

How Much Can Taxes Help Selfish Routing?

Richard Cole^{*}
Department of Computer
Science
New York University
New York, NY 10012
cole@cs.nyu.edu

Yevgeniy Dodis[†]
Department of Computer
Science
New York University
New York, NY 10012
dodis@cs.nyu.edu

Tim Roughgarden[‡]
Department of Computer
Science
Cornell University
Ithaca, NY 14853
timr@cs.cornell.edu

ABSTRACT

We study economic incentives for influencing selfish behavior in networks. We consider a model of selfish routing in which the latency experienced by network traffic on an edge of the network is a function of the edge congestion, and network users are assumed to selfishly route traffic on minimum-latency paths. The quality of a routing of traffic is historically measured by the sum of all travel times, also called the *total latency*.

It is well known that the outcome of selfish routing (a *Nash equilibrium*) does not minimize the total latency and can be improved upon with coordination, and that *marginal cost pricing*—charging each network user for the congestion effects caused by its presence—eliminates the inefficiency of selfish routing. However, the principle of marginal cost pricing assumes that (possibly very large) taxes cause no disutility to network users; this is only appropriate when collected taxes can be feasibly returned (directly or indirectly) to the users, for example via a lump-sum refund. If this assumption does not hold and we wish to minimize the total user disutility (latency plus taxes paid)—the total *cost*—how should we price the network edges? Intuition may suggest that taxes should never be able to improve the cost of a Nash equilibrium, but the famous *Braess’s Paradox* shows this intuition to be incorrect.

We consider strategies for pricing network edges to reduce the cost of a Nash equilibrium. Since levying a sufficiently large tax on an edge effectively removes it from the network, our study generalizes previous work on network design [22]. In this paper, we prove the following results.

- In a large class of networks—including all networks with linear latency functions—marginal cost taxes do not improve the cost of the Nash equilibrium.

^{*}Supported in part by NSF grant CCR0105678.

[†]Supported in part by an NSF CAREER Award.

[‡]Supported by ONR grant N00014-98-1-0589. Research performed in part while visiting NYU.

- The best-possible benefit from arbitrary taxes does not exceed that of edge removal. In *every* network with linear latency functions, taxes cannot improve over removing edges. There are networks with nonlinear latency functions, however, in which taxes are radically more powerful than edge removal.
- Strong inapproximability results known for network design [22] carry over to the problem of computing optimal taxes, proving this problem computationally intractable.

1. INTRODUCTION

Selfish Routing and Marginal Cost Pricing

We study economic incentives for influencing selfish behavior in networks. We focus on a simple model of *selfish routing*, defined by Wardrop [28] and first studied in the theoretical computer science literature by Roughgarden and Tardos [26]. In this model, we are given a directed network in which each edge possesses a latency function describing the common delay experienced by all traffic on the edge as a function of the edge congestion. There is a fixed amount of traffic wishing to travel from a source vertex s to a sink vertex t ; as in earlier works, we assume that the traffic comprises a very large population of users, so that the actions of a single individual have negligible effect on network congestion. The quality of an assignment of traffic to s - t paths is historically measured by the resulting sum of all travel times—the *total latency*. We assume that each network user, when left to its own devices, acts selfishly and routes itself on a minimum-latency path, given the network congestion due to the other users. In general such a “selfishly motivated” assignment of traffic to paths (a *Nash equilibrium*) does not minimize the total latency; put differently, the outcome of selfish behavior can be improved upon with coordination.

The inefficiency of selfish routing motivates the introduction of economic incentives to ensure that selfish behavior results in a socially desirable routing of traffic. An ancient idea—discussed informally as early as 1920 [20]—is to use *marginal cost pricing*. The principle of marginal cost pricing asserts that on each edge, each network user on the edge should pay a tax equal to the additional delay its presence causes for the other users on the edge. Assuming that all network users choose routes to minimize the sum of the latency experienced and the taxes paid, this principle ensures that the resulting Nash equilibrium achieves the minimum-possible total latency [3]. Briefly, the inefficiency of selfish routing can always be eradicated by pricing network edges appropriately.




Figure 1: Braess's Paradox

The following observation motivates our work: the principle of marginal cost pricing is single-minded in its pursuit of a minimum-latency flow, and ignores the disutility to network users due to (possibly very large) taxes. This assumption is only appropriate when collected taxes can be feasibly returned (directly or indirectly) to the network users, for example by refunding taxes equally to all users (a “lump-sum refund”). In this paper, we are interested in settings where this assumption is not reasonable. For example, refunding the collected taxes to network users could be logistically or economically infeasible, or taxes could represent quantities of a non-monetary, non-refundable good such as time delays.

If we wish to minimize the total user disutility (latency plus taxes paid)—a quantity we call *cost*—rather than merely the total latency, how should we price the network edges?

Braess's Paradox and the Power of Taxes

Intuition may suggest that taxes, which can only increase the disutility incurred by users along an s - t path, could never improve the cost of a Nash equilibrium. A famous example that we describe next, called *Braess's Paradox* and shown in Figure 1, shows this intuition to be incorrect [5].

Each edge of the network in Figure 1 is labeled with its latency function, giving the delay incurred by traffic on the link as a function of the fraction of the traffic that uses the link. In the (unique) Nash equilibrium, all traffic uses route $s \rightarrow v \rightarrow w \rightarrow t$ and experiences two units of latency. On the other hand, if half a unit of tax is levied on edge (v, w) , then in the Nash equilibrium half of the traffic uses each of the routes $s \rightarrow v \rightarrow t$ and $s \rightarrow w \rightarrow t$. In this Nash equilibrium, everyone experiences latency $\frac{3}{2}$ and no taxes are paid; this outcome has cost $\frac{3}{2}$ and is clearly superior to the original Nash equilibrium in the absence of taxes.

It is illuminating to study the edge taxes that are dictated by the principle of marginal cost pricing for the network of Figure 1. As it turns out, these taxes are $\frac{1}{2}$ on edges (s, v) and (w, t) and 0 on the other three edges. With these taxes, we obtain the same Nash equilibrium as in our previous solution with taxes (and thus all traffic experiences latency $\frac{3}{2}$); however, all traffic must also pay $\frac{1}{2}$ unit of tax. This solution thus has a cost of 2, and it seems evident that the previous solution should be preferred over marginal cost taxes in this example.

The potential power of taxes in decreasing the disutility experienced by selfish network users, together with the failure of marginal

cost pricing to minimize this disutility, motivate the central questions that we study in this paper:

- (1) Are marginal cost taxes ever a good idea for minimizing the cost of a Nash equilibrium?
- (2) A sufficiently large edge tax effectively removes the edge from the network, and the power of removing edges to improve a Nash equilibrium is well understood in the context of network design [22]. Taxation is thus at least as powerful as edge removal; when it is strictly more powerful?
- (3) Can we compute or approximate optimal taxes efficiently?

Our Results

Our contributions on these three questions are as follows.

- While marginal cost taxes can dramatically improve the cost of a Nash equilibrium in specific examples, we show that in a large class of networks—including all networks with linear latency functions—such taxes do not improve the cost of the Nash equilibrium.
- We prove that the maximum-possible benefit due to edge taxation is no more than what can be achieved, in the best case, via edge removal.
- For networks with linear latency functions, we show the following stronger statement: in every such network—not merely worst-case examples—taxes cannot decrease the cost of a Nash equilibrium beyond what can be achieved by removing edges from the network. By contrast, we show that in general there are networks in which taxes can radically improve over the best subgraph solution.
- We show that most of the very strong inapproximability results known for network design [22] carry over to the problem of computing optimal taxes. As a consequence, essentially the best-possible polynomial-time heuristic for the problem, unless $P = NP$, is to assign no taxes at all.

These results are summarized in Table 1. Determining the maximum-possible benefit of taxes and the (in)approximability of the problem of computing optimal taxes requires only reasonably straightforward extensions of previous work on network design [22]; all other results of this paper require new constructions and proof approaches.

Related Work

The model of selfish routing studied in this paper was first defined in the 1950's by Wardrop [28], and has been extensively studied ever since. It has been generalized in multiple directions that we do not explore here, such as in the literature on congestion games [21, 25], potential games [17], and summarization games [14]. In addition, several recent papers have proposed new game-theoretic ideas to better model distributed and asynchronous networks in which players have little information about their strategic environment, such as the Internet [10, 18, 27]. For further discussion and references on these topics, see [24].

We make no attempt at a detailed survey of previous work on the optimal pricing of resources, and instead refer the reader to [6, 13] for recent collections of surveys on pricing by the transportation

Problem Studied	Linear	General
Can marginal cost taxes help?	no	yes
Maximum benefit of taxes	4/3	$n/2$
Taxes better than network design?	no	yes
Approximability of optimal taxes	4/3	$[\omega(n^{1-\epsilon}), n/2]$

Table 1: Contributions of this paper. All results depend on the class of allowable edge latency functions; for simplicity, we list only our results for networks with linear latency functions and for networks with general latency functions. The benefit of taxes is measured by the ratio in Nash flow cost before and after taxes are levied. Approximability refers to the best-possible approximation ratio achievable by a polynomial-time algorithm, assuming $P \neq NP$.

science literature, and to [16] for an overview on approaches to optimal pricing in welfare economics. The transportation science community has long realized that marginal cost taxes may cause users to pay more than is necessary (see e.g. [4, 12]), but has apparently neglected the idea of incorporating the amount of taxes paid into the definition of social welfare. The more general problem of using taxation to optimize the sum of user utilities has been extensively studied by economists (see e.g. [16]), although not with the same emphasis on worst- and best-case performance guarantees and on computational efficiency that characterizes our work here. In the networking literature there are numerous papers on pricing networks (see e.g. [15]), but this body of work has not studied the model of selfish routing considered in this paper. Finally, our approach is similar in spirit to work of Archer and Tardos [1, 2] on “frugal mechanisms”—mechanisms (such as auctions) that solve an optimization problem in an incentive-compatible way, and also make use of only moderate incentives.

The paper closest to the present work is that of Roughgarden on designing networks for selfish users [22]. Indeed, the central questions of [22], which concern the maximum-possible benefit from and the algorithmic complexity of removing edges from a network with selfish routing, can be viewed as a special case of some of the problems considered here (with edge taxes restricted to be either 0 or $+\infty$). Our work in this paper extends some of the results from [22] to the setting of more general taxes, and in addition tackles some problems that have not been previously considered, such as determining the benefits of marginal cost taxes and separating the power of taxes from that of edge removal.

Organization

In Section 2 we formally define our traffic routing model and review useful results from past works. In Section 3 we ask if marginal cost pricing can improve the cost of selfish routing; we resolve this question in the negative for networks with linear latency functions, and in the affirmative for general networks. In Sections 4 and 5 we ask if the power of taxes exceeds that of edge removal. In Section 4 we adapt previous work on network design [22] to show that the best-possible improvement due to taxes cannot exceed that of edge removal. In Section 5 we show that taxes never improve over the best solution achievable by removing edges in networks with linear latency functions, but can be radically more powerful in general networks. Finally, in Section 6 we ask whether optimal taxes can be efficiently computed; we answer this question in the negative by extending known inapproximability results for network design [22] to the problem of computing optimal taxes.

2. PRELIMINARIES

2.1 The Model

We follow the notation and conventions of Roughgarden and Tardos [26]. We study a single-commodity flow network, described by a directed graph $G = (V, E)$ with source vertex s and sink vertex t . We allow parallel edges but have no use for self-loops. We denote the set of simple s - t paths by \mathcal{P} , and we assume that this set is nonempty. A flow is a function $f : \mathcal{P} \rightarrow \mathcal{R}^+$; for a fixed flow f we define $f_e = \sum_{P: e \in P} f_P$ as the amount of traffic using edge e en route from s to t . With respect to a finite and positive traffic rate r , a flow f is said to be *feasible* if $\sum_{P \in \mathcal{P}} f_P = r$.

The network G suffers from congestion effects; to model this, we assume each edge e possesses a nonnegative, continuous, nondecreasing latency function ℓ_e that describes the delay incurred by traffic on e as a function of the edge congestion f_e . The latency of a path P in G with respect to a flow f is then given by $\ell_P(f) = \sum_{e \in P} \ell_e(f_e)$. The quality of a flow is historically measured by its total latency $L(f)$, defined by

$$L(f) \equiv \sum_{P \in \mathcal{P}} \ell_P(f) f_P = \sum_{e \in E} \ell_e(f_e) f_e,$$

with the equality proved by writing $\ell_P(f)$ as a sum over edges and reversing the order of summation. We will call a flow minimizing $L(\cdot)$ *optimal* or *minimum-latency*.

Finally, we will allow a set of nonnegative taxes $\{\tau_e\}_{e \in E}$ to be placed on the edges of a network G . We will write $\tau_P = \sum_{e \in P} \tau_e$ for the total taxes on a path P . The cost $C(f, \tau)$ of a flow f in a network with taxes τ is the total disutility caused to network users, accounting for disutility due to both latency and taxes:

$$C(f, \tau) \equiv \sum_{P \in \mathcal{P}} [\ell_P(f) + \tau_P] f_P = \sum_{e \in E} [\ell_e(f_e) + \tau_e] f_e.$$

The functions $L(\cdot)$ and $C(\cdot, \tau)$ coincide if and only if $\tau = 0$.¹ We will call a triple (G, r, ℓ) an *instance*, and will use the notation $(G, r, \ell + \tau)$ to denote an instance in which taxes τ have been levied on the edges of G .

2.2 Flows at Nash Equilibrium

We assume that noncooperative behavior results in a Nash equilibrium—a “stable point” in which no traffic has an incentive to unilaterally alter its strategy (i.e., its route from s to t). We also assume that all agents seek to minimize the sum of the latency experienced and the tax paid.² We therefore expect that, in a flow at Nash equilibrium, all traffic is routed on paths with minimum-possible latency plus tax. Formally, we have the following definition.

Definition 2.1 A flow f feasible for $(G, r, \ell + \tau)$ is at *Nash equilibrium* or is a *Nash flow* if for all $P_1, P_2 \in \mathcal{P}$ with $f_{P_1} > 0$,

$$\ell_{P_1}(f) + \tau_{P_1} \leq \ell_{P_2}(f) + \tau_{P_2}.$$

¹In several previous papers on selfish routing (without taxes), such as [23, 26], the terms total latency and cost were used synonymously.

²This assumption, while classical, is obviously quite strong. In general, we expect different agents to have different objective functions and to trade off time and money in different ways; this objection raises several interesting issues that are the subject of our ongoing work [7].

In Section 5, we will make use of the following alternative definition of Nash flows.

Proposition 2.2 ([22, 24]) *Let f be a flow feasible for $(G, r, \ell + \tau)$. For a vertex v in G , let $d(v)$ denote the length, with respect to edge lengths $\ell_e(f_e) + \tau_e$, of a shortest s - v path in G . Then f is at Nash equilibrium if and only if, for all edges $e = (v, w)$, either $f_e = 0$ or $d(w) - d(v) = \ell_e(f_e) + \tau_e$.*

An immediate consequence of Proposition 2.2 is that the property of being at Nash equilibrium is a property only of the flow vector on edges $\{f_e\}$ induced by a flow f , rather than on the particular path decomposition.

Corollary 2.3 *Suppose f, \tilde{f} are flows for $(G, r, \ell + \tau)$ with $f_e = \tilde{f}_e$ for all edges e . Then f is at Nash equilibrium if and only if \tilde{f} is at Nash equilibrium.*

Flows at Nash equilibrium enjoy several nice properties. We first consider the existence and uniqueness of Nash flows.

Proposition 2.4 ([3, 9, 26]) *Every instance $(G, r, \ell + \tau)$ admits a flow at Nash equilibrium.*

Proposition 2.5 ([3, 9, 26]) *If f, \tilde{f} are flows at Nash equilibrium for $(G, r, \ell + \tau)$, then:*

- (a) $\ell_e(f_e) = \ell_e(\tilde{f}_e)$ for all edges e ;
- (b) $C(f, \tau) = C(\tilde{f}, \tau)$.

The next proposition states that all paths used by a Nash flow have the same combined latency and tax (this is immediate from Definition 2.1), and that the cost of a Nash flow is therefore expressible in a very simple form.

Proposition 2.6 *If f is at Nash equilibrium for $(G, r, \ell + \tau)$, then there is a constant $c \geq 0$ with $\ell_P(f) + \tau_P = c$ whenever $f_P > 0$. Moreover, $C(f, \tau) = r \cdot c$.*

We adopt the notation $c(G, r, \ell + \tau)$ for the value of the constant c in Proposition 2.6 for instance $(G, r, \ell + \tau)$, which is uniquely defined by Propositions 2.4 and 2.5. Our final proposition states that, for fixed G, ℓ , and τ , the value of $c(G, r, \ell + \tau)$ is nondecreasing in r ; this fact was proved by Hall [11] and we will use it in Section 3.

Proposition 2.7 ([11]) *The value $c(G, r, \ell + \tau)$ is nondecreasing in r .*

3. WHEN DO MARGINAL COST TAXES HELP?

In this section we study the cost of applying the principle of marginal cost pricing. Recall that this principle posits that each user should pay a tax equal to the additional delay other users experience because of its presence. Mathematically, this principle asserts that in a flow f feasible for instance (G, r, ℓ) , the tax τ_e assigned to edge e should be $\tau_e = f_e \cdot \ell'_e(f_e)$ where ℓ'_e denotes the derivative of ℓ_e (assume for simplicity that the latency functions are differentiable).

The term $\ell'_e(f_e)$ corresponds to the marginal increase in latency caused by one user on the edge, and the term f_e is the amount of traffic that suffers from this increase. Marginal cost taxes come with the following guarantee, which is classical (see [3, 9, 24] for proofs).

Proposition 3.1 *Let (G, r, ℓ) be an instance with differentiable latency functions admitting minimum-latency flow f^* . Let $\tau_e = f_e^* \cdot \ell'_e(f_e^*)$ denote the marginal cost tax for edge e . Then f^* is at Nash equilibrium for $(G, r, \ell + \tau)$.*

In words, marginal cost taxes induce the minimum-latency flow as a flow at Nash equilibrium.

Proposition 3.1 gives an effective way to minimize the total latency of a Nash flow with edge taxes. But how effective are marginal cost taxes when we also account for the disutility to traffic due to taxes? Our next theorem identifies a reasonably large class of networks—networks in which all latency functions are linear, with the form $\ell(x) = ax + b$ —in which marginal cost taxes are *guaranteed* to be unnecessary, if not detrimental. While the theorem does not hold in general networks (see Remark 3.4 below), we trust that it illustrates the dangers of marginal cost pricing in settings where minimizing the total latency is not the sole goal of taxation.

Theorem 3.2 *Let (G, r, ℓ) be an instance with linear latency functions and τ the corresponding marginal cost taxes. Let f and f^τ be Nash flows for (G, r, ℓ) and $(G, r, \ell + \tau)$, respectively. Then,*

$$C(f, 0) \leq C(f^\tau, \tau).$$

PROOF. Let (G, r, ℓ) be an instance with linear latency functions, admitting minimum-latency flow f^* and Nash flow f . Write $\ell_e(x) = a_e x + b_e$ for each edge e . The principle of marginal cost pricing dictates that $\tau_e = f_e^* \cdot \ell'_e(f_e^*) = a_e f_e^*$.

Define the modified latency function ℓ_e^* by $\ell_e^*(x) = 2a_e x + b_e$. The functions $\ell + \tau$ and ℓ^* are not identically equal, but the identity $\ell_e^*(f_e^*) = 2a_e f_e^* + b_e = \ell_e(f_e^*) + \tau_e$ holds on every edge. It follows that f^τ is at Nash equilibrium not only for the instance $(G, r, \ell + \tau)$, but also for the instance (G, r, ℓ^*) . Moreover, in the notation of Proposition 2.6, $c(G, r, \ell^*) = c(G, r, \ell + \tau)$.

We next claim that $f/2$ is at Nash equilibrium for $(G, r/2, \ell^*)$ with $c(G, r/2, \ell^*) = c(G, r, \ell)$.³ To see why, we note that since $\ell_e(x) = a_e x + b_e$ and $\ell_e^*(x) = 2a_e x + b_e$, edge and path latencies with respect to $f/2$ in $(G, r/2, \ell^*)$ and with respect to f in (G, r, ℓ) are identical. That f is a Nash flow for (G, r, ℓ) then implies that $f/2$ is at Nash equilibrium for $(G, r/2, \ell^*)$, with $c(G, r/2, \ell^*) = c(G, r, \ell)$.

Combining the two equalities of the previous paragraphs with an application of Proposition 2.7, we obtain

$$c(G, r, \ell) = c(G, r/2, \ell^*) \leq c(G, r, \ell^*) = c(G, r, \ell + \tau);$$

the theorem now follows from Proposition 2.6. \square

Remark 3.3 Theorem 3.2 concerns only networks with linear latency functions, but it is easily extended to networks in which, for some fixed $p \geq 0$, every edge e has a latency function of the form

³This claim was in essence proved in [26].

$a_e x^p + b_e$ with $a_e, b_e \geq 0$. It is not clear if Theorem 3.2 extends to networks with arbitrary polynomial latency functions.

Remark 3.4 There are n -node networks with more exotic latency functions in which marginal cost taxes decrease the cost of a Nash flow by as much as an $n/2 - \epsilon$ factor, for arbitrary $\epsilon > 0$. This construction makes use of the Braess graphs of [22] (see also Subsection 5.2), and is omitted.

4. HOW POWERFUL ARE ARBITRARY TAXES?

In this section we study the following question: how much can the cost of a Nash flow improve after levying taxes on the edges? As we will see, a precise answer to this question follows easily from previous work on network design with selfish routing [22].

The maximum-possible benefit from taxes will depend crucially on the type of network latency functions that we allow. This dependence is characteristic of much of the work on selfish routing [24]; indeed, we have already seen a glimpse of such a dependence in the previous section, where marginal cost pricing can help in general, but not in networks with linear latency functions.

We will give two types of upper bounds on the maximum-possible reduction in cost due to taxes, depending on the generality of the network latency functions. The first two upper bounds are corollaries of previous work bounding the *price of anarchy* [19], defined as the largest possible ratio between the total latency of a Nash flow and of a minimum-latency flow. The price of anarchy is a function of the class of allowable latency functions, and this dependence is by now well understood. For example, the following statements are known (see e.g. [23] for further examples).

Proposition 4.1 ([26]) *The price of anarchy in networks with linear latency functions is $4/3$.*

Proposition 4.2 ([23]) *The price of anarchy in networks with latency functions that are polynomials with degree at most p and nonnegative coefficients is asymptotically $\Theta(p/\log p)$ as $p \rightarrow \infty$.*

Upper bounds on the price of anarchy translate directly to upper bounds on the largest decrease in cost achievable with taxes: at best, taxes replace the Nash flow in the original network with the minimum-latency flow, while causing no additional disutility to network users. We therefore have the following corollaries.

Corollary 4.3 *Let (G, r, ℓ) be an instance with linear latency functions and τ a tax on edges. Let f and f^τ be Nash flows for (G, r, ℓ) and $(G, r, \ell + \tau)$, respectively. Then $L(f) \leq \frac{4}{3} \cdot C(f^\tau, \tau)$.*

Corollary 4.4 *For some constant $c_1 > 0$ the following statement holds for $p \geq 2$: If (G, r, ℓ) is an instance with polynomial latency functions with degree at most p and nonnegative coefficients, τ is a tax on edges, and f and f^τ are Nash flows for (G, r, ℓ) and $(G, r, \ell + \tau)$, respectively, then $L(f) \leq c_1 \frac{p}{\log p} \cdot C(f^\tau, \tau)$.*

Roughgarden [22] gave a family of networks in which deleting edges can dramatically improve the total latency of a Nash flow. As we have noted, sufficiently large taxes can always simulate edge deletions and the examples of [22] thus translate directly into our

setting. In particular, these examples imply that the bound of Corollary 4.3 is best possible, and that Corollary 4.4 is sharp up to a constant factor.

In networks with more general latency functions, the price of anarchy is unbounded, even restricting attention to two-node, two-link networks. However, while a bounded price of anarchy is *sufficient* for a bound on the largest-possible benefit of taxes, it is not *necessary*. Indeed, with no assumptions whatsoever on network latency functions (other than continuity and monotonicity), we can still obtain an upper bound that is a function of the network size.

Theorem 4.5 *Let (G, r, ℓ) be an instance with n vertices and τ a tax on edges. Let f and f^τ be Nash flows for (G, r, ℓ) and $(G, r, \ell + \tau)$, respectively. Then*

$$L(f) \leq \left\lfloor \frac{n}{2} \right\rfloor \cdot C(f^\tau, \tau).$$

The proof of Theorem 4.5 is a straightforward extension of an argument from [22, Thm 4.1] proving the weaker statement that deleting edges from a network can improve the total latency of a Nash flow by at most an $\lfloor n/2 \rfloor$ factor, and we omit it. Examples from [22] again show that the bound of Theorem 4.5 is best possible.

5. ARE TAXES MORE POWERFUL THAN EDGE REMOVAL?

In Section 4, we saw that there is a strong connection between the power of taxation and the power of edge removal. Specifically, we found that for several natural classes of networks (with polynomial latency functions, or with general latency functions), the maximum-possible reduction in cost achievable by levying taxes on edges is the same as that from removing edges from the network. We have therefore resolved in the negative, for these classes of networks, the “extremal” version of the following question: are taxes strictly more powerful than edge removal? However, we have not resolved whether taxes can *ever* improve over solutions obtainable by removing edges. We study this issue in this section. In Subsection 5.1 we show that in *every* network with linear latency functions, taxes cannot improve the Nash flow more than the removal of some set of edges. By contrast, in Subsection 5.2 we show that in general networks, taxes can improve the Nash flow far beyond what is achievable by merely deleting edges from the network.

5.1 The Power of Edge Removal for Networks with Linear Latency Functions

In this subsection we consider only networks with linear latency functions (latency functions of the form $\ell(x) = ax + b$), and show that taxes are never more powerful than edge removal for these networks. To state this formally, we will say that a set τ of taxes for the instance (G, r, ℓ) is $0/\infty$ if, for some Nash flow f^τ for $(G, r, \ell + \tau)$, $\tau_e = 0$ or $f_e^\tau = 0$ for each edge e . We note that $0/\infty$ taxes are no more powerful than removing edges from the network, since if τ is $0/\infty$ then $c(G, r, \ell + \tau) = c(H, r, \ell)$, where H is the subgraph of G consisting of the edges with zero tax. A tax is *optimal* for an instance (G, r, ℓ) if it minimizes $c(G, r, \ell + \tau)$ over all nonnegative tax vectors τ .

Theorem 5.1 *An instance with linear latency functions admits an optimal set of taxes that is $0/\infty$.*

PROOF. The proof of this result is somewhat involved. To avoid considering an arbitrary network with linear latency functions, our proof argues by contradiction and studies a minimal counterexample; as we will see, the minimal counterexample must possess several convenient properties that facilitate the argument.

Suppose the theorem is false, and choose a counterexample instance $(G, r, \ell + \tau)$ having as few edges as possible, admitting Nash flow f^τ with $C(f^\tau, \tau)$ smaller than the cost of any Nash flow with a set of $0/\infty$ taxes. We next claim that this minimal counterexample must possess several nice properties.

Claim: *The following properties hold for the minimal counterexample $(G, r, \ell + \tau)$:*

- (1) $f_e^\tau > 0$ for all edges $e \in E$.
- (2) There is a constant c^τ such that $\ell_P(f^\tau) + \tau_P = c^\tau$ for all s - t paths P of G .
- (3) A flow f is at Nash equilibrium for $(G, r, \ell + \tau)$ if and only if $f_e^\tau = f_e$ for all edges e .
- (4) The graph G is directed acyclic.
- (5) For every pair v, w of vertices in G , there is at most one v - w path of G comprising only edges with constant latency functions.

Proof Sketch of Claim:

- (1) If $f_e^\tau = 0$, then deleting e from G yields a counterexample with fewer edges.
- (2) By property (1), we can choose a path decomposition f of $\{f_e^\tau\}$ with $f_P > 0$ for all paths $P \in \mathcal{P}$; the claim now follows from Corollary 2.3 and Proposition 2.6.
- (3) The “if” direction is Corollary 2.3. For the “only if” direction, suppose f is another Nash flow for (G, r, ℓ) with $\{f_e^\tau\} \neq \{f_e\}$. Since all latency functions are linear, Proposition 2.5(a) implies that f^τ and f differ only on edges with a constant latency function. By Proposition 2.2, every affine combination of $\{f_e^\tau\}$ and $\{f_e\}$ that yields a nonnegative vector on edges corresponds (after a path decomposition) to some Nash flow for (G, r, ℓ) . An appropriate affine combination yields a Nash flow \tilde{f} with $\tilde{f}_e = 0$ for some edge e ; as in (1), this contradicts the minimality of (G, r, ℓ) .
- (4) Since every instance admits an acyclic Nash flow [22], this follows immediately from properties (1) and (3).
- (5) If there are two v - w paths P_1 and P_2 containing only edges with constant latency functions, then by property (1) and Proposition 2.2 we can transfer flow from one to the other to obtain a Nash flow \hat{f} with $\{\hat{f}_e\} \neq \{f_e^\tau\}$; this contradicts property (3).

□

We now return to the proof of Theorem 5.1. Among all optimal taxes for the minimal counterexample (G, r, ℓ) , choose a tax τ minimizing $\sum_{e \in E} \tau_e$. It is not obvious that such a minimum exists; this requires a technical argument, making use of properties

(1)–(5), that we defer to Appendix A (see Proposition A.1). Since (G, r, ℓ) is a counterexample instance, $\sum_e \tau_e > 0$; we will show that this leads to a contradiction by a perturbation argument.

Write $\ell_e(x) = a_e x + b_e$ for each edge e of G . The equations

$$\sum_{e \in P} [a_e f_e + b_e + \tau_e] = c$$

for all $P \in \mathcal{P}$, together with the standard flow conservation constraints for f , form a system of equations linear in the $m + 1$ variables $\{f_e\}$ and c (where m is the number of edges in G). By property (3) of (G, r, ℓ) , $(\{f_e^\tau\}, c^\tau)$ is the unique solution to this system, with flow nonnegativity constraints automatically satisfied (indeed, strictly). Choosing $m + 1$ linearly independent constraints with at least one constraint corresponding to a path with nonzero tax, there is a square linear system

$$A \begin{bmatrix} f \\ c \end{bmatrix} = d$$

for which $\{f_e^\tau\}, c^\tau$ is the unique solution, namely $A^{-1}d$.

Let $\tau_{\tilde{e}} > 0$ be a positive edge tax appearing in this linear system (necessarily as part of the right-hand side d). Consider perturbing this tax by subtracting a small number $\epsilon > 0$. This translates to a perturbation of adding ϵ to the right-hand side of all constraints corresponding to paths that include the edge \tilde{e} . Since $f_e^\tau > 0$ for all edges e , for perturbations sufficiently close to 0 the new Nash flow is given by $(\{\tilde{f}_e\}, \tilde{c}) = A^{-1}\tilde{d}$, where \tilde{d} is the old right-hand side d following the perturbation. Since τ is an optimal tax minimizing $\sum_e \tau_e$, subtracting $\epsilon > 0$ from $\tau_{\tilde{e}}$ must produce a non-optimal tax. By Proposition 2.6, it follows that $c^\tau < \tilde{c}$. By linearity, however, the opposite perturbation of adding ϵ to $\tau_{\tilde{e}}$ has the opposite effect, producing a tax $\hat{\tau}$ inducing a solution $(\{\hat{f}_e\}, \hat{c})$ with $\hat{c} < c^\tau$. This contradicts the optimality of c^τ , and the proof is complete. □

5.2 The Power of Taxes in General Networks

In the previous subsection we saw that, in networks with linear latency functions, taxes never improve over the best solution attainable by removing edges from the network. In this subsection we will demonstrate that this fails in general: in fact, for each value of $n \geq 2$, there is an n -node network in which general taxes can improve upon $0/\infty$ taxes by an $\lfloor n/2 \rfloor$ factor. With this result, we will have a good understanding of the relationship between taxes and edge removal in general networks. Briefly, removing edges can improve the Nash flow by an $\lfloor n/2 \rfloor$ factor [22, Prop 4.2], taxes can improve a Nash flow by an $\lfloor n/2 \rfloor$ factor beyond what is achievable by removing edges (Theorem 5.2 below), but their combined power (or equivalently, the power of taxes) cannot improve a Nash flow by more than an $\lfloor n/2 \rfloor$ factor (Theorem 4.5).

Theorem 5.2 *For each integer $n \geq 2$, there is an instance (G, r, ℓ) with $c(H, r, \ell) = \lfloor n/2 \rfloor$ for all subgraphs H of G but $c(G, r, \ell + \tau) = 1$ for some tax $\tau \geq 0$.*

PROOF. The construction is similar to that in an inapproximability result for network design [22, Thm 4.3]. We can assume that n is even (for n odd, add an isolated vertex or subdivide an edge) and at least 4. We will make use of the *Braess graphs*, introduced in [22, §4.2]. The k th Braess graph $B^k = (V^k, E^k)$ is defined as follows: start with a set $V^k = \{s, v_1, \dots, v_k, w_1, \dots, w_k, t\}$ of $2k + 2$ vertices and define E^k by $\{(s, v_i), (v_i, w_i), (w_i, t) : 1 \leq$





Figure 2: The second and third Braess graphs

$i \leq k\} \cup \{(v_i, w_{i-1}) : 2 \leq i \leq k\} \cup \{(v_1, t)\} \cup \{(s, w_k)\}$ (see Figure 2). We will work with the Braess graph B^k for which $2k + 2 = n$.

We define latency functions ℓ on the graph B^k as follows.

- (A) Edges of the form (v_i, w_i) receive the latency function $\ell(x) = 0$.
- (B) Edges of the form $(v_i, w_{i-1}), (v_1, t), (s, w_k)$ receive a latency function ℓ satisfying $\ell(x) = 1$ for $x \leq 1/(k + 1)$ and $\ell(x) = n/2$ for $x \geq 1/(k + 1) + \epsilon$, where $\epsilon > 0$ is a sufficiently small constant.
- (C) Edges of the form (w_i, t) or (s, v_{k-i+1}) receive a latency function ℓ satisfying $\ell(x) = 0$ for $x \leq 1 + 1/(k + 1)$, $\ell(1 + 1/k) = i$, and $\ell(x) = n/2$ for $x \geq 1 + 1/k + \epsilon$.

We will call edges of the form (v_i, w_i) *type A edges*, and so forth. If a type B edge carries at least $1/(k + 1) + \epsilon$ units of flow or a type C edge carries at least $1 + 1/k + \epsilon$ units of flow, we will say that the edge is *oversaturated*. A simple but important observation is that if a Nash flow oversaturates an edge in (H, r, ℓ) for some traffic rate r and subgraph H of B^k , then $c(H, r, \ell) \geq n/2$.

To describe flows in B^k in a convenient way, we require further notation. For $i = 1, \dots, k$, let P_i denote the path $s \rightarrow v_i \rightarrow w_i \rightarrow t$. For $i = 2, \dots, k$, let Q_i denote the path $s \rightarrow v_i \rightarrow$

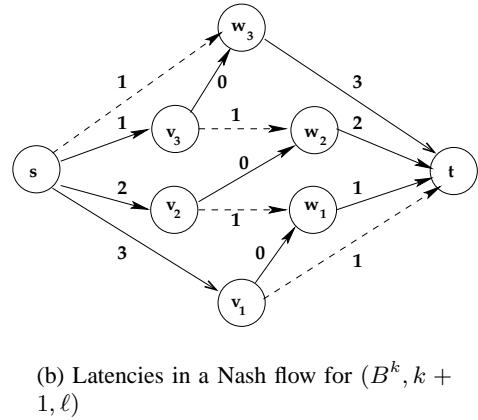
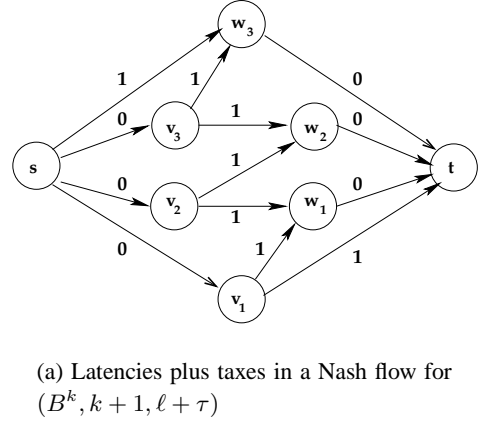


Figure 3: Proof of Theorem 5.2 when $k = 3$. Solid edges carry flow in the flow at Nash equilibrium, dashed edges do not. Edges are labeled with their cost (sum of latency and tax), where latencies are with respect to flows at Nash equilibrium.

$w_{i-1} \rightarrow t$; define Q_1 to be the path $s \rightarrow v_1 \rightarrow t$ and Q_{k+1} the path $s \rightarrow w_k \rightarrow t$.

If we put a tax τ of 1 on each edge of the form (v_i, w_i) and 0 elsewhere, then the following flow is at Nash equilibrium for $(B^k, k + 1, \ell + \tau)$: 1 unit of flow is assigned to each of P_1, P_2, \dots, P_k , and $\frac{1}{k+1}$ units of flow are assigned to each of Q_1, Q_2, \dots, Q_{k+1} (see Figure 3(a)). This Nash flow proves that $c(B^k, k + 1, \ell + \tau) = 1$.

To finish the proof, we must show that $c(H, k + 1, \ell) \geq n/2$ for every subgraph H of B^k ; this requires a bit of case analysis. If H is all of B^k , then $c(H, k + 1, \ell) = n/2$ because routing $1 + 1/k$ units of traffic on each of P_1, P_2, \dots, P_k provides a flow at Nash equilibrium (see Figure 3(b)). This is equally true if H only omits edges eschewed by this Nash flow—the type B edges.

So suppose H omits some type C edge. If edge (s, v_i) is not in H , then any flow in H must oversaturate some edge incident to s , provided $\epsilon > 0$ is sufficiently small. Thus $c(H, k + 1, \ell) \geq n/2$ if H omits a type C edge incident to s . A symmetric argument applies to subgraphs H that omit a type C edge incident to t .

Finally, suppose H omits some type A arc, say (v_i, w_i) . The vertex v_i then has at most one outgoing edge in H , which must be a type B edge. If this edge is not oversaturated, then the edge (s, v_i) carries at most $1/(k+1) + \epsilon$ units of flow; as in the previous paragraph, this implies that some edge incident to s is oversaturated. In either case, $c(H, k+1, \ell) \geq n/2$ and the proof is complete. \square

6. THE COMPLEXITY OF COMPUTING OPTIMAL TAXES

In this section, we study the optimization problem of computing the optimal taxes for a network—taxes that minimize the cost of the ensuing Nash flow. We will see that, as in Section 4, the (in)approximability of this problem is easily resolved by extending previous work on network design [22]. Unfortunately, we will discover that this optimization problem is intractable in a strong sense.

The maximum-possible benefit achievable with taxes, as determined in Section 4, has immediate consequences for the performance guarantee of the *trivial algorithm*—the algorithm that assigns all edges zero tax.

Corollary 6.1 *The trivial algorithm is a:*

- $\frac{4}{3}$ -approximation algorithm for the problem of computing the optimal tax in networks with linear latency functions
- $\Theta(p/\log p)$ -approximation algorithm for the problem of computing the optimal tax in networks with latency functions that are polynomials with degree p and nonnegative coefficients
- $\lfloor n/2 \rfloor$ -approximation algorithm for the problem of computing the optimal tax in n -node networks with arbitrary latency functions.

Roughgarden [22] gave several inapproximability results for the problem of computing the best $0/\infty$ taxes (equivalently, computing the subgraph admitting the best Nash flow). Using the same constructions of [22] and somewhat more involved arguments, most of these inapproximability results can be extended to the problem of computing optimal taxes.

Theorem 6.2 *Unless $P = NP$, for $\epsilon > 0$ there is no*

- $(\frac{4}{3} - \epsilon)$ -approximation algorithm for the problem of computing the optimal tax in networks with linear latency functions
- $o(p/\log p)$ -approximation algorithm for the problem of computing the optimal tax in networks with latency functions that are polynomials with degree p and nonnegative coefficients
- $O(n^{1-\epsilon})$ -approximation algorithm for the problem of computing the optimal tax in n -node networks with arbitrary latency functions.

Remark 6.3 It is interesting to note that the construction of the strongest inapproximability result known for network design [22, Thm 4.3], stating that optimal $0/\infty$ taxes are hard to approximate to within an $\lfloor n/2 \rfloor - \epsilon$ factor in networks with general latency functions, cannot be used to prove negative results for computing more general taxes. We leave open the question of whether some variant of this construction, or some other proof approach, can be used to prove a lower bound of $\lfloor n/2 \rfloor$ on the approximability of computing optimal taxes in general networks.

Acknowledgements

We thank the anonymous referees for their comments.

7. REFERENCES

- [1] A. Archer and É. Tardos. Truthful mechanisms for one-parameter agents. In *Proceedings of the 42nd Annual Symposium on Foundations of Computer Science*, pages 482–491, 2001.
- [2] A. Archer and É. Tardos. Frugal path mechanisms. In *Proceedings of the 13th Annual Symposium on Discrete Algorithms*, pages 991–999, 2002.
- [3] M. Beckmann, C. B. McGuire, and C. B. Winsten. *Studies in the Economics of Transportation*. Yale University Press, 1956.
- [4] P. Bergendorff, D. W. Hearn, and M. V. Ramana. Congestion toll pricing of traffic networks. In P. M. Pardalos, D. W. Hearn, and W. W. Hager, editors, *Network Optimization*, pages 51–71. Springer-Verlag, 1997.
- [5] D. Braess. Über ein paradoxon der verkehrsplanung. *Unternehmensforschung*, 12:258–268, 1968. Available from <http://homepage.ruhr-uni-bochum.de/Dietrich.Braess/>.
- [6] K. J. Button and E. T. Verhoef. *Road Pricing, Traffic Congestion and the Environment*. Edward Elgar, 1998.
- [7] R. Cole, Y. Dodis, and T. Roughgarden. Pricing network edges for heterogeneous selfish users. In *Proceedings of the 35th Annual ACM Symposium on the Theory of Computing*, 2003.
- [8] S. C. Dafermos and A. Nagurney. Sensitivity analysis for the asymmetric network equilibrium problem. *Mathematical Programming*, 28:174–184, 1984.
- [9] S. C. Dafermos and F. T. Sparrow. The traffic assignment problem for a general network. *Journal of Research of the National Bureau of Standards, Series B*, 73B(2):91–118, 1969.
- [10] E. J. Friedman and S. Shenker. Learning and implementation on the Internet. Working paper, 1997.
- [11] M. A. Hall. Properties of the equilibrium state in transportation networks. *Transportation Science*, 12(3):208–216, 1978.
- [12] D. W. Hearn and M. V. Ramana. Solving congestion toll pricing models. In P. Marcotte and S. Nguyen, editors, *Equilibrium and Advanced Transportation Modeling*, pages 109–124. Kluwer Academic Publishers, 1998.
- [13] B. Johansson and L. Mattsson. *Road Pricing: Theory, Empirical Assessment and Policy*. Kluwer, 1995.
- [14] M. Kearns and Y. Mansour. Efficient Nash computation in large population games with bounded influence. In *18th Conference on Uncertainty in Artificial Intelligence*, 2002.
- [15] L. W. McKnight and J. P. Bailey. *Internet Economics*. MIT Press, 1998.

- [16] J. A. Mirrlees. The theory of optimal taxation. In K. J. Arrow and M. D. Intriligator, editors, *Handbook of Mathematical Economics, Volume III*, chapter 24, pages 1197–1249. North-Holland, 1986.
- [17] D. Monderer and L. S. Shapley. Potential games. *Games and Economic Behavior*, 14:124–143, 1996.
- [18] D. Monderer and M. Tennenholtz. Distributed games. *Games and Economic Behavior*, 28:55–72, 1999.
- [19] C. H. Papadimitriou. Algorithms, games, and the Internet. In *Proceedings of the 33rd Annual ACM Symposium on the Theory of Computing*, pages 749–753, 2001.
- [20] A. C. Pigou. *The economics of welfare*. Macmillan, 1920.
- [21] R. W. Rosenthal. A class of games possessing pure-strategy Nash equilibria. *International Journal of Game Theory*, 2:65–67, 1973.
- [22] T. Roughgarden. Designing networks for selfish users is hard. In *Proceedings of the 42nd Annual Symposium on Foundations of Computer Science*, pages 472–481, 2001. Full version available from <http://www.cs.cornell.edu/timr>.
- [23] T. Roughgarden. The price of anarchy is independent of the network topology. In *Proceedings of the 34th Annual ACM Symposium on the Theory of Computing*, pages 428–437, 2002. Full version to appear in *Journal of Computer and System Sciences*.
- [24] T. Roughgarden. *Selfish Routing*. PhD thesis, Cornell University, 2002.
- [25] T. Roughgarden and É. Tardos. Bounding the inefficiency of equilibria in nonatomic congestion games. Technical Report TR2002-1866, Cornell University, 2002.
- [26] T. Roughgarden and É. Tardos. How bad is selfish routing? *Journal of the ACM*, 49(2):236–259, 2002. Preliminary version in *FOCS '00*.
- [27] M. Tennenholtz. Competitive safety analysis: Robust decision-making in multi-agent systems. *Journal of Artificial Intelligence Research*, 17:363–378, 2002.
- [28] J. G. Wardrop. Some theoretical aspects of road traffic research. In *Proceedings of the Institute of Civil Engineers, Pt. II*, volume 1, pages 325–378, 1952.

APPENDIX

A. PROOF OF THEOREM 5.1

In this appendix we complete the proof of Theorem 5.1. It will be convenient to abuse notation and, with respect to an instance (G, r, ℓ) , to write $C(\tau)$ for $C(f^\tau, \tau)$, where f^τ is at Nash equilibrium for $(G, r, \ell + \tau)$. We note that $C(\tau)$ is well defined by Proposition 2.5(b).

To complete the proof of Theorem 5.1, it suffices to prove the following proposition. By *the minimal instance*, we mean the minimal counterexample from the proof of Theorem 5.1. Recall that an optimal tax is one that minimizes $C(\tau)$.

Proposition A.1 *Let O denote the set of optimal taxes for the minimal instance. Then the infimum $\inf_{\tau \in O} \sum_e \tau_e$ is attained by some tax τ .*

To prove Proposition A.1, we require two technical lemmas. The first states that bounded taxes suffice to minimize the cost of a Nash flow.

Lemma A.2 *Let (G, r, ℓ) be the minimal instance. Let G have n vertices and define $\ell_{max} = \max_{e \in E} \ell_e(r)$. If $\tau \geq 0$ is a tax, then there is a tax $\tilde{\tau}$ with $\max_e \tilde{\tau}_e \leq n\ell_{max}$ and $C(\tilde{\tau}) \leq C(\tau)$.*

PROOF. Let τ be any tax for (G, r, ℓ) , with f^τ at Nash equilibrium for $(G, r, \ell + \tau)$. We next show how to decrease taxes while leaving f^τ at Nash equilibrium.

By property (4) of the minimal instance, G is directed acyclic and we can therefore topologically sort the vertices of G so that all edges of G travel forward. Since $f_e^\tau > 0$ on all edges of G (property (1) of the minimal instance), s is the first vertex in the ordering and t is the last. Beginning with the penultimate vertex and proceeding backwards in the vertex ordering, we perform the following operation for each vertex $v \neq s$: if $\tau_v \geq 0$ is the minimum tax on any edge with tail v , subtract τ_v from the tax on each edge with tail v and add τ_v to the tax on each edge with head v . This operation does not affect the total tax on any s - t path nor the cost of any flow and leaves f^τ at Nash equilibrium. When the source s is reached, subtract τ_s from the tax on all edges with tail s ; f^τ remains at Nash equilibrium and its cost can only decrease. Call the new set of taxes $\tilde{\tau}$. We have already argued that $C(\tilde{\tau}) \leq C(\tau)$; it remains to show that $\max_e \tau_e \leq n\ell_{max}$.

We first observe that the tax-reducing operations ensure the following property: every vertex other than t has an outgoing edge with zero $\tilde{\tau}$ -tax. From this, it follows that an entire s - t path, say P_0 , has zero $\tilde{\tau}$ -tax. By property (2) of the minimal instance,

$$\ell_P(f^\tau) + \tilde{\tau}_P = \ell_{P_0}(f^\tau)$$

for all paths $P \in \mathcal{P}$. Since $\ell_{P_0}(f^\tau) \leq n\ell_{max}$, no edge tax in $\tilde{\tau}$ can exceed $n\ell_{max}$, and the proof is complete. \square

The second technical lemma asserts continuity of the map $\tau \mapsto C(\tau)$.

Lemma A.3 *The map $\tau \mapsto C(\tau)$, with respect to the minimal instance (G, r, ℓ) , is continuous.*

PROOF. Property (3) of the minimal instance ensures that the map $\tau \mapsto \{f_e^\tau\}$ is uniquely defined. Property (5) and a result of Dafermos and Nagurny [8, Thm 3.1] imply that this map is continuous. The lemma then follows easily. \square

Proof of Proposition A.1. We first claim that O , the set of optimal taxes for the minimal instance, is not empty. This is tantamount to showing that the infimum $\inf_{\tau \geq 0} C(\tau)$ is attained. Let B denote the taxes with all components bounded by $n\ell_{max}$, where n is the number of vertices of G and $\ell_{max} = \max_{e \in E} \ell_e(r)$. By Lemma A.2, $\inf_{\tau \in B} C(\tau) = \inf_{\tau \geq 0} C(\tau)$. Since B is a compact subset of \mathcal{R}^E and C is continuous by Lemma A.3, the infimum in $\inf_{\tau \in B} C(\tau)$, and hence in $\inf_{\tau \geq 0} C(\tau)$, is attained.

It remains to show that the infimum $\inf_{\tau \in O} \sum_e \tau_e$ is attained. By Lemma A.2, there is some optimal tax $\tau \in O$ with $\sum_{e \in E} \tau_e \leq mn\ell_{max}$ where m is the number of edges of G . Since O is the inverse image of a closed set under a continuous map (write $O = C^{-1}(C(\tau))$ for some $\tau \in O$), it is a closed subset of \mathcal{R}^E . Restricting O to taxes with sum of all components at most $mn\ell_{max}$, we obtain a nonempty compact subset S of optimal taxes. Since $\tau \mapsto \sum_e \tau_e$ is a continuous function, it attains a minimum on S ; this must also be its minimum on O , and the proof is complete. ■