17].

1.1 Background and Previous Work

A basic method for resource allocation is the *proportional allocation mechanism* of Kelly [9]. In the context of networks, it operates as follows: each potential consumer submits a bid b_q; bandwidth is then allocated to the consumers in proportion to their bids. This simple idea has also been studied within economics by Shapley and Shubik [14] as a model for understanding pricing in market economies. In a groundbreaking result, Johari and Tsitsiklis [5] showed that the welfare loss incurred by this mechanism is at most 25% of optimal.

Observe that Kelly's is a scalar-paramterized mechanism for a one-sided market: every participant is a consumer. Johari and Tsitsiklis [7] also examined one-sided markets with supply-side competition only. There, under a corresponding single-parameter mechanism, the welfare loss tends to zero as the level of competition increases. We remark that we cannot simply analyze supply-side competition by trying to model suppliers as demand-side consumers [3].

Of course, competition in markets typically occurs on *both* sides. Consequently, understanding the efficiency of *two-sided market*[2] mechanisms is an important problem. In this work, we analyze the price of anarchy in a mechanism for a two-sided market in which consumers and producers compete simultaneously to determine the production and allocation of goods. This mechanism was first proposed by Neumayer [10] and is the natural generalization of both the demand-side model of Kelly [9] and the supply-side model of Johari and Tsitsiklis [7].

In order to examine how the generalized proportional allocation mechanism performs in a two-sided market, it is important to note that there are three primary causes of welfare loss. First, the underlying allocation problem may be computationally hard. In other words, in some settings (such as combinatorial auctions, for example), it may be hard to compute the optimal allocation even when the players' utilities are known. Secondly, even if the allocation problem is computationally simple, the mechanism itself may still be insufficiently sophisticated to solve it. Thirdly, the mechanism may be susceptible to gaming; namely, the mechanism may incentivize selfish agents to behave in a manner that produces a poor overall outcome. As we will see in Section 4, the first two causes do not arise here: as long as the users do not behave strategically, the proportional allocation mechanism can quickly find optimal allocations in two-sided markets. Thus, we are concerned only with the third factor: how adversely is the proportional allocation mechanism affected by gaming agents? That is, the mechanism may be capable of producing an optimal solution, but how will the agents' selfish behaviour affect social welfare at the resultant equilibria?

In this paper we prove that the proportional allocation mechanism does perform well in two-sided markets. Specifically, under quite general assumptions, the mechanism admits a constant factor price of anarchy guarantee. Moreover, there exists a large family of mechanisms among which the proportional allocation mechanism uniquely achieves the best possible price of anarchy guarantee. We state our exact results in Section 3, after we have described the model and our assumptions.

[2] It should be noted that "two-sided market" often has a different meaning in the economics literature than the one we use here. There it refers to a specific class of markets where externalities occur between groups on the two sides of the market.

2 The Model

2.1 The two-sided proportional allocation mechanism

We now formally present the two-sided proportional allocation mechanism due to Neumayer [10]. There are Q consumers and R suppliers in the market. Each consumer q has a valuation function $V_q(d_q)$, where d_q is the amount of the resource allocated to consumer q, and each supplier r has a cost function $C_r(s_r)$, where s_r is the amount produced by supplier r. Consumers and suppliers respectively input bids b_q and b_r to the mechanism. Doing so, consumers are implicitly selecting b_q-parametrized demand functions of the form $D(b_q, p) = \frac{b_q}{p}$, and suppliers are selecting b_r-parametrized supply functions of the form $S(b_r, p) = 1 - \frac{b_r}{p}$. We can also interpret a high consumer bid as an indicator of high willingness to pay for the product, and a low supplier bid as an indicator of a high willingness to supply (alternatively, a high bid indicates a high cost supplier). The actual choice of constant used for the supply functions does not affect our results, and so we choose it to be 1.

Observe that the parametrized demand functions are identical to the ones in the demand-side mechanism of Kelly [9], and the supply functions are identical to the ones in the supply-side mechanism of Johari and Tsitsiklis [7]. The peculiar form of the supply functions comes from the interesting fact that for most scalar-parametrized mechanisms, in order to have a non-zero welfare ratio, the supply functions have to be bounded from above. In other words, suppliers' strategies must necessarily be constrained in order to obtain high welfare; see the full version of the paper for the precise statement of this fact. This rules out, for instance, Cournot-style mechanisms where suppliers directly submit the quantities they wish to produce.

More detailed justifications for this choice of model can be found in [10], as well as in [9] and [7]. Further justification for the mechanism will be provided by our results. Specifically, the proportional allocation mechanism generally produces high welfare allocations and, in addition, it is the optimal mechanism amongst a class of single-parameter mechanisms for two-sided markets.

Given the bids, the mechanism sets a price $p(\mathbf{b})$ that clears the market; i.e. that satisfies the supply equals demand equation: $\sum_{q=1}^{Q} \frac{b_q}{p} = \sum_{r=1}^{R}(1 - \frac{b_r}{p})$. The price therefore gets set to $p(\mathbf{b}) = \frac{\sum_q b_q + \sum_r b_r}{R}$. Consumer q then receives d_q units of the resource, and pays pd_q, while supplier r produces s_r units and receives a payment of ps_r. In the game induced by this mechanism, the payoff (or utility) to consumer q placing a bid b_q is defined to be

$$\Pi_q(b_q) = \begin{cases} V_q\left(\frac{b_q}{\sum_{q \in Q} b_q + \sum_r b_{r \in R}} R\right) - b_q & \text{if } b_q > 0 \\ V_q(0) & \text{if } b_q = 0 \end{cases}$$

and the payoff to supplier r placing a bid b_r is defined as

$$\Pi_r(b_r) = \begin{cases} \frac{\sum_{q \in Q} b_q + \sum_{r \in R} b_r}{R} - b_r - C_r\left(1 - \frac{b_r}{\sum_{q \in Q} b_q + \sum_{r \in R} b_r} R\right) & \text{if } b_r > 0 \\ \frac{\sum_{q \neq r} b_q + \sum_{r \in R} b_r}{R} - C_r(1) & \text{if } b_r = 0 \end{cases}$$

2.2 The Welfare Ratio

Given a vector of bids $\mathbf{b}$, the *social welfare* at the resulting mechanism allocation is defined to be

$$\mathcal{W}(\mathbf{b}) = \sum_{q=1}^{Q} V_q(d_q(\mathbf{b})) - \sum_{r=1}^{R} C_r(s_r(\mathbf{b}))$$

If the agents do not strategically anticipate the effects of their actions on the price, that is if they act as "price-takers", we show in Section 4 that the mechanism maximizes social welfare. However, since the price is a function of their bid, each agent is a "price-maker". If agents attempt to exploit this market power, then a welfare loss may occur at a Nash equilibrium. Consequently we are interested in maximizing (over all equilibria) the *welfare ratio*, more commonly known as the *price of anarchy*, $\frac{\mathcal{W}^{\text{NE}}}{\mathcal{W}^{\text{OPT}}}$. Equivalently, we wish to minimize the *welfare loss*, $1 - \frac{\mathcal{W}^{\text{NE}}}{\mathcal{W}^{\text{OPT}}}$.

2.3 Assumptions

We make the following assumption on the valuation and cost functions.

Assumption 1 *For each consumer q, the valuation function $V_q(d_q) : \mathbb{R}^+ \to \mathbb{R}^+$ is strictly increasing and concave. For each supplier r, the cost function $C_r(s_r) : \mathbb{R}^+ \to \mathbb{R}^+$ is strictly increasing and convex.*

Assumption 1 corresponds to decreasing marginal valuations and increasing marginal costs. The assumption is standard in the literature. It certainly may not hold in every market[3], but without it there will be a natural incentive for the number of agents to decline on both sides of the market. In this paper, we will also assume that our functions are differentiable over their entire domain; this property is assumed primarily for clarity and is not essential.

Assumption 1, however, is not sufficient to ensure a large welfare ratio. In fact, the welfare ratio depends upon the curvature of the *marginal cost functions*. Specifically, if the marginal cost functions are convex, then we show in Section 4 that the welfare ratio is at least 0.58. Concave marginal cost functions also exhibit constant welfare ratios, provided the corresponding total cost function is sufficiently non-linear. However, in the limit as the total cost functions become linear, the welfare ratio degrades to zero (see Section 5 for more details).

Our main result thus concerns convex marginal cost functions. Formally, for most of the paper, we assume that

Assumption 2 *For each supplier r, the marginal cost function $C_r'(s_r)$ is convex. Furthermore, we assume that $C_r(0) = C_r'(0) = 0$.*

Convex marginal cost functions are extremely common in both the theoretical and the practical literature on industrial theory [18], so this assumption is not particularly restrictive. In Assumption 2 we also set $C_r'(0) = 0$, but as we show in the full version of the paper, constant welfare ratios still arise whenever $C_r'(0)$ is bounded below one (it cannot be higher than one or the firm is uncompetitive).

We also remark that Assumption 2 was used in Johari, Mannor and Tsitsiklis [4] in their analysis of the demand-side proportional allocation mechanism with elastic supply. Most of the results of Johari and Tsitsiklis [6] and Tobias and Harks [2] on demand-side Cournot competition with elastic supply also hold under the assumption of convex marginal costs.

3 Our Results

Our first results are concerned with the performance of the mechanism when the users act as price-takers. Under Assumption 1, we prove that:

Theorem 1. *A unique competitive equilibrium exists for the two-sided proportional allocation mechanism. The social welfare attained at the competitive equilibrium is optimal.*

This property was exhibited by Kelly's original proportional allocation mechanism, and has been a feature of all subsequent generalizations by Johari and Tsitsiklis. It is very appealing from a practical point of view, as in actual networks, users are likely to have little information about each other, making it difficult to manipulate the system.

In many other settings however, users will be incentivized to act strategically. In that case, we need to use the stronger solution concept of a Nash equilibrium to analyze the resulting game. Our second result establishes the existence and uniqueness of such equilibria under Assumption 1.

[3] For example, in markets exhibiting economies of scale.

Theorem 2. *The two-sided proportional allocation mechanism has a unique Nash equilibrium for $R \geq 2$.*

Our main result measures the loss of welfare at that unique Nash equilibrium under Assumption 2.

Theorem 3. *The worst case welfare ratio for the mechanism involving $R \geq 2$ suppliers equals*

$$\frac{s^2((R-1)^2 + 4(R-1)s + 2s^2)}{(R-1)(R-1+2s)}$$

where s is the unique positive root of the quartic polynomial $\gamma(s) = 16s^4 + (R-1)s^2(49s-24) + 10(R-1)^2 s(3s-2) + (R-1)^3(5s-4)$. Furthermore, this bound is tight.

It follows that the mechanism admits a constant bound on the price of anarchy. Moreover, Theorem 3 allows us to measure the effects of market competition on social welfare. The following two corollaries are concerned with that relationship.

Corollary 1. *The worst possible price of anarchy is achieved when the supply side is a duopoly ($R = 2$). It evaluates numerically to about 0.588727.*

Corollary 2. *When the supply side is fully competitive ($R \to \infty$) the price of anarchy equals precisely 0.64.*

Consequently, as supply-side competition increases, the welfare ratio improves. In contrast, the welfare ratio *decreases* as demand-side competition increases. Although this fact may seem surprising at first, it turns out to have a simple intuitive explanation. The optimal demand-side allocation consists in giving the entire production to the user which derives from it the highest utility. When more consumers are present in the market, they selfishly request more of the resource for themselves, leaving less for the most needy user and reducing the overall social welfare.

The best welfare ratios thus arise when there is only one consumer ($Q = 1$), that is, in the case of a *monopsony*. In the two-sided proportional allocation mechanism, the best possible price of anarchy over all possible values of Q and R is given by the next corollary.

Corollary 3. *In a market in which a monopsonist faces a fully competitive supply side, the price of anarchy equals $\sqrt{3} - 1$, which is about 0.7321.*

Recall that in the one-sided proportional allocation mechanism for suppliers facing a fixed demand, the welfare loss tends to zero when the supply side is fully competitive [7]. In contrast, Corollary 3 implies that in two-sided markets, that result no longer holds and that full efficiency cannot be achieved.

So far, our results assumed the convexity of marginal costs. Dropping that assumption, we find that the welfare ratio equals zero when the providers' total cost functions are linear. However, the price of anarchy remains bounded for a class of concave marginal cost functions, and degrades smoothy to zero as the total costs become linear.

Corollary 4. *The welfare ratio for cost functions $C_r(s_r) = c_r s_r^{1+\frac{1}{d}}$ where $c_r > 0$ and $d \geq 1$ is $\Omega(\frac{1}{d^2})$.*

Like its one-sided versions, the two-sided mechanism can be generalized to multi-resource markets. An important multi-resource setting is that of bandwidth shared on a network of links. The same guarantees as in the single-resource setting hold for the network version of our market, as well as for more general multi-resource markets (see the full version of the paper for more details).

Theorem 4. *The welfare ratio in networks equals that of the single-resource model.*

Theorem 5. *The welfare guarantees hold for more general multi-resource markets.*

Finally, we show that the proportional allocation mechanism is optimal in the following way:

Theorem 6. *In two-sided markets, the proportional allocation mechanism provides the best welfare ratio amongst a class of single-parameter market-clearing mechanisms.*

Our proof techniques are inspired by the approaches and techniques developed to analyze single-sided markets by Johari [3], Johari and Tsitsiklis ([5], [8] and [7]), Johari, Mannor and Tsitsiklis [4], Tobias and Harks [2], and Roughgarden [12]. Due to space limitations, most of our results will be deferred to the full version of the paper. Here, we will focus upon the proof of Theorem 3.

4 Optimization in Eight Steps

The proof of the main result, Theorem 3, is presented below in eight steps. We formulate the efficiency loss problem as an optimization program in Step III. To be able to formulate this we first need to understand the structure of optimal solutions and of equilibria under this mechanism. This we do in Steps I and II, where we give necessary and sufficient conditions for optimal solutions and for equilibrium. This leads us to an optimization problem that initially appears slightly formidable, so we then attempt to simplify it. In Steps IV and V, we show how to simplify the demand constraints in the program, and in Steps VI and VII, we simplify the supply constraints. This produces an optimization program in a form more amenable to quantitive analysis; we perform this analysis in Step VIII.

Step I: Optimality Conditions. The best possible allocation is the solution to the system:

$$\text{(OPT)} \qquad \begin{aligned} \max \ & \sum_{q=1}^{Q} V_q(d_q^{\text{OPT}}) - \sum_{r=1}^{R} C_r(s_r^{\text{OPT}}) \\ \text{s.t} \ & \sum_{q=1}^{Q} d_q^{\text{OPT}} = \sum_{r=1}^{R} s_r^{\text{OPT}} \\ & 0 \leq s_r^{\text{OPT}} \leq 1 \\ & d_q^{\text{OPT}} \geq 0 \end{aligned}$$

Since the constraints are linear, there exists an optimal solution at which the Karush-Kuhn-Tucker (KKT) conditions hold. As the objective function is concave, the following first order conditions are both necessary and sufficient:

$$\begin{aligned} C_r'\left(s_r^{\text{OPT}}\right) &\leq \lambda && \text{if } 0 < s_r^{\text{OPT}} \leq 1 \\ C_r'\left(s_r^{\text{OPT}}\right) &\geq \lambda && \text{if } 0 \leq s_r^{\text{OPT}} < 1 \\ V_q'\left(d_q^{\text{OPT}}\right) &\leq \lambda && \text{if } d_q^{\text{OPT}} = 0 \\ V_q'\left(d_q^{\text{OPT}}\right) &= \lambda && \text{if } d_q^{\text{OPT}} > 0 \end{aligned}$$

We have used λ to denote the dual variable corresponding to the equality constraint.

Step II: Equilibria Conditions. Here we describe necessary and sufficient conditions for a set of bids $\mathbf{b}$ to form Nash equilibrium.

First, observe that there must be at least two suppliers, that is $R \geq 2$. If not, then we have a monopolist k whose payoff is is strictly increasing in b_k. Specifically,

$$\Pi_k(b_k, b_{-k}) = \sum_q b_q - C_k\left(1 - \frac{b_k}{b_k + \sum_q b_q}\right) = \sum_q b_q - C_k\left(\frac{\sum_q b_q}{b_k + \sum_q b_q}\right)$$

Next, we show that if $\mathbf{b}$ is a Nash equilibrium, then at least two bids must be positive. Suppose for a contradiction that we have a supplier k and $\sum_{r\neq k} b_r = \sum_q b_q = 0$. Then $\Pi_k(0) = -C_k(1)$, and $\Pi_k(b_k) = -\frac{R-1}{R} b_k$ when $b_k > 0$. For the second expression, we used the fact that $C_k(x) = 0$ for any $x \leq 0$. Observe that if $b_k = 0$ then the firm can profitably deviate by increasing b_k infinitesimally; on the other hand, if $b_k > 0$ then the firm should infinitesimally decrease b_k. Thus, there is no equilibrium in which either all bids are zero, or a single supplier is the only agent to make a positive bid. Thus there must be at least two positive bids at equilibrium.

Since at least two bids are positive, the payoffs Π_k are differentiable and concave, and the following conditions are necessary and sufficient for the existence of a Nash equilibrium. For the suppliers,

$$C_r'(s_r)\left(1+\frac{s_r^{\mathrm{NE}}}{R-1}\right) \geq p \quad \text{if } 0 < b_r \leq p \qquad\qquad C_r'(s_r)\left(1+\frac{s_r^{\mathrm{NE}}}{R-1}\right) \leq p \quad \text{if } 0 \leq b_r < p$$

For the consumers, $V_q'(0) \leq p$ and $V_q'(d_q^{\mathrm{NE}})\left(1-\frac{d_q}{R}\right) = p$ if $d_q^{\mathrm{NE}} > 0$.

Step III: An optimization problem. We can now formulate the welfare ratio as an optimization problem.

$$\min \quad \frac{\sum_{q=1}^{Q} V_q(d_q^{\mathrm{NE}}) - \sum_{r=1}^{R} C_r(s_r^{\mathrm{NE}})}{\sum_{q=1}^{Q} V_q(d_q^{\mathrm{OPT}}) - \sum_{r=1}^{R} C_r(s_r^{\mathrm{OPT}})} \tag{1}$$

$$\text{s.t.} \quad V_q'(d_q^{\mathrm{NE}})\left(1-\frac{d_q^{\mathrm{NE}}}{R}\right) \geq p \quad \forall q \text{ s.t. } d_q^{\mathrm{NE}} > 0 \tag{2}$$

$$V_q'(d_q^{\mathrm{NE}})\left(1-\frac{d_q^{\mathrm{NE}}}{R}\right) \leq p \quad \forall q \tag{3}$$

$$C_r'(s_r^{\mathrm{NE}})\left(1+\frac{s_r^{\mathrm{NE}}}{R-1}\right) \leq p \quad \forall r \text{ s.t. } 0 < s_r^{\mathrm{NE}} \leq 1 \tag{4}$$

$$C_r'(s_r^{\mathrm{NE}})\left(1+\frac{s_r^{\mathrm{NE}}}{R-1}\right) \geq p \quad \forall r \text{ s.t. } 0 \leq s_r^{\mathrm{NE}} < 1 \tag{5}$$

$$\sum_{q=1}^{Q} d_q^{\mathrm{NE}} = \sum_{r=1}^{R} s_r^{\mathrm{NE}} \tag{6}$$

$$C_r'(s_r^{\mathrm{OPT}}) \leq \lambda \quad \forall r \text{ s.t. } 0 < s_r^{\mathrm{OPT}} \leq 1 \tag{7}$$

$$C_r'(s_r^{\mathrm{OPT}}) \geq \lambda \quad \forall r \text{ s.t. } 0 \leq s_r^{\mathrm{OPT}} < 1 \tag{8}$$

$$V_q'(d_q^{\mathrm{OPT}}) \leq \lambda \quad \forall q \text{ s.t. } d_q^{\mathrm{OPT}} = 0 \tag{9}$$

$$V_q'(d_q^{\mathrm{OPT}}) = \lambda \quad \forall q \text{ s.t. } d_q^{\mathrm{OPT}} > 0 \tag{10}$$

$$\sum_{q=1}^{Q} d_q^{\mathrm{OPT}} = \sum_{r=1}^{R} s_r^{\mathrm{OPT}} \tag{11}$$

$$d_q^{\mathrm{OPT}}, d_q^{\mathrm{NE}} \geq 0 \quad \forall q \tag{12}$$

$$0 \leq s_r^{\mathrm{NE}}, s_r^{\mathrm{OPT}} \leq 1 \quad \forall q, r \tag{13}$$

$$p, \lambda \geq 0 \tag{14}$$

Given the cost and valuation functions, the constraints (2)-(6) are necessary and sufficient conditions for a Nash equilibrium by Step II, and constraints (7)-(11) are the optimality conditions from Step I. We now want to find the worst-case cost and valuation functions for the mechanism.

Step IV: Linear Valuation Functions. To evaluate this intimidating looking program we attempt to simplify it. First, efficiency loss is worst when each consumer has a linear valuation function. This is simple to show using a standard trick (see, for example, [5]). Thus, we restrict ourselves to linear functions of the form $V_q(d_q) = \alpha_q d_q$. Without loss of generality, we may assume that $\alpha_1 \geq \alpha_2 \geq ... \geq \alpha_Q$ and that $\max_q \alpha_q = 1$ after we normalize the functions by $1/\max_q \alpha_q$. Observe that this implies that $d_1^{\mathrm{OPT}} = \sum_r s_r^{\mathrm{OPT}}$ and $d_q^{\mathrm{OPT}} = 0$ for $q > 1$. As a result the objective function becomes $\left(d_1^{\mathrm{NE}} + \sum_{q=2}^{Q} \alpha_q d_q^{\mathrm{NE}} - \sum_{r=1}^{R} C_r(s_r^{\mathrm{NE}})\right) / \left(\sum_{r=1}^{R} s_r^{\mathrm{OPT}} - \sum_{r=1}^{R} C_r(s_r^{\mathrm{OPT}})\right)$, and the optimality constraints become $C_r'(s_r^{\mathrm{OPT}}) \leq 1, \forall r$ s.t. $0 < s_r^{\mathrm{OPT}} \leq 1$ and $C_r'(s_r^{\mathrm{OPT}}) \geq 1, \forall r$ s.t. $0 \leq s_r^{\mathrm{OPT}} < 1$. With linear valuations, the new optimality constraints ensure s_r^{OPT} is optimal by setting the marginal cost of each supplier to the marginal valuation, $\alpha_1 = 1$, of the first consumer.

Step V: Eliminating the Demand Constraints. In this step, we describe how to eliminate the demand constraints from the program. First we show that we can transform constraint (14) into $0 \leq p < 1$. Since $\alpha_q \leq 1, \forall q$, we see that constraint (2) implies that $p \leq 1$. Furthermore, if $p = 1$, then (2) can never be satisfied, and so we must have $d_q^{\mathrm{NE}} = 0, \forall q$. The supply equals demand constraint (6) then gives $s_r^{\mathrm{NE}} = 0, \forall r$. This gives a contradiction as the resulting allocation is not a Nash equilibrium: any supplier can increase its profits by providing a bid slightly smaller than p (remember that $C_r'(0) = 0$ by Assumption 2). Thus $p < 1$. This, in turn, implies that $d_1^{\mathrm{NE}} > 0$. To see this, note that if $d_1^{\mathrm{NE}} = 0$ then (3) cannot be satisfied for $q = 1$. Consequently, constraints (2) and (3) must hold with equality for $q = 1$. In fact, without loss of generality, constraints (2) and (3) hold with equality for $q > 1$. If constraint (2) does not hold with equality, we can reduce α_q, and this does not increase the value of the objective function. If $d_q^{\mathrm{NE}} = 0$ and constraint (3) does not hold with equality, we can set $\alpha_q = p$ and the objective function will be unaffected. So, $\alpha_q = \frac{p}{1-d_q^{\mathrm{NE}}/R}$ for all q. Substituting into the objective function:

$$\min \quad \frac{d_1^{\mathrm{NE}} + p\sum_{q=2}^{Q} \frac{d_q^{\mathrm{NE}}}{1-d_q^{\mathrm{NE}}/R} - \sum_{r=1}^{R} C_r(s_r^{\mathrm{NE}})}{\sum_{r=1}^{R} s_r^{\mathrm{OPT}} - \sum_{r=1}^{R} C_r(s_r^{\mathrm{OPT}})} \tag{15}$$

$$\text{s.t.} \quad \left(1 - \frac{d_1^{\mathrm{NE}}}{R}\right) = p \tag{16}$$

$$C_r'(s_r^{\mathrm{NE}})\left(1 + \frac{s_r^{\mathrm{NE}}}{R-1}\right) \leq p \qquad \forall r \text{ s.t. } 0 < s_r^{\mathrm{NE}} \leq 1 \tag{17}$$

$$C_r'(s_r^{\mathrm{NE}})\left(1 + \frac{s_r^{\mathrm{NE}}}{R-1}\right) \geq p \qquad \forall r \text{ s.t. } 0 \leq s_r^{\mathrm{NE}} < 1 \tag{18}$$

$$\sum_{q=1}^{Q} d_q^{\mathrm{NE}} = \sum_{r=1}^{R} s_r^{\mathrm{NE}} \tag{19}$$

$$C_r'(s_r^{\mathrm{OPT}}) \leq 1 \qquad \forall r \text{ s.t. } 0 < s_r^{\mathrm{OPT}} \leq 1 \tag{20}$$

$$C_r'(s_r^{\mathrm{OPT}}) \geq 1 \qquad \forall r \text{ s.t. } 0 \leq s_r^{\mathrm{OPT}} < 1 \tag{21}$$

$$d_q^{\mathrm{NE}} \geq 0 \qquad \forall q \geq 2 \tag{22}$$

$$d_1^{\mathrm{NE}} > 0 \tag{23}$$

$$0 \leq s_r^{\mathrm{NE}}, s_r^{\mathrm{OPT}} \leq 1 \qquad \forall r \tag{24}$$

$$0 \leq p < 1 \tag{25}$$

Now, observe that the objective function is convex and symmetric in the variables $d_2, ..., d_Q$, when all the other variables are held fixed. Convexity holds because our function is a sum of functions $\frac{d_q^{\mathrm{NE}}}{1-d_q^{\mathrm{NE}}/R}$, $q = 2, ..., Q$, that are convex on the range $[0, R]$; note that $d_q^{\mathrm{NE}} \leq R$ by (6), (12) and (13).

Therefore, for any given fixed assignment to the other variables, we must have $d_2 = ... = d_Q := x$. Otherwise, we could reshuffle the variable labels and obtain a second minimum, which is impossible by the convexity of the objective function. So, after replacing every d_q by x, constraint (19) becomes $x = \left(\sum_{r=1}^{R} s_r^{\mathrm{NE}} - d_1^{\mathrm{NE}}\right) / (Q-1)$. After inserting constraint (16) and the new constraint (19), the numerator of the objective function (15) becomes

$$
\begin{aligned}
&(1-p)R + p(Q-1)\frac{x}{1-x/R} - \sum_{r=1}^{R} C_r(s_r^{\mathrm{NE}}) \\
&= (1-p)R + p(Q-1)\frac{\left(\sum_{r=1}^{R} s_r^{\mathrm{NE}} - d_1^{\mathrm{NE}}\right) / (Q-1)}{1 - \frac{1}{R}\left(\sum_{r=1}^{R} s_r^{\mathrm{NE}} - d_1^{\mathrm{NE}}\right) / (Q-1)} - \sum_{r=1}^{R} C_r(s_r^{\mathrm{NE}}) \\
&= (1-p)R + p\frac{\sum_{r=1}^{R} s_r^{\mathrm{NE}} - (1-p)R}{1 - \left(\sum_{r=1}^{R} s_r^{\mathrm{NE}} - d_1^{\mathrm{NE}}\right) / R\,(Q-1)} - \sum_{r=1}^{R} C_r(s_r^{\mathrm{NE}})
\end{aligned}
$$

Finally, observe that if we increase Q by one, the objective function (1) cannot increase, since we can set $d_{Q+1} = 0$ and at least keep the same objective function value as before. Therefore, without loss of generality, we can take the limit as $Q \to \infty$. Note that this only changes the objective function, as all the constraints that contained Q have been inserted into the function and can be eliminated. After these changes, the optimization problem becomes

$$
\begin{aligned}
\min \quad & \frac{(1-p)^2 R + p\sum_{r=1}^{R} s_r^{\mathrm{NE}} - \sum_{r=1}^{R} C_r(s_r^{\mathrm{NE}})}{\sum_{r=1}^{R} s_r^{\mathrm{OPT}} - \sum_{r=1}^{R} C_r(s_r^{\mathrm{OPT}})} && (26) \\
\text{s.t.} \quad & C_r'(s_r^{\mathrm{NE}})\left(1 + \frac{s_r^{\mathrm{NE}}}{R-1}\right) \le p \quad \forall r \text{ s.t. } 0 < s_r^{\mathrm{NE}} \le 1 && (27) \\
& C_r'(s_r^{\mathrm{NE}})\left(1 + \frac{s_r^{\mathrm{NE}}}{R-1}\right) \ge p \quad \forall r \text{ s.t. } 0 \le s_r^{\mathrm{NE}} < 1 && (28) \\
& C_r'(s_r^{\mathrm{OPT}}) \le 1 \quad \forall r \text{ s.t. } 0 < s_r^{\mathrm{OPT}} \le 1 && (29) \\
& C_r'(s_r^{\mathrm{OPT}}) \ge 1 \quad \forall r \text{ s.t. } 0 \le s_r^{\mathrm{OPT}} < 1 && (30) \\
& 0 \le s_r^{\mathrm{NE}}, s_r^{\mathrm{OPT}} \le 1 \quad \forall r && (31) \\
& 0 \le p < 1 && (32)
\end{aligned}
$$

Hence, we have achieved our goal and completely eliminated the demand side of the optimization problem. Specifically, all the demand constraints have been replaced with an expression that is a function of the supply-side allocation. Now we must find the worst such allocation.

Step VI: Linear Marginal Cost Functions The next step is to show that, in searching for a worst case allocation, we can restrict our attention to linear marginal cost functions of the form $C_r'(s_r) = \beta_r s_r$ where $\beta_r > 0$. In this section, we briefly sketch the proof of this fact and defer the full treatment to the full version of the paper. Our proof technique is based on the work of Johari, Mannor and Tsitsiklis on demand-side markets with elastic supply [4], [6].

The proof consists in exhibiting, for any family of cost functions $C_r(s_r), r \in R$, two new families $\hat{C}_r()$ and $\bar{C}_r()$ with the property that the C_r have a better performance ratio than the $\bar{C}_r$ which, in turn, have a a better performance ratio than the $\hat{C}_r$. Furthermore, the $\hat{C}_r$ will be a family with linear marginal costs, as desired. The cost functions are defined as

$$\bar{C}'_r(s_r) = \begin{cases} C'_r(s_r) & \text{if } s_r < s_r^{\text{NE}} \\ \dfrac{C'_r(s_r^{\text{NE}})}{s_r^{\text{NE}}}\, s_r & \text{if } s_r \geq s_r^{\text{NE}} \end{cases} \qquad \text{and} \qquad \hat{C}'_r(s_r) = \frac{C'_r(s_r^{\text{NE}})}{s_r^{\text{NE}}}\, s_r$$

where s_r^{NE} is the Nash equilibrium allocation to supplier r when the cost functions are $C_r(s_r)$. Observe that the s_r^{NE} still satisfy the Nash equilibrium conditions (27) and (28) for both $\bar{C}_r$ and $\hat{C}_r$. Thus $\bar{s}_r^{\text{NE}} = \hat{s}_r^{\text{NE}} = s_r^{\text{NE}}$. The heart of the proof consists in showing that the optimal welfare can only improve when going from one family to the next.

Step VII: Eliminating the Supply Constraints. Assuming linear marginal cost functions, the optimization problem (26)-(32) becomes

$$\min \quad \frac{(1-p)^2 R + p\sum_{r=1}^R s_r^{\text{NE}} - \frac{1}{2}\sum_{r=1}^R \beta_r (s_r^{\text{NE}})^2}{\sum_{r=1}^R s_r^{\text{OPT}} - \frac{1}{2}\sum_{r=1}^R \beta_r (s_r^{\text{OPT}})^2} \tag{33}$$

$$\text{s.t.} \quad \beta_r s_r^{\text{NE}}\left(1 + \frac{s_r^{\text{NE}}}{R-1}\right) \leq p \quad \forall r \text{ s.t. } 0 < s_r^{\text{NE}} \leq 1 \tag{34}$$

$$\beta_r s_r^{\text{NE}}\left(1 + \frac{s_r^{\text{NE}}}{R-1}\right) \geq p \quad \forall r \text{ s.t. } 0 \leq s_r^{\text{NE}} < 1 \tag{35}$$

$$\beta_r s_r^{\text{OPT}} \leq 1 \quad \forall r \text{ s.t. } 0 < s_r^{\text{OPT}} \leq 1 \tag{36}$$

$$\beta_r s_r^{\text{OPT}} \geq 1 \quad \forall r \text{ s.t. } 0 \leq s_r^{\text{OPT}} < 1 \tag{37}$$

$$0 \leq s_r^{\text{NE}}, s_r^{\text{OPT}} \leq 1 \quad \forall r \tag{38}$$

$$\beta_r > 0 \quad \forall r \tag{39}$$

$$0 \leq p < 1 \tag{40}$$

with the new variables β_r, $r = 1, ..., R$. From $s_r^{\text{OPT}} \geq s_r^{\text{NE}}$, we can then deduce that (34) and (35) hold with equality. Suppose they don't for some r. Then $s_r^{\text{NE}} = s_r^{\text{OPT}} = 1$. Constraint (34) is $\beta_r < \frac{p}{1+1/(R-1)} < p < 1$. Hence, $\beta_r = \frac{p}{1+1/(R-1)}$ will be a feasible solution (i.e. constraint (36) will still be satisfied). Furthermore, increasing β_r to $\frac{p}{1+1/(R-1)}$ will only decrease the objective function since this is equivalent to subtracting a positive number from the numerator and the denominator. We can further simplify the system by replacing constraints (36) and (37) with $s_r^{\text{OPT}} = \min(1/\beta_r, 1)$. It is easy to see that s_r^{OPT} and β_r satisfy the equation above if and only if they satisfy (36) and (37). The reduced optimization problem now becomes:

$$\min \quad \frac{(1-p)^2 R + p\sum_{r=1}^R s_r^{\text{NE}} - \frac{1}{2}\sum_{r=1}^R \beta_r (s_r^{\text{NE}})^2}{\sum_{r=1}^R s_r^{\text{OPT}} - \frac{1}{2}\sum_{r=1}^R \beta_r (s_r^{\text{OPT}})^2} \tag{41}$$

$$\text{s.t.} \quad \beta_r s_r^{\text{NE}}\left(1 + \frac{s_r^{\text{NE}}}{R-1}\right) = p \quad \forall r \tag{42}$$

$$s_r^{\text{OPT}} = \min(1/\beta_r, 1) \quad \forall r \tag{43}$$

$$0 < s_r^{\text{NE}}, s_r^{\text{OPT}} \leq 1 \quad \forall r \tag{44}$$

$$\beta_r > 0 \quad \forall r \tag{45}$$

$$0 \leq p < 1 \tag{46}$$

We can insert the equality constraints (42) and (43) into the objective function (41) to obtain:

$$\min \quad \frac{(1-p)^2 R + p\sum_{r=1}^{R} s_r^{\mathrm{NE}} - \frac{p}{2}\sum_{r=1}^{R} \frac{s_r^{\mathrm{NE}}}{1+s_r^{\mathrm{NE}}/(R-1)}}{\sum_{r=1}^{R} \min(1/\beta_r, 1) - \frac{p}{2}\sum_{r=1}^{R} \frac{\min(1/\beta_r,1)^2}{s_r^{\mathrm{NE}}(1+s_r^{\mathrm{NE}}/(R-1))}} \tag{47}$$

$$\text{s.t.} \quad 0 < s_r^{\mathrm{NE}} \le 1 \qquad \forall r \tag{48}$$

$$\beta_r = \frac{p}{s_r^{\mathrm{NE}}\left(1+s_r^{\mathrm{NE}}/(R-1)\right)} \qquad \forall r \tag{49}$$

$$0 \le p < 1 \tag{50}$$

The objective function (47) can be rewritten as:

$$\frac{\sum_{r=1}^{R}\left((1-p)^2 + p s_r^{\mathrm{NE}} - \frac{p}{2}\frac{s_r^{\mathrm{NE}}}{1+s_r^{\mathrm{NE}}/(R-1)}\right)}{\sum_{r=1}^{R}\left(\min(1/\beta_r,1) - \frac{p}{2}\frac{\min(1/\beta_r,1)^2}{s_r^{\mathrm{NE}}(1+s_r^{\mathrm{NE}}/(R-1))}\right)}$$

Consequently, the minimum of the optimization problem (47)-(50) is greater than or equal to

$$\min \quad \frac{(1-p)^2 + ps - \frac{p}{2}\frac{s}{1+s/(R-1)}}{\min(\frac{s(1+s/(R-1))}{p}, 1) - \frac{p}{2s(1+s/(R-1))}\min(\frac{s(1+s/(R-1))}{p}, 1)^2} \tag{51}$$

$$\text{s.t.} \quad 0 < s \le 1 \tag{52}$$

$$0 \le p < 1 \tag{53}$$

We have now reduced the system (33)-(40) to a two-dimensional minimization problem. The next step is to try to explicitly find the minimum.

Step VIII: Computing the Worst Case Welfare Ratio. To obtain Theorem 3 we need to solve the optimization problem (51)-(53) with R as a parameter. We show how to do this in the full version of the paper. Thus we have proved our main result. It has several ramifications. Firstly, the worst case welfare ratio occurs with duopolies, that is when $R = 2$. There we obtain $s = 0.566812\cdots$ which gives a worst case welfare ratio of $0.588727\cdots$. Moreover, observe that this bound is tight. Our proof is essentially constructive; costs and valuations can be defined to to create an instance that produces the bound. Secondly, the welfare ratio improves as the number of supplies increases. Specifically as $R \to \infty$, the bound tends to $\frac{16}{25}$. Thus we obtain Corollaries 1 and 2.

So, as supply-side competition increases, the welfare ratio does improves. The opposite occurs as demand-side competition increases. Specifically, adapting our approach gives Corollary 3.

5 Concave Marginal Cost Functions.

The welfare ratio tends to zero if the cost function is linear, that is if the marginal cost function is a constant; for an example see the full version of the paper. We can get some idea of how the welfare ratio tends to zero for concave marginal cost functions by considering a class of polynomial cost functions with degree $1 + \frac{1}{d}$. These functions give a welfare ratio of $\Omega(\frac{1}{d^2})$, for any constant d. A proof of this (Corollary 4) is given in the full version of the paper. See Neumayer [10] for another example of inefficiency in the presence of linear cost functions.

6 Extensions to Networks and Arbitrary Markets.

We can generalize our results for bandwidth markets over a single network connection to the case where bandwidth is shared over an entire network. In that model, each consumer q is associated with a source-sink pair, and providers at associated with edges of the network at which they can offer bandwidth. A consumer's payoff is a function of the maximum (s_q, t_q)-flow it can obtain using the bandwidth it has purchased in the network.

The welfare guarantees for the network model are the same as for the single-link case. A formal description of the network model and a proof of Theorem 4 is given in the full version of the paper. Moreover, if we identify links $e \in E$ with arbitrary resources, then our results extend to a general class of markets with any number of resources. The exact definition of these markets and a proof of Theorem 5 are also given in the full version of the paper.

7 Smooth Market-Clearing Mechanisms

It was shown in [3] and [8] that in one-sided markets, the proportional allocation mechanism uniquely achieves the best possible welfare ratio within a broad class of so-called *smooth market-clearing mechanisms.* This family has a natural extension to the case of two-sided mechanisms, and we show that, given a symmetry condition, the two-sided proportional allocation mechanism is optimal amongst that class of single-parameter mechanisms. A description of smooth market-clearing mechanisms and a proof of Theorem 6 is given in the full version of the paper.

References

1. Gupta, A., Stahl, D.O., Whinston, A.B.: The economics of network management. Commun. ACM 42(9), 57–63 (1999)
2. Harks, T., Miller, K.: Efficiency and stability of nash equilibria in resource allocation games. In: GameNets'09: Proceedings of the First ICST international conference on Game Theory for Networks. pp. 393–402. IEEE Press, Piscataway, NJ, USA (2009)
3. Johari, R.: Efficiency loss in market mechanisms for resource allocation. Ph.D. thesis, MIT (2004)
4. Johari, R., Mannor, S., Tsitsiklis, J.N.: Efficiency loss in a network resource allocation game: The case of elastic supply. Mathematics of Operations Research 29, 407–435 (2005)
5. Johari, R., Tsitsiklis, J.N.: Efficiency Loss in a Network Resource Allocation Game. Mathematics of Operations Research 29(3), 407–435 (2004)
6. Johari, R., Tsitsiklis, J.N.: A scalable network resource allocation mechanism with bounded efficiency loss. IEEE Journal on Selected Areas in Communications 24(5), 992–999 (2005)
7. Johari, R., Tsitsiklis, J.N.: Parameterized supply function bidding: Equilibrium and welfare (2006)
8. Johari, R., Tsitsiklis, J.N.: Efficiency of scalar-parameterized mechanisms. Oper. Res. 57(4), 823–839 (2009)
9. Kelly, F.: Charging and rate control for elastic traffic. European Transactions on Telecommunications 8, 33–37 (1997)
10. Neumayer, S.: Efficiency Loss in a Class of Two-Sided Mechanisms. Master's thesis, MIT (2007)
11. Rosen, J.B.: Existence and uniqueness of equilibrium points for concave n-person games. Econometrica 33(3), 520–534 (1965)
12. Roughgarden, T.: Potential functions and the inefficiency of equilibria. Proceedings of the International Congress of Mathematicians (ICM) 3, 1071–1094 (2006)
13. Semret, N.: Market mechanisms for network resource sharing. Ph.D. thesis, Columbia University, New York, NY, USA (1999)
14. Shapley, L.S., Shubik, M.: Trade using one commodity as a means of payment. Journal of Political Economy 85(5), 937–68 (October 1977)
15. Shenker, S., Clark, D., Estrin, D., Herzog, S.: Pricing in computer networks: reshaping the research agenda. SIGCOMM Comput. Commun. Rev. 26(2), 19–43 (1996)
16. Songhurst, D.J.: Charging communication networks: from theory to practice. Elsevier Science Inc., New York, NY, USA (1999)
17. Stoft, S.: Power system economics : designing markets for electricity. IEEE Press, Piscataway, NJ (2002)
18. Tirole, J.: The Theory of Industrial Organization, vol. 1. The MIT Press, 1 edn. (1988)