Operations Research
Vol. 57, No. 6, November–December 2009, pp. 1421–1437
ISSN 0030-364X | EISSN 1526-5463 | 09 | 5706 | 1421

informs®
DOI 10.1287/opre.1080.0653

The Impact of Oligopolistic Competition in Networks

Roberto Cominetti
Departamento de Ingeniería Matemática, Universidad de Chile, Santiago, Chile, rcominet@dim.uchile.cl

José R. Correa
Departamento de Ingeniería Industrial, Universidad de Chile, Santiago, Chile, jcorrea@dii.uchile.cl

Nicolás E. Stier-Moses
Graduate School of Business, Columbia University, New York, New York 10027, stier@gsb.columbia.edu

In the traffic assignment problem, first proposed by Wardrop in 1952, commuters select the shortest available path to travel from their origins to their destinations. We study a generalization of this problem in which competitors, who may control a nonnegligible fraction of the total flow, ship goods across a network. This type of games, usually referred to as atomic games, readily applies to situations in which the competing freight companies have market power. Other applications include intelligent transportation systems, competition among telecommunication network service providers, and scheduling with flexible machines.

Our goal is to determine to what extent these systems can benefit from some form of coordination or regulation. We measure the quality of the outcome of the game without centralized control by computing the worst-case inefficiency of Nash equilibria. The main conclusion is that although self-interested competitors will not achieve a fully efficient solution from the system's point of view, the loss is not too severe. We show how to compute several bounds for the worst-case inefficiency that depend on the characteristics of cost functions and on the market structure in the game. In addition, building upon the work of Catoni and Pallotino, we show examples in which market aggregation (or collusion) adversely impacts the aggregated competitors, even though their market power increases. For example, Nash equilibria of atomic network games may be less efficient than the corresponding Wardrop equilibria. When competitors are completely symmetric, we provide a characterization of the Nash equilibrium using a potential function, and prove that this counterintuitive phenomenon does not arise. Finally, we study a pricing mechanism that elicits more coordination from the players by reducing the worst-case inefficiency of Nash equilibria.

Subject classifications: networks/graphs: multicommodity, theory; games/group decisions: noncooperative, atomic; transportation: models, network; programming: complementarity.
Area of review: Optimization.
History: Received July 2007; revision received January 2008; accepted June 2008. Published online in *Articles in Advance* June 3, 2009.

1. Introduction

Logistic and freight companies routinely transport goods between different points in the world to serve their clients. They make use of trucks, trains, ships, and planes to deliver goods from their points of origin to their destinations. Companies that provide this service compete in at least two dimensions: the price they charge for shipping and the service level they provide. To improve their competitive advantage, these companies need to be strategic in how they deliver the goods and minimize costs and delivery times. The main operative decision is choosing the routes to be used in the actual shipments, where each route consists of a sequence of basic segments that may combine various modes.

Although these companies may own and operate some of the resources needed to move and sort the goods, frequently they also subcontract other multimodal freight transportation companies, and pay to receive services and to use resources. Some examples consist of leasing additional planes at a peak time such as the holiday season, paying landing and take-off fees to airports, and sending a shipment to an area not covered by the company's network. From the perspective of the company that provides service to the freight company, standard economic arguments imply that additional demand is associated with a price increase. For example, if an airport becomes very popular, it may decide to increase landing and take-off fees because the demand supports it. A potential increase may lead some freight companies to consider other routes that were too expensive with the previous fees. Besides the economic consideration, another aspect is congestion. For many facilities, including airports, an increase in demand causes an increase in service time, which may impact the feasibility of a given route.

We focus on the freight companies, and model companies that provide service to them implicitly. As we described above, the cost and the delay generated by a given resource depend on which freight companies use it and the quantity of goods they ship through it. These negative externalities generate an interdependence between freight companies, and leads to a competitive situation between them. We study this competition from the perspective of noncooperative game theory, and use the Nash equilibrium as the solution concept of the game. An equilibrium is a solution under which no competitor has any regret after seeing what all competitors have done.

The main goal of this article is to provide a methodology to understand under what conditions equilibria are efficient or, at least, not extremely inefficient. We measure the efficiency of a solution, which represents the collective decisions made by the competitors, using a *social cost function*. A low or high social cost, in itself, does not imply that an equilibrium is good or bad, because there may be instances that are intrinsically more expensive than others; instead, we compare the social cost of an equilibrium to a reference point provided by the socially optimal solution. This solution encodes what would happen if a single decision maker controlled all the freight companies and minimized the social cost. To quantify the efficiency loss arising from the self-mindedness of freight companies, one can compute the worst-case ratio of the social cost of an equilibrium to that of the social optimum. If the ratio, generally referred to as the *price of anarchy*, happens to be close to one, equilibria are rather efficient, which suggests that companies are better off by making routing decisions on their own. Even if the competitors could get together and coordinate themselves, they would not be able to lower the social cost significantly. Moreover, the extra effort to coordinate has a cost (e.g., deploying systems to collect information, computing optimal coordinated solutions, disseminating information back to the coalition members, and enforcing that companies do as told) that could offset the reduction arising from the additional coordination. This does not even consider that participants lose their free will, which raises the cost as well. Conversely, when the ratio is significantly higher than one, competitors may benefit from some kind of coordination. This coordination can take the form of regulations for the market structure (e.g., no company can hold more than a given market share), regulations for the network (e.g., trucks are not allowed to circulate in certain roads), incentives (e.g., subsidies or taxes for some resources), etc. In this article, we look at pricing mechanisms that can approximate the resulting equilibria to the socially optimal solution.

1.1. A Model of Competition in Freight Transportation

We represent the different segments and resources, henceforth referred to as arcs, by a directed network $G=(V,A)$. For example, an arc may represent a route from Hong Kong to New York by sea going through the Panama Canal, landing in the Atlanta airport, or leasing a truck from San Francisco to Denver. Note that we do not mean a specific boat, plane, train, or truck; we mean that the shipment uses that particular arc along its route. We denote the set of all freight companies, henceforth referred to as players, by $[K]=\{1,\dots,K\}$. We assume that player $k\in[K]$ has to send d_k units of freight from node s_k to node t_k (§6 describes a generalization to multiple sources and destinations). We refer to (s_k,t_k) as an origin-destination (OD) pair. Each player executes its contracts by selecting how much freight to send along each possible route connecting the corresponding OD pair. This decision is encoded by a flow that specifies shipments along each arc, and that satisfies flow conservation constraints at every node. Summarizing, each player $k\in[K]$ chooses a flow $x^k\in\mathbb{R}_+^A$ that routes d_k units of flow from s_k to t_k. We refer collectively to the flows for all players by $\vec{x}:=(x^1,\dots,x^K)\in\mathbb{R}_+^{A\times K}$. In addition, to simplify notation we henceforth let $x:=\sum_{k\in[K]}x^k$ be the aggregate flow induced by all K players. Hence, we denote the flow that player k ships through arc a by x_a^k, and the total flow by x_a.

As we described previously, arcs are subject to congestion and to competition. More demand for a carrier increases its delay and its price. Because both negatively affect the cost incurred by a given company, the standard way of modeling this is to merge delay and price in a single *cost function*. This modeling simplification can be achieved by expressing delays in currency units. For example, one can assign a penalty (usually called the *value of time*) to each unit of time a product is late. This penalty reflects the customer goodwill that is lost from the delay, or the cost of having capital tied in the form of inventory for one extra period of time. Formally, we associate a cost function $c_a(\cdot)\colon\mathbb{R}_+\to\mathbb{R}_+$ to every arc. These functions map the total flow x_a on arc a to its per-unit cost $c_a(x_a)$, which is equal for all players. Notice that the cost function depends on the aggregated flow x_a, but not on the decomposition $\vec{x}_a$. Cost functions are assumed to be increasing, differentiable, and convex, although for some of our results the convexity assumption can be slightly relaxed. In addition, in this article we only consider separable cost functions, meaning that the cost in one arc only depends on the flow in the same arc. Two commonly used cost functions are polynomials of small degree (e.g., the Bureau of Public Roads 1964 uses the well-known BPR cost functions to measure delay in road segments; these functions are polynomials of degree 4) and delay functions of queues (e.g., $(c_a-x_a)^{-1}$, where c_a is the capacity of the queue). Of course, the choice of cost function in a given arc will ultimately depend on the role of that arc in the logistic network. We assume that cost functions are taken from a set of allowable cost functions $\mathscr{C}$.

The goal of player k is to send its total demand d_k minimizing its own cost $C^k(\vec{x}):=\sum_{a\in A}x_a^k c_a(x_a)$. Note that players can divide their flows among many paths if they think it is convenient for them. Indeed, in some situations it is advantageous to send a fraction of the goods along a more expensive route to lower the cost in a bottleneck. Because of the convexity of cost functions, this translates into savings for most of the freight, which can ultimately reduce the total cost of the shipment.

Because the above cost function heavily depends of the decisions made by other players, the natural solution concept is that of an equilibrium. Thus, a solution for all players $\vec{x}^{\text{NE}}$ is called a Nash equilibrium if each flow $x^{\text{NE},k}$ minimizes $C^k(x^k, x^{\text{NE},-k})$ among all flows x^k that are feasible for player k, assuming that flows of other players $x^{\text{NE},-k}$ are fixed. Because the goal is quantifying the quality of equilibria, we introduce a social cost function, given by $C(\vec{x}) = C(x) := \sum_{k\in[K]} C^k(\vec{x}) = \sum_{a\in A} x_a c_a(x_a)$. Notice that the total cost does not depend on $\vec{x}$ directly, but on the aggregated flow x. A socially optimal flow is a solution $\vec{x}^{\text{OPT}}$ that minimizes $C(x)$ among all feasible solutions $\vec{x}$. Such a solution may not be stable because players could have an incentive to deviate from it. Notice that our assumptions guarantee that the system optimum is unique because the objective function is strictly convex and it is minimized over a polytope. It is well known that a Nash equilibrium can be inefficient with respect to a social optimum (Pigou 1920, Dubey 1986); actually, it may even be worse for *all* players (Braess 1968). In addition, it need not even be Pareto optimal.

An instance of the game introduced above is defined by the network topology, the set of players with their corresponding OD pairs and demands, and the cost functions associated to arcs. We consider a set of allowed instances $\mathcal{I}$ and denote an arbitrary Nash equilibrium and the social optimum of a given instance $I \in \mathcal{I}$ by $\vec{x}^{\text{NE}}(I)$ and $\vec{x}^{\text{OPT}}(I)$, respectively.

Koutsoupias and Papadimitriou (1999) proposed to use the worst-case ratio of the social cost of equilibria and that of a socially optimal solution as a way to quantify the impact of not being able to coordinate the players of a game. This quantity, which became known as the *price of anarchy* (Papadimitriou 2001), can be computed by solving

$$\sup_{I\in\mathcal{I}} \frac{C(\vec{x}^{\text{NE}}(I))}{C(\vec{x}^{\text{OPT}}(I))}. \tag{1}$$

As any worst-case measure, the price of anarchy tends to be pessimistic when considered broadly. For example, if $\mathcal{I}$ contains all possible instances, the supremum is unbounded (Roughgarden and Tardos 2002). To get a more realistic estimation of the efficiency loss, we consider smaller sets $\mathcal{I}$. Past work and different parts of this paper restrict either the cost functions, the OD pairs, or the demands to have certain characteristics.

1.2. Our Contributions and Related Literature

The game presented in the previous section is generally called a *network game*, although it also belongs to the more general class of *congestion games* introduced by Rosenthal (1973). The distinctive characteristic of these games is that the per-unit cost of a resource (arc in this case) depends only on the number of players that selected the resource (total flow in this case), not on the identities of those players. Although we concentrate on network games to simplify the presentation, everything holds true for the more general class as well. Section 6 provides further details and an application of congestion games in our context.

Most of the previous work on network games considers that there are infinitely many players, and none of them substantially controls the market. For this reason, players cannot influence prices unilaterally, causing them to be price taking. In this situation, we say that the game is *nonatomic*. A common application is given by a transportation network in which players represent drivers that commute in the network. Here, players are small compared to the scale of the whole system, and cannot modify the congestion level on a given road by themselves. An equilibrium is an assignment of commuters to routes such that everybody is simultaneously taking a shortest path under the prevailing conditions. This solution is commonly referred to as a *Wardrop equilibrium*, due to the seminal paper about road traffic modeling by Wardrop (1952).

Roughgarden and Tardos (2002) initiated the study of the price of anarchy in nonatomic network games. They showed that Wardrop equilibria can be arbitrarily inefficient compared to social optima if one considers all possible instances. For that reason, it is relevant to compute the worst-case inefficiency, parameterized with the class $\mathcal{C}$ of cost functions that are allowed to appear in the network. For example, $\mathcal{C}$ can be the set of affine functions, the polynomials of degree smaller than a fixed constant, or the M/M/1 delay functions. For affine cost functions, Roughgarden and Tardos (2002) showed that the price of anarchy is 4/3, which implies that the efficiency loss that arises from the self-mindedness of players is at most 33%. Following their work, a series of papers generalized the initial results by considering more general assumptions. Roughgarden (2003) considered a general class $\mathcal{C}$ of (nonlinear) functions and established that an instance achieving the supremum in (1) always has a simple structure, which facilitates the computation of the price of anarchy. Exploiting that, he found that the price of anarchy is 1.626 for quadratic functions, 1.9 for cubic ones, and it grows as $b/\ln b$ if $\mathcal{C}$ contains nonnegative polynomials of degree at most b. Correa et al. (2004) introduced the use of variational inequalities in this setting, which allowed them to add side constraints to the problem without increasing the price of anarchy, and to drop assumptions made previously for technical reasons. Chau and Sim (2003) extended the analysis to allow for symmetric, nonseparable cost functions and elastic demands, whereas Perakis (2007) considered asymmetric, nonseparable ones. Because nonseparable cost functions depend on the flow on *all* arcs in the network, they can represent more general congestion and competition effects. Finally, Roughgarden and Tardos (2004) and Correa et al. (2008) generalized earlier results from network to congestion games. Furthermore, the latter reference provides a graphical interpretation of the inefficiency of equilibria.

An *atomic game* represents situations in which some of the players have a significant market power. In this case, players control arbitrary demands, as opposed to an infinitesimal amount of flow, as was assumed in a nonatomic game. Depending on the specific application, flow can or cannot be divided along different routes. For example, for freight transportation the shipper can normally send different packets along different routes if it deems it convenient. In some cases, companies may require shipments to follow a single route to minimize the likelihood of losing or misplacing items. For telecommunications networks, depending on the protocols in use, traffic is or is not required to follow the same route (e.g., IP traffic versus ATM traffic). Roughgarden and Tardos (2002) also presented some results for the case of unsplittable demand. Later, Fotakis et al. (2005) studied special classes of networks, and Awerbuch et al. (2005) and Christodoulou and Koutsoupias (2005) independently proved that if cost functions are linear, the price of anarchy is bounded by a small constant.

The splittable case, which is the model we focus on, was first considered by Orda et al. (1993), who noted that existence of equilibria follows directly from the classical result about concave games of Rosen (1965). That article and one by Altman et al. (2002), among others, have obtained uniqueness results for some special cases; nevertheless, uniqueness does not hold in general as shown recently by Bhaskar et al. (2009). Although Roughgarden and Tardos (2002), Roughgarden (2005), and Correa et al. (2005) studied the inefficiency of equilibria in a similar model, unfortunately there are some problems with those results, as we describe below. For this reason, this article presents the first upper bounds on the price of anarchy of network games with atomic players and splittable flow. Our main conclusion is that although Nash equilibria may be strictly worse than a socially optimal solution, the gap between the two is not too large. On the negative side, a Nash equilibrium of an atomic game may be worse than a Wardrop equilibrium of the corresponding nonatomic game, implying that market power can have a negative effect on the quality of solutions. Additionally, a counterintuitive phenomenon may arise: If some firms collude and aggregate their demands, one would expect that their collective efficiency improves. We provide examples that show that this need not happen. Moreover, firms outside the cartel may find that the cost they incur is lower when they compete with the cartel than when they compete with the individual companies.

After looking at general market structures, we consider assumptions that allow us to provide stronger results. Specifically, when competitors ship from a common origin to a common destination, we find a bound that depends on the Herfindahl index (Tirole 1988)—a standard measure of the industry concentration that is used by the U.S. Federal Trade Commission to evaluate mergers and acquisitions. Our bound shows that the price of anarchy decreases when going from oligopolies with few companies that dominate the market to instances in which companies' market shares are similar. Finally, assuming that market shares are exactly equal allows us to characterize equilibria using a potential function. In this case, all players are completely symmetric because they also ship from a common origin to a common destination. Besides simplifying the calculation of equilibria, the characterization also implies that when more companies compete, equilibria become less efficient because companies are more difficult to coordinate. This rules out paradoxes like the one previously mentioned and implies that equilibria with atomic players are at least as efficient as the Wardrop equilibrium of the corresponding nonatomic instance. Independently of this work, Hayrapetyan et al. (2006) study the effect of collusion in network games and reach a similar conclusion for networks with parallel arcs and splittable demands. Our results consider more restrictive assumptions on the players, but are valid for arbitrary networks.

When players of a game internalize the negative externalities they generate (by paying a tax or toll), the resulting outcome is socially optimal. Inspired by this insight, several researchers in different domains have designed payment mechanisms that provide the incentives to the participants to make decisions that are optimal from the system's perspective. This has important regulatory and operational consequences because solutions that the system designer has in mind can be enforced without introducing explicit coordination among the users. For example, in the area of transportation networks, this concept has been called *congestion pricing* (Vickrey 1969, Johnson and Mattson 1992). This mechanism assigns tolls to certain arcs of the network, which are charged to users that take routes through them. Congestion pricing has been used in cities such as Singapore and London as a measure to help relieve the ever-increasing congestion. The best possible set of tolls—called *optimal tolls*—is one for which an equilibrium of the system with tolls is socially optimal in the original network. Beckmann et al. (1956) proved that charging users the difference between the marginal and the real cost makes them internalize the externality they generate, thus showing that optimal tolls exist for nonatomic network games with homogeneous users. More recently, Cole et al. (2003) considered the more realistic situation in which players are heterogeneous in their valuation of time. Their main result says that when all users share the same origin and destination there is an optimal set of tolls. Yang and Huang (2004), and later Fleischer et al. (2004) and Karakostas and Kolliopoulos (2004) proved that there are optimal tolls for heterogeneous users even in general networks. For the case of atomic players, we provide a mechanism that charges uniform prices under which the Nash equilibrium is closer to the social optimum, thus reducing the price of anarchy.

Although our mechanism improves the efficiency of equilibria, they do not lead to fully efficient solutions. In the case of symmetric games, tolls similar to those proposed by Beckmann et al. (1956) are optimal.

Most research in the area of supply chain management focuses on the design of mechanisms to achieve full efficiency instead of analyzing the status quo to decide if something actually needs to be done. For example, Taylor (2002), Cachon and Netessine (2004), Bernstein and Federgruen (2005), and Golany and Rothblum (2006) study contracts and pricing mechanisms that can induce efficient Nash equilibria. Nevertheless, Parlakturk and Kumar (2004) and Perakis and Roels (2007) compute the worst-case inefficiency of equilibria in service systems and in various types of supply chains, respectively. Switching to other application domains, a series of papers study competition and pricing in telecommunication networks. Johari and Tsitsiklis (2004), building on Kelly (1997), prove that auctioning capacity in networks leads to an efficiency loss of at most 33%. Sanghavi and Hajek (2004) and Maheswaran and Başar (2004) study extensions of the basic model. Finally, Acemoglu and Ozdaglar (2007) study competition when network providers compete for traffic by setting prices, and prove a tight worst-case bound for the efficiency loss.

Structure of This Paper. In §2, we present a variational inequality characterization of Nash equilibria for the atomic case, and work out a simple example. Then, §3 shows an upper bound on the price of anarchy for arbitrary networks when cost functions belong to a set given a priori. In addition, we provide a lower bound that arises from a particularly bad instance. Section 4 concentrates in games with a single origin and destination. We provide a bound on the price of anarchy that depends on the variability of the market power across players and a bound for the case of symmetric players. Section 5 discusses pricing mechanisms that reduce the price of anarchy. We conclude in §6 by presenting additional applications of our model, and some directions in which our results can be generalized.

2. Characterization of Nash Equilibria in the Atomic Network Game

Recall that a solution $\vec{x}^{\text{NE}}$ for all players is a Nash equilibrium if it is a best reply strategy for each player. For that to be the case, the flow $x^{\text{NE},k}$ for player $k \in [K]$ needs to be a solution to the following optimization problem in which flows $x^{\text{NE},i}$ are fixed for $i \neq k$. For ease of notation, we introduce a reverse arc with zero cost between t^k and s^k for each k:

$$
\begin{aligned}
(\text{NE}^k)\quad & \min\ C^k(x^{\text{NE},1},\dots,x^{\text{NE},k-1},x^k,x^{\text{NE},k+1},\dots,x^{\text{NE},K})\\
& \text{s.t.}\ \sum_{(u,v)\in A} x^k_{(u,v)} - \sum_{(v,w)\in A} x^k_{(v,w)} = 0 \quad \text{for all } v\in V,\\
& x^k_{(t^k,s^k)} = d_k,\\
& x^k_a \geqslant 0 \quad \text{for all } a\in A.
\end{aligned}
$$

Note that our assumptions guarantee that these optimization problems are convex. Then, an equilibrium always exists (Rosen 1965). Using the convexity of $C^k(\vec{x})$ and the first-order optimality conditions of problem (NE^k), we can characterize equilibria with a variational inequality. Indeed, $\vec{x}^{\text{NE}}$ is at equilibrium if and only if, for all $k \in [K]$, $x^{\text{NE},k}$ solves

$$
\sum_{a\in A} c^k_a(\vec{x}^{\text{NE}}_a)(x^k_a - x^{\text{NE},k}_a) \geqslant 0 \quad \text{for any feasible flow } x^k \text{ for player } k. \tag{2}
$$

Here, the modified cost function $c^k_a(\vec{x}_a) := c_a(x_a) + x^k_a c'_a(x_a)$ is the derivative with respect to x^k_a of the term $x^k_a c_a(x_a)$ in $C^k(\vec{x})$. Intuitively, the second term accounts for player k's ability to affect prices.

At times, we will consider the Wardrop equilibria x^{WE} of an instance where each atomic player is replaced by nonatomic ones controlling the same total flow. Under a Wardrop equilibrium, all used paths serving the same OD pair need to have the same cost with respect to $c_a(x^{\text{WE}}_a)$. In addition, following Harker (1988), we consider situations in which some OD pairs are controlled by atomic players, whereas others are controlled by infinitely many nonatomic players. These games can be viewed as limits of games in which the number of players tends to infinity, but some of them retain market power to set prices, whereas others are relegated to be price takers. Harker (1988) referred to the equilibria of those games as *mixed behavior equilibria*, and he showed how to characterize them using a set of variational inequalities similar to (2). Except where otherwise stated, all results in this paper are valid for the three classes of equilibria that we introduced (Nash, Wardrop, and mixed) because we work with arbitrary market powers.

We remind the reader that the quality of equilibria for a given set of allowable instances is determined by solving the problem shown in (1). In this article, we work with arbitrary network topologies, we assume that cost functions belong to an arbitrary but fixed set of functions $\mathscr{C}$, and we consider alternative assumptions for the structure of the players and their OD pairs. For example, market structures can be arbitrary (§3), have a single OD pair with arbitrary demands (§4.1), or have a single OD pair with symmetric demands (§4.2).

2.1. A Simple Example with Linear Costs

In this section, we provide a simple example with linear costs (i.e., the cost equals a constant times the flow) to illustrate that price-setting players can hurt the system. Although in this case Wardrop equilibria are known to be optimal, Nash equilibria may be inefficient when players are atomic. (Dafermos and Sparrow 1969 showed that when cost functions are of the form $c_a(x_a) = r_a x^b_a$ for a fixed b, a flow is a Wardrop equilibrium if and only if it has minimal total cost.)

The example we present in Figure 1 is inspired from a discussion on traffic paradoxes by Catoni and Pallotino

Figure 1. An instance of a network game with atomic players and splittable flow.

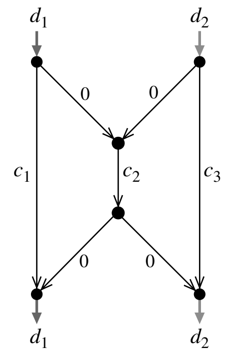

(1991) and will be the basis for other instances in which Nash equilibria exhibit a peculiar behavior. There are two atomic players with demands equal to $d_1 := 2$ and $d_2 := 3$, and cost functions equal to $c_1(x_1) := x_1$, $c_2(x_2) := x_2$, and $c_3(x_3) := 2x_3$. In this example, the unique Wardrop equilibrium x^{WE} (and therefore also the social optimum) routes all d_1 units of flow along arc 1 making $x_1^{\text{WE}} = 2$, whereas the d_2 units are split between $x_2^{\text{WE}} = 2$ and $x_3^{\text{WE}} = 1$. Notice that $c_1(2) = c_2(2) = c_3(1) = 2$ as expected, and the total cost equals 10.

Nash equilibria of atomic games with linear cost functions generally do not minimize the total cost. The main difference compared to a Wardrop equilibrium is that when competitors have market power they give relatively less importance to others. Using (2), it can be verified that the unique Nash equilibrium is the flow that routes $x_1^{\text{NE}} \approx 1.48$, $x_2^{\text{NE}} \approx 0.52 + 1.91$, and $x_3^{\text{NE}} \approx 1.09$, where x_2^{NE} is the sum of the flows coming from the two players. Indeed, all used paths have the same length with respect to the modified cost functions $c_a^k(\vec{x}_a) = r_a(x_a + x_a^k)$. Its total cost is approximately equal to 10.47, and therefore the efficiency loss in this instance amounts to 4.7%. The degradation arises because both players give less importance to the competitor's flow and they load arc 2 too much. A carefully constructed instance with linear cost functions implies that the price of anarchy is at least 1.17. For the case of a single OD pair, Altman et al. (2002) show that Wardrop equilibria, Nash equilibria, and social optima all coincide when cost functions are monomials of a fixed degree.

3. Atomic Games with General Players

In this section, we study the price of anarchy for atomic games with arbitrary networks and demand configurations. For example, players with arbitrary market power may coexist with price-taking players. The most important conclusion that can be drawn from the results of this section is that if marginal costs do not increase too steeply, then the total cost at a Nash equilibrium is not too large when compared to the cost of a socially optimal solution. Even if players could collude and distribute the benefits fairly, the total savings will not be significant.

As a warm-up exercise and before considering arbitrary cost functions, we derive a bound on the price of anarchy for the case in which $\mathscr{C}$ is the set of affine cost functions. To this end, we define an optimization problem whose first-order optimality conditions correspond to the equilibrium conditions. In particular, this optimization problem implies that the game is potential (Monderer and Shapley 1996) and that there is an essentially unique equilibrium. (Boulogne 2004 pointed out that this conclusion is implied by the results of Rosen 1965; "essentially unique" means that if there are multiple equilibria, they are indistinguishable from the players' perspective.)

Consider an affine cost function of the form $c_a(x_a) = q_a x_a + r_a$. Let us define a modified cost function $\hat{c}_a$: $\mathbb{R}_+^K \to \mathbb{R}_+$ by $\hat{c}_a(\vec{x}_a) := q_a \sum_{i \leqslant j \in [K]^2} x_a^i x_a^j + r_a \sum_{k \in [K]} x_a^k$. Note that $2\sum_{i \leqslant j \in [K]^2} x_a^i x_a^j = (\sum_{k \in [K]} x_a^k)^2 + \sum_{k \in [K]} (x_a^k)^2$, which implies that $\hat{c}_a(\vec{x}_a)$ is convex, or strictly convex when costs are strictly increasing. We define problem (NLP-NE) as the minimization of the potential function $\widehat{C}(\vec{x}) := \sum_{a \in A} \hat{c}_a(\vec{x}_a)$ among all feasible flows $\vec{x}$. Strict convexity implies that there is a single solution to the previous problem. Because its first-order optimality conditions coincide with the conditions that characterize a Nash equilibrium, the latter has to be unique. In addition, problem (NLP-NE) can be used to approximate a Nash equilibrium up to a fixed additive term in polynomial time (Potra and Ye 1993). One cannot expect to do better than an additive approximation because an equilibrium may require irrational numbers. This potential function can be used to derive bounds on the price of anarchy (e.g., see Roughgarden and Tardos 2002, Johari and Tsitsiklis 2004). However, this bound is looser than that of Proposition 3.2.

Proposition 3.1. *Consider an atomic congestion game with K players and affine cost functions. Let $\vec{x}^{\text{NE}}$ be a Nash equilibrium and $\vec{x}^{\text{OPT}}$ be a social optimum. Then, $C(x^{\text{NE}}) \leqslant (2K/(K+1))C(x^{\text{OPT}})$.*

Proof. To compare $\widehat{C}(\cdot)$ and $C(\cdot)$, notice that $x_a c_a(x_a) \leqslant (2K/(K+1))\hat{c}_a(\vec{x}_a)$ for an arbitrary decomposition of x_a into $\vec{x}_a$. Summing over the arcs,

$$C(x^{\text{NE}}) \leqslant \frac{2K}{K+1}\widehat{C}(\vec{x}^{\text{NE}}) \leqslant \frac{2K}{K+1}\widehat{C}(\vec{x}^{\text{OPT}}) \leqslant \frac{2K}{K+1}C(x^{\text{OPT}}),$$

where the middle inequality holds because $\vec{x}^{\text{NE}}$ minimizes problem (NLP-NE). □

This approach can be easily extended to games with a mix of atomic and nonatomic players (Harker 1988). We need only to correct the cost function introduced above to take into account that some users are price setters, whereas the rest are price takers. Denoting the former by $[\widetilde{K}]$, the resulting cost function is $\hat{c}_a(\vec{x}_a) := (q_a/2)(\sum_{k \in [K]} x_a^k)^2 + (q_a/2)\sum_{k \in [\widetilde{K}]} (x_a^k)^2 + r_a \sum_{k \in [K]} x_a^k$.

3.1. The Price of Anarchy for General Cost Functions

Using the variational inequality displayed in (2), we can prove a stronger upper bound on the price of anarchy for atomic congestion games. The upper bound we provide below originates in Roughgarden (2005), using ideas from Correa et al. (2004). We define

$$\beta^K(c) := \sup_{\vec{x},\,\vec{y}\in\mathbb{R}_+^K} \frac{\sum_{k\in[K]}\{(c^k(\vec{x})-c(y))y^k+(c(x)-c^k(\vec{x}))x^k\}}{xc(x)},$$

and $\beta^K(\mathscr{C}) := \sup_{c\in\mathscr{C}} \beta^K(c)$. Notice that in the previous definition (and later on when working with β) we overloaded notation slightly because x, $\vec{x}$, y, and $\vec{y}$ represent flows on a single arc instead of representing flows for the entire instance, as usual. Also, in this supremum we implicitly assume that $x=\sum_{k\in[K]} x^k$ and $y=\sum_{k\in[K]} y^k$. For this definition and the ones below to work, we shall assume that $0/0=0$. It is straightforward to see that $\beta^K(\mathscr{C})\geqslant 0$. The magnitude of β is related to the steepness of cost functions. For a geometric interpretation of this calculation in the setting of nonatomic games, we refer the reader to Correa et al. (2008).

We now give a bound on the price of anarchy for games with k players that depends on $\beta^K(\mathscr{C})$. To simplify notation, we will not explicitly distinguish the case of $\beta^K(\mathscr{C})\geqslant 1$ and assume that $(1-\beta^K(\mathscr{C}))^{-1}=+\infty$ in such a case.

Proposition 3.2 (Roughgarden 2005). *Consider an atomic congestion game with K players and separable cost functions drawn from $\mathscr{C}$. Let $\vec{x}^{\text{NE}}$ be a Nash equilibrium and $\vec{x}^{\text{OPT}}$ be a social optimum. Then, $C(x^{\text{NE}})\leqslant (1-\beta^K(\mathscr{C}))^{-1}\,C(x^{\text{OPT}})$.*

Proof. Using (2) and the definition of $\beta^K(\mathscr{C})$ in order, we get that

$$\begin{aligned} C(x^{\text{NE}}) &= \sum_{a\in A}\sum_{k\in[K]} \left\{(c_a(x_a^{\text{NE}})-c_a^k(\vec{x}_a^{\text{NE}}))x_a^{\text{NE},k}+c_a^k(\vec{x}_a^{\text{NE}})x_a^{\text{NE},k}\right\}\\ &\leqslant \sum_{a\in A}\sum_{k\in[K]} \left\{(c_a(x_a^{\text{NE}})-c_a^k(\vec{x}_a^{\text{NE}}))x_a^{\text{NE},k}+c_a^k(\vec{x}_a^{\text{NE}})y_a^k\right\}\\ &\leqslant \beta^K(\mathscr{C})C(x^{\text{NE}})+C(y)\end{aligned}$$

for any solution $\vec{y}$. We finish by setting $\vec{y}=\vec{x}^{\text{OPT}}$. □

Note that this bound on the price of anarchy is also valid for the mixed atomic and nonatomic games. This is because those games can be seen as the limit of atomic games when the number of players goes to infinity. Roughgarden (2005) proved that the price of anarchy under the same situation equals $\alpha^K(\mathscr{C}) := \sup_{c\in\mathscr{C}} \alpha^K(c)$, where

$$\alpha^K(c) := \sup_{\vec{x},\,\vec{y}\in\mathbb{R}_+^K} \frac{xc(x)}{yc(y)+\sum_{k\in[K]}(x^k-y^k)c^k(\vec{x})} \tag{3}$$

for $c\in\mathscr{C}$. The two bounds match because $\alpha^K(\mathscr{C}) = (1-\beta^K(\mathscr{C}))^{-1}$ when $\beta^K(\mathscr{C})\leqslant 1$.

Although Roughgarden (2005) and Correa et al. (2005) independently claimed that the price of anarchy for atomic network games cannot exceed that in nonatomic ones, Figure 2 presents a counterexample. The OD pair on the left is nonatomic, whereas that on the right is controlled by a single player. At Nash equilibrium, the common arc has 1 and 0.9 units of demand coming from the left and right OD pairs, respectively, and the total cost is 3.89. Under the social optimum, the common arc has 0 and 1 units of demand and the total cost is 2.9. Dividing, we get a price of anarchy of approximately 1.341, which is larger than 4/3 (the price of anarchy when players are nonatomic and cost functions are affine).

As a side remark, a Wardrop equilibrium for the same instance incurs a smaller total cost than the Nash equilibrium. Furthermore, the same happens considering only the cost paid by the atomic OD pair. One could argue that the corresponding player may anticipate the response of the nonatomic players and simulate a Wardrop equilibrium to her advantage. This reasoning fails because this behavior for the atomic player is not allowed in a Nash equilibrium because the atomic player has to select a best response to the flow that the nonatomic players choose. In other words, a Wardrop equilibrium is actually better for the atomic player, but it is not a Nash equilibrium of the atomic game.

Figure 2. Example of an instance with affine costs and price of anarchy larger than 4/3.

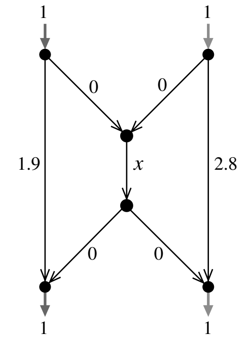

The following remark implies that atomic games are provably harder to coordinate than nonatomic ones. Indeed, the affine case discussed previously is not an anomaly. The price of anarchy for atomic games grows by at least a factor of $\ln b$ faster than that of nonatomic ones, where b denotes the maximum degree of the polynomials that appear as cost functions.

Remark 3.1. Using a structure similar to the previous example, let us consider an instance that consists of two arcs with constant cost, a common arc with cost function equal to x^b, and nonatomic and atomic OD pairs, both with unit demands. Optimizing over the two constant costs, we get lower bounds on the price of anarchy of 1.343, 1.67, 1.981, and 2.287 for polynomials of degree one to four, respectively. Asymptotically, the price of anarchy grows as $\Omega(b)$, in contrast to the price of anarchy for nonatomic games, which grows as $\Theta(b/\ln b)$.

To conclude the example, let us add that it is not necessary to use a nonatomic OD pair. We could have constructed a similar example with a finite number of players.

That would require replacing the nonatomic OD pair by $K-1$ atomic players, each controlling $1/(K-1)$ units of demand. If K is large, both equilibria are similar by continuity (e.g., Haurie and Marcotte 1985 proved that equilibria in atomic games converge to those in nonatomic games when players lose market power). For example, for the affine instance with price of anarchy equal to 1.343, if the nonatomic demand is controlled by 94 or more identical players, the price of anarchy is already larger than 4/3.

We now compute a concrete expression for the price of anarchy under specific sets of cost functions. The key is to first obtain a simpler expression for $\beta^K(c)$.

Theorem 3.1. *The constant $\beta^K(c)$ is at most*

$$\sup_{x,\, y\in\mathbb{R}^2_+} \frac{yc(x)-yc(y)+c'(x)(y^2/4-(x-y/2)^2/K)}{xc(x)}.$$

Proof. Starting from the definition of $\beta^K(c)$, we get

$$\begin{aligned}\beta^K(c) &= \sup_{\vec{x},\,\vec{y}\in\mathbb{R}^K_+} \frac{xc(x)-yc(y)+\sum_{k\in[K]} c^k(\vec{x})(y^k-x^k)}{xc(x)}\\ &= \sup_{\vec{x},\,\vec{y}\in\mathbb{R}^K_+} \frac{yc(x)-yc(y)+c'(x)(\sum_{k\in[K]} y^k x^k-\sum_{k\in[K]}(x^k)^2)}{xc(x)}.\end{aligned} \tag{4}$$

As c is nondecreasing, $c'(x)\geqslant 0$. Thus, assuming w.l.o.g. that $x^1\geqslant x^k$ for all $k\in[K]$, to make (4) as big as possible we have to set $(y^1,\dots,y^K)$ to $(y,0,\dots,0)$. It follows that

$$\beta^K(c)=\sup_{\substack{\vec{x}\in\mathbb{R}^K_+;\, y\in\mathbb{R}_+\\ x=\sum_{k\in[K]}x^k,\\ x^1=\max(\vec{x})}} \frac{yc(x)-yc(y)+c'(x)(x^1y-\sum_{k\in[K]}(x^k)^2)}{xc(x)}. \tag{5}$$

To find the best choice of $\vec{x}$, we fix the total flow x and compute the optimal decomposition. It is enough to solve $\max\{x^1y-\sum_{k\in[K]}(x^k)^2\colon \vec{x}\in\mathbb{R}^K_+,\, x^1=\max(\vec{x})\}$. By symmetry, an optimal solution to this problem satisfies $x^2=\cdots=x^K$. Therefore, we replace x^1 by u and $x^2,\dots,x^K$ by $(x-u)/(K-1)$, and solve

$$\max_{x\geqslant u\geqslant x/K} uy-u^2-\frac{(x-u)^2}{K-1}. \tag{6}$$

This is a concave program so we can conclude that the optimal solution is $u^*=\min\{x, x/K+y(K-1)/2K\}$. Plugging in $x^1=u^*$ and $x^k=\max\{x/K-y/2K,0\}$ for $k=2,\dots,K$ in (5), we have that

$$\begin{aligned}&\beta^K(c)\\ &=\max\left\{\sup_{0\leqslant y\leqslant 2x}\frac{yc(x)-yc(y)+c'(x)(y^2/4-(x-y/2)^2/K)}{xc(x)},\right.\\ &\qquad\left.\sup_{0\leqslant 2x\leqslant y}\frac{yc(x)-yc(y)+c'(x)(xy-x^2)}{xc(x)}\right\}\\ &\leqslant \sup_{x,\,y\geqslant 0}\frac{yc(x)-yc(y)+c'(x)(y^2/4-(x-y/2)^2/K)}{xc(x)}.\end{aligned}$$

The last inequality follows by removing the constraint $x\geqslant u$ from (6). □

Corollary 3.1. *If $xc(x)$ is a convex function, then*

$$\beta^K(c)=\sup_{0\leqslant y\leqslant x}\frac{yc(x)-yc(y)+c'(x)(y^2/4-(x-y/2)^2/K)}{xc(x)}.$$

Moreover, defining

$$\beta^\infty(c):=\sup_{0\leqslant y\leqslant x}\left(c(x)-c(y)+\frac{c'(x)y}{4}\right)\frac{y}{xc(x)},$$

we have that $\beta^K(c)\leqslant\beta^\infty(c)$.

Proof. Consider the function $h(y)$ defined as the numerator inside the supremum of (5). To prove that the solution satisfies $y\leqslant x$, we will show that $h'(y)\leqslant 0$ if $y\geqslant x$. Because $h'(y)=c(x)-c(y)-yc'(y)+x^1c'(x)$, the derivative is negative if and only if $c(x)+x^1c'(x)\leqslant c(y)+yc'(y)$. Because $xc(x)$ is convex, its derivative is increasing, implying that $c(x)+xc'(x)\leqslant c(y)+yc'(y)$ for $x\leqslant y$. The first claim follows because $x^1\leqslant x$. The second claim is straightforward because a square is nonnegative. □

The definition of the constant $\beta^\infty(\mathscr{C})$ is very similar to that of

$$\beta(\mathscr{C}):=\sup_{0\leqslant y\leqslant x}\frac{y(c(x)-c(y))}{xc(x)},$$

which provides a bound on the price of anarchy for nonatomic games (Correa et al. 2004). The only difference between the two expressions is the last term in the numerator of $\beta^\infty(\mathscr{C})$, which penalizes equilibria in the case of atomic players.

3.2. Computing the Price of Anarchy

In this section, we show how to evaluate $\beta^\infty(\mathscr{C})$. We start with a rough estimate of $\beta^\infty(\mathscr{C})$, and we continue by computing it exactly. In particular, that allows us to conclude that the price of anarchy is at most 3/2, 2.464, and 7.826, for affine, quadratic, and cubic cost functions, respectively.

Roughgarden (2002) and Correa et al. (2007) used the constant

$$\gamma(\mathscr{C}):=\sup\{1+c'(x)x/c(x)\colon c\in\mathscr{C},\ x\in\mathbb{R}_+\}$$

to bound the unfairness of a socially optimal solution. Notice that $\gamma(\mathscr{C})$ can be easily computed for given sets of cost functions. For example, $\gamma(\text{degree-}b\text{ polynomials})=b+1$.

Corollary 3.2. *The constant $\beta^\infty(\mathscr{C})$ satisfies that*

$$\frac{\gamma(\mathscr{C})-1}{4}\leqslant\beta^\infty(\mathscr{C})\leqslant\frac{\gamma(\mathscr{C})-1}{4}+\beta(\mathscr{C}).$$

Proof. The lower bound arises from evaluating $\beta^\infty(\mathscr{C})$ in $x = y$. The upper bound holds because the supremum of a sum is bounded by the sum of the suprema. $\square$

The following bound provides a tighter expression for common sets of allowable cost functions such as polynomials of a fixed degree.

Proposition 3.3. *Let $\mathscr{C}$ be a family of continuous and nondecreasing cost functions c that satisfy that $xc(x)$ is convex. Furthermore, assume that $c(rx) \geqslant s(r)c(x)$ for all $r \in [0, 1]$, where s: $[0, 1] \to [0, 1]$ is a differentiable function satisfying $s(1) = 1$. Then,*

$$\beta^\infty(c) \leqslant \max_{0\leqslant u\leqslant 1} u\left(1 - s(u) + s'(1)\frac{u}{4}\right).$$

Proof. Let us first bound $c'(x)$ using that for any two values $z \geqslant z'$, $c(z') = c((z'/z)z) \geqslant s(z'/z)c(z)$. We have that

$$c'(x) = \lim_{\varepsilon\downarrow 0} \frac{c(x+\varepsilon) - c(x)}{\varepsilon} \leqslant c(x) \lim_{\varepsilon\downarrow 0} \frac{1 - s(x/(x+\varepsilon))}{\varepsilon} = c(x)\frac{s'(1)}{x},$$

where the last equality follows by applying l'Hôpital's rule. Therefore,

$$\begin{aligned}\beta^\infty(c) &\leqslant \sup_{0\leqslant y\leqslant x} \frac{yc(x)(1 - c(y)/c(x) + s'(1)y/(4x))}{xc(x)}\\ &\leqslant \sup_{0\leqslant y\leqslant x} \frac{y(1 - s(y/x) + s'(1)y/(4x))}{x}\\ &= \max_{0\leqslant u\leqslant 1} u\left(1 - s(u) + \frac{s'(1)u}{4}\right). \quad \square\end{aligned}$$

Corollary 3.3. *If $\mathscr{C}$ only contains polynomials of degree at most b, the price of anarchy is at most*

$$\left(1 - \max_{0\leqslant u\leqslant 1} u(1 - u^b + bu/4)\right)^{-1}.$$

Proof. The assumption of Proposition 3.3 is now satisfied with $s(x) = x^b$. Therefore, $s'(1) = b$ and $\beta^\infty(c) \leqslant \max_{0\leqslant u\leqslant 1} u(1 - u^b + bu/4)$. $\square$

Using Corollary 3.3, we can determine that our bound on the price of anarchy when cost functions are affine equals $3/2$. For the case of quadratic or cubic polynomials, it is approximately 2.564 or 7.826, respectively. For polynomials of degree 4, our bound evaluates to infinity. This stands in contrast with the situation of nonatomic games in which the price of anarchy grows as $b/\ln b$, where b is the degree of the polynomials. Using Corollary 3.2, it is straightforward to see that $b/4 \leqslant \beta^\infty(\text{degree-}b \text{ polynomials}) \leqslant 1 + b/4$. This follows from the definition of $\gamma(\mathscr{C})$ and $0 \leqslant \beta(\mathscr{C}) \leqslant 1$. Notice that although there is a gap between the lower and upper bound for the price of anarchy, the value of $\beta^\infty(\mathscr{C})$ that we computed is exact. The following proposition provides a bound that depends on the range that the derivative of the cost functions is allowed to take.

Proposition 3.4. *Suppose that*

$$\max_{x\in\mathbb{R}_+} c'(x) \leqslant \gamma \min_{x\in\mathbb{R}_+} c'(x)$$

for a given $\gamma \geqslant 1$. Then, $\beta^\infty(c) \leqslant \gamma/3$.

Proof. Denote $\min_{x\in\mathbb{R}_+} c'(x)$ by $\underline{c}$ and $\max_{x\in\mathbb{R}_+} c'(x)$ by $\bar{c}$. Using the Mean Value Theorem, and the convexity and nonnegativity of c, it is easy to see that $c(x) - c(y) \leqslant (x - y)\bar{c}$ and that $c(x) \geqslant x\underline{c}$. Replacing the values in the bound from Corollary 3.1, we get that

$$\begin{aligned}\beta^\infty(c) &\leqslant \sup_{0\leqslant y\leqslant x} \frac{y((x-y)\bar{c} + \bar{c}y/4)}{x^2\underline{c}}\\ &\leqslant \gamma \sup_{0\leqslant y\leqslant x} \frac{yx - 3y^2/4}{x^2} = \gamma/3. \quad \square\end{aligned}$$

As an example, the last proposition can be used to show that if, for all $c \in \mathscr{C}$, the largest derivative is not bigger than twice the smallest derivative, then the price of anarchy is bounded by three. It also provides another proof of Corollary 3.3 in the case of affine cost functions because the ratio of the derivatives is equal to one.

Now we analyze the dependence of the price of anarchy on the number of players K. For $K = 1$, the single player computes a social optimum, and therefore one should expect that if the bound provided by Proposition 3.2 is tight, then $\alpha^1(\mathscr{C})$ should be one. This is the case, in fact, when cost functions are convex. The following proposition establishes the price of anarchy for K players and affine cost functions. We do not include a proof because it simply consists of technical calculations.

Proposition 3.5. *If there are K players and the allowable cost functions are affine, the price of anarchy is bounded from above by $\alpha^K(\text{affine}) = (3K + 1)/(2K + 2)$.*

The previous proposition implies that when $K \to \infty$, $\alpha^K(\text{affine}) \to \alpha^\infty(\text{affine}) = 3/2$. In particular, for $K > 5$, the upper bound is larger than $\alpha(\text{affine}) = 4/3$, the price of anarchy for nonatomic games. Recall that for the example that we presented before, we needed approximately 100 players to achieve a price of anarchy larger than $4/3$.

3.3. Pseudoapproximations

We now concentrate on pseudoapproximation results (also known as bicreteria results) that compare the Nash equilibrium to a social optimum in an instance with expanded demands (Roughgarden and Tardos 2002). The main motivation of this comparison is that a large coordinating power is required to achieve a social optimum, but an equilibrium arises naturally and without any coordination. To balance this difference, we impose more costs to the socially optimal solution by increasing its demand. We measure the quality of equilibria by determining how much more demand is needed to make the social costs of the two solutions equal. If a large expansion factor for the demand is

Figure 3. Example with pseudoapproximation guarantee larger than 2.

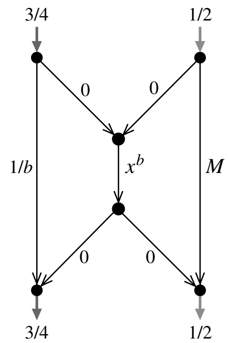

needed, it means that equilibria are inefficient. These results can also be interpreted as a way to compute the expansion of capacity that is needed to offset the lack of coordination in the network. Roughgarden and Tardos (2002) proved that the social cost of a Wardrop equilibrium is bounded by that of a social optimum of a game with demands doubled. This means that instead of trying to coordinate the network, one can double the capacity of all arcs and achieve a similar social cost.

Roughgarden and Tardos (2002) extended the pseudoapproximation bound to atomic games. This extension was based on a characterization of equilibria of atomic congestion games that they proposed. Unfortunately, this characterization is not correct, and hence the results in Correa et al. (2005) that used it are not valid. Figure 3 presents an example for which the Nash equilibrium is more costly than the social optimum with demands doubled. The OD pair on the left is nonatomic, whereas the one on the right is atomic. Consider $M := (1-\varepsilon)^b + b(1/4-\varepsilon)(1-\varepsilon)^{b-1}$, where ε is such that $(1-\varepsilon)^b < 1/b$. The parameters M and ε are chosen so that the Nash equilibrium is the flow in which the nonatomic demand routes all its 3/4 units of flow in the middle arc and the atomic player splits its flow in $1/4-\varepsilon$ along the middle arc and the rest in the other. The social cost of the equilibrium equals $(1-\varepsilon)^{b+1} + (1/4+\varepsilon)M$. Consider the flow that routes twice the demand in which 3/2 units of flow take the left arc, $1-\varepsilon$ units take the middle arc, and ε units take the right arc. Therefore, the social cost of the social optimum is at most $\varepsilon M + (1-\varepsilon)^{b+1} + 3/(2b)$. Comparing the two costs, we conclude that to find a counterexample, we need to find b and ε such that $b(1-\varepsilon)^b < 1$ and $Mb/6 > 1$. This is achieved by taking $\varepsilon = 0.1$ and $b = 34$. Modifying the example slightly, we can obtain a counterexample with polynomials of degree 26. On the other hand, if we allow polynomials of arbitrary degree, it can be seen that the cost of the Nash equilibrium can be made arbitrarily higher than that of the social optimum with demands doubled.

In addition, one cannot expect to prove a theorem of this type with a constant expansion factor if arbitrary cost functions are allowed. To see this, consider the same example as in Figure 3 and a parameter $0 < \delta < 1$. The nonatomic demand is $1-\delta$, the demand of the atomic player is 2δ, and the cost functions, from left to right, are 0, a step function that is 0 for $x \leqslant 1$ and 1 otherwise, and 2. It can be seen that there is one equilibrium with total cost equal to 2δ, whereas the social optimum when the demand is amplified by $1/(2\delta)$ has zero cost. The example can be worked out for polynomial cost functions (of arbitrary high degree). The previous discussion leads us to the following result.

Proposition 3.6. *Let $\vec{x}^{\text{NE}}$ be a Nash equilibrium and, for an arbitrary $\alpha > 1$, let $\vec{x}^{\text{OPT}}$ be a social optimum of the game when demands are multiplied by α. Then, there exists an instance of the atomic network game with convex and increasing cost functions such that $C(x^{\text{NE}}) > C(x^{\text{OPT}})$.*

In view of the previous negative results, we now prove a pseudoapproximation result for atomic games that hinges on ideas of Correa et al. (2008). The following proposition provides a bound that depends on the allowable cost functions $\mathscr{C}$. For example, in the case of affine cost functions, an expansion factor equal to 4/3 makes the social cost of an equilibrium be bounded by that of the expanded social optimum.

Proposition 3.7. *Let $\vec{x}^{\text{NE}}$ be a Nash equilibrium of an atomic congestion game with K players and with separable cost functions drawn from $\mathscr{C}$. If $\vec{x}^{\text{OPT}}$ denotes a social optimum of the game with demands multiplied by $1+\beta^K(\mathscr{C})$, then $C(x^{\text{NE}}) \leqslant C(x^{\text{OPT}})$.*

Proof. Let $\vec{y}$ be a flow that routes $(1+\beta^K(\mathscr{C}))d_k$ units of demand from s_k to t_k for $k \in [K]$. Then,

$$\begin{aligned} C(x^{\text{NE}}) &= (1+\beta^K(\mathscr{C})) \sum_{a\in A}\sum_{k\in[K]} \{(c_a(x_a^{\text{NE}}) - c_a^k(\vec{x}_a^{\text{NE}}))x_a^{\text{NE},k} \\ &\qquad + c_a^k(\vec{x}_a^{\text{NE}})x_a^{\text{NE},k}\} - \beta^K(\mathscr{C})C(x^{\text{NE}}) \\ &\leqslant (1+\beta^K(\mathscr{C})) \sum_{a\in A}\sum_{k\in[K]} \Bigg\{(c_a(x_a^{\text{NE}}) - c_a^k(\vec{x}_a^{\text{NE}}))x_a^{\text{NE},k} \\ &\qquad + \frac{c_a^k(\vec{x}_a^{\text{NE}})\,y_a^k}{1+\beta^K(\mathscr{C})}\Bigg\} - \beta^K(\mathscr{C})C(x^{\text{NE}}), \end{aligned}$$

where the inequality follows using (2) with $y_a^k/(1+\beta^K(\mathscr{C}))$. As $c_a(x_a^{\text{NE}}) - c_a^k(\vec{x}_a^{\text{NE}}) \leqslant 0$,

$$\begin{aligned} C(x^{\text{NE}}) &\leqslant \sum_{a\in A}\sum_{k\in[K]} \{(c_a(x_a^{\text{NE}}) - c_a^k(\vec{x}_a^{\text{NE}}))x_a^{\text{NE},k} + c_a^k(\vec{x}_a^{\text{NE}})y_a^k\} \\ &\qquad - \beta^K(\mathscr{C})C(x^{\text{NE}}) \\ &\leqslant \beta^K(\mathscr{C})C(x^{\text{NE}}) + C(y) - \beta^K(\mathscr{C})C(x^{\text{NE}}) = C(y). \end{aligned}$$

The proof follows by evaluating in $\vec{y} = \vec{x}^{\text{OPT}}$. $\square$

Remark 3.2. Note that the example above shows that $\beta^\infty(\mathscr{C})$ is unbounded for general cost functions (continuous and convex).

4. Atomic Games with a Single OD Pair

In this section, we concentrate on atomic games played on networks with arbitrary topology in which all K players share the same source s and sink t. This context is particularly relevant in settings in which goods are produced in one place (e.g., Asia) and consumers are located in another one (e.g., the United States). In another example that we present in §6, we consider a flexible manufacturing environment. Here, production involves a series of operations starting from raw materials that are assembled into finished goods by different producers. This process can be modeled using a congestion game with units that start from a state corresponding to raw material and that evolve into a state corresponding to finished goods.

Single-source single-sink instances are easier to analyze because the same set of paths is available to all players. This fact will allow us to provide improved results compared to the general case. We consider two alternatives: either players control arbitrary fractions of the market, or players are symmetric. This type of games and the two alternatives have also been considered by Orda et al. (1993), although they restricted the network topology to be of parallel arcs.

The presentation is divided into two sections: In the first, we consider the case in which different players control different amounts of demand, resulting in different market shares. We prove an upper bound on the price of anarchy that depends on the variability of the market power of the different players. To the best of our knowledge, this is the first known bound of this type. In the second part, we consider the case of symmetric players in which all players have the same demand to route through the network. This symmetry allows us to provide improved results and to rule out counterintuitive phenomena such as a paradox described by Catoni and Pallotino (1991). We provide improved results for the particular case of affine cost functions in the appendix.

4.1. Variable Market Power

We consider the case in which different players control different amounts of demand, leading to different levels of market power. Our main result is a bound on the price of anarchy that depends on the variability of market power across players. To that extent, we use the *Herfindahl index* which is a standard measure of industry concentration. We denote it by $H := \sum_{k\in[K]}(d_k/D)^2$, where $D := \sum_{k\in[K]} d_k$ is the total demand. This index is a number between $1/K$ and 1. A higher index means that the market is less competitive, and the case of $H = 1$ corresponds to a monopoly. The case in which $H = 1/K$ corresponds to instances with symmetric players (see the next section).

The following proposition combines Theorem 3.1 and Corollary 3.1 and achieves a better bound by reinterpreting the definition of $\beta^K(c)$.

Proposition 4.1. *Consider an instance with a single OD pair and Herfindahl index equal to H. Letting*

$$\tilde{\beta}(c, H) := \sup_{0\leqslant y\leqslant x} \frac{y(c(x) - c(y) + c'(x)yH/4)}{xc(x)},$$

and $\tilde{\beta}(\mathscr{C}, H) := \sup_{c\in\mathscr{C}} \tilde{\beta}(c, H)$, we have that $C(x^{\mathrm{NE}}) \leqslant (1 - \tilde{\beta}(\mathscr{C}, H))^{-1} C(x^{\mathrm{OPT}})$.

Proof. Looking at the proof of Proposition 3.2, the constant $\beta^K(c)$ can be interpreted as the minimum number for which the last inequality of the proof holds. Because that inequality does not depend at all on the decomposition of y into y^k, we can set y^k in the way that is most convenient instead of using its worst realization, as we have done in (5). The only restriction in setting y_a^k is that when we sum the inequalities derived from each arc, y^k has to be a feasible flow for player k. This can be easily done in the case of a single OD pair by decomposing y proportionally to the demand of each player, i.e., as $(d_k y/D)_{k\in[K]}$. The claim follows after solving the supremum in (4) with the new decomposition of y, and redoing the proof of Corollary 3.1. □

Providing bounds of this type for multiple OD pairs is an interesting question that our work leaves open. Our techniques do not easily extend to multiple OD pairs because it is not clear how to create a feasible flow arc by arc. Nevertheless, §6 outlines a generalization in that direction.

The difference compared to the expression provided by Corollary 3.1 is the factor H in the last term of the numerator. Observe that as $H \leqslant 1$, this result can only reduce the price of anarchy. Moreover, if each player controls at most a fraction $\phi(K)$ of the demand such that $\phi(K) \to 0$ when $K \to \infty$, the price of anarchy is asymptotically equal to that in the nonatomic game. Indeed, the worst case for the market power variability is that there are $1/\phi(K)$ players, each controlling a fraction $\phi(K)$ of the demand, whereas the rest of the players control an infinitesimal. In that case, $H \leqslant (1/\phi(K))\phi(K)^2 = \phi(K) \to 0$. For example, in an oligopoly with K players that control a total demand equal to K, but in which $K/\ln K$ players control $\ln K$ units of demand each and the rest of the players do not have market power, the analysis above shows that this oligopoly approaches the nonatomic game when K grows.

Proposition A1 in the appendix shows that the price of anarchy in the case of affine cost functions is at most $(4-H)/(3-H)$. This generalizes that the price of anarchy is equal to $4/3$ for nonatomic games ($H = 0$) and at most $3/2$ in general (arbitrary H). Nevertheless, we know that when $H = 1$ the price of anarchy equals 1. By perturbing the monopolistic case, we can show that the price of anarchy for the case of a single OD pair is strictly less than $3/2$. However, this analysis is quite technical, and it is unlikely to provide a bound that is tight.

4.2. Symmetric Players

When all players have the same demand d to route through the network, Orda et al. (1993) showed that there is a unique Nash equilibrium. Our first contribution in this section is to provide a convex optimization problem whose optimum is the unique equilibrium. This implies that the game with symmetric players is a potential game. To facilitate notation, we add a reverse arc between t and s with zero cost:

$$\text{(SNE)}\quad \min \sum_{a\in A} x_a c_a(x_a) + (K-1)\sum_{a\in A}\int_0^{x_a} c_a(\tau)\,d\tau$$

$$\text{s.t.}\quad \sum_{(u,v)\in A} x_{(u,v)} - \sum_{(v,w)\in A} x_{(v,w)} = 0 \quad \text{for all } v\in V,$$

$$x_{(t,s)} = dK,$$

$$x_a \geqslant 0 \quad \text{for all } a\in A.$$

Interestingly, problem (SNE) consists of finding a feasible flow that minimizes a convex combination between the objective functions of the problems used to compute a social optimum and a Nash equilibrium of a nonatomic game. When there is a single player, the second part vanishes, leaving only the social cost. Instead, when there are many players the second part is dominant and the social cost becomes negligible. The next result shows that a solution is optimal for problem (SNE) if and only if it is a Nash equilibrium. Therefore, if the cost functions are strictly increasing, there is exactly one Nash equilibrium. Additionally, we make use of the potential function to derive results on the efficiency of equilibria.

Theorem 4.1. *If x solves problem (SNE), then $\vec{x}^{\text{NE}} = (x/K, \ldots, x/K)$ is a Nash equilibrium of the symmetric game with atomic players.*

Proof. Because problem (SNE) is a convex program, the Karush-Kuhn-Tucker conditions say that x is an optimal solution if and only if it is a feasible flow satisfying

$$Kc_a(x_a) + x_a c_a'(x_a) = \lambda_u - \lambda_v + \mu_a \quad \text{for all } a=(u,v)\in A,$$

$$0 = \lambda_t - \lambda_s + \lambda_{(t,s)},$$

$$\mu_a x_a = 0 \quad \text{for all } a\in A,$$

$$\mu \geqslant 0.$$

By letting $x^{\text{NE},k} = x/K$, $\lambda^k = \lambda/K$ and $\mu^k = \mu/K$, and by dividing all previous equations by K, we obtain that $x^{\text{NE},k}$ is feasible for problem (NE^k), and it satisfies

$$c_a(x_a^{\text{NE}}) + x_a^{\text{NE},k} c_a'(x_a^{\text{NE}}) = \lambda_u^k - \lambda_v^k + \mu_a^k \quad \text{for all } a=(u,v)\in A,$$

$$0 = \lambda_t^k - \lambda_s^k + \lambda_{(t,s)}^k,$$

$$\mu_a^k x_a^{\text{NE},k} = 0 \quad \text{for all } a\in A,$$

$$\mu^k \geqslant 0,$$

which are exactly the Karush-Kuhn-Tucker conditions corresponding to problem (NE^k). □

We now use Theorem 4.1 to derive results on the efficiency of equilibria for symmetric network games with atomic players.

Proposition 4.2. *Let $\vec{x}\in\mathbb{R}_+^K$ be a Nash equilibrium in an atomic game with K players who control d units of flow each, and let $\vec{y}\in\mathbb{R}_+^{\tilde{K}}$ be a Nash equilibrium in an atomic game with $\tilde{K} < K$ players who control $dK/\tilde{K}$ units of flow each. Then, $C(y)\leqslant C(x)$.*

Proof. Using the optimality of x and y in their respective problems as before,

$$\begin{aligned}
&\sum_{a\in A} x_a c_a(x_a) + (K-1)\sum_{a\in A}\int_0^{x_a} c_a(\tau)\,d\tau \\
&\quad\leqslant \sum_{a\in A} y_a c_a(y_a) + (K-1)\sum_{a\in A}\int_0^{y_a} c_a(\tau)\,d\tau \\
&\quad\leqslant \sum_{a\in A} x_a c_a(x_a) + (\tilde{K}-1)\sum_{a\in A}\int_0^{x_a} c_a(\tau)\,d\tau \\
&\qquad + (K-\tilde{K})\sum_{a\in A}\int_0^{y_a} c_a(\tau)\,d\tau.
\end{aligned}$$

Thus, $\sum_{a\in A}\int_0^{x_a} c_a(\tau)\,d\tau \leqslant \sum_{a\in A}\int_0^{y_a} c_a(\tau)\,d\tau$, which implies that

$$\begin{aligned}
&\sum_{a\in A} y_a c_a(y_a) + (\tilde{K}-1)\sum_{a\in A}\int_0^{y_a} c_a(\tau)\,d\tau \\
&\quad\leqslant \sum_{a\in A} x_a c_a(x_a) + (\tilde{K}-1)\sum_{a\in A}\int_0^{x_a} c_a(\tau)\,d\tau \\
&\quad\leqslant \sum_{a\in A} x_a c_a(x_a) + (\tilde{K}-1)\sum_{a\in A}\int_0^{y_a} c_a(\tau)\,d\tau. \quad \square
\end{aligned}$$

The previous proposition implies that the price of anarchy in symmetric games with K players increases as the number of players increases. Going to the limit when $K\to\infty$, we get the following corollary.

Corollary 4.1. *The social cost of a Nash equilibrium in an atomic game is bounded by that of the Wardrop equilibrium in the corresponding nonatomic game.*

Hence, when the number of players goes to infinity, the price of anarchy approaches that in the nonatomic case. The conclusion is that when players are completely symmetric, the ability to set prices does not degrade the quality of equilibria with respect to price-taking players. This stands in clear contrast to the case of atomic asymmetric games whose price of anarchy is larger than that of nonatomic games. Proposition A2 in the appendix evaluates the price of anarchy of symmetric games as a function of the number of players K when cost functions are affine. As expected, it tends to 4/3 as the number of players grows.

The results for symmetric players can be generalized to the asymmetric case with a single OD pair if we assume that for each arc either all players have a positive flow on it, or no player uses it. Orda et al. (1993, p. 518) referred to this assumption by "all-positive flows," and proved that in this case there is a unique Nash equilibrium. Specifically, if we consider the KKT conditions of each of the player problems under an equilibrium and sum them up, we get the KKT conditions of problem (SNE). This implies that, on every arc, the total flow induced by an equilibrium matches the optimal solution of (SNE), and thus the results we have presented in this section carry over to this setting. Note, however, that the equilibrium flows do not necessarily consist of the proportional decomposition of the total flow, as was the case for symmetric players.

5. Pricing Mechanisms

In this section, we consider that players are charged a per-unit price when sending flow through taxed arcs. These charges are levied by the system administrator, designer or regulator with the sole purpose of encouraging coordination among the participants of the game. This approach extends previous work in the setting of nonatomic congestion games. Because prices are payments inside the system, they do not change the functional form of the social cost.

If it is possible to charge different prices to different users (price discrimination), it is always feasible to achieve a fully efficient solution. It is enough to compute a social optimum and charge player $k \in [K]$ the difference between its marginal cost and the real cost. Indeed, consider a social optimum x^{OPT} with an arbitrary flow decomposition by player. Because x^{OPT} satisfies the first-order optimality conditions of its optimization problem, the following inequality holds for all players:

$$\sum_{a\in A}[c_a(x_a^{\text{OPT}}) + x_a^{\text{OPT}} c_a'(x_a^{\text{OPT}})](y_a^k - x_a^{\text{OPT},k}) \geqslant 0 \qquad \text{for any feasible flow } y^k \text{ for player } k. \quad (7)$$

Notice the resemblance between this inequality and (2). Actually, if player k in arc a is charged a nonnegative unit price of $\rho_a^k := (x_a^{\text{OPT}} - x_a^{\text{OPT},k})c_a'(x_a^{\text{OPT}})$, player k's total unit cost in arc a will be $c_a(x_a) + \rho_a^k$. Plugging that cost function into (2), we see that the inequality coincides with (7), implying that $\vec{x}^{\text{OPT}}$ is at equilibrium.

We now present some other cases in which we can guarantee full efficiency without price discrimination. The case of a nonatomic game, which was first analyzed by Beckmann et al. (1956), is included in the analysis above. Indeed, the flow controlled by any player is infinitesimally small, implying that $\rho_a := x_a^{\text{OPT}} c_a'(x_a^{\text{OPT}})$. This case results in uniform prices because the term that discriminates players vanishes. The case of a symmetric game can be handled in a similar way: Using the flow decomposition $x^{\text{OPT},k} = x^{\text{OPT}}/K$ for all $k \in [K]$, we get that $\rho_a := (1-1/K)x_a^{\text{OPT}} c_a'(x_a^{\text{OPT}})$.

If one is not allowed to price discriminate, it is not clear that achieving an optimal solution is possible. For a general network topology and market structure, we find a set of prices for each arc that reduces the price of anarchy. Denoting the tax that we add to the cost function on each arc a by ρ_a, we want to find the taxes ρ_a^* that minimize the price of anarchy. To this end, we need to redefine $\beta^K(c)$ as follows:

$$\beta^K(c,\rho) := \sup_{\vec{x},\vec{y}\in\mathbb{R}_+^K} \frac{\sum_{k\in[K]}\{(c^k(\vec{x}) - c(y) + \rho)y^k + (c(x) - c^k(\vec{x}) - \rho)x^k\}}{xc(x)},$$

and $\beta^K(\mathscr{C},\rho) := \sup_{c\in\mathscr{C}} \beta^K(c,\rho)$. With this, we have a new version of Proposition 3.2.

Proposition 5.1. *Consider an atomic congestion game with K players, separable cost functions c_a, and prices ρ_a. Let $\vec{x}^{\text{NE}}$ be a Nash equilibrium with tolls and $\vec{x}^{\text{OPT}}$ be a social optimum without tolls. Then,*

$$C(x^{\text{NE}}) \leqslant \big(1 - \max_{a\in A} \beta^K(c_a,\rho_a)\big)^{-1} C(x^{\text{OPT}}).$$

Proof. Using (2) and the definition of $\beta^K(\mathscr{C})$ in order, we get that

$$\begin{aligned}
C(x^{\text{NE}}) &= \sum_{a\in A}\sum_{k\in[K]}\big\{(c_a(x_a^{\text{NE}}) - (c_a^k(\vec{x}_a^{\text{NE}}) + \rho_a))x_a^{\text{NE},k} \\
&\qquad + (c_a^k(\vec{x}_a^{\text{NE}}) + \rho_a)x_a^{\text{NE},k}\big\} \\
&\leqslant \sum_{a\in A}\sum_{k\in[K]}\big\{(c_a(x_a^{\text{NE}}) - c_a^k(\vec{x}_a^{\text{NE}}) - \rho_a)x_a^{\text{NE},k} \\
&\qquad + (c_a^k(\vec{x}_a^{\text{NE}}) + \rho_a)y_a^k\big\} \\
&\leqslant \max_{a\in A} \beta^K(c_a,\rho_a) C(x^{\text{NE}}) + C(y),
\end{aligned}$$

for any solution $\vec{y}$. We finish by setting $\vec{y} = \vec{x}^{\text{OPT}}$. □

Proceeding as in Proposition 3.7, we can also get the following:

Proposition 5.2. *Let $\vec{x}^{\text{NE}}$ be a Nash equilibrium of an atomic congestion game with K players, separable cost functions c_a, and prices ρ_a. If $\vec{x}^{\text{OPT}}$ denotes a social optimum of the game with demands multiplied by $1 + \max_{a\in A}\beta^K(c_a,\rho_a)$, then $C(x^{\text{NE}}) \leqslant C(x^{\text{OPT}})$.*

Motivated by the symmetric case, we propose that the price charged in every link is $\rho_a(\eta) := \eta x_a^{\text{OPT}} c_a'(x_a^{\text{OPT}})$, where $\eta \in [0,1]$ is a constant that is going to be chosen to minimize the price of anarchy.

Proposition 5.3. *If $\mathscr{C}$ is the set of nonnegative polynomials of degree at most b, the optimal price for this mechanism is $\eta^* = \min\{1/2, 1/b\}$. The corresponding price of anarchy is bounded by $\beta^K(\mathscr{C},\rho(\eta^*)) \leqslant \max_{0\leqslant u\leqslant 1} u(1 - u^{b-1} + bu/4)$.*

Proof. Starting with a general η and proceeding as in §3.1, we have that for $c \in \mathscr{C}$:

$$\begin{aligned}
&\beta^K(c,\rho(\eta)) \\
&\quad \leqslant \beta^\infty(c,\rho(\eta)) \\
&\quad := \sup_{x,y\in\mathbb{R}_+} \frac{y(c(x) - c(y) + c'(x)y/4 + \eta c'(y)(y-x))}{xc(x)}. \quad (8)
\end{aligned}$$

We can assume without loss of generality that cost functions are monomials of degree at most b (otherwise, subdivide each arc into multiple arcs with a monomial each). Hence, we only need to compute $\beta^\infty(ax^b, \rho(\eta))$ because lower-degree monomials have a smaller price of anarchy. Using (8) and the change of variables $u := y/x$, $\beta^K(\mathscr{C}, \rho(\eta)) \leqslant \sup_{u\in\mathbb{R}_+} u(1+(1-\eta b)u^b - \eta b u^{b-1} + bu/4)$. For convenience, let us call the argument of the supremum $h_\eta(u)$. A requirement for the supremum to be bounded is that the highest-degree monomial of $h_\eta(\cdot)$ has a nonpositive coefficient. Therefore, optimal prices are achieved at $\eta^* \leqslant \min\{3/4, 1/b\}$.

We will concentrate first on the case $b > 1$, which implies that $\eta^* \leqslant 1/2$. Evaluating, the derivative $h'_\eta(1)$ is $b(\eta - 1/2) \leqslant 0$, and the second derivative $h''_\eta(u)$ is $(b+1)b(\eta b - 1)u^{b-1} - \eta b^2(b-1)u^{b-2} + b/2$. This implies that $h'_\eta(u)$ decreases for $u \geqslant 1$, and thus the optimal value for u verifies $0 \leqslant u \leqslant 1$. Finally, as $\partial h_\eta(u)/\partial\eta = bu^b(u-1) \leqslant 0$ in $0 \leqslant u \leqslant 1$, we conclude that for larger values of η, $h_\eta(u)$ decreases. Hence, we should set η^* to its largest possible value $1/b$. Plugging this value into $h_\eta(\cdot)$, we get the claim.

The only case we have not yet considered is the affine one. If $0 \leqslant \eta \leqslant 1/2$, reasoning as before, we conclude that $\eta^* = 1/2$ and the optimal value of u satisfies $0 \leqslant u \leqslant 1$. When $1/2 < \eta \leqslant 3/4$, as $h_\eta(1) = b/4$ and $h'_\eta(1) > 0$, the supremum is strictly higher than that with $\eta = 1/2$, so this case does not provide an optimal value for η. □

In Table 1, we compute values of β and bounds for the price of anarchy with optimal prices and without pricing. Notice that the pricing mechanism reduces both values. Moreover, in the affine case the pricing mechanism is able to reduce the price of anarchy to the level of nonatomic games.

6. Concluding Remarks

We now discuss some possible extensions of the model we have presented. Although we have assumed that each player routes flow only from a single origin to a single destination, this can be relaxed. Equation (2) still holds when each player has to route flow from multiple origins to multiple destinations, implying that our bounds on the price of anarchy hold too. In §4, this means that instead of a single OD pair, there may be many OD pairs. Section 4.1 requires that market share of player k in OD pair s-t (i.e., $d^k_{st}/\sum_{j\in[K]} d^j_{st}$, where d^k_{st} is the demand controlled by player k in OD pair s-t) is the same throughout all OD pairs. Instead, §4.2 requires that $d^k_{st} = d^j_{st}$ for all players k and j and all OD pairs.

After the publication of a preliminary version of this article (Cominetti et al. 2006), there has been some work related to atomic network games with splittable flow. Harks (2008) strengthened our upper bounds on the price of anarchy for general networks and for nonlinear polynomials of bounded degree. His improvement arises from introducing and optimizing upon another free variable in the definition of $\beta(\mathscr{C})$. In addition, using the framework developed by Fleischer et al. (2004), Swamy (2007) and Yang and Zhang (2008) proved that tolls that induce a socially optimal flow always exist and can be computed efficiently.

6.1. Further Applications

We conclude by presenting other examples that fit the abstract model we have presented, beyond competition in the setting of freight transportation. Intelligent transportation systems (ITS) provide users with information about travel options and allow them to make informed travel decisions. Eventually, ITS could be used to provide route guidance services to users. These services will not only provide information about traffic network conditions, but may also provide a user with detailed guidance from her current position to her final destination. This situation naturally fits our model. Route guidance service providers are atomic players (they have market power as they control a nonnegligible fraction of the cars) who strive to minimize the overall travel time of their clients. The rest of the users in the transportation network are nonatomic players because they make their decisions independently. Although route assignments that achieve minimal travel time may route some users on excessively long paths, the overall adverse effect from the user perspective is small because users are assigned to paths randomly (see also Jahn et al. 2005 for a route guidance model that specifically addresses this issue). Hence, users will improve the travel time in expectation because the likelihood of being assigned to a long path is insignificant. Interestingly, it has been frequently mentioned that ITS promises to improve the usage of the existing road network infrastructure and to help manage congestion. Our results seem to indicate that route guidance systems that minimize total delay may not always improve users' performance. It

Table 1. Bounds on the price of anarchy for polynomials of degree up to b with and without pricing.

		Optimal pricing		No pricing		Nonatomic	
b	η^*	$\beta^\infty(\mathscr{C}, \rho(\eta^*))$	$(1-\beta^\infty(\mathscr{C}, \rho(\eta^*)))^{-1}$	$\beta^\infty(\mathscr{C}, 0)$	$(1-\beta^\infty(\mathscr{C}, 0))^{-1}$	$\beta(\mathscr{C})$	$(1-\beta(\mathscr{C}))^{-1}$
1	1/2	1/4	4/3	1/3	3/2	1/4	4/3
2	1/2	1/2	2	0.61…	2.56…	0.38…	1.63…
3	1/3	0.78…	4.53…	0.87…	7.83…	0.47…	1.90…
4	1/4	1.05…	∞	1.13…	∞	0.53…	2.15…
5	1/5	1.32…	∞	1.38…	∞	0.58…	2.39…

may happen that by joining a route guidance service, a user may end up being worse than when she selected the path with shortest delay. Nevertheless, this probably happens very rarely. Although the instances we provide were carefully constructed to show that degradation may occur, in real-world instances delays would typically be lower when guidance is available.

Another application domain that fits this model is that of telecommunication networks. Internet service providers (ISPs) route packets using an underlying physical network whose arcs are owned by network providers. To deliver their packets, ISPs must subcontract capacity from link owners and determine routing patterns to deliver packets across the network. Competition arises from the fact that network providers can sell their capacity to different ISPs. As before, the two performance measures of interest are cost paid to link owners and an indicator of service level, such as routing delay. Mapping both measures into cost, we can model this game with the help of price curves that map demand to unit prices.

Even though the network structure arises naturally in applications that involve networks such as those presented above, this model can also represent other situations. For example, in a flexible manufacturing environment, producing goods involves a series of operations that are performed by different flexible machines. There may be more than one sequence of machines that can process parts to manufacture the finished product. In this setting, the flexible machines are owned by different agents who sell capacity to contractors. Contractors, in turn, compete in the two dimensions previously mentioned: If more contractors buy capacity of a particular machine, its price goes up on the one hand, but on the other hand its waiting time, and thus the completion time, increases. Naturally, a contractor would like to produce at minimum cost—the reason for which competition arises—and hence has to determine which machine sequences to use for production. However, such an environment cannot be modeled directly by network games; instead, one needs to use an atomic congestion game. In this class of games there are resources (machines in this case), and players (contractors) have to pay to use them. Here, payments depend only on how many players select the resources, and not on their identities. Because, our results did not make use of the network structure, but relied on variational inequalities to characterize Nash equilibria, all results we have presented extend to congestion games without modification.

Appendix. Games with One OD Pair and Affine Costs

We prove additional results for networks with a single OD pair and affine costs. Namely, we compute the price of anarchy as a function on the Herfindahl index, and as a function on the number of players for symmetric games.

Proposition A1. *Consider an atomic game with a single OD pair and affine cost functions. The price of anarchy is bounded from above by* $(4-H)/(3-H)$, *where* H *is the Herfindahl index.*

Proof. Using the extension of Proposition 3.3 to this setting, we have that

$$\beta^{\infty}(c) \leqslant \max_{0\leqslant u\leqslant 1} u(1-u+Hu/4) = \frac{1}{4-H},$$

from which the claim follows. □

Proposition A2. *Consider a symmetric atomic game with* K *players and affine cost functions. The social cost of an equilibrium is bounded from above by* $4K^2/(K+1)(3K-1)$ *times that of a social optimum.*

Proof. Let $c_a(x_a) = q_a x_a + r_a$. By Theorem 4.1, we know that equilibria are of the form $(x/K, \dots, x/K)$, where x solves problem (SNE). The optimality condition of Problem (SNE) can be expressed as

$$\sum_{a\in A}((K+1)q_a x_a^2 + K r_a x_a) \leqslant \sum_{a\in A}((K+1)q_a x_a y_a + K r_a y_a) \quad \text{for any feasible flow } y.$$

By reorganizing the terms in the above inequality, we obtain

$$\begin{aligned}
&(K+1)\sum_{a\in A} x_a(q_a x_a + r_a)\\
&\quad\leqslant (K+1)\sum_{a\in A} q_a x_a y_a + K\sum_{a\in A} r_a y_a + \sum_{a\in A} r_a x_a\\
&\quad\leqslant K\sum_{a\in A} y_a(q_a y_a + r_a) + \frac{(K+1)^2}{4K}\sum_{a\in A} q_a x_a^2 + \sum_{a\in A} r_a x_a\\
&\quad\leqslant K\sum_{a\in A} y_a(q_a y_a + r_a) + \frac{(K+1)^2}{4K}\sum_{a\in A} x_a(q_a x_a + r_a).
\end{aligned}$$

The second inequality follows from $((x/2)\sqrt{(K+1)/K} - y\sqrt{K/(K+1)})^2 \geqslant 0$, whereas the third holds as $(K+1)^2 \geqslant 4K$, for all $K \geqslant 1$. We conclude that

$$\sum_{a\in A} x_a(q_a x_a + r_a) \leqslant \frac{4K^2}{(K+1)(3K-1)}\sum_{a\in A} y_a(q_a y_a + r_a)$$

for any feasible flow y; in particular, it holds for the social optimum. □

Let us prove that the bound provided by Proposition A2 is tight. To this end, consider a two-node two-arc network with K unit-demand players. The cost of arc 1 is $c_1(x) = 1$, whereas the cost of arc 2 is given by $c_2(x) = x/(K+1)$. It is easy to see that the optimal flow pattern routes $(K-1)/2$ and $(K+1)/2$ on arcs 1 and 2, respectively, leading to a total cost of $(3K-1)/4$. On the other hand, at equilibrium each player routes all its flow along arc 2, and therefore the social cost of the equilibrium is $K^2/(K+1)$. The resulting ratio is exactly as previously stated.

Acknowledgments

The authors thank two anonymous referees and Nicolás Figueroa, Andreas Schulz, and Gabriel Weintraub for useful comments that helped them to improve the paper. The research of the first author was partially supported by FONDAP in Applied Mathematics (CONICYT-Chile) and Instituto Milenio SCI (P05-004F). The research of the second author was partially supported by CONICYT-Chile through grants Anillo en Redes ACT08 and FONDEYT 1060035, and the research of the third author was partially supported by the Center for International Business Education and Research (CIBER) and by the Center for Excellence in E-Business (CEBIZ), both at Columbia University. An extended abstract of this article appears in ICALP 2006 (Cominetti et al. 2006).

References

Acemoglu, D., A. Ozdaglar. 2007. Competition and efficiency in congested markets. *Math. Oper. Res.* **32**(1) 1–31.

Altman, E., T. Başar, T. Jimenez, N. Shimkin. 2002. Competitive routing in networks with polynomial costs. *IEEE Trans. Automatic Control* **47**(1) 92–96.

Awerbuch, B., Y. Azar, A. Epstein. 2005. The price of routing unsplittable flow. *Proc. 37th Annual ACM Sympos. Theory Comput.* (*STOC*). ACM Press, New York, 57–66.

Beckmann, M. J., C. B. McGuire, C. B. Winsten. 1956. *Studies in the Economics of Transportation*. Yale University Press, New Haven, CT.

Bernstein, F., A. Federgruen. 2005. Decentralized supply chains with competing retailers under demand uncertainty. *Management Sci.* **51**(1) 18–29.

Bhaskar, U., L. Fleischer, D. Hoy, C.-C. Huang. 2009. Equilibria of atomic flow games are not unique. *Proc. 19th Annual ACM-SIAM Sympos. Discrete Algorithms* (*SODA*), New York. SIAM, Philadelphia, 748–757.

Boulogne, T. 2004. Nonatomic strategic games and applications to networks. Ph.D. thesis, Université Paris 6, Paris.

Braess, D. 1968. Über ein Paradoxon aus der Verkehrsplanung. *Unternehmensforschung* **12** 258–268. [English translation: Braess, D., A. Nagurney, T. Wakolbinger. 2005. On a paradox of traffic planning. *Transportation Sci.* 2005. **39**(4) 446–450.]

Bureau of Public Roads. 1964. Traffic assignment manual. U.S. Department of Commerce, Urban Planning Division, Washington, D.C.

Cachon, G., S. Netessine. 2004. Game theory in supply chain analysis. D. Simchi-Levi, S. Wu, Z. Shen, eds. *Handbook of Quantitative Supply Chain Analysis*: *Modeling in the E-business Era*. Kluwer, Boston, 13–66.

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** 327–334.

Christodoulou, G., E. Koutsoupias. 2005. The price of anarchy of finite congestion games. *Proc. 37th Annual ACM Sympos. Theory Comput.* (*STOC*). ACM Press, New York, 67–73.

Cole, R., Y. Dodis, T. Roughgarden. 2003. Pricing network edges for heterogeneous selfish users. *Proc.* 35*th Annual ACM Sympos. Theory Comput.* (*STOC*). ACM Press, New York, 521–530.

Cominetti, R., J. R. Correa, N. E. Stier-Moses. 2006. Network games with atomic players. M. Bugliesi, B. Preneel, V. Sassone, I. Wegener, eds. *Automata, Languages and Programming. Proc. 33rd Internat. Colloquium* (*ICALP*). Vol. 4051, *Lecture Notes in Computer Science*, Part 1. Springer, Heidelberg, 525–536.

Correa, J. R., A. S. Schulz, N. E. Stier-Moses. 2004. Selfish routing in capacitated networks. *Math. Oper. Res.* **29**(4) 961–976.

Correa, J. R., A. S. Schulz, N. E. Stier-Moses. 2005. On the inefficiency of equilibria in congestion games. M. Jünger, V. Kaibel, eds. *Proc. 11th Internat. Integer Programming Combin. Optim. Conf.* (*IPCO*), Vol. 3509, *Lecture Notes in Computer Science*. Springer, Heidelberg, 167–181.

Correa, J. R., A. S. Schulz, N. E. Stier-Moses. 2007. Fast, fair, and efficient flows in networks. *Oper. Res.* **55**(2) 215–225.

Correa, J. R., A. S. Schulz, N. E. Stier-Moses. 2008. A geometric approach to the price of anarchy in nonatomic congestion games. *Games Econom. Behav.* **64**(2) 457–469.

Dafermos, S. C., F. T. Sparrow. 1969. The traffic assignment problem for a general network. *J. Res. U.S. National Bureau of Standards* **73B** 91–118.

Dubey, P. 1986. Inefficiency of Nash equilibria. *Math. Oper. Res.* **11**(1) 1–8.

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.* (*FOCS*). IEEE Computer Society Press, Los Alamitos, CA, 277–285.

Fotakis, D., S. C. Kontogiannis, P. G. Spirakis. 2005. Selfish unsplittable flows. *Theoret. Comput. Sci.* **348** 226–239.

Golany, B., U. G. Rothblum. 2006. Inducing coordination in supply chains through linear reward schemes. *Naval Res. Logist.* **53** 1–15.

Harker, P. T. 1988. Multiple equilibrium behaviors of networks. *Transportation Sci.* **22**(1) 39–46.

Harks, T. 2008. Stackelberg strategies and collusion in network games with splittable flow. E. Bampis, M. Skutella, eds. *Sixth Workshop on Approximation and Online Algorithms* (*WAOA*) *Lecture Notes in Computer Science*. Springer, Heidelberg, 133–146.

Haurie, A., P. Marcotte. 1985. On the relationship between Nash-Cournot and Wardrop equilibria. *Networks* **15** 295–308.

Hayrapetyan, A., É. Tardos, T. Wexler. 2006. The effect of collusion in congestion games. *Proc. 38th Annual ACM Sympos. Theory Comput.* (*STOC*). ACM Press, New York, 89–98.

Jahn, O., R. H. Möhring, A. S. Schulz, N. E. Stier-Moses. 2005. System-optimal routing of traffic flows with user constraints in networks with congestion. *Oper. Res.* **53**(4) 600–616.

Johari, R., J. N. Tsitsiklis. 2004. Network resource allocation and a congestion game. *Math. Oper. Res.* **29**(3) 407–435.

Johnson, B., L.-G. Mattson. 1992. *Road Pricing*: *Theory, Empirical Assessment and Policy*. Kluwer Academic Press, Dordrecht, The Netherlands.

Karakostas, G., S. G. Kolliopoulos. 2004. Edge pricing of multicommodity networks for heterogeneous users. *Proc. 45th Annual IEEE Sympos. Foundations Comput. Sci.* (*FOCS*). IEEE Computer Society Press, Los Alamitos, CA, 268–276.

Kelly, F. P. 1997. Charging and rate control for elastic traffic. *Eur. Trans. Telecomm.* **8** 33–37.

Koutsoupias, E., C. H. Papadimitriou. 1999. Worst-case equilibria. C. Meinel, S. Tison, eds. *Proc. 16th Annual Sympos. Theoret. Aspects Comput. Sci.* (*STACS*), Vol. 1563, *Lecture Notes in Computer Science*. Springer, Heidelberg, 404–413.

Maheswaran, R. T., T. Başar. 2004. Social welfare of selfish agents: Motivating efficiency for divisible resources. *43rd IEEE Conf. Decision and Control* (*CDC*), Vol. 2. Atlantis, Paradise Island, Bahamas, 1550–1555.

Monderer, D., L. S. Shapley. 1996. Potential games. *Games Econom. Behav.* **14** 124–143.

Orda, A., R. Rom, N. Shimkin. 1993. Competitive routing in multiuser communication networks. *IEEE/ACM Trans. Networking* **1**(5) 510–521.

Papadimitriou, C. H. 2001. Algorithms, games, and the Internet. *Proc. 33rd Annual ACM Sympos. Theory Comput.* (*STOC*). ACM Press, New York, 749–753.

Parlakturk, A., S. Kumar. 2004. Self-interested routing in queueing networks. *Management Sci.* **50**(7) 949–967.

Perakis, G. 2007. The "price of anarchy" under nonlinear and asymmetric costs. *Math. Oper. Res.* **32**(3) 614–628.

Perakis, G., G. Roels. 2007. The price of anarchy in supply chains: Quantifying the efficiency of price-only contracts. *Management Sci.* **53**(8) 1249–1268.

Pigou, A. C. 1920. *The Economics of Welfare*. Macmillan, London.

Potra, F., Y. Ye 1993. A quadratically convergent polynomial algorithm for solving entropy optimization problems. *SIAM J. Optim.* **3**(4) 843–860.

Rosen, J. B. 1965. Existence and uniqueness of equilibrium points for concave N-person games. *Econometrica* **33**(3) 520–534.

Rosenthal, R. W. 1973. A class of games possessing pure-strategy Nash equilibria. *Internat. J. Game Theory* **2** 65–67.

Roughgarden, T. 2002. How unfair is optimal routing? *Proc. 13th Annual ACM-SIAM Sympos. Discrete Algorithms* (*SODA*). SIAM, Philadelphia, 203–204.

Roughgarden, T. 2003. The price of anarchy is independent of the network topology. *J. Comput. System Sci.* **67** 341–364.

Roughgarden, T. 2005. Selfish routing with atomic players. *Proc. 16th Annual ACM-SIAM Sympos. Discrete Algorithms* (*SODA*). SIAM, Philadelphia, 973–974.

Roughgarden, T., É. Tardos. 2002. How bad is selfish routing? *J. ACM* **49** 236–259.

Roughgarden, T., É. Tardos. 2004. Bounding the inefficiency of equilibria in nonatomic congestion games. *Games Econom. Behav.* **47**(2) 389–403.

Sanghavi, S., B. Hajek. 2004. Optimal allocation of a divisible good to strategic buyers. *43rd IEEE Conf. Decision and Control* (*CDC*), Vol. 3. Atlantis, Paradise Island, Bahamas, 2748–2753.

Swamy, C. 2007. The effectiveness of Stackelberg strategies and tolls for network congestion games. *Proc. 18th Annual ACM-SIAM Sympos. Discrete Algorithms* (*SODA*). SIAM, Philadelphia, 1133–1142.

Taylor, T. 2002. Supply chain coordination under channel rebates with sales effort effects. *Management Sci.* **48** 992–1007.

Tirole, J. 1988. *The Theory of Industrial Organization*. MIT Press, Cambridge, MA.

Vickrey, W. S. 1969. Congestion theory and transport investment. *Amer. Econom. Rev.* **59** 251–261.

Wardrop, J. G. 1952. Some theoretical aspects of road traffic research. *Proc. Institution of Civil Engineers*, Part II, Vol. 1. 325–378.

Yang, H., H.-J. Huang. 2004. The multi-class, multi-criteria traffic network equilibrium and systems optimum problem. *Transportation Res. Part B* **38** 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* **42** 99–112.