

The Worst-Case Efficiency of Cost Sharing Methods in Resource Allocation Games

Tobias Harks

Department of Quantitative Economics, Maastricht University, 6200 MD Maastricht, The Netherlands,
t.harks@maastrichtuniversity.nl

Konstantin Miller

Telecommunication Networks Group, Technische Universität Berlin, 10587 Berlin, Germany,
miller@tkn.tu-berlin.de

Resource allocation problems play a key role in many applications, including traffic networks, telecommunication networks, and economics. In most applications, the allocation of resources is determined by a finite number of independent players, each optimizing an individual objective function. An important question in all these applications is the degree of suboptimality caused by selfish resource allocation. We consider the worst-case efficiency of cost sharing methods in resource allocation games in terms of the ratio of the minimum guaranteed surplus of a Nash equilibrium and the maximal surplus. Our main technical result is an upper bound on the efficiency loss that depends on the class of allowable cost functions and the class of allowable cost sharing methods. We demonstrate the power of this bound by evaluating the worst-case efficiency loss for three well-known cost sharing methods: incremental cost sharing, marginal cost pricing, and average cost sharing.

Subject classifications: resource allocation games; cost sharing methods; network routing games; marginal cost pricing; inefficiency of equilibria.

Area of review: Optimization.

History: Received September 2009; revisions received July 2010, October 2010, February 2011, accepted April 2011.

1. Introduction

Resource allocation problems play a key role in many applications. Whenever a set of resources needs to be matched to a set of demands, the goal is to find the most profitable or least costly allocation of the resources to the demands. Examples of such applications come from a wide range of areas, including traffic networks (Beckmann et al. 1956, Knight 1924, Roughgarden 2005, Smith 1979, Wardrop 1952), telecommunication networks (Johari and Tsitsiklis 2006, Kelly et al. 1998, Srikant 2003), and economics (Mas-Colell et al. 1995, Moulin 2008, Novshek 1985). In most of the above applications, the allocation of resources is determined by a finite number of independent players, each optimizing an individual objective function. A natural framework for analyzing such noncooperative games are *congestion games* as introduced by Rosenthal (1973). Congestion games model the interaction of a finite set of strategic players that compete over a finite set of resources. A pure strategy of a player consists of a subset of resources, and the payoff of a player depends only on the number of players choosing the same or overlapping strategies.

An important variant of congestion games is known as *resource allocation games* in which each player assigns a nonnegative demand to each of its subsets available. The

payoff for a player is defined as the difference between the utility associated with the sum of the demands and the costs associated with the resources used. A prominent example of such a game is the traffic routing game of Haurie and Marcotte (1985), which builds upon the classical model of Wardrop (1952): the arcs in a given network represent the resources, the different origin-destination pairs correspond to the players, and the subsets of resources are the paths available in the network for each origin-destination pair. A strategy of a player is a distribution of traffic flow over its available paths. The latency that a player experiences traversing an arc is given by a (nondecreasing) function of the total flow on that arc. The cost for a player on an arc is given by the product of the latency and the player's flow contribution on the arc.

Resource allocation games also play a key role in telecommunication networks, where users want to route packets from their source node to some sink node in the network. In this type of application it is frequently assumed that each user receives a nonnegative utility from transmitting at a certain packet rate and that each link (resource) determines a congestion-dependent price per unit flow that is charged to its users; see Kelly et al. (1998) and Srikant (2003). In Kelly et al. it is assumed that every link has a total cost function (modeling total delay or packet loss)

and the price per unit flow is defined by the marginal cost function.

The above two examples can be cast in the light of cost sharing methods: every resource incurs a cost that is passed on to its users by charging every user a cost share. In the terminology of the cost sharing literature, the prevailing cost sharing method in transportation networks is *average cost sharing*, because the cost of a resource is the total delay, while every user pays the product of the current latency and its flow contribution. In telecommunication networks (see Kelly et al. 1998), every user is charged the *marginal cost* per unit of resource, which corresponds to *marginal cost pricing*. Note that in both cases the cost sharing method charges a *single price per unit of resource*. This property is considered desirable and indispensable for large-scale networks, because every resource needs only to pass a one-dimensional information to its users; see also the motivation given in Johari and Tsitsiklis (2009), Kelly et al. (1998), and Srikant (2003).

An important question in all these areas is the degree of suboptimality caused by selfish resource allocation. Because this suboptimality crucially depends on the specific cost sharing method used, we first have to define the design space of cost sharing methods. To this end, we define the following five properties listed below, which are defined more formally in §3.

1. *Separability*: The cost sharing method of a resource is a function only of the consumption of the considered resource.

2. *Cost-covering*: The cost of a resource is covered by the cost shares collected from the users.

3. *No charge for zero demand*: The cost share for every player is zero on resources not used by her.

4. *Nash-inducing*: The cost sharing method is a nonnegative, nondecreasing, differentiable, and convex function in the resource consumption of every player.

5. *Scalability*: The cost sharing method charges a single price per unit of resource.

We briefly discuss the above five requirements. The first assumption requires that the cost share of a resource depends only on the vector of its consumption by the players. This implies that the cost shares of a resource are independent of the usage of other resources and thus precludes any coordination between different resources. While this property seems restrictive, it is crucial for practical applications in which cost sharing methods have only local information about their own usage (see for instance the TCP/IP protocol design, where routers drop packets based on some function of the number of packets in the queue; see Srikant 2003). Assumptions (2) and (3) are standard in the economics literature and are the least controversial. The fourth assumption gives a sufficient condition on the existence of a pure Nash equilibrium of the induced resource allocation game and is frequently used in the economics literature; see Moulin (2008). The last requirement, certainly the most restrictive one, is motivated by focusing on cost sharing

methods that are applicable in the context of large-scale networks; see the discussion above. In the following, we will call a cost sharing method *basic* if it satisfies Assumptions (1)–(4). We will call a cost sharing method *scalable* if it satisfies Assumptions (1)–(5).

1.1. Our Results

We study the efficiency loss of Nash equilibria in the context of resource allocation games with basic and scalable cost sharing methods. Given a class of cost functions \mathcal{C} and a class of basic cost sharing methods \mathcal{D}_n for n players, we develop a general lower bound on the worst-case efficiency of Nash equilibria that depends only on \mathcal{C} and \mathcal{D}_n but not on the player's private utilities. We show that among all basic cost sharing mechanisms, there is an optimal mechanism (incremental cost sharing) that achieves full efficiency. Because the incremental cost sharing method is not scalable, we analyze the worst-case efficiency of two well-known scalable cost sharing methods: *marginal cost pricing* and *average cost sharing*. By applying our generic lower bound to marginal cost pricing and average cost sharing, we obtain the following results that are summarized below.

Results for Marginal Cost Pricing. For differentiable, nondecreasing, and convex marginal cost functions, we prove a lower bound of $4/(3 + \sqrt{5 + 4n})$ on the worst-case efficiency. In particular, this bound carries over to practically relevant $M/M/1$ functions that model queuing delays with arc-capacities. We complement this bound by presenting an asymptotically matching upper bound of $2(n - \sqrt{n})/(\sqrt{n}(n - 1))$, leaving only a gap for small n . We completely characterize the worst-case efficiency for polynomial cost functions with nonnegative coefficients (previous results covered only affine marginal costs). For symmetric games (players have equal utility functions and equal strategy space), we present a series of results showing that the worst-case efficiency of Nash equilibria significantly improves. In particular, we prove a lower bound of $2n/(2n + 1)$ for differentiable, nondecreasing, and convex marginal cost functions. For polynomial cost functions with nonnegative coefficients we prove a tight bound of $3/4$.

Results for Average Cost Sharing. For differentiable, nondecreasing, and convex cost functions, we prove a lower bound of $1/n$ on the worst-case efficiency. If we further assume that the average cost functions are convex (e.g., polynomials with nonnegative coefficients), we present a tight bound of $4/(n + 3)$. For symmetric games this bound improves to $4n/(n + 1)^2$.

1.2. Significance and Techniques Used

Our main technical contribution is a general template to derive an upper bound on the efficiency loss of basic cost sharing methods in resource allocation games. This generality stems from two aspects: On one hand, the restriction to *basic* cost sharing methods requires only mild

assumptions on the feasible design space; see also the discussion in Moulin (2008). On the other hand, our template works for general resource allocation games, including the single resource case as in Moulin (2008) as well as multi-commodity network variants considered in Johari and Tsitsiklis (2006). We see this as a nontrivial generalization of previous works as, for instance, in Johari and Tsitsiklis (2006), the network structure is explicitly used to prove bounds on the price of anarchy (essentially through max-flow computations).

Our proof technique is quite simple and different from Johari and Tsitsiklis (2006) and Moulin (2008). In Johari and Tsitsiklis (2006), the authors consider marginal cost pricing and explicitly identify the worst possible game by analytically solving a sequence of quadratic optimization problems (assuming linear marginal cost functions). The resulting optimization problem explicitly involves the coefficients a and b of an affine marginal cost function $c(x) = ax + b$. Hence, this approach becomes increasingly technical if this optimization problem involves, e.g., polynomial cost functions of higher degree. For general convex marginal cost functions it is not clear whether the approach of Johari and Tsitsiklis (2006) gives an optimization problem that is structured enough to be solved.

Moulin (2008) derives lower bounds on the worst-case efficiency (using a different measure of efficiency) of three cost sharing methods: *average cost sharing*, *serial cost sharing*, and *incremental cost sharing*. His bounds are valid for resource allocation games with a single resource. Clearly, this assumption simplifies the subsequent analysis. From a technical point of view, Moulin proves an upper bound on the efficiency loss for each of the three cost sharing methods separately. Our approach gives a unified bound on the efficiency loss for an entire class of cost sharing methods, including those considered in Moulin (2008) and Johari and Tsitsiklis (2006).

Key to our approach is the use of variational inequalities, which allow us to relate the surplus of a Nash equilibrium to that of an optimal profile. Because variational inequalities do not rely on the specific combinatorial structure of the strategy spaces, this approach is applicable to general resource allocation games, which contain games with network structure as a special case. We note here that variational inequalities have been used before for bounding the efficiency loss of Nash equilibria; see Cominetti et al. (2009), Correa et al. (2004), Roughgarden (2002), and Yang et al. (2008).

1.3. Outline

The remainder of this paper is structured as follows. After reviewing the related work in §2, we introduce in §3 the fundamentals of a resource allocation game consisting of a congestion model and a cost sharing method. For the class of basic cost sharing methods, we develop in §4 a general lower bound on the worst-case efficiency of Nash equilibria that depends only on the used cost functions and cost

sharing methods but not on the player's private utilities. We use this general bound to show that the incremental cost sharing method is optimal. Because the incremental cost sharing method is basic but not scalable, we focus in the rest of the paper on two scalable cost sharing methods: marginal cost pricing and average cost sharing. In §5, we apply our general lower bound to marginal cost pricing, and we derive several lower and upper bounds on the worst-case efficiency of Nash equilibria depending on the used cost functions. In §6, we subsequently apply our generic bound to average cost sharing. We conclude the paper in §7 with a brief summary of our results and a discussion of open problems. Appendix A provides a table of notation. All missing proofs can be found in the e-companion to this paper. An electronic companion to this paper is available as part of the online version that can be found at <http://or.journal.informs.org/>.

2. Related Work

Network Resource Allocation Games. Kelly (1997) and Kelly et al. (1998) studied network resource allocation games and proposed a pricing mechanism termed proportionally fair pricing in which every resource charges a price per unit resource equal to *marginal cost*. Despite the simplicity and scalability of this mechanism, Kelly et al. showed that an optimal solution can be achieved as an equilibrium if players are *price takers*, that is, if they do not anticipate the consequence of price change in response to a change of their flow.

Johari et al. (2005) and Johari and Tsitsiklis (2004) studied network resource allocation games, where players submit a bid to each resource in the network and resources are allocated to the players according to Kelly's proportionally fair allocation mechanism. For this mechanism they established a bounded efficiency loss of the marginal pricing scheme with fixed and elastic resource capacities. However, the proposed mechanism is not scalable since each player has to submit an individual bid to each resource. If, instead, players can submit only a single bid per path, it was proved that the efficiency can be arbitrarily low for the case of hard capacities by Yang and Hajek (2006) and for the case of elastic capacities by Johari (2004).

Johari and Tsitsiklis (2005, 2006) studied network resource allocation games with marginal cost pricing. On the negative side, they showed that for nondifferentiable marginal cost functions, the price of anarchy is unbounded even for games with two players. For the special case of linear marginal cost functions, Johari and Tsitsiklis (2006) showed that the efficiency loss is bounded by $2/3$. Remarkably, this result holds for an arbitrary collection of concave utility functions and arbitrary networks. For a game with one resource and n players having equal utility functions, Johari and Tsitsiklis (2005) proved a bound of $2n/(2n+1)$ for convex marginal cost functions.

Chen and Zhang (2010) recently presented a class of pricing mechanisms for network resource allocation

games satisfying four axioms that are considered desirable. In particular, their mechanisms are characterized by the axioms *rescaling*, *additivity*, *positivity*, and *weak consistency*, which have been proposed by Samet and Tauman (1982). This family of price mechanisms includes marginal cost pricing, Aumann-Shapley pricing, and average cost pricing. The main objective of Chen and Zhang (2010) is to find among all mechanisms that satisfy the four axioms an *optimal mechanism*, i.e., one that minimizes the induced price of anarchy. Their main result states that for affine cost functions, the optimal mechanism is obtained by an affine transformation of marginal cost prices, and that marginal cost pricing itself is nearly optimal (achieving a slightly better efficiency guarantee (0.686) than the bound $(2/3)$).

Cost Sharing in Cournot Games. Cournot's oligopoly model clearly is one of the cornerstones of economic theory; see Mas-Colell et al. (1995), Owen (1982) for an overview of classical work in this area. Johari and Tsitsiklis (2005) proved that Cournot oligopoly games are basically equivalent (in terms of the worst-case efficiency of Nash equilibria) to resource allocation games with a single resource (which are termed Cournot oligopsonies in Johari and Tsitsiklis 2005). Moulin (2008) studied the price of anarchy for resource allocation games on a single resource with three different pricing mechanisms: average cost sharing, incremental cost sharing, and serial cost sharing. An important difference between our approach and that of Moulin is the definition of the efficiency loss of a cost-sharing method. The total surplus of a Nash equilibrium in Moulin (2008) is defined as the sum of the player's payoffs, which inevitably involve the cost shares collected. In our model (and that of Johari and Tsitsiklis 2005, 2006), we assume that the collected cost shares are internalized, so that we count only the player's utilities for using the resources minus the *actual* cost of using the resources. Only for exactly balanced cost sharing methods (such as average cost sharing) this difference vanishes as the actual costs and the collected cost shares coincide. In fact, it turns out that two of our results for average cost sharing (the bound $1/n$ in Theorem 6.1 and $4/(3+n)$ in Theorem 6.2) coincide with Moulin's bounds.

Guo and Yang (2005) studied Cournot oligopoly models and derived bounds on the price of anarchy for marginal cost pricing. In a Cournot oligopoly game, there is a set of players that each produce quantities so as to satisfy an elastic demand. The production cost for every player is modeled by a convex cost function, and the market price is modeled by a decreasing function in the total supplied quantity. The goal of every player is to maximize revenue. Guo and Yang (2005) derived, among other results, a lower bound of the worst-case efficiency of $4/(\sqrt{4n+5}+3)$ for concave marginal price functions. Using the equivalence result of Johari and Tsitsiklis (2005), this bound translates to the case of resource allocations games with a single resource, marginal cost pricing, and convex marginal cost

functions. We show in Theorem 5.1 that the same bound holds even for general resource allocation games.

Nonatomic Network Routing. In nonatomic network routing games, Roughgarden and Tardos (2002) showed that the price of anarchy for network routing games with nonatomic players and linear latency functions is $4/3$. The case of more general families of latency functions has been studied by Roughgarden (2002) and Correa et al. 2004. (For an overview of related results, we refer to the book by Roughgarden 2005 and the survey by Altman et al. 2006.) Despite these bounds for specific classes of latency functions, it is known that the price of anarchy in routing games with general latency functions is unbounded even on simple parallel-arc networks (Roughgarden and Tardos 2002). Chau and Sim (2003) studied the price of anarchy for nonatomic network games with elastic demands and general cost functions. They obtain bounds for the more general case of separable cost functions and elastic demands. The case of asymmetric cost functions has been studied by Perakis (2007). Yang et al. (2010) presented a detailed study on the worst-case efficiency loss of different variants of marginal cost pricing for the case of nonatomic users with fixed and elastic demands, respectively.

Atomic Splittable Network Routing. In atomic splittable network routing games there is a finite number of players who can split the flow along available paths; see Altman et al. (2002), Cominetti et al. (2009), Harks (2011), Haurie and Marcotte (1985), Hayrapetyan et al. (2006), Yang et al. (2008). Haurie and Marcotte presented a general framework for studying atomic splittable network games with elastic demands. Haurie and Marcotte, however, do not study the efficiency of Nash equilibria with respect to an optimal solution. Along similar lines as Haurie and Marcotte (1985), Harker (1988) considers games with atomic players and nonatomic players at the same time. Harker referred to the equilibria of those games as mixed behavior equilibria and gave a characterization of these equilibria by means of variational inequalities.

Hayrapetyan et al. (2006) studied congestion games with colluding players. Their goal is to investigate the price of collusion: the factor by which the quality of Nash equilibria can deteriorate when coalitions form. Altman et al. (2002) and Cominetti et al. (2009) studied the atomic splittable selfish routing model. Altman et al. bounded the price of anarchy for monomial latency functions (plus a constant). They also derived conditions under which a Nash equilibrium is unique. Uniqueness of Nash equilibria has been further studied by Bhaskar et al. (2009), Orda et al. (1993), and Yang and Zhang (2008). Cominetti et al. observed that the price of anarchy of the atomic splittable game might exceed that of the standard nonatomic selfish routing game. Based on the work of Catoni and Pallotino (1991), they presented an instance with affine latency functions where the price of anarchy is 1.34. For affine latencies, they presented an upper bound of 1.5 on the price of anarchy. In Harks (2011), a general upper bound on the price of anarchy is derived

that depends on the class of latency functions. This bound is tight, as shown in Roughgarden and Schoppmann (2011).

An important difference between our model and that of Hayrapetyan et al. (2006) and Cominetti et al. (2009) is that our model involves *elastic* demands that are varied by players. As a result, in our model the payoff of players is a linear combination of utility (derived from sending flow) and associated costs.

Tolls in Network Games. A large body of work in the area of transportation networks is concerned with congestion toll pricing; see, for example, Knight (1924), Beckmann et al. (1956), Smith (1979), and Hearn and Ramana (1998). This mechanism assigns tolls to certain arcs of the network that are charged to those users that decide to take routes through them. The toll mechanism has the desirable property that every user is charged a single price per unit resource.

Beckmann et al. (1956) showed that for the Wardrop model with homogeneous users, charging the difference between the marginal cost and the real cost in the socially optimal solution (marginal cost pricing) leads to an equilibrium flow that is optimal. Cole et al. (2003) considered the case of heterogeneous users, that is, users value latency relative to monetary cost differently. For single-commodity networks, the authors showed the existence of tolls that induce an optimal flow as Nash flow. Fleischer et al. (2004), Karakostas and Kolliopoulos (2004), and Yang and Huang (2004) proved that there are tolls inducing an optimal flow for heterogeneous users even in general networks. Swamy (2007) and Yang and Zhang (2008) proved the existence of optimal tolls for the atomic splittable model with fixed demands. Note that for computing the corresponding tolls, the works by Cole et al. (2003), Fleischer et al. (2004), Karakostas and Kolliopoulos (2004), Swamy (2007), and Yang and Huang (2004), and Yang and Zhang (2008) use a mathematical programming approach that requires central knowledge about the users, including their locations, private utility functions, and demands. In this sense, the toll mechanisms are not scalable because the underlying cost sharing method is a function of these private values.

Finally, Acemoglu and Ozdaglar (2007) and Ozdaglar (2008) study a model of parallel arc networks in which the arcs are owned by service providers that compete for the available traffic by setting prices. For this model they prove a tight worst-case bound for the efficiency loss of equilibria.

3. The Model

In this section, we introduce resource allocation games as natural generalizations of variants of congestion games. As the two building blocks of a resource allocation game, we first define a congestion model and then introduce the notion of a cost sharing method.

3.1. Congestion Model

DEFINITION 3.1 (CONGESTION MODEL). A tuple $\mathcal{M} = (N, R, \{X_i\}_{i \in N}, \{C_r\}_{r \in R})$ is called a *congestion model* if $N = \{1, \dots, n\}$ is a nonempty, finite set of players; $R = \{1, \dots, m\}$ is a nonempty, finite set of resources; and for each player $i \in N$, her collection of accessible sets $X_i = \{R_{i1}, \dots, R_{im_i}\}$, $m_i \in \mathbb{N}$, is a nonempty, finite set of subsets of R . We will use the shorthand notation $M_i = \{1, \dots, m_i\}$. Every resource $r \in R$ has a cost function $C_r: \mathbb{R}_+ \rightarrow \mathbb{R}_+$.

ASSUMPTION 3.1. *Cost functions $C_r: \mathbb{R}_+ \rightarrow \mathbb{R}_+$, $r \in R$, are differentiable, convex, nondecreasing functions, with $\lim_{x \rightarrow \infty} C_r(x)/x = \infty$.*

Given a congestion model $\mathcal{M} = (N, R, \{X_i\}_{i \in N}, \{C_r\}_{r \in R})$, we derive a corresponding *resource allocation model* $\mathcal{RM} = (N, R, \{X_i\}_{i \in N}, \Phi, \{C_r\}_{r \in R})$, where $\Phi = \times_{i \in N} \Phi_i$, and $\Phi_i = \mathbb{R}_+^{m_i}$ defines the strategy space for player i . A strategy profile $\varphi_i = (\varphi_{i1}, \dots, \varphi_{im_i})$ of player i can be interpreted as a distribution of non-negative *demands* over the elements in X_i . The total demand of player i is defined by $d_i(\varphi) = \sum_{j=1}^{m_i} \varphi_{ij}$. For $i \in N$, $\Phi_{-i} = \Phi_1 \times \dots \times \Phi_{i-1} \times \Phi_{i+1} \times \dots \times \Phi_n$ denotes the strategy space of all players except for player i . With a slight abuse of notation we will sometimes write a strategy profile as $\varphi = (\varphi_i, \varphi_{-i})$, meaning that $\varphi_i \in \Phi_i$ and $\varphi_{-i} \in \Phi_{-i}$. For a given profile φ , the load generated by player $i \in N$ on resource $r \in R$ is defined by $\varphi_i^r = \sum_{j \in M_i: r \in R_{ij}} \varphi_{ij}$. We denote by $\varphi^r = (\varphi_i^r, i \in N)$ the load vector of resource $r \in R$. The total load on resource $r \in R$ is defined by $l_r(\varphi) = \sum_{i=1}^n \varphi_i^r$. We will give two examples of a resource allocation model.

EXAMPLE 3.1 (NETWORK RESOURCE ALLOCATION). A resource allocation model $\mathcal{RM} = (N, R, \{X_i\}_{i \in N}, \Phi, \{C_r\}_{r \in R})$ is called a *network resource allocation model* if the set of resources correspond to the set of arcs of a directed or undirected graph G , every player i corresponds to a commodity having two distinguished vertices (s_i, t_i) (s_i is the source and t_i the terminal vertex in G , respectively), and the collection of player i 's accessible sets (X_i) is the set of corresponding (s_i, t_i) -paths. Thus, a strategy for player i corresponds to sending a nonnegative demand along the available (s_i, t_i) -paths.

EXAMPLE 3.2 (MATROID RESOURCE ALLOCATION). A resource allocation model $\mathcal{RM} = (N, R, \{X_i\}_{i \in N}, \Phi, \{C_r\}_{r \in R})$ is called *matroid resource allocation model* if for every $i \in N$, there is a matroid $M_i = (R, \mathcal{F}_i)$ (note that \mathcal{F}_i refers to an independence system in R ; see Schrijver 2003 for an introduction to matroids) such that X_i equals the set of bases of M_i . A prominent example of a *matroid resource allocation model* arises if the resources form a graph and the set of bases correspond to the set of spanning trees in G . In this case, a strategy for player i corresponds to sending a nonnegative demand along the available spanning trees of G .

3.2. Cost Sharing Methods

We define a cost sharing method as a collection of functions, one for each resource that takes as input the vector of the players' loads on the resource and outputs a vector of cost shares for each player. We restrict the set of feasible cost sharing methods as defined below.

DEFINITION 3.2. Given a resource allocation model $\mathcal{RM} = (N, R, \{X_i\}_{i \in N}, \Phi, \{C_r\}_{r \in R})$, a cost sharing method for a resource $r \in R$ is a mapping $\xi^r: \mathbb{R}_+^n \rightarrow \mathbb{R}_+^n$. We define the following conditions:

1. **Cost-covering:** $\sum_{i=1}^n \xi_i^r(\varphi^r) \geq C_r(l_r(\varphi))$ for all $\varphi \in \Phi$;

2. **Nash-inducing:** $\xi_i^r(\varphi^r)$ is nondecreasing, differentiable, and convex in φ_i^r for all $i \in N$;

3. **No charge for zero demand:** $\xi_i^r(\varphi^r) = 0$ for all $\varphi \in \Phi$ with $\varphi_i^r = 0$, for all $i \in N$;

4. **Scalability:** $\xi_i^r(\varphi^r)/\varphi_i^r = (\xi_j^r(\varphi^r))/\varphi_j^r$ for all $i, j \in N$, and all $\varphi \in \Phi$ with $\varphi_i^r, \varphi_j^r > 0$.

A cost sharing method is called *basic* if it satisfies Assumptions 1–3, and it is called *scalable* if it satisfies the Assumptions 1–4. Note that a basic cost sharing method is automatically *separable* in the sense of condition (1) in §1 because every ξ^r has only φ^r as argument.

We next discuss the above assumptions in detail. The first assumption is standard in the economics literature and the least critical: the cost of using a resource is passed to its users. The second assumption ensures the existence of a pure Nash equilibrium of the induced resource allocation game. Moreover, a positive charge for zero resource consumption prevents users from participation and is thus considered undesirable; see Moulin (2008). Assumption 4 stating that the price per unit resource consumption must be equal for all players is perhaps the most restrictive and controversial one. In the context of large-scale networks (e.g., the TCP/IP protocol suite used in the Internet), this property is considered desirable and indispensable because every resource needs only to pass a one-dimensional information to its users. For a detailed discussion on this subject, we refer the reader to Johari and Tsitsiklis (2009), Kelly et al. (1998), and Srikant (2003). We give in the following three examples of cost sharing methods that we will analyze throughout this paper.

EXAMPLE 3.3 (VERAGE COST SHARING). In average cost sharing, the cost share for player i on resource r under profile φ is defined as $\xi_i^r(\varphi^r) = \varphi_i^r \cdot C_r(l_r(\varphi))/l_r(\varphi)$. This cost sharing method is frequently used in the transportation literature (cf. Beckmann et al. 1956, Haurie and Marcotte 1985) for modeling the experienced travel time, where the term $c_r(l_r(\varphi)) := C_r(l_r(\varphi))/l_r(\varphi)$ models the load-dependent latency function on r . Note that average cost sharing is a scalable cost sharing method.

PROPOSITION 3.1. *The only cost sharing method that is exactly budget balanced and fulfills Assumption 4 (single price per unit) in Definition 3.2 is average cost sharing.*

EXAMPLE 3.4 (MARGINAL COST PRICING). In marginal cost pricing, the cost share for player i on resource r under profile φ is defined as $\xi_i^r(\varphi^r) = \varphi_i^r \cdot C_r'(l_r(\varphi))$. Note that marginal cost pricing is a scalable cost sharing method.

EXAMPLE 3.5 (INCREMENTAL COST SHARING). In incremental cost sharing, the cost share for player i on resource r under profile φ is defined as $\xi_i^r(\varphi^r) = C_r(l_r(\varphi)) - C_r(l_r(0, \varphi_{-i}))$. One can easily show that incremental cost sharing is not scalable.

REMARK 3.1. While the incremental cost sharing method is not scalable, it still satisfies the symmetry condition: $\xi_i^r(\varphi^r) = \xi_j^r(\varphi^r)$ for all $i, j \in N$ and $\varphi \in \Phi$ with $\varphi_i^r = \varphi_j^r$. The above property is considered desirable in the economics literature and refers to the notion of fairness between resource consumers: if two players have an equal resource consumption, their cost share must be equal.

3.3. Resource Allocation Games

We are now ready to formally define a resource allocation game. By choosing a strategy φ_i , player i receives a certain benefit measured by a utility function $U_i(d_i(\varphi))$. We assume that utility functions satisfy the following conditions.

ASSUMPTION 3.2. *Each utility function $U_i: \mathbb{R}_+ \rightarrow \mathbb{R}_+$ is differentiable, strictly increasing, and concave.*

DEFINITION 3.3 (RESOURCE ALLOCATION GAME). Given a resource allocation model $\mathcal{RM} = (N, R, \{X_i\}_{i \in N}, \Phi, \{C_r\}_{r \in R})$, the corresponding *resource allocation game* is the strategic game $G(\mathcal{RM}) = (N, \Phi, \pi)$, where the payoff $\pi = (\pi_1, \dots, \pi_n)$ is defined as $\pi_i(\varphi) := U_i(d_i(\varphi)) - \sum_{r \in R} \xi_i^r(\varphi^r)$, where $\xi_i^r: \mathbb{R}_+^n \rightarrow \mathbb{R}_+^n$ is the cost sharing method of resource $r \in R$.

For the remainder of this paper, we will write G instead of $G(\mathcal{RM})$.

REMARK 3.2. Assumptions 3.1, 3.2, and Definition 3.2 imply $\lim_{\|\varphi_i\| \rightarrow \infty} \pi_i(\varphi_i; \varphi_{-i}) = -\infty$, hence we can effectively restrict the strategy space for every player to a compact set. Because the payoff functions are concave, a pure Nash equilibrium exists; see the result of Rosen (1965).

The *total surplus* of a profile φ is defined as $\mathcal{U}(\varphi) := \sum_{i=1}^n U_i(d_i(\varphi)) - C(\varphi)$, where $C(\varphi) = \sum_{r \in R} C_r(l_r(\varphi))$ is the total cost function for the profile φ . A profile of maximum total surplus is called *optimal*. We define the following functions: $\hat{\xi}_i^r(\varphi^r) := \partial \xi_i^r(\varphi^r)/\partial \varphi_i^r$ and $\hat{\xi}_{ij}^r(\varphi) := \sum_{r \in R_{ij}} \hat{\xi}_i^r(\varphi^r)$. The next lemma establishes necessary and sufficient conditions for a profile to be optimal and a Nash equilibrium, respectively.

LEMMA 3.1. *Consider a resource allocation game G with basic cost sharing methods $(\xi_r, r \in R)$. The profiles ϑ and*

ψ are a Nash equilibrium and an optimal profile, respectively, if and only if for all players i the following conditions hold:

$$\nabla \pi_i(\vartheta_i; \vartheta_{-i}) \cdot (\varphi_i - \vartheta_i) \leq 0, \quad \text{for all } \varphi_i \in \Phi_i, \quad (1)$$

$$\begin{aligned} U'_i(d_i(\vartheta)) &= \hat{\xi}_{ij}(\vartheta), \quad \text{for all } j \in M_i \text{ with } \vartheta_{ij} > 0, \\ U'_i(d_i(\vartheta)) &\leq \hat{\xi}_{ij}(\vartheta), \quad \text{for all } j \in M_i \text{ with } \vartheta_{ij} = 0, \end{aligned} \quad (2)$$

$$\begin{aligned} U'_i(d_i(\psi)) &= \sum_{r \in R_{ij}} C'_r(l_r(\psi)), \\ &\quad \text{for all } j \in M_i \text{ with } \psi_{ij} > 0, \\ U'_i(d_i(\psi)) &\leq \sum_{r \in R_{ij}} C'_r(l_r(\psi)), \\ &\quad \text{for all } j \in M_i \text{ with } \psi_{ij} = 0. \end{aligned} \quad (3)$$

In the following sections we analyze the worst-case efficiency of Nash equilibria for several cost sharing methods. In §4 we first develop a general lower bound for basic cost sharing methods. We proceed by studying marginal cost pricing and average cost sharing in §§5 and 6, respectively.

4. Worst-Case Efficiency of Basic Cost Sharing Methods

In the following, we will study the worst-case efficiency loss for a class of basic cost sharing methods. Throughout the analysis we assume that cost functions satisfy Assumption 3.1 and utility functions satisfy Assumption 3.2. Before we give a formal definition of the worst-case efficiency loss, we prove an auxiliary lemma, showing that basic cost sharing methods always guarantee a nonnegative total surplus for every Nash equilibrium.

LEMMA 4.1. *Let G be a resource allocation game with n players, cost functions in \mathcal{C} , and basic cost sharing methods $\xi_r \in \mathcal{D}_n$ for all $r \in R$. Let Θ_G be the set of Nash equilibria. Then, $\mathcal{U}(\vartheta) \geq 0$ for all $\vartheta \in \Theta_G$.*

Next, we provide a formal definition of the worst-case efficiency loss.

DEFINITION 4.1. Let \mathcal{C} be a class of cost functions. Let $\mathcal{G}_n(\mathcal{C}, \mathcal{D}_n)$ be the set of all resource allocation games with n players, cost functions in \mathcal{C} , and basic cost sharing methods $\xi_r \in \mathcal{D}_n$ for all $r \in R$. For $G \in \mathcal{G}_n(\mathcal{C}, \mathcal{D}_n)$, let ψ_G be an optimal profile and let Θ_G be the set of pure Nash equilibria, respectively. Then the *worst-case efficiency* is defined by

$$\rho_n(\mathcal{C}, \mathcal{D}_n) = \begin{cases} \inf_{G \in \mathcal{G}_n(\mathcal{C}, \mathcal{D}_n)} \inf_{\vartheta \in \Theta_G} \frac{\mathcal{U}_G(\vartheta)}{\mathcal{U}_G(\psi_G)}, & \text{if } \mathcal{U}_G(\psi_G) > 0, \\ 1, & \text{otherwise.} \end{cases}$$

Here, \mathcal{U}_G denotes the total surplus function for game G . Conversely, $1 - \rho_n(\mathcal{C}, \mathcal{D}_n)$ defines the *worst-case efficiency loss* or *price of anarchy*.

REMARK 4.1. Note that by Lemma 4.1 the case $\mathcal{U}_G(\psi_G) = 0$ implies that for basic cost sharing methods, every Nash equilibrium is optimal. Therefore, we can assume without loss of generality that every optimal profile recovers a strictly positive total surplus, i.e., $\mathcal{U}_G(\psi_G) > 0$.

We show next that for bounding the worst-case efficiency of basic cost sharing methods it is sufficient to consider games with only linear utility functions. The next lemma can be proved using ideas of Johari and Tsitsiklis (2006), Moulin (2008), and Chen and Zhang (2010).

LEMMA 4.2. *Let $\mathcal{G}_n(\mathcal{C}, \mathcal{D}_n)$ be the set of all resource allocation games with n players, cost functions in \mathcal{C} , and basic cost sharing methods $\xi_r \in \mathcal{D}_n$ for all $r \in R$. Then, for bounding the worst-case efficiency it is enough to consider resource allocation games in which all utility functions are linear.*

We proceed by calculating the surplus of an optimal solution and that of a Nash equilibrium in terms of the cost functions and cost sharing methods involved, respectively.

LEMMA 4.3. *Consider a game G with basic cost sharing methods and linear utility functions, that is, $U_i(x) = u_i \cdot x$, $u_i \geq 0$, $i \in N$. Let ψ be an optimal profile and ϑ be a Nash equilibrium. Then ψ and ϑ generate a total surplus of $\mathcal{U}(\psi) = \sum_{r \in R} (l_r(\psi) \cdot C'_r(l_r(\psi)) - C_r(l_r(\psi)))$ and $\mathcal{U}(\vartheta) = \sum_{i=1}^n \sum_{j=1}^{m_i} \xi_{ij}(\vartheta) \vartheta_{ij} - C(\vartheta)$.*

We provide in this section a general proof template that enables us to derive a bound on the worst-case efficiency for a resource allocation game with basic cost sharing methods. The main idea for proving such bounds is an application of the variational inequality. Let ψ and ϑ be an optimal and a Nash profile, respectively. Observe that for any λ , the following inequality holds: $\mathcal{U}(\psi) \leq \lambda \mathcal{U}(\vartheta) + \mathcal{U}(\psi) + \sum_{i=1}^n \nabla \pi_i(\vartheta_i; \vartheta_{-i}) \cdot (\vartheta_i - \psi_i) - \lambda \mathcal{U}(\vartheta)$. If we can derive an inequality of the form $(\mathcal{U}(\psi) + \sum_{i=1}^n \nabla \pi_i(\vartheta_i; \vartheta_{-i}) \cdot (\vartheta_i - \psi_i) - \lambda \mathcal{U}(\vartheta)) / \mathcal{U}(\psi) \leq \omega(\lambda)$ for some $\omega(\lambda) < 1$, we would obtain the inequality $\mathcal{U}(\psi) \leq \lambda \mathcal{U}(\vartheta) + \omega(\lambda) \mathcal{U}(\psi)$, which yields a bound on the worst-case efficiency of $(1 - \omega(\lambda)) / \lambda$. As a consequence, we could then optimize over λ (which of course involves $\omega(\lambda)$) so as to derive the best possible bound. This technique (λ -technique) has been previously applied to bound the price of anarchy in atomic splittable congestion games; see Harks (2011).

In the following, we denote by \mathcal{D}_n a class of basic cost sharing methods for n players. For a cost function C , a cost sharing method $\xi \in \mathcal{D}_n$, and a parameter $\lambda > 0$, we define the following value:

$$\begin{aligned} \omega_n(C, \xi, \lambda) &:= \sup_{x, y \in \mathbb{R}_+^n} \frac{\sum_{i=1}^n \hat{\xi}_i(x)(y_i - \lambda x_i) + \lambda C(l(x)) - C(l(y))}{C'(l(y)) \cdot l(y) - C(l(y))}, \quad (4) \end{aligned}$$

where $l(x) = \sum_{i=1}^n x_i$. For a class of cost functions \mathcal{C} and a class of basic cost sharing methods \mathcal{D}_n , we define

$\omega_n(\mathcal{C}, \mathcal{D}_n, \lambda) := \sup_{\xi \in \mathcal{D}_n} \sup_{C \in \mathcal{C}} \omega_n(C, \xi, \lambda)$. We define the feasible λ -region as $\Lambda(\mathcal{C}, \mathcal{D}_n) := \{\lambda > 0 \mid \omega_n(\mathcal{C}, \mathcal{D}_n, \lambda) < 1\}$.

THEOREM 4.1. *Consider the set $\mathcal{G}_n(\mathcal{C}, \mathcal{D}_n)$ of resource allocation games with basic cost sharing methods $\xi^r \in \mathcal{D}_n$, $r \in R$, and cost functions in \mathcal{C} . Then the worst-case efficiency is at least*

$$\rho(\mathcal{C}, \mathcal{D}_n) \geq \sup_{\lambda \in \Lambda(\mathcal{C}, \mathcal{D}_n)} \left[\frac{1 - \omega_n(\mathcal{C}, \mathcal{D}_n, \lambda)}{\lambda} \right].$$

PROOF. Let $G \in \mathcal{G}_n(\mathcal{C}, \mathcal{D}_n)$. Using Lemma 4.2 we may assume that utility functions are linear. Let ψ and ϑ be an optimal and a Nash profile, respectively. Observe that for any λ , the following inequalities hold:

$$\begin{aligned} \mathcal{U}(\psi) &\leq \lambda \mathcal{U}(\vartheta) + \mathcal{U}(\psi) + \sum_{i=1}^n \nabla \pi_i(\vartheta_i; \vartheta_{-i}) \\ &\quad \cdot (\vartheta_i - \psi_i) - \lambda \mathcal{U}(\vartheta) \end{aligned} \quad (5)$$

$$\begin{aligned} &= \lambda \mathcal{U}(\vartheta) - C(\psi) + \sum_{i=1}^n u_i \cdot d_i(\vartheta) \\ &\quad + \sum_{r \in R} \sum_{i=1}^n \hat{\xi}_i^r(\vartheta^r) (\psi_i^r - \vartheta_i^r) - \lambda \mathcal{U}(\vartheta) \\ &= \lambda \mathcal{U}(\vartheta) - C(\psi) + \sum_{r \in R} \sum_{i=1}^n \hat{\xi}_i^r(\vartheta^r) \psi_i^r - \lambda \mathcal{U}(\vartheta) \end{aligned} \quad (6)$$

$$\begin{aligned} &= \lambda \mathcal{U}(\vartheta) - C(\psi) \\ &\quad + \sum_{r \in R} \sum_{i=1}^n \hat{\xi}_i^r(\vartheta^r) (\psi_i^r - \lambda \vartheta_i^r) + \lambda C(\vartheta). \end{aligned} \quad (7)$$

Here, (5) and (6) follow from Lemma 3.1, while (7) follows from Lemma 4.3. To complete the proof, we need to show that

$$\begin{aligned} \sum_{r \in R} \sum_{i=1}^n \hat{\xi}_i^r(\vartheta^r) (\psi_i^r - \lambda \vartheta_i^r) + \lambda C(\vartheta) - C(\psi) \\ \leq \omega_n(\mathcal{C}, \mathcal{D}_n, \lambda) \cdot \mathcal{U}(\psi). \end{aligned} \quad (8)$$

By definition of $\omega_n(\mathcal{C}, \mathcal{D}_n, \lambda)$, we have

$$\begin{aligned} \frac{\sum_{i=1}^n \hat{\xi}_i^r(\vartheta^r) (\psi_i^r - \lambda \vartheta_i^r) + \lambda C_r(l_r(\vartheta)) - C_r(l_r(\psi))}{l_r(\psi) \cdot C_r(l_r(\psi)) - C_r(l_r(\vartheta))} \\ \leq \omega_n(\mathcal{C}, \mathcal{D}_n, \lambda) \end{aligned}$$

for all $r \in R$. Multiplying this inequality by $l_r(\psi) \cdot C_r(l_r(\psi)) - C_r(l_r(\vartheta))$, summing up over $r \in R$, and using Lemma 4.3, we obtain (8). \square

We briefly pause here to discuss implications of the above result. Theorem 4.1 provides a lower bound on the worst-case efficiency of Nash equilibria that depends only on \mathcal{C} and \mathcal{D}_n but neither on the player's private utilities nor on the strategy space. If the sets \mathcal{C} and \mathcal{D}_n have a specific form (e.g., convex cost functions and marginal

cost pricing), then evaluating the concrete bound in Theorem 4.1 amounts to solving a highly structured optimization problem. In the remainder of the paper we will actually solve this optimization problem for three specific cost sharing methods (incremental cost sharing, marginal cost pricing, and average cost sharing) and different classes of cost functions. We will first apply Theorem 4.1 to prove that the incremental cost sharing method is actually an optimal mechanism among all basic mechanisms.

PROPOSITION 4.1. *For incremental cost sharing, every Nash equilibrium is optimal.*

Moulin (2008) showed that incremental cost sharing is optimal for resource allocation games with a single resource. Proposition 4.1 generalizes Moulin's result to hold for general resource allocation games.

5. The Worst-Case Efficiency of Marginal Cost Pricing

In the previous section, we showed that among all basic cost sharing mechanisms, there is an optimal mechanism (incremental cost sharing) that achieves full efficiency. Because the incremental cost sharing method is not scalable, we will focus in this section on *marginal cost pricing* which is a well-known scalable cost sharing method. More precisely, we will study the price of anarchy in games where all resources use marginal cost pricing as cost sharing method. We thus have $\xi_i^r(\phi^r) = \phi_i^r \cdot C_r'(l_r(\phi))$ for all $r \in R$, $i \in N$, and $\phi \in \Phi$. We will call $C_r'(\cdot)$ the *marginal cost function* of resource $r \in R$. For the rest of this section, we assume that all cost functions $C_r(\cdot)$ are twice differentiable for all $r \in R$. Instead of $\mathcal{G}_n(\mathcal{C}, \mathcal{D}_n)$, we will use the shorthand $\mathcal{G}_n(\mathcal{C})$ assuming that \mathcal{D}_n corresponds to marginal cost pricing. In Lemma 4.3, we represented the total surplus of a Nash equilibrium and that of an optimal profile in terms of the involved cost functions for a general cost sharing method. The following lemma is a special case of this result for marginal cost pricing.

LEMMA 5.1. *Consider a game G with marginal cost pricing and linear utility functions, that is, $U_i(x) = u_i \cdot x$, $u_i \geq 0$, $i \in N$. Let ϑ be a Nash equilibrium. Then, ϑ generates total surplus of*

$$\begin{aligned} \mathcal{U}(\vartheta) &= \sum_{r \in R} \left(l_r(\vartheta) \cdot C_r'(l_r(\vartheta)) \right. \\ &\quad \left. + \sum_{i=1}^n (\vartheta_i^r)^2 \cdot C_r''(l_r(\vartheta)) - C_r(l_r(\vartheta)) \right). \end{aligned} \quad (9)$$

The proof follows from Lemma 4.3. We proceed by deriving an upper bound for $\omega_n(C, \xi, \lambda)$ using that ξ corresponds to marginal cost pricing.

LEMMA 5.2. Let ξ be a marginal cost pricing method for n players. Then, $\omega_n(C, \xi, \lambda) \leq \omega_n^{\text{mcp}}(C, \lambda)$, where

$$\begin{aligned} \omega_n^{\text{mcp}}(C, \lambda) &:= \sup_{\substack{x, y \in \mathbb{R}_+ \\ \mu \in \{0\} \cup [1/n, 1]}} (C'(x)y + C''(x)\mu xy - \lambda(C'(x)x \\ &+ (\mu^2 + (1 - \mu)^2/(n - 1))C''(x)x^2) + \lambda C(x) - C(y)) \\ &\cdot (C'(y) \cdot y - C(y))^{-1}. \end{aligned} \quad (10)$$

An essential element of the definition of $\omega_n^{\text{mcp}}(C, \lambda)$ is the parameter μ defined as the largest ratio of the load of a single player and the overall load on a resource. We note that this ratio has been used before in the context of bounding the price of anarchy in atomic splittable network games with fixed demands; see Cominetti et al. (2009), Harks (2011), and Yang et al. (2008).

5.1. Cost Functions with a Convex Derivative

We start by applying Lemma 5.2 to convex marginal cost functions, that is, we consider cost functions with a convex derivative.

THEOREM 5.1. Let $\mathcal{E}^{\text{convD}}$ be a class of cost function that have a convex derivative. Consider the set $\mathcal{G}_n(\mathcal{E}^{\text{convD}})$ of games with at most n players. Then, $\rho_n(\mathcal{E}^{\text{convD}}) \geq 4/(3 + \sqrt{5 + 4n})$.

PROOF. We define $\lambda = (3 + \sqrt{5 + 4n})/4$ and prove the claim by showing that $\omega_n^{\text{mcp}}(C; \lambda) \leq 0$ for all $C \in \mathcal{E}^{\text{convD}}$. We bound the nominator of (10) by a case distinction. First, we assume $x \geq y$. We get

$$\begin{aligned} &C'(x)y + C''(x)\mu xy - \lambda(C'(x)x \\ &+ (\mu^2 + \frac{(1 - \mu)^2}{n - 1})C''(x)x^2) + \lambda C(x) - C(y) \\ &\leq C''(x) \left(\mu xy - \lambda \left(\mu^2 + \frac{(1 - \mu)^2}{n - 1} \right) x^2 \right) \\ &\leq C''(x)x^2 \left(\mu - \lambda \left(\mu^2 + \frac{(1 - \mu)^2}{n - 1} \right) \right). \end{aligned}$$


For the first inequality, we used that $C'(x)y - \lambda C'(x)x + \lambda C(x) - C(y) \leq 0$, because $y \leq x$, $\lambda \geq 1$, and $C'(\cdot)$ is convex. The second inequality follows from $y \leq x$ and $C''(x) \geq 0$ (since $C(\cdot)$ convex). Then, $\lambda = (3 + \sqrt{5 + 4n})/4$ yields $\omega_n^{\text{mcp}}(C; \lambda) \leq 0$, because

$$\max_{\mu \in \{0\} \cup [1/n, 1]} \mu - \left(\frac{3 + \sqrt{5 + 4n}}{4} \right) \left(\mu^2 + \frac{(1 - \mu)^2}{n - 1} \right) \leq 0.$$

Now we consider the case $x < y$. We define $\beta := x/y \in [0, 1)$. Observe that $C(y) - \lambda C(\beta y) = C(y) - C(\beta y) - (\lambda - 1)C(\beta y)$. Then we use the following inequality, which is illustrated in Figure 1.

$$C(y) - C(\beta y) \geq (y - \beta y)C'(\beta y) + \frac{(y - \beta y)^2}{2}C''(\beta y). \quad (11)$$

Figure 1. Illustration of the inequality (11) in the proof of Theorem 5.1.



Note. The shaded area illustrates the term $C(y) - C(\beta y) = \Delta_1 + \Delta_2$. The linear approximation $L_{\beta y}(\cdot)$ of the convex function $C'(\cdot)$ bounds $C'(z)$ from below, i.e., $L_{\beta y}(z) \leq C'(z)$. Then we have $\Delta_1 = (y - \beta y)C'(\beta y)$ and $\Delta_2 \geq ((y - \beta y)^2/2)C''(\beta y)$.

Together with $(\lambda - 1)C(\beta y) \leq (\lambda - 1)C'(\beta y)\beta y$, we obtain

$$\begin{aligned} \omega_n^{\text{mcp}}(C; \lambda) &\leq \sup_{\substack{\beta \in [0, 1), y \in \mathbb{R}_+ \\ \mu \in \{0\} \cup [1/n, 1]}} (C''(\beta y)y^2(\beta\mu - \lambda(\mu^2 + (1 - \mu)^2/(n - 1))) \\ &\cdot \beta^2 - ((1 - \beta)^2/2)) \cdot (C'(y)y - C(y))^{-1}. \end{aligned}$$

We then use $\max_{\beta \in [0, 1), \mu \in [0, 1]} (\beta\mu - \lambda(\mu^2 + (1 - \mu)^2/(n - 1))\beta^2 - (1 - \beta)^2/2) \leq (n - 4\lambda^2 + 6\lambda - 1)/4\lambda(n + 2\lambda)$, where $\beta^* = (n + 1)/(n + 2\lambda)$ and $\mu^* = (n - 1 + 4\lambda)/2\lambda(n + 1)$ are the unique maximizer. The value of λ solves $n - 4\lambda^2 + 6\lambda - 1 = 0$, thus, we obtain $\omega_n^{\text{mcp}}(C; \lambda) \leq 0$. Applying Lemma 5.2 for both cases proves the claim. \square

REMARK 5.1. The bound of Theorem 5.1 has been established before by Guo and Yang (2005) for the case of a single resource.

The above result gives a bound on the efficiency loss for differentiable and convex marginal cost functions scaling with the number of players. This result complements a negative result of Johari and Tsitsiklis (2005) for two-player games with *nondifferentiable* convex marginal cost functions, where the efficiency loss might be arbitrarily high. For nondifferentiable marginal cost functions, a Nash equilibrium can be characterized by optimality conditions expressed by the left and right directional derivatives of the marginal cost function. The key ingredient of the instance in Johari and Tsitsiklis (2005) is to increase the difference between two such values (for a point of non-differentiability) giving rise to a Nash equilibrium with low total surplus. In contrast, if marginal cost functions are differentiable, then by Lemma 5.1 the total surplus of an arbitrary Nash equilibrium can be expressed in terms of the involved cost functions and their *well-defined*

derivatives ruling out the instance constructed in Johari and Tsitsiklis (2005). Note that in many applications the considered marginal cost functions are differentiable, e.g., polynomial delay functions considered in transportation networks (Branston 1976) and $M/M/1$ functions modeling queuing delays in telecommunication networks (Srikant 2003).

REMARK 5.2. In the next section (Proposition 5.1), we present an asymptotically matching upper bound for the efficiency loss of $\rho_n(\mathcal{C}^{\text{convD}}) \leq (2(n - \sqrt{n}))/\sqrt{n}(n - 1)$, which grows as $O(1/\sqrt{n})$.

5.2. Polynomial Cost Functions

In practice, the most frequently used functions modeling delay are polynomials whose degrees and coefficients are determined from real-world data through statistical evaluation methods; see Patriksson (1994) and Branston (1976). Thus, we will explicitly calculate the price of anarchy for the class $\mathcal{C}_d := \{C(z) = \sum_{j=0}^d a_j z^j, a_j \geq 0\}$, $d \in \{2, 3, \dots\}$. Note that we have to demand $d \geq 2$ because otherwise Assumption 3.1 would be violated and a Nash equilibrium might not exist.

To simplify the analysis, we focus on the general case $n \in \mathbb{N}^* \cup \{\infty\}$. Let us define $\omega_\infty^{\text{mcp}}(C; \lambda) := \lim_{n \rightarrow \infty} \omega_n^{\text{mcp}}(C; \lambda)$. It is easy to see that $\omega_\infty^{\text{mcp}}(C; \lambda) \geq \omega_n^{\text{mcp}}(C; \lambda)$ for any $n \in \mathbb{N}^*$, implying $\rho_\infty(C) \leq \rho_n(C)$.

REMARK 5.3. We observe that for marginal cost pricing, the payoff functions $\pi_i(\cdot)$ are affine linear in each of the cost functions $C_r(\cdot)$. We can reduce the analysis to monomial cost functions subdividing each resource r into $d + 1$ resources r_0, \dots, r_d with monomial cost functions $C_{r_s}(l_r(\varphi)) = a_{r_s} \cdot (l_r(\varphi))^s$ for $s \in \{0, 1, \dots, d\}$. By extending the accessible sets of every player accordingly, we obtain a transformed game in which the set of Nash equilibria, optimal profiles, and corresponding surplus values coincide.

We present in the next lemma an upper bound for the value $\omega_\infty^{\text{mcp}}(\mathcal{M}_d; \lambda)$.

LEMMA 5.3. Consider the class $\mathcal{M}_d := \{C(z) = a_d z^d, a_d \geq 0, d \in \{2, 3, \dots\}\}$. Then, it holds that

$$\omega_\infty^{\text{mcp}}(\mathcal{M}_d; \lambda) \leq \left(\frac{1 + \mu(d-1)}{\lambda(1 + \mu^2 d)} \right)^{d-1} \left(\frac{d}{d-1} + \mu - 1 \right) - \frac{1}{d-1}, \quad \text{where } \mu(d) = \frac{1}{\sqrt{d-1} + 1}.$$

Given the above upper bound on $\omega_\infty^{\text{mcp}}(\mathcal{M}_d; \lambda)$ we now give a precise bound for $\rho(\mathcal{C}_d)$.

THEOREM 5.2. Let \mathcal{C}_d be the class of polynomial cost functions with nonnegative coefficients and maximum degree $d \in \{2, 3, \dots\}$. Then $\rho(\mathcal{C}_d) = (1 + \mu(d)^2 d)/(1 + \mu(d) \cdot (d-1))^{1+1/(d-1)}$, where $\mu(d) = 1/(\sqrt{d-1} + 1)$.

PROOF. We define $\lambda = ((1 + \mu(d)(d-1))^{1+1/(d-1)})/(1 + \mu(d)^2 d)$. Then Lemma 5.3 implies $\omega_\infty(\mathcal{M}_d; \lambda) \leq 0$ for all $d' < d$ and $\omega_\infty(\mathcal{M}_d; \lambda) = 0$. Thus, using Lemma 5.2 and Theorem 4.1, we have $\rho(\mathcal{C}_d) \geq (1 + \mu(d)^2 d)/(1 + \mu(d)(d-1))^{1+1/(d-1)}$.

Now we prove the upper bound. Consider a game with one resource having the cost function $C(x) = (1/d)x^d$ for some $d \in \{2, 3, \dots\}$. Assume we have n players, where player 1 has the utility function $U_1(\varphi_1) = \varphi_1$ while the remaining $n-1$ players have utility functions $U_k(\varphi_k) = b\varphi_k$ for some $b \in [0, 1]$ specified later. Consider a Nash equilibrium $\vartheta(n)$ in this game. Without loss of generality, we can assume $l(\vartheta(n)) > 0$. Using Lemma 3.1, we obtain $\vartheta_1(n) = (1 - l(\vartheta(n))^{d-1})/((d-1)l(\vartheta(n))^{d-1})$ for player 1 and $\vartheta_k(n) = (b - l(\vartheta(n))^{d-1})/((d-1)l(\vartheta(n))^{d-1})$ for players $k = 2, \dots, n$. Summing all demands, we get

$$l(\vartheta(n)) = \frac{1 - l(\vartheta(n))^{d-1}}{(d-1)l(\vartheta(n))^{d-1}} + (n-1) \frac{b - l(\vartheta(n))^{d-1}}{(d-1)l(\vartheta(n))^{d-1}} \\ \Leftrightarrow l(\vartheta(n)) = \left(\frac{1 + b(n-1)}{d-1+n} \right)^{1/(d-1)}.$$

In the limit, we get

$$\lim_{n \rightarrow \infty} l(\vartheta(n)) = b^{1/(d-1)}, \quad \lim_{n \rightarrow \infty} \vartheta_1(n) = \frac{b^{1/(d-1)}(1-b)}{b(d-1)},$$

and

$$\lim_{n \rightarrow \infty} b(n-1)\vartheta_k(n) = \frac{b^{1/(d-1)}(bd-1)}{d-1}.$$

Thus, we get as limit for the total surplus of the Nash equilibrium $\vartheta(n)$

$$\lim_{n \rightarrow \infty} \mathcal{U}(\vartheta(n)) = \frac{b^{1/(d-1)}(1-b)}{b(d-1)} + \frac{b^{1/(d-1)}(bd-1)}{d-1} - \frac{b^{1/(d-1)}b}{d}.$$

An optimal solution is given by $\psi = (1, 0, \dots, 0)$ with total surplus of $\mathcal{U}(\psi) = 1 - 1/d$. Now choosing $b = (1 + (d-1)^{3/2})/(d^2 - d + 2)$ the ratio $\mathcal{U}(\vartheta)/\mathcal{U}(\psi)$ coincides with the lower bound of the theorem. \square

REMARK 5.4. The worst-case efficiency for cost functions in \mathcal{C}_d is asymptotically bounded from below by $\Omega(1/\sqrt{d-1})$.

Note that the example used in the previous proof can also be used to construct an upper bound for $\rho_n(\mathcal{C}^{\text{convD}})$ complementing Theorem 5.1.

PROPOSITION 5.1. Let $\mathcal{C}^{\text{convD}}$ be a class of cost functions with a convex derivative. Consider the set $\mathcal{G}_n(\mathcal{C}^{\text{convD}})$ of games with at most $n \in \mathbb{N}^*$ players. Then $\rho_n(\mathcal{C}^{\text{convD}}) \leq 2(n - \sqrt{n})/\sqrt{n}(n - 1)$, which grows as $O(1/\sqrt{n})$.

5.3. Symmetric Games

In this section, we consider *symmetric games* in which all players have the same utility function $U_i = U_j$ for all $i, j \in N$ and the same strategy space, that is, $\Phi_i = \Phi_j$ for all $i, j \in N$. Symmetric resource allocation games have been considered before in the context of single-commodity network games with atomic players, unit demands, and splittable flows, see Cominetti et al. (2009) and Altman et al. (2002). Another example of a symmetric resource allocation game arises in *scheduling games* in which there are m machines used by n players having the same utility function. The strategy of every player is simply a distribution of her workload over the machines. We prove that the symmetry assumption implies improved bounds on the worst-case efficiency loss. Consider a symmetric game with n players. Then, there exists a symmetric optimal profile ψ in the sense that $\psi_i^r = (1/n)\psi^r$ for all $i \in N$. We obtain the following bound for $\omega_n(C, \xi, \lambda)$.

LEMMA 5.4. *Consider a symmetric game with ξ being marginal cost pricing. Then it holds that $\omega_n(C, \xi, \lambda) \leq \omega_n^{\text{mcp, sym}}(C, \lambda)$, where*

$$\begin{aligned} \omega_n^{\text{mcp, sym}}(C; \lambda) &:= \sup_{x, y \in \mathbb{R}_+} (C'(x)y + C''(x)(xy/n) \\ &\quad - \lambda(C'(x)x + C''(x)((x^2)/n) - C(x)) - C(y)) \\ &\quad \cdot (C'(y) \cdot y - C(y))^{-1}. \end{aligned}$$

Note that the above Lemma uses only the existence of a symmetric optimal profile but does not rely on symmetry of every Nash equilibrium. The following result for cost functions with a convex derivative has been obtained previously by Johari and Tsitsiklis (2005) for the special case of games with a single resource. We present here a more general result (arbitrary symmetric strategy space) with a simpler proof.

PROPOSITION 5.2. *Let $\mathcal{C}^{\text{convD}}$ be the class of cost functions with a convex derivative. Consider the set $\mathcal{G}_n(\mathcal{C}^{\text{convD}})$ of symmetric games with at most $n \in \mathbb{N}^*$ players. Then $\rho_n(\mathcal{C}^{\text{convD}}) \geq 2n/(2n+1)$.*

For polynomials with nonnegative coefficients and arbitrary degree $d \in \{2, 3, \dots\}$, we prove the following.

THEOREM 5.3. *Let \mathcal{C}_d be the class of polynomial cost function with nonnegative coefficients and arbitrary degree $d \in \{2, 3, \dots\}$. Consider the set $\mathcal{G}_n(\mathcal{C}_d)$ of symmetric games with n players and cost functions in \mathcal{C}_d . Then $\rho_n(\mathcal{C}_d) = \frac{3}{4}$.*

Note that the above bound on the worst-case efficiency does neither depend on the maximum degree d of the polynomial nor on the number of players n .

6. The Worst-Case Efficiency of Average Cost Sharing

In this section, we derive lower bounds on the worst-case efficiency of average cost sharing, which is the prevailing cost sharing method in transportation networks (cf. Beckmann et al. 1956, Haurie and Marcotte 1985). In the context of transportation networks, there is a load-dependent latency function $c_r(l_r(\varphi))$ on every resource and the cost of resource r under profile φ is defined as $C_r(l_r(\varphi)) = c_r(l_r(\varphi))l_r(\varphi)$, while the cost share for user i on resource r is determined as $\xi_i^r(\varphi) = c_r(l_r(\varphi))\varphi_r^i = (C_r(l_r(\varphi))/l_r(\varphi))\varphi_r^i$. Note that average cost sharing is a scalable cost sharing method. Given a cost function C_r , we define the *per-unit cost function* by $c_r(l_r(\varphi)) = (C_r(l_r(\varphi)))/l_r(\varphi)$. Instead of $\mathcal{G}_n(\mathcal{C}, \mathcal{D}_n)$, we will use the shorthand $\mathcal{G}_n(\mathcal{C})$ assuming that \mathcal{D}_n corresponds to average cost sharing.

Similar to the analysis of the marginal cost sharing method, we study the worst-case efficiency of average cost sharing for three types of cost functions. First, we consider general convex cost functions and derive a lower bound for the worst-case efficiency of $1/n$. Second, we consider cost functions with convex per-unit costs and characterize the price of anarchy for this case. Finally, we conclude the section by analyzing average cost sharing in symmetric games.

THEOREM 6.1. *Let $\mathcal{C}^{\text{conv}}$ be a class of convex cost functions. Consider the set $\mathcal{G}_n(\mathcal{C}^{\text{conv}})$ of games with at most $n \in \mathbb{N}$ players. Then $\rho_n(\mathcal{C}^{\text{conv}}) \geq 1/n$.*

We proceed by considering average cost functions, where the per-unit cost function is convex, that is, the functions $c_r, r \in R$ are convex.

THEOREM 6.2. *Let $\mathcal{C}^{\text{convU}}$ be a class of cost functions with convex unit costs. Then, $\rho(\mathcal{C}^{\text{convU}}) = 4/(n+3)$.*

REMARK 6.1. The bounds of Theorems 6.1 and 6.2 have been established before by Moulin (2008) for the case of a single resource.

We close this section by analyzing the efficiency loss of average cost sharing in symmetric games.

THEOREM 6.3. *Let $\mathcal{C}^{\text{convU}}$ be a class of cost functions with convex unit costs. Consider the set $\mathcal{G}_n(\mathcal{C}^{\text{convU}})$ of symmetric games with at most $n \in \mathbb{N}^*$ players. Then $\rho(\mathcal{C}^{\text{convU}}) = 4n/(n+1)^2$.*

7. Conclusions and Future Work

In this work, we studied the worst-case efficiency of Nash equilibria in resource allocation games for different cost-sharing methods. We derived various new results about the efficiency loss for marginal cost pricing and average cost sharing depending on the structure of allowable cost functions. In particular, we were able to prove tight bounds for the worst-case efficiency loss for average cost sharing and marginal cost pricing involving polynomial costs with

nonnegative coefficients. Because this class of functions is quite rich and widely used for modeling, for instance, queuing delays at resources, we see our results as an important contribution toward the applicability of these cost sharing methods in practice. While we proved that the incremental cost sharing method is optimal among all basic cost sharing methods, such a strong result is not known for the class of scalable cost sharing methods. In light of the high practical relevance of the scalability property of cost sharing methods, we see the design of an *optimal* cost sharing method among all scalable mechanisms for differentiable and convex cost functions as the most important open problem.

Appendix A

This appendix provides a table of notation.

A. Table of Notation

$N = \{1, \dots, n\}, R = \{1, \dots, m\}$	Set of players, set of resources
In the following:	$i = 1, \dots, n$ and $r = 1, \dots, m$
$R_{ij} \subset R$	j th accessible set of player i
$X_i = \{R_{i1}, \dots, R_{im_i}\}, m_i \in \mathbb{N}$	Set of accessible sets of player i
$M_i = \{1, \dots, m_i\}$	Set of indices of accessible sets of player i
$\mathcal{M} = (N, R, \{X_i\}_{i \in N}, \{C_r\}_{r \in R})$	Congestion model
$\Phi_i = \mathbb{R}_+^{m_i}, \Phi_{-i}$	Strategy space of player i and all players except i
$\phi_i = (\phi_{i1}, \dots, \phi_{im_i}) \in \Phi_i, \phi = (\phi_i, i \in N)$	Strategy of player i , strategy profile
$\Phi = \times_{i \in N} \Phi_i$	Space of strategy profiles
$d_i(\phi) = \sum_{j=1}^{m_i} \phi_{ij}$	Total demand of player i
$\mathcal{R}\mathcal{M} = (N, R, \{X_i\}_{i \in N}, \Phi, \{C_r\}_{r \in R})$	Resource allocation model
$C_R: \mathbb{R}_+ \rightarrow \mathbb{R}_+$	Cost function of resource r
$C(\phi) = \sum_{r \in R} C_r(l_r(\phi)), \mathcal{C}$	Total cost for ϕ , class of cost functions
$\phi_i^r = \sum_{j \in M_i: r \in R_{ij}} \phi_{ij}$	Load of player i on resource r
$\phi^r = (\phi_i^r, i \in N), l_r(\phi) = \sum_{i=1}^n \phi_i^r$	Load vector of resource r , total load of r
$\xi^r: \mathbb{R}_+^n \rightarrow \mathbb{R}_+^n$	Cost sharing method for a resource r
$\hat{\xi}_i^r(\phi^r) = \frac{\partial \xi_i^r(\phi^r)}{\partial \phi_i^r}, \hat{\xi}_{ij}(\phi) = \sum_{r \in R_{ij}} \hat{\xi}_i^r(\phi^r)$	Short notation
\mathcal{D}_n	Class of cost sharing methods for n players
$U_i: \mathbb{R}_+ \rightarrow \mathbb{R}_+$	Utility function of player i
$\mathcal{U}(\phi) = \sum_{i=1}^n U_i(d_i(\phi)) - C(\phi)$	Total surplus of a profile ϕ
$\pi_i(\phi) = U_i(d_i(\phi)) - \sum_{r \in R} \xi_i^r(\phi^r)$	Payoff of player i
$\pi = (\pi_1, \dots, \pi_n)$	Payoff vector of all players
$G(\mathcal{R}\mathcal{M}) = (N, \Phi, \pi)$ or G	Resource allocation game
$\vartheta, \Theta_G, \psi_G$	Nash equilibrium, set of Nash equilibria, optimal profile
$\mathcal{G}_n(\mathcal{C}, \mathcal{D}_n), \rho_n(\mathcal{C}, \mathcal{D}_n)$	Set of games with n players, worst-case efficiency

8. Electronic Companion

An electronic companion to this paper is available as part of the online version that can be found at <http://or.journal.informs.org/>.

Acknowledgments

The authors are grateful to the two anonymous reviewers for their detailed suggestions for extending the scope of this paper. They also thank Felix G. König, Max Klimm, and Rolf H. Möhring for their thorough proofreading. Some of the results of §5 appeared as an extended abstract in Harks and Miller (2009).

References

- Acemoglu, D., A. E. Ozdaglar. 2007. Competition and efficiency in congested markets. *Math. Oper. Res.* **32**(1) 1–31.
- Altman, E., T. Basar, T. Jimenez, N. Shimkin. 2002. Competitive routing in networks with polynomial costs. *IEEE Trans. Automatic Control* **47**(1) 92–96.
- Altman, E., T. Boulogne, R. El-Azouzi, T. Jiménez, L. Wynter. 2006. A survey on networking games in telecommunications. *Comput. Oper. Res.* **33**(2) 286–311.
- Beckmann, M., C. McGuire, C. Winsten. 1956. *Studies in the Economics and Transportation*. Yale University Press, New Haven, CT.
- Bhaskar, U., L. Fleischer, D. Hoy, C. Huang. 2009. Equilibria of atomic flow games are not unique. *Proc. 20th Annual ACM-SIAM Sympos. Discrete Algorithms*. Society for Industrial and Applied Mathematics, Philadelphia, 748–757.
- Branstón, D. 1976. Link capacity functions. *Transportation Res.* **10**(4) 223–236.
- Catoni, S., S. Pallotino. 1991. Traffic equilibrium paradoxes. *Transportation Sci.* **25**(3) 240–244.
- Chau, C. K., K. M. Sim. 2003. The price of anarchy for non-atomic congestion games with symmetric cost maps and elastic demands. *Oper. Res. Lett.* **31**(3) 327–334.
- Chen, Y.-J., J. Zhang. 2010. Design of price mechanisms for network resource allocation via price of anarchy. *Math. Programming*, ePub ahead of print May 12.
- Cole, R., Y. Dodis, T. Roughgarden. 2003. Pricing network edges for heterogeneous selfish users. *Proc. 35th Annual ACM Sympos. Theory Comput.* ACM, New York, 521–530.
- Cominetti, R., J. R. Correa, N. E. Stier-Moses. 2009. The impact of oligopolistic competition in networks. *Oper. Res.* **57**(6) 1421–1437.
- Correa, J. R., A. S. Schulz, N. E. Stier-Moses. 2004. Selfish routing in capacitated networks. *Math. Oper. Res.* **29**(4) 961–976.
- Fleischer, L., K. Jain, M. Mahdian. 2004. Tolls for heterogeneous selfish users in multicommodity networks and generalized congestion games. *Proc. 45th Annual IEEE Sympos. Foundations Comput. Sci.*, IEEE Computer Society, Washington, DC, 277–285.
- Guo, X., H. Yang. 2005. The price of anarchy of Cournot oligopoly. Deng, X., Y. Ye, eds. *Proc. 1st Internat. Workshop Internet Network Econom. (WINE)*, Lecture Notes in Computer Science, Vol. 3828, Springer, Berlin, 246–257.
- Harker, P. T. 1988. Multiple equilibrium behaviors on networks. *Transportation Sci.* **22**(1) 39–46.
- Harks, T. 2011. Stackelberg strategies and collusion in network games with splittable flow. *Theory Comput. Systems* **48**(4) 781–802.
- Harks, T., K. Miller. 2009. Efficiency and stability of Nash equilibria in resource allocation games. *Proc. First ICST Internat. Conf. Game Theory for Networks. GameNets '09*, IEEE Press, Piscataway, NJ, 398–402.
- Haurie, A., P. Marcotte. 1985. On the relationship between Nash-Cournot and Wardrop equilibria. *Networks* **15**(3) 295–308.
- Hayrapetyan, A., E. Tardos, T. Wexler. 2006. The effect of collusion in congestion games. *Proc. 38th Annual ACM Sympos. Theory Comput.*, ACM Press, New York, 89–98.
- Hearn, D., M. Ramana. 1998. Solving congestion toll pricing models. P. Marcotte, S. Nguyen, eds. *Equilibrium and Advanced Transportation Modeling*. Kluwer Academic Publishers, Boston, 109–124.
- Johari, R. 2004. Efficiency loss in market mechanisms for resource allocation. Ph.D. thesis, Massachusetts Institute of Technology, Cambridge, MA.
- Johari, R., J. N. Tsitsiklis. 2004. Efficiency loss in a network resource allocation game. *Math. Oper. Res.* **29**(3) 407–435.
- Johari, R., J. Tsitsiklis. 2005. Efficiency loss in Cournot games. Technical report, LIDS-P-2639, Laboratory for Information and Decision Systems, Massachusetts Institute of Technology, Cambridge, MA.
- Johari, R., J. Tsitsiklis. 2006. A scalable network resource allocation mechanism with bounded efficiency loss. *IEEE J. Selected Areas Comm.* **24**(5) 992–999.
- Johari, R., J. N. Tsitsiklis. 2009. Efficiency of scalar-parameterized mechanisms. *Oper. Res.* **57**(4) 823–839.
- Johari, R., S. Mannor, J. N. Tsitsiklis. 2005. Efficiency loss in a network resource allocation game: The case of elastic supply. *IEEE Trans. Automatic Control* **50**(11) 1712–1724.
- Karakostas, G., S. G. Kolliopoulos. 2004. Edge pricing of multicommodity networks for heterogeneous selfish users. *Proc. 45th Annual IEEE Sympos. Foundation Comput. Sci.*, IEEE Computer Society, Washington, DC, 268–276.
- Kelly, F. 1997. Charging and rate control for elastic traffic. *Eur. Trans. Telecomm.* **8**(1) 33–37.
- Kelly, F., A. Maulloo, D. Tan. 1998. Rate control in communication networks: Shadow prices, proportional fairness, and stability. *J. Oper. Res. Soc.* **49**(3) 237–252.
- Knight, F. H. 1924. Some fallacies in the interpretation of social cost. *Quart. J. Econom.* **38**(4) 582–606.
- Mas-Colell, A., M. D. Whinston, J. R. Green. 1995. *Microeconomic Theory*. Oxford University Press, New York.
- Moulin, H. 2008. The price of anarchy of serial, average and incremental cost sharing. *Econom. Theory* **36**(3) 379–405.
- Novshek, W. 1985. On the existence of Cournot equilibrium. *Rev. Econom. Stud.* **52**(1) 85–98.
- Orda, A., R. Rom, N. Shimkin. 1993. Competitive routing in multi-user communication networks. *IEEE/ACM Trans. Networking* **1**(5) 510–521.
- Owen, G. 1982. *Game Theory*, 2nd ed. Academic Press, New York.
- Ozdaglar, A. E. 2008. Price competition with elastic traffic. *Networks* **52**(3) 141–155.
- Patriksson, M. 1994. *Traffic Assignment Problems: Models and Methods*. Topics in Transportation. V.S.P. International Science, Utrecht, The Netherlands.
- Perakis, G. 2007. The “price of anarchy” under nonlinear and asymmetric costs. *Math. Oper. Res.* **32**(3) 614–628.
- Rosen, J. 1965. Existence and uniqueness of equilibrium points in concave n -player games. *Econometrica* **33**(3) 520–534.
- Rosenthal, R. 1973. A class of games possessing pure-strategy Nash equilibria. *Internat. J. Game Theory* **2**(1) 65–67.
- Roughgarden, T. 2002. The price of anarchy is independent of the network topology. *J. Comput. System Sci.* **67**(2) 341–364.
- Roughgarden, T. 2005. *Selfish Routing and the Price of Anarchy*. MIT Press, Cambridge, MA.
- Roughgarden, T., F. Schoppmann. 2011. Local smoothness and the price of anarchy in atomic splittable congestion games. *Proc. 22nd Annual ACM-SIAM Sympos. Discrete Algorithms*, Society for Industrial and Applied Mathematics, Philadelphia, 255–267.
- Roughgarden, T., É. Tardos. 2002. How bad is selfish routing? *J. ACM* **49**(2) 236–259.
- Samet, D., Y. Tauman. 1982. The determination of marginal cost prices under a set of axioms. *Econometrica* **50**(4) 895–909.
- Schrijver, A. 2003. *Combinatorial Optimization: Polyhedra and Efficiency*. Springer-Verlag, Berlin.
- Smith, M. 1979. The marginal cost taxation of a transportation network. *Transportation Res. Part B: Methodological* **13**(3) 237–242.
- Srikant, R. 2003. *The Mathematics of Internet Congestion Control*. Birkhäuser, Boston.
- Swamy, C. 2007. The effectiveness of Stackelberg strategies and tolls for network congestion games. *Proc. 18th ACM-SIAM Symp. Discrete Algorithms*, Society for Industrial and Applied Mathematics, Philadelphia, 1133–1142.
- Wardrop, J. G. 1952. Some theoretical aspects of road traffic research. *Proc. Institution Civil Engineers, Part II* **1**(36) 352–362.
- Yang, H., H. Huang. 2004. The multi-class, multi-criteria traffic network equilibrium and systems optimum problem. *Transportation Res. Part B: Methodological* **38**(1) 1–15.
- Yang, H., X. Zhang. 2008. Existence of anonymous link tolls for system optimum on networks with mixed equilibrium behaviors. *Transportation Res. Part B: Methodological* **42**(2) 99–112.
- Yang, H., D. Han, H. Lo. 2008. Efficiency of atomic splittable selfish routing with polynomial cost functions. *Networks Spatial Econom.* **8**(4) 443–451.
- Yang, H., W. Xu, B. Heydecker. 2010. Bounding the efficiency of road pricing. *Transportation Res. Part E: Logistics and Transportation Rev.* **46**(1) 90–108.
- Yang, S., B. Hajek. 2007. VCG-Kelly mechanisms for allocation of divisible goods: Adapting VCG mechanisms to one-dimensional signals. *IEEE J. Selected Areas Comm.* **25**(6) 1237–1243.