

Optimal Mechanisms for Robust Coordination in Congestion Games

Philip N. Brown and Jason R. Marden

Abstract—Uninfluenced social systems often exhibit sub-optimal performance; a common mitigation technique is to charge agents specially-designed taxes, influencing the agents’ choices and thereby bringing aggregate social behavior closer to optimal. In general, the efficiency guaranteed by a particular taxation methodology is limited by the quality of information available to the tax-designer. If the tax-designer possesses a perfect characterization of the system, it is often straightforward to design taxes which perfectly optimize the behavior of the agent population. In this paper, we investigate situations in which the tax-designer lacks such a perfect characterization and must design taxes that are robust to a variety of model imperfections. Specifically, we study the application of taxes to a network-routing game, and we assume that the tax-designer knows neither the network topology nor the tax-sensitivities and demands of the agents. Nonetheless, we show that it is possible to design taxes that guarantee that network flows are arbitrarily close to optimal flows, despite the fact that agents’ tax-sensitivities are unknown to us. We term these taxes “universal,” since they enforce optimal behavior in any routing game without *a priori* knowledge of the specific game parameters. In general, these taxes may be very high; accordingly, for affine-cost parallel-network routing games, we explicitly derive the optimal bounded tolls and the best-possible efficiency guarantee as a function of a toll upper-bound.

I. INTRODUCTION

It is well-known that in systems that are driven by social behavior, agents’ self-interested behavior can lead to significant system-level inefficiencies. This inefficiency is commonly referred to as the *price of anarchy*; defined as the ratio between the worst-case social welfare resulting from selfish behavior and the optimal social welfare [1]. This inefficiency due to selfish behavior has been the the subject of research in the areas of network resource allocation [2], distributed control [3], traffic congestion [4]–[6], and others. As a result, there is a growing body of research geared at influencing social behavior to improve system performance [7]–[13].

To study the issues surrounding the problem of influencing selfish social behavior, we turn to a simple model of traffic routing: a unit mass of traffic needs to be routed across a network in such a way that minimizes the average network transit time. If a central planner has the ability to direct traffic explicitly, it is straightforward to compute the routing profile that minimizes total congestion. However, in real systems, it

may not be possible to implement such direct centralized control: for example, if the network represents a city’s road network, individual drivers make their own routing choices in response to their own personal objectives.

Accordingly, we model this routing problem as a non-atomic congestion game, where the traffic can be viewed as a collection of infinitely-many users, each controlling an infinitesimally-small amount of traffic and seeking to minimize its own experienced transit time. We use the concept of a *Nash flow* (defined as a routing profile in which no user can switch to a different path and decrease her transit delay) to characterize the routing profile resulting from such self-interested behavior. It is widely known that Nash flows can be significantly less efficient than optimal flows; an important result in this setting states that a Nash flow on a network with linear-affine latency functions can be up to 33% worse than the optimal flow; that is, the price of anarchy in this setting is $4/3$. For networks with general latency functions, the price of anarchy can be unbounded [14].

A natural approach to mitigating this inefficiency is to charge monetary taxes for the use of network links, thereby modifying the users’ costs and inducing a new, more efficient Nash flow. Existing research has shown that it is possible to design such optimal taxes given that the tax-designer has access to certain information regarding the system. Typically, these types of results have strict informational requirements; for example, in [15]–[17] it is shown that optimal “fixed” taxes can be computed for any routing game, but the computation requires precise characterizations of the network topology, user demands, and user tax-sensitivities. In contrast, [18], [19] derive optimal taxes known as “marginal-cost taxes” which require no knowledge of the network topology or user demands, but require that all users share a common tax-sensitivity. In Section III, we survey these existing results in greater detail; see the first two rows of Table I for a side-by-side comparison of the particular design constraints and informational dependencies of these two taxation mechanisms.

In this paper, we ask if it is possible to compute optimal taxes with no information about the system. Our main contribution is the derivation of a universal taxation mechanism that guarantees arbitrarily-high efficiency for any routing game without requiring *a priori* knowledge of the specific network or distribution of user sensitivities or demands. This result holds for networks with very general latency functions and any topology. In the third row of Table I, we summarize this contribution in context with some existing results.

Since very high tolls may be impossible (or politically unpalatable) to implement, our second contribution is to

This research was supported by AFOSR grant #FA9550-12-1-0359, ONR grant #N00014-12-1-0643, and NSF Grant #ECCS-1351866.

P. N. Brown is a graduate research assistant with the Department of Electrical, Computer, and Energy Engineering, University of Colorado, Boulder, CO 80309, philip.brown@colorado.edu. Corresponding author.

J. R. Marden is with the Department of Electrical, Computer, and Energy Engineering, University of Colorado, Boulder, CO 80309, jason.marden@colorado.edu.

TABLE I

Toll Type	Design Constraints		Informational Dependencies			Efficiency Guarantee
	Constant Functions	Boundedness	Topology	Demands	Sensitivities	
Fixed [16], [17]	✓		✓	✓	✓	100%
Marginal-Cost [18], [19]					✓	100%
Universal (Theorem 1)						100%
Bounded Affine (Theorem 2)		✓		—		Function of toll upper bound

Design constraints, informational dependencies and efficiency guarantees of several taxation methodologies. Fixed tolls are simple constant functions of flow, but to to guarantee optimality, they depend heavily on a precise system characterization. Marginal-cost tolls, though flow-varying, guarantee optimality while only requiring knowledge of the (homogeneous) user-sensitivities. In this paper, Theorem 1 defines tolls which require none of the above information, but are flow-varying and may be arbitrarily large. By contrast, Theorem 2 derives the optimal bounded tolls for a sub-class of networks, and guarantees efficiency that is increasing in the toll upper-bound.

explore the effect of an upper bound on the allowable tolling functions. To that end, for parallel networks with linear-affine latency functions, we derive the optimal tolling functions that minimize worst-case efficiency losses for any unknown distribution of user sensitivities and toll upper bound, again requiring no *a priori* knowledge of the network topology. Surprisingly, these optimal tolls are simple affine functions of flow. We show that for parallel networks with linear-affine cost functions and simple user demands, the worst-case efficiency losses decrease monotonically with the toll upper bound, illustrating the concept that we can compensate for a poor characterization of user sensitivities by charging higher tolls. The last row of Table I pertains to this result.

II. MODEL AND PERFORMANCE METRICS

A. Routing Game

Consider a network routing problem in which a unit mass of traffic needs to be routed across a network (V, E) , which consists of a vertex set V and edge set $E \subseteq (V \times V)$. We call a source/destination vertex pair $(s_c, t_c) \in (V \times V)$ a *commodity*, and the set of all commodities \mathcal{C} . We assume that for each $c \in \mathcal{C}$, there is a mass of traffic $r_c > 0$ that needs to be routed from s_c to t_c . We write $\mathcal{P}_c \subset 2^E$ to denote the set of *paths* available to traffic in commodity c , where each path $p \in \mathcal{P}_c$ consists of a set of edges connecting s_c to t_c . Let $\mathcal{P} = \cup \{\mathcal{P}_c\}$.

A *feasible flow* $f \in \mathbb{R}^{|\mathcal{P}|}$ is an assignment of traffic to various paths such that for each commodity, $\sum_{p \in \mathcal{P}_c} f_p = r_c$, where $f_p \geq 0$ denotes the mass of traffic on path p . Without loss of generality, we assume that $\sum_{c \in \mathcal{C}} r_c = 1$.

Given a flow f , the flow on edge e is given by $f_e = \sum_{p: e \in p} f_p$. To characterize transit delay as a function of traffic flow, each edge $e \in E$ is associated with a specific latency function $\ell_e : [0, 1] \rightarrow [0, \infty)$. We assume that latency functions are nondecreasing, continuously differentiable, and convex. We measure the efficiency of a flow f by the *total latency*, given by

$$\mathcal{L}(f) = \sum_{e \in E} f_e \cdot \ell_e(f_e) = \sum_{p \in \mathcal{P}} f_p \cdot \ell_p(f_p), \quad (1)$$

where $\ell_p(f) = \sum_{e \in p} \ell_e(f_e)$ denotes the latency on path p . We denote the flow that minimizes the total latency by

$$f^* \in \underset{f \text{ is feasible}}{\operatorname{argmin}} \mathcal{L}(f). \quad (2)$$

Due to the convexity of ℓ_e , $\mathcal{L}(f^*)$ is unique.

A *routing problem* is given by the tuple $G = (V, E, \mathcal{C}, \{\ell_e\})$. We write the set of all such routing problems as \mathcal{G} . We will often use shorthand notation such as $e \in \mathcal{G}$ to denote $(e \in G : G \in \mathcal{G})$.

In this paper we study taxation mechanisms for influencing the emergent collective behavior resulting from self-interested price-sensitive users. To that end, we model the above routing problem as a non-atomic congestion game. We assign each edge $e \in E$ a flow-dependent, nondecreasing taxation function $\tau_e : \mathbb{R}^+ \rightarrow \mathbb{R}^+$. We characterize the taxation sensitivities of the users in commodity c in the following way: let each user $x \in [0, r_c]$ have a taxation sensitivity $s_x^c \in [S_L, S_U] \subseteq \mathbb{R}^+$ where $S_U \geq S_L \geq 0$ denote upper and lower sensitivity bounds, respectively. Given a flow f , the cost that user x experiences for using path $\tilde{p} \in \mathcal{P}_c$ is of the form

$$J_x(f) = \sum_{e \in \tilde{p}} [\ell_e(f_e) + s_x^c \tau_e(f_e)], \quad (3)$$

and we assume that each user selects the lowest-cost path from the available source-destination paths. We call a flow f a *Nash flow* if for all commodities $c \in \mathcal{C}$ and all users $x \in [0, r_c]$ we have

$$J_x(f) = \min_{p \in \mathcal{P}_c} \left\{ \sum_{e \in p} [\ell_e(f_e) + s_x^c \tau_e(f_e)] \right\}. \quad (4)$$

It is well-known that a Nash flow exists for any non-atomic congestion game of the above form [20].

In our analysis, we assume that the sensitivity distribution function s is unknown; for a given routing problem G and $S_U \geq S_L \geq 0$ we define the set of possible sensitivity distributions as the set of Lebesgue-measurable functions $\mathcal{S} = \{s^c : [0, r_c] \rightarrow [S_L, S_U]\}_{c \in \mathcal{C}}$.

For a given routing problem $G \in \mathcal{G}$, we gauge the efficacy of a collection of taxation functions $\tau = \{\tau_e\}_{e \in E}$ by comparing the total latency of the resulting Nash flow and the total latency associated with the optimal flow, and then performing a worst-case analysis over all possible sensitivity distributions. Let $\mathcal{L}^*(G)$ denote the total latency associated with the optimal flow, and $\mathcal{L}^{\text{nf}}(G, s, \tau)$ denote the total latency of the Nash flow resulting from taxation functions τ and sensitivity distributions s . The worst-case efficiency loss associated with this specific instance is captured by the *price of anarchy* which is of the form

$$\text{PoA}(G, \tau) = \sup_{s \in \mathcal{S}} \left\{ \frac{\mathcal{L}^{\text{nf}}(G, s, \tau)}{\mathcal{L}^*(G)} \right\} \geq 1.$$

B. Summary of Our Contributions

In Theorem 1 we prove that if each edge’s taxation function is given by

$$\tau_e(f_e) = \kappa \left(\ell_e(f_e) + f_e \frac{d}{df_e} \ell_e(f_e) \right), \quad (5)$$

for $\kappa \in \mathbb{R}^+$, the price of anarchy converges to 1 as κ approaches infinity, for any user sensitivities and network topology. Thus, the toll designer can enforce arbitrarily-high efficiency simply by charging these tolls with sufficiently high κ . Note that these tolls are universal in the sense that they have no dependence on the specific network or sensitivity distribution.

However, in some situations it may be impractical to charge very high tolls; for example, it may be politically unpalatable, or there may be a degree of elasticity in network demand. Accordingly, in Theorem 2, we investigate the effect of an upper bound T on allowable tolling functions for single-commodity parallel networks in which each ℓ_e is linear-affine. Though one might expect that the optimal tolling functions in this situation would be equal to the toll presented in Theorem 1 for some value of κ , this is not generally the case.

Theorem 2 shows that there exist functions $\kappa_1(\mathcal{G}, \mathcal{S}, T)$ and $\kappa_2(\mathcal{G}, \mathcal{S}, T)$ such that if an edge’s latency function is $\ell_e(f_e) = a_e f_e + b_e$, the optimal tolling function is given by

$$\tau_e(f_e) = \kappa_1(\mathcal{G}, \mathcal{S}, T) a_e f_e + \kappa_2(\mathcal{G}, \mathcal{S}, T) b_e, \quad (6)$$

and we derive expressions for the price of anarchy when using this tolling methodology. Since $\kappa_1(\cdot)$ and $\kappa_2(\cdot)$ do not depend on instance-specific parameters, these tolls can be applied without *a priori* knowledge of the specific routing instance. Thus, these efficiency guarantees are robust to a wide variety of mischaracterizations of the routing scenario.

For the simple 2-link network known as Pigou’s Example depicted in Figure 1, we plot the price of anarchy resulting from the the taxation mechanisms proposed in Theorems 1 and 2 with respect to a toll upper bound. Note that though both curves converge to 1 (i.e., they both guarantee perfectly optimal flows in the large-toll limit), the tolls from Theorem 2 converge much more quickly. This shows that the universal guarantees made by Theorem 1 come at a price: if we have additional information about the specific class of networks, we may be able to guarantee significantly higher efficiency for a given upper bound.

III. RELATED WORK

There has been significant research geared towards developing taxation mechanisms to eradicate the inefficiency caused by users’ self-interested routing choices. A taxation mechanism simply computes edge tolls as a function of some set of information about the system; here we survey the informational dependencies of several taxation approaches in the literature.

– *Omniscient taxation mechanisms*: These taxation mechanisms are assumed to have access to complete information regarding the routing game. For edge $e \in G$, with sensitivity




Fig. 1. Price of Anarchy plot contrasting the Universal toll result in Theorem 1 (dashed line) with the optimal toll result from Theorem 2 (solid line). Note that the price of anarchy of either taxation mechanism converges to 1 as the upper bound increases, but the solid line converges much more quickly. This is because Theorem 2 gives the optimal tolls for a specific class of networks, but the universal tolls from Theorem 1 are designed to work on *all* classes of networks.

distribution $s \in \mathcal{S}$, the edge tolling function takes the following form: $\tau_e(f_e; G, s)$. That is, each edge’s taxation function can depend on the entire routing problem G and the sensitivity distribution s . Recent results have identified taxation mechanisms of this form that assign fixed tolls (i.e., for any $e \in G$, $\tau_e(f_e) = q_e$ for some $q_e \geq 0$) that guarantee a price of anarchy of 1 [16], [17].

However, the robustness of these mechanisms to variations or mischaracterizations of network topology is unknown, and in [21], the authors show that fixed tolls can never guarantee a price of anarchy of 1 if the user sensitivities are unknown.

– *Network-agnostic taxation mechanisms*: This type of taxation mechanism is agnostic to network specifications. Here, a system designer essentially commits to a taxation function for each potential edge $e \in \mathcal{G}$, and any network realization $(G, s) \in \mathcal{G} \times \mathcal{S}$ merely employs a subset of these pre-defined taxation functions. An edge’s toll cannot depend on any *other* edge’s congestion properties or location in the network.

A commonly-studied network-agnostic taxation mechanism is the marginal-cost (or Pigovian) taxation mechanism τ^{mc} , which is of the following form: for any $e \in \mathcal{G}$ with latency function ℓ_e , the accompanying taxation function is

$$\tau^{\text{mc}}(f_e) = f_e \cdot \frac{d}{df_e} \ell_e(f_e), \quad \forall f_e \geq 0. \quad (7)$$

In [18] the author shows that for any $G \in \mathcal{G}$ we have $\mathcal{L}^*(G) = \mathcal{L}^{\text{nf}}(G, s, \tau^{\text{mc}})$ provided that all users have a sensitivity exactly equal to 1. Hence, irrespective of the underlying network structure, a marginal-cost taxation mechanism always ensures the optimality of the resulting Nash flow, provided that all users share a common known sensitivity.

Finally, in [21], the authors show that marginal-cost taxes scaled by $\sqrt{S_L S_U}$ do possess a degree of robustness to mischaracterizations of user sensitivities, but can no longer guarantee a price of anarchy of 1.

IV. THEOREM 1: A UNIVERSAL TAXATION MECHANISM

In this paper, we prove that network-agnostic tolls exist which can drive the price of anarchy to 1 for general networks and latency functions. We term these “universal” because they take the same form and provide the same efficiency guarantee regardless of which particular routing scenario they are applied to. Using this taxation mechanism, we show in Theorem 1 that for all networks, regardless of

network topology, user demands, or price-sensitivity functions, the price of anarchy can be made arbitrarily close to 1 if we allow edge tolls to be sufficiently high.

Theorem 1: For any network edge $e \in \mathcal{G}$ with convex, nondecreasing, continuously differentiable latency function ℓ_e , define the generalized Pigovian taxation function on edge e as

$$\tau^{\text{gpt}}(\kappa) = \kappa \left(\ell_e(f_e) + f_e \cdot \frac{d}{df_e} \ell_e(f_e) \right). \quad (8)$$

Then for any routing problem $G \in \mathcal{G}$ and any $S_U \geq S_L > 0$,

$$\lim_{\kappa \rightarrow \infty} \text{PoA}(G, \tau^{\text{gpt}}(\kappa)) = 1. \quad (9)$$

That is, on *any* network being used by *any* population of users, the total latency can be made arbitrarily close to the optimal latency, and each individual link toll is a simple continuous function of that link's flow. The reason for this is that as κ increases, the original latency function has a smaller and smaller relative effect on the users' cost functions; in the large-toll limit, the only cost experienced by the users is the tolling function itself which is specifically designed to induce optimal Nash flows.

Proof: Using a sequence of tolls, we construct a sequence of Nash flows that converges to an optimal flow. Let κ_n be an unbounded, increasing sequence of tolling coefficients.

For any routing problem $G \in \mathcal{G}$ and price-sensitivities $s \in \mathcal{S}$, let $f^n = (f_p^n)_{p \in \mathcal{P}}$ denote the Nash flow resulting from the tolling coefficient κ_n . For each commodity c , let $\mathcal{P}_c^n \subseteq \mathcal{P}_c$ denote the set of paths that have positive flow in f^n . For any $p \in \mathcal{P}_c^n$, there must be some user $x \in [0, r_c]$ using p ; suppose this user has sensitivity s_x^c , then the cost experienced by this user is given by

$$J_x(f^n) = \sum_{e \in p} \left[\ell_e(f_e) + \kappa_n s_x^c \left(\ell_e(f_e) + f_e \cdot \frac{d}{df_e} \ell_e(f_e) \right) \right].$$

Define

$$\gamma_{n,x} = \frac{\kappa_n s_x^c}{1 + \kappa_n s_x^c}.$$

Let $\ell_e^*(f_e) = f_e \cdot \frac{d}{df_e} \ell_e(f_e)$; then for any other path $p' \in \mathcal{P}_c \setminus p$, user x must experience a lower cost on p than on p' , or

$$\sum_{e \in p} \ell_e(f_e) - \sum_{e \in p'} \ell_e(f_e) \leq \gamma_{n,x} \left[\sum_{e \in p'} \ell_e^*(f_e) - \sum_{e \in p} \ell_e^*(f_e) \right]. \quad (10)$$

Therefore, for any $n \geq 1$, f^n must satisfy some set of inequalities defined by (10). Note that for all $c \in \mathcal{C}$ and any $x \in [0, r_c]$, $\lim_{\kappa_n \rightarrow \infty} \gamma_{n,x} = 1$, so because all the functions in (10) are continuous, f^n converges to a set F^* of feasible flows that satisfy

$$\sum_{e \in p} \ell_e(f_e) - \sum_{e \in p'} \ell_e(f_e) \leq \left[\sum_{e \in p'} \ell_e^*(f_e) - \sum_{e \in p} \ell_e^*(f_e) \right] \quad (11)$$

for all c , all $p \in \mathcal{P}_c^*$, and $p' \in \mathcal{P}_c$, where $\mathcal{P}_c^* \subseteq \mathcal{P}_c$ is some subset of paths. But inequalities (11) (combined with




Fig. 2. Figure for Example 1. The plot shows the Nash flows and price of anarchy as a function of κ . The bold curves represent the possible edge-1 Nash flows as a function of κ : the gray-shaded area highlights all Nash flows that could result from some sensitivity distribution in $[1, 10]$. Note that if $\kappa = 0$, the figure shows a Nash flow with $f_1 = 1$, but that the Nash flows in the shaded area converge to $f_1 = 1/2$ as $\kappa \rightarrow \infty$. The dotted horizontal lines show the price of anarchy that results from a flow at that level. Note that the price of anarchy decreases rapidly with κ , and by the time κ is greater than 10, the price of anarchy is already well below 1.01.

the feasibility constraints on f) also specify a Nash flow for G for a unit-sensitivity population with marginal-cost taxes as defined in (7); any such Nash flow must be optimal [18]; that is, any $f \in F^*$ is a minimum-latency flow for G . Thus, since $\mathcal{L}(f)$ is a continuous function of f ,

$$\lim_{n \rightarrow \infty} \mathcal{L}(f^n) = \mathcal{L}^*(G), \quad (12)$$

obtaining the proof of the theorem. \blacksquare

Example 1 [An Application of Theorem 1] Consider again the simple two-link network depicted on the left in Figure 1; this is the canonical network known as ‘‘Pigou’s Example.’’ An un-tolled Nash flow on this network has all traffic using the upper congestion-sensitive link (with a total latency of 1), while the optimal flow has the traffic split evenly between link 1 and link 2 (with a total latency of 0.75), for a price of anarchy of $4/3$.

Suppose we only know the toll-sensitivities of the user population to within 10%, or $S_L = 1$ and $S_U = 10$, and we wish to design tolls that reduce the price of anarchy as close to 1 as possible. On this network, Theorem 1 assigns tolling functions $\tau_1(f_1) = 2\kappa f_1$ and $\tau_2(f_2) = \kappa$; we simply need to set κ high enough to achieve our desired performance.

Figure 2 shows plots of the Nash flows and price of anarchy as a function of κ . Note that for a two-link network, a network flow is uniquely determined by the flow on a single edge. Thus, the bold curves represent the possible edge-1 Nash flows as a function of κ ; the gray-shaded area highlights all Nash flows that could result from some sensitivity distribution in $[1, 10]$. Note that if $\kappa = 0$, the figure shows a Nash flow with $f_1 = 1$, but that any sequence of Nash flows in the shaded area converges to $f_1 = 1/2$ as $\kappa \rightarrow \infty$. The dotted horizontal lines show the price of anarchy that results from a flow at that level. Note that the price of anarchy decreases rapidly with κ , and by the time κ is greater than 10, the price of anarchy is already well below 1.01.

V. THEOREM 2: OPTIMAL BOUNDED TOLLS

Of course, it may be impractical or politically infeasible to charge extremely high tolls. Therefore, in Theorem 2, we analyze the effect of placing an upper bound on the allowable tolling functions. If tolling functions are bounded, we show that the price of anarchy is strictly decreasing in the toll upper bound, and analytically characterize the effect of this upper bound for single-commodity parallel-network routing games. Additionally, we show that for routing games with affine costs, linear-affine tolling functions are sufficient to achieve the optimal price of anarchy given a toll upper bound. That is, we have no need to consider more complicated classes of tolling functions.

For parallel networks with affine cost functions in which every edge has positive flow in an un-tolled Nash flow, we explicitly derive the optimal bounded taxation mechanism, and then provide an expression for the price of anarchy. To this end, we say a taxation mechanism is *bounded* if it never assigns taxation functions that exceed some upper bound:

Definition 1: Taxation mechanism τ is *bounded* by T on a class of routing problems $\bar{\mathcal{G}}$ if for every edge $e \in \bar{\mathcal{G}}$, τ assigns a tolling function that satisfies

$$\tau_e : [0, 1] \rightarrow [0, T]. \quad (13)$$

We write the set of taxation mechanisms bounded by T on $\bar{\mathcal{G}}$ as $\mathcal{T}(T, \bar{\mathcal{G}})$.

For the following results, let $\mathcal{G}^p \subseteq \mathcal{G}$ represent the class of all single-commodity, parallel-link routing problems with affine latency functions. That is, for all $e \in \mathcal{G}^p$, the latency function satisfies

$$\ell_e(f_e) = a_e f_e + b_e \quad (14)$$

where a_e and b_e are non-negative edge-specific constants. By ‘‘single-commodity,’’ we mean that all traffic has access to all network edges. Furthermore, we assume that every edge has positive flow in an un-tolled Nash flow. It will be necessary to describe classes of networks with bounded latency functions; to this end, we define $\mathcal{G}(\bar{a}, \bar{b}) \subset \mathcal{G}^p$ as the set of parallel, affine-cost networks such that for every $e \in \mathcal{G}(\bar{a}, \bar{b})$, the latency function coefficients satisfy $a_e \leq \bar{a}$ and $b_e \leq \bar{b}$.

Definition 2: For every edge $e \in \mathcal{G}$ with latency function ℓ_e a *network-agnostic* taxation mechanism is a mapping $\tau^{\text{na}} : [0, 1] \times \{\ell_e\}_{e \in \mathcal{G}} \rightarrow \{\tau_e\}$ that assigns the following flow-dependent taxation function to edge e :

$$\tau_e(f_e) = \tau^{\text{na}}(f_e; \ell_e) \quad (15)$$

where $\tau^{\text{na}}(f, \ell)$ satisfies the following additivity condition:¹ for all $e, e' \in \mathcal{G}$ and $f \in [0, 1]$,

$$\tau^{\text{na}}(f; \ell_e + \ell_{e'}) = \tau^{\text{na}}(f; \ell_e) + \tau^{\text{na}}(f; \ell_{e'}). \quad (16)$$

Note that by this definition, the universal taxation mechanism we defined in Theorem 1 is network-agnostic.

¹The additivity condition in Definition 2 is a natural assumption which simply ensures that two edges connected in series will be assigned the same taxation function as if they were replaced by a single edge whose latency function is the sum of the underlying latency functions.

Our goal is to derive the bounded network-agnostic taxation mechanism that minimizes the worst-case selfish routing on \mathcal{G}^p . We define the price of anarchy with respect to class of routing problems \mathcal{G} and bound T as the best price of anarchy we can achieve on \mathcal{G} with a taxation mechanism bounded by T :

$$\text{PoA}(\mathcal{G}, T) = \inf_{\tau \in \mathcal{T}(T, \mathcal{G})} \left\{ \sup_{G \in \mathcal{G}} \text{PoA}(G, \tau) \right\}. \quad (17)$$

Theorem 2: Let $\mathcal{G}(\bar{a}, \bar{b}) \subset \mathcal{G}^p$ be some subset of parallel, affine-cost networks with finite \bar{a} and \bar{b} . For any toll bound T and $S_U \geq S_L > 0$, define the set of universal parameters by the tuple $U = (S_L, S_U, \bar{a}, \bar{b}, T)$. Then there exist functions $\kappa_1(U)$ and $\kappa_2(U)$ such that the optimal network-agnostic taxation mechanism bounded by T on $\mathcal{G}(\bar{a}, \bar{b})$ assigns tolling functions

$$\tau_e(f_e) = \kappa_1(U) a_e f_e + \kappa_2(U) b_e. \quad (18)$$

Furthermore, the price of anarchy $\text{PoA}(\mathcal{G}(\bar{a}, \bar{b}), T)$ is given by the following:

$$\begin{cases} \frac{4}{3} \left(1 - \frac{\kappa_1(U) S_L}{(1 + \kappa_1(U) S_L)^2} \right) & \text{if } \kappa_1(U) < \frac{1}{\sqrt{S_L S_U}} \\ \frac{4}{3} \left(1 - \frac{(1 + \kappa_1(U) S_L) \left(\frac{S_L}{S_U} + \kappa_1(U) S_L \right)}{(1 + 2\kappa_1(U) S_L + \frac{S_L}{S_U})^2} \right) & \text{if } \kappa_1(U) \geq \frac{1}{\sqrt{S_L S_U}}. \end{cases} \quad (19)$$

For the reader’s convenience, we include a closed-form expression for $\kappa_1(\cdot)$ in the appendix as (35), and for $\kappa_2(\cdot)$ in the proof of Theorem 2 as (27). It is evident from these expressions that $\kappa_1(\cdot)$ and $\kappa_2(\cdot)$ are both nondecreasing and unbounded in T ; among other things, this implies that $\lim_{T \rightarrow \infty} \text{PoA}(\mathcal{G}(\bar{a}, \bar{b}), T) = 1$. Qualitatively, it is important to note that they depend only on parameters that are common to all network edges. Thus, the above price of anarchy expression is universal in the sense that it applies to all networks in the class $\mathcal{G}(\bar{a}, \bar{b})$.

We now proceed with the proof of Theorem 2, which relies on two supporting lemmas. For our first milestone, we restrict attention to simple affine tolling functions:

Lemma 2.1: Let $\tau^A(\kappa_1, \kappa_2)$ denote an affine taxation mechanism that assigns tolling functions $\tau_e(f_e) = \kappa_1 a_e f_e + \kappa_2 b_e$. For any $\kappa_{\max} \geq 0$, the optimal coefficients κ_1^* and κ_2^* satisfying

$$(\kappa_1^*, \kappa_2^*) \in \arg \min_{\kappa_1, \kappa_2 \leq \kappa_{\max}} \left\{ \sup_{G \in \mathcal{G}^p} \text{PoA}(G, \tau^A(\kappa_1, \kappa_2)) \right\} \quad (20)$$

are given by

$$\kappa_1^* = \kappa_{\max}, \quad (21)$$

$$\kappa_2^* = \max \left\{ 0, \frac{\kappa_{\max}^2 S_L S_U - 1}{S_L + S_U + 2\kappa_{\max} S_L S_U} \right\}. \quad (22)$$

Furthermore, the price of anarchy $\text{PoA}(\mathcal{G}, \tau^A(\kappa_1^*, \kappa_2^*))$ is given by the following expression:

$$\begin{cases} \frac{4}{3} \left(1 - \frac{\kappa_{\max} S_L}{(1 + \kappa_{\max} S_L)^2} \right) & \text{if } \kappa_{\max} < \frac{1}{\sqrt{S_L S_U}} \\ \frac{4}{3} \left(1 - \frac{(1 + \kappa_{\max} S_L) \left(\frac{S_L}{S_U} + \kappa_{\max} S_L \right)}{(1 + 2\kappa_{\max} S_L + \frac{S_L}{S_U})^2} \right) & \text{if } \kappa_{\max} \geq \frac{1}{\sqrt{S_L S_U}}. \end{cases} \quad (23)$$

See the Appendix for the proof of Lemma 2.1.

Next, in Lemma 2.2, we investigate the possibility that some other taxation mechanism could perform better than the affine $\tau^A(\kappa_1^*, \kappa_2^*)$ while still respecting the bound T . To that end, we assume that some arbitrary taxation mechanism outperforms affine tolls, and deduce various properties of these hypothetical tolls. We show that this hypothetical “better” taxation mechanism must universally charge higher tolls than our optimal affine tolls.

Lemma 2.2: Let τ^* be any network-agnostic taxation mechanism such that for $\kappa_{\max} \geq 0$

$$\text{PoA}(\mathcal{G}^p, \tau^*) < \text{PoA}(\mathcal{G}^p, \tau^A(\kappa_1^*, \kappa_2^*)). \quad (24)$$

Then τ^* must charge strictly higher tolls than $\tau^A(\kappa_1^*, \kappa_2^*)$ on every edge in every network:

$$\forall e \in \mathcal{G}^p, \forall f_e \in [0, 1], \tau_e^*(f_e) > \tau_e^A(f_e). \quad (25)$$

The proof of Lemma 2.2 appears in the Appendix.

Proof: [Theorem 2] For any non-negative κ_1 and κ_2 , $\tau^A(\kappa_1, \kappa_2)$ is tightly bounded by $(\kappa_1 \bar{a} + \kappa_2 \bar{b})$ on $\mathcal{G}(\bar{a}, \bar{b})$. Note that for κ_1^* and κ_2^* as defined in Lemma 2.1, $(\kappa_1^* \bar{a} + \kappa_2^* \bar{b})$ is a strictly increasing, continuous function of κ_{\max} . Thus, for any $T \geq 0$, there is a unique $\kappa_{\max}^* \geq 0$ for which $\tau^A(\kappa_1^*, \kappa_2^*)$ is tightly bounded by T on $\mathcal{G}(\bar{a}, \bar{b})$. We define the function $\kappa_1(U)$ as the maximal κ_{\max}^* for any $T \geq 0$, given S_L, S_U, \bar{a} and \bar{b} . That is, we define $\kappa_1(U)$ implicitly as the unique function satisfying

$$\kappa_1(U) \bar{a} + \max \left\{ 0, \frac{(\kappa_1^2(U) S_L S_U - 1) \bar{b}}{S_L + S_U + 2\kappa_1(U) S_L S_U} \right\} = T. \quad (26)$$

For completeness, in the appendix we include a closed-form expression for $\kappa_1(U)$ as (35). We define $\kappa_2(U)$ as

$$\kappa_2(U) = \max \left\{ 0, \frac{\kappa_1^2(U) S_L S_U - 1}{S_L + S_U + 2\kappa_1(U) S_L S_U} \right\}. \quad (27)$$

Let $e' \in \bar{\mathcal{G}}$ be an edge with latency function $\ell_{e'}(f_{e'}) = \bar{a} f_{e'} + \bar{b}$. By construction, the tolling function assigned by $\tau^A(\kappa_1(U), \kappa_2(U))$ to e' satisfies bound T with equality: $\tau_{e'}^A(1) = T$.

Now let τ^* be any taxation mechanism with a strictly lower price of anarchy than $\tau^A(\kappa_1(U), \kappa_2(U))$. By Lemma 2.2, τ^* assigns higher tolling functions than $\tau^A(\kappa_1(U), \kappa_2(U))$ on every edge for every flow rate. In particular, on edge e' , $\tau_{e'}^*(1) > \tau_{e'}^A(1) = T$, violating bound T and proving the optimality of $\tau^A(\kappa_1(U), \kappa_2(U))$ over the space of all network-agnostic taxation mechanisms bounded by T . By substituting $\kappa_1(U)$ for κ_{\max} in expression (23), we obtain the complete price of anarchy expression (19). ■

VI. CONCLUSION

In this paper we have explored several avenues for influencing social behavior when aspects of the underlying system are unknown. We showed in Theorem 1 that in theory, it is possible to charge tolls that induce arbitrarily-efficient Nash flows without requiring knowledge of the network topology, user demands, or user sensitivities, but that the required tolls may be very high. To make this more realistic, in

Theorem 2 we investigated the effect of an upper bound on the allowable tolling functions for affine-cost parallel networks. We showed that affine tolls are sufficient to achieve the lowest price of anarchy over the space of all possible tolling functions, and derived the price of anarchy as an explicit function of the upper bound on tolling coefficients. This neatly demonstrated the principle that the more we can charge, the higher efficiency we can guarantee.

This work is part of a growing body of research on applying incentive mechanisms to uncertain situations. Here, we investigate simple affine-latency congestion games; future work will focus on extending the class of applicable networks and latency functions. The setting studied in this paper assumed that user demands were inelastic; an interesting extension would be to model a degree of elasticity, allowing users to simply “stay home” if the network travel cost is too high.

REFERENCES

- [1] C. Papadimitriou, “Algorithms, Games, and the Internet,” in *Proc. of the 28th International Colloquium on Automata, Languages and Programming*, 2001.
- [2] R. Johari and J. N. Tsitsiklis, “Efficiency Loss in a Network Resource Allocation Game,” *Mathematics of Operations Research*, vol. 29, pp. 407–435, Aug. 2004.
- [3] J. R. Marden and J. Shamma, “Game Theory and Distributed Control,” in *Handbook of Game Theory Vol. 4* (H. Young and S. Zamir, eds.), Elsevier Science, 2014.
- [4] H. Youn, M. Gastner, and H. Jeong, “Price of Anarchy in Transportation Networks: Efficiency and Optimality Control,” *Physical Review Letters*, vol. 101, p. 128701, Sept. 2008.
- [5] G. Piliouras, E. Nikolova, and J. S. Shamma, “Risk sensitivity of price of anarchy under uncertainty,” in *Proceedings of the fourteenth ACM conference on Electronic commerce - EC '13*, vol. 9, (New York, New York, USA), pp. 715–732, ACM Press, 2013.
- [6] N. Nisan, T. Roughgarden, E. Tardos, and V. Vazirani, eds., *Algorithmic Game Theory*. Cambridge University Press, 2007.
- [7] P. N. Brown and J. R. Marden, “Social Coordination in Unknown Price-Sensitive Populations,” in *52nd IEEE Conference on Decision and Control*, pp. 1168 – 1173, 2013.
- [8] M. Gairing, “Covering games: Approximation through non-cooperation,” in *Lecture Notes in Computer Science*, vol. 5929 LNCS, pp. 184–195, 2009.
- [9] H.-L. Chen, T. Roughgarden, and G. Valiant, “Designing Network Protocols for Good Equilibria,” 2010.
- [10] J. R. Marden and A. Wierman, “Distributed Welfare Games,” *Operations Research*, vol. 61, pp. 155–168, 2013.
- [11] R. Gopalakrishnan, J. R. Marden, and A. Wierman, “Potential games are necessary to ensure pure nash equilibria in cost sharing games,” *Proceedings of the fourteenth ACM conference on Electronic commerce - EC '13*, pp. 563–564, 2013.
- [12] W. Sandholm, “Negative Externalities and Evolutionary Implementation,” *The Review of Economic Studies*, vol. 72, pp. 885–915, July 2005.
- [13] F. Farokhi and K. H. Johansson, “A Study of Truck Platooning Incentives Using a Congestion Game,” *IEEE Transactions on Intelligent Transportation Systems (in press)*.
- [14] T. Roughgarden, *Selfish Routing and the Price of Anarchy*. MIT Press, 2005.
- [15] R. Cole, Y. Dodis, and T. Roughgarden, “Pricing network edges for heterogeneous selfish users,” in *Proc. of the 35th ACM symp. on Theory of computing*, (New York, New York, USA), pp. 521–530, ACM Press, 2003.
- [16] L. Fleischer, K. Jain, and M. Mahdian, “Tolls for Heterogeneous Selfish Users in Multicommodity Networks and Generalized Congestion Games,” *45th Annu. IEEE Symp. on Foundations of Computer Science*, pp. 277–285, 2004.
- [17] G. Karakostas and S. Kolliopoulos, “Edge pricing of multicommodity networks for heterogeneous selfish users,” *45th Annual IEEE Symposium on Foundations of Computer Science*, pp. 268–276, 2004.

- [18] M. Beckman, C. McGuire, and C. B. Winsten, "Studies in the Economics of Transportation," 1956.
- [19] W. Sandholm, "Evolutionary Implementation and Congestion Pricing," *The Review of Economic Studies*, vol. 69, no. 3, pp. 667–689, 2002.
- [20] J. Wardrop, "Some Theoretical Aspects of Road Traffic Research," in *Proc. of the Institution of Civil Engineers, Part II, Vol. 1, No. 36*, pp. 352–362, 1952.
- [21] P. N. Brown and J. R. Marden, "Robust Taxation Mechanisms in Affine Congestion Games with Price-Sensitive Users," (*Under Review*) <http://rintintin.colorado.edu/phbr4549/papers/scaledmctolls.pdf>, 2015.

APPENDIX: PROOFS OF SUPPORTING LEMMAS

In the proof of Lemma 2.1, we show that a Nash flow on a network with affine tolling coefficients κ_1 and κ_2 for some sensitivity distribution s is identical to a Nash flow on the same network with scaled marginal-cost tolls with $\kappa = \kappa_1$ for some other sensitivity distribution s' . We can then use known analytical techniques for scaled marginal-cost tolls to derive the optimal κ_1 and κ_2 .

Definition 3 (Brown and Marden, [21]): The scaled marginal-cost taxation mechanism assigns the following tolls to any edge $e \in \mathcal{G}^p$ for any $\kappa \geq 0$:

$$\tau_e^{\text{smc}}(f_e) = \kappa a_e f_e. \quad (28)$$

Proof of Lemma 2.1

Let $G \in \mathcal{G}^p$, and $\kappa_1 \geq \kappa_2 \geq 0$.² For user $x \in [0, 1]$ with sensitivity $s_x \in [S_L, S_U]$, the cost of edge $e \in G$ given flow f under affine tolls is given by

$$J_x^e(f) = (1 + \kappa_1 s_x) a_e f_e + (1 + \kappa_2 s_x) b_e.$$

Note that we may scale $J_x^e(f)$ by any factor without changing the underlying preferences of agent x , provided that the scale factor is the same for all edges. Thus, without loss of generality, we may write

$$J_x^e(f) = \frac{1 + \kappa_1 s_x}{1 + \kappa_2 s_x} a_e f_e + b_e. \quad (29)$$

Now, define sensitivity distribution s' by the following: for any $x \in [0, 1]$, s'_x satisfies

$$s'_x = \frac{s_x(\kappa_1 - \kappa_2)}{\kappa_1(1 + \kappa_2 s_x)}. \quad (30)$$

By a series of algebraic manipulations, we may combine (29) and (30) to obtain

$$J_x^e(f) = (1 + \kappa_1 s'_x) a_e f_e + b_e, \quad (31)$$

which is simply the cost resulting from scaled marginal-cost tolls (28). Thus, for any sensitivity distribution s , we may model a Nash flow resulting from affine tolls with coefficients κ_1 and κ_2 as a Nash flow for sensitivity distribution s' resulting from scaled marginal-cost tolls with $\kappa = \kappa_1$.

In [21], the authors show that in this class of networks, the optimal value of κ for scaled marginal-cost tolls is $\frac{1}{\sqrt{S_L S_U}}$.

²Here, the requirement that $\kappa_1 \geq \kappa_2$ is without loss of generality; later analysis shows that $\kappa_2 > \kappa_1$ would always result in a Nash flow with higher congestion than the un-tolled case.

Therefore, assuming at first that κ_{\max} is sufficiently high, our optimal choice of κ_1 is that which satisfies

$$\kappa_1 = \frac{1}{\sqrt{S'_L S'_U}}, \quad (32)$$

where S'_L and S'_U are computed according to (30).

We may combine (32) and (30) to obtain the following characterization of the optimal κ_2 with respect to κ_1 , for $\kappa_{\max} \geq (S_L S_U)^{-1/2}$:

$$\kappa_2 = \frac{\kappa_1^2 S_L S_U - 1}{S_L + S_U + 2\kappa_1 S_L S_U}. \quad (33)$$

In [21], the authors derive the following expression for the price of anarchy with respect to the sensitivity ratio $q = S_L/S_U$ for $\kappa^* = (S_L S_U)^{-1/2}$:

$$\text{PoA}(G, \tau^{\text{smc}}(\kappa^*)) = \frac{4}{3} \left(1 - \frac{\sqrt{q}}{(1 + \sqrt{q})^2} \right). \quad (34)$$

We may use this expression evaluated at $q = S'_L/S'_U$ to compute the price of anarchy resulting from optimal affine tolls for this high- κ_{\max} case, obtaining the following for $\text{PoA}(G, \tau^A(\kappa_1^*, \kappa_2^*))$:

$$\frac{4}{3} \cdot \left(1 - \frac{(1 + \kappa_{\max} S_L) \left(\frac{S_L}{S'_U} + \kappa_{\max} S_L \right)}{\left(1 + 2\kappa_{\max} S_L + \frac{S_L}{S'_U} \right)^2} \right).$$

Finally, we must consider the case when $\kappa_{\max} < (S_L S_U)^{-1/2}$. Now, (33) prescribes a negative value for κ_2 , so the optimal choice is to let κ_2 saturate at 0. Now, we are precisely applying scaled marginal-cost tolls with $\kappa = \kappa_1$, so we apply the fact shown in Lemma 1.2 of [21] that on this class of networks, if $\kappa \leq (S_L S_U)^{-1/2}$, the worst-case total latency of a Nash flow always occurs for the extreme low-sensitivity homogeneous sensitivity distribution given by $s_x \equiv S_L$ for all $x \in [0, 1]$.

Equation (46) in [21] gives the total latency of a Nash flow for a homogeneous population with sensitivity S_L as

$$\mathcal{L}^{\text{nf}}(G, S_L, \kappa) = L_R - \frac{\kappa S_L}{(1 + \kappa S_L)^2} \Theta,$$

where L_R and Θ are positive constants depending only on G , satisfying $\Theta \leq L_R$. It is easy to verify that the above expression is minimized on a subset of $[0, (S_L S_U)^{-1/2}]$ by maximizing κ , and using the fact that $\Theta \leq L_R$, we may compute the price of anarchy for $\kappa_{\max} < (S_L S_U)^{-1/2}$

$$\text{PoA}(G, \tau^A(\kappa_1^*, \kappa_2^*)) = \frac{4}{3} \left(1 - \frac{\kappa_{\max} S_L}{(1 + \kappa_{\max} S_L)^2} \right),$$

obtaining the proof of Lemma 2.1. \blacksquare

We now proceed with the proof of Lemma 2.2, in which we derive properties of any taxation mechanism that outperforms $\tau^A(\kappa_1^*, \kappa_2^*)$.

$$\kappa_1(U) = \min \left\{ \frac{T}{\bar{a}}, \frac{2TS_L S_U - (S_L + S_U)\bar{a} + \sqrt{((S_L + S_U)\bar{a} + 2TS_L S_U)^2 + 4\bar{b}S_L S_U (2\bar{a} + \bar{b} + T(S_L + S_U))}}{2S_L S_U (2\bar{a} + \bar{b})} \right\} \quad (35)$$

Fig. 3. Closed-form expression for $\kappa_1(U)$ used in Theorem 2. Note that it is a continuous, unbounded, strictly increasing function of T .

Proof of Lemma 2.2

We define the set of routing problems \mathcal{G}^0 as follows: $G \in \mathcal{G}^0$ is a parallel network consisting of two edges, with $\ell_1(f_1) = cf_1$ and $\ell_2(f_2) = c$.

Let $G \in \mathcal{G}^0$. For any c , the optimal flow on G is $(f_1^*, f_2^*) = (1/2, 1/2)$ and the optimal total latency is $\mathcal{L}^*(G) = 3c/4$, but the un-tolled Nash flow has a total latency of $\mathcal{L}^{\text{nf}}(G, s, \emptyset) = c$, so the un-tolled price of anarchy is $4/3$. It is straightforward to show furthermore that if $S_U > S_L \geq 0$, for any $\kappa_{\max} > 0$, the price of anarchy of this particular network equals the expression given in (19); i.e., $\text{PoA}(G, \tau^A(\kappa_1^*, \kappa_2^*)) = \text{PoA}(G, \tau^A(\kappa_1^*, \kappa_2^*))$. Thus, if our hypothetical τ^* outperforms τ^A in general, it must specifically outperform τ^A on any network $G \in \mathcal{G}^0$, or

$$\text{PoA}(G, \tau^*) < \text{PoA}(G, \tau^A(\kappa_1^*, \kappa_2^*)).$$

Next, we investigate the performance of τ^* on networks in \mathcal{G}^0 . Given a network $G \in \mathcal{G}^0$, the hypothetical tolling mechanism τ^* assigns edge tolling functions $\tau_1^*(f_1)$ and $\tau_2^*(f_2)$. Recall that since τ^* is network-agnostic, there is some function $\tau^*(f; a, b)$ such that an edge $e \in E$ with latency function $\ell_e(f_e) = a_e f_e + b_e$ is assigned tolling function $\tau^*(f_e; a_e, b_e)$. By analyzing networks in \mathcal{G}^0 , we can deduce properties of the function with the 2nd and 3rd arguments set to 0, since $\tau_1^*(f_1) = \tau^*(f_1; c, 0)$ and $\tau_2^*(f_2) = \tau^*(f_2; 0, c)$.

Now we show that τ^* must assign higher tolls than $\tau^A(\kappa_1^*, \kappa_2^*)$. Let $S_U > S_L$. By design, the worst-case Nash flows resulting from $\tau^A(\kappa_1^*, \kappa_2^*)$ occur for homogeneous populations with $s = S_L$ and $s = S_U$. Since any network $G \in \mathcal{G}^0$ has only 2 links, we can characterize a Nash flow simply by the flow on edge 1; accordingly, let $f_L(c)$ denote the flow as a function of c on edge 1 in the Nash flow resulting from sensitivity distribution $s = S_L$, and $f_H(c)$ the corresponding edge 1 flow for $s = S_U$. These flows are the solutions to the following equations:

$$cf_L(c)(1 + \kappa_1^* S_L) = c(1 + \kappa_2^* S_L), \quad (36)$$

$$cf_H(c)(1 + \kappa_1^* S_U) = c(1 + \kappa_2^* S_U). \quad (37)$$

We may combine and rearrange the above in the following way:

$$\kappa_1^*(f_L(c) - f_H(c)) = \frac{f_H(c)}{S_U} - \frac{f_L(c)}{S_L} + \frac{1}{S_L} - \frac{1}{S_U}. \quad (38)$$

It is always true that $f_H(c) < f_L(c)$. By design, $\mathcal{L}(f_L(c)) = \mathcal{L}(f_H(c))$. Note that \mathcal{L} is simply a concave-up parabola in the flow on edge 1.

Now, let $f_L^*(c)$ and $f_H^*(c)$ be similarly defined as the Nash flows resulting from τ^* for a given value of c ; i.e., the

solutions to

$$cf_L^*(c) + \tau_1^*(f_L^*(c))S_L = c + \tau_2^*(1 - f_L^*(c))S_L, \quad (39)$$

$$cf_H^*(c) + \tau_1^*(f_H^*(c))S_U = c + \tau_2^*(1 - f_H^*(c))S_U. \quad (40)$$

Since τ^* guarantees better efficiency than $\tau^A(\kappa_1^*, \kappa_2^*)$, it must do so in particular for these homogeneous sensitivity distributions $s = 1$ and $s = S_U$. Since \mathcal{L} is a parabola, this means that for any c , $f_H(c) < f_H^*(c) < f_L^*(c) < f_L(c)$.

Define the nondecreasing function $\Delta^*(f) = \tau_2^*(f) - \tau_1^*(1 - f)$ (which is implicitly also a function of c), so equations (39) and (40) can be combined and rearranged to show

$$\begin{aligned} \Delta^*(f_L^*(c)) - \Delta^*(f_H^*(c)) &= c \left[\frac{f_H^*(c)}{S_U} - \frac{f_L^*(c)}{S_L} + \frac{1}{S_L} - \frac{1}{S_U} \right] \\ &> c \left[\frac{f_H(c)}{S_U} - \frac{f_L(c)}{S_L} + \frac{1}{S_L} - \frac{1}{S_U} \right] \\ &= \kappa_1^* c (f_L(c) - f_H(c)) \end{aligned} \quad (41)$$

We can loosen the above inequality even further by replacing $f_L^*(c)$ with $f_L(c)$ and $f_H^*(c)$ with $f_H(c)$, and substituting from (38) and rearranging, we finally obtain

$$\frac{\Delta^*(f_L(c)) - \Delta^*(f_H(c))}{f_L(c) - f_H(c)} > \kappa_1^* c. \quad (42)$$

Since this must be true for any $c > 0$, the average slope of $\Delta^*(f)$ must be greater than $\kappa_1^* c$ for all $f > 0$. Since $\tau_2^*(f) \geq 0$ this implies that $\tau_1^*(f) > \kappa_1^* cf$ for all $f > 0$, or that

$$\tau^*(f; a, 0) > \tau^A(f; a, 0) \quad (43)$$

for all positive f and a .

Now consider the following rearrangement of (40):

$$\begin{aligned} \tau_2^*(1 - f_H^*(c)) &= [cf_H^*(c) + \tau_1^*(f_H^*(c)) - cS_U] \cdot \frac{1}{S_U} \\ &> c[(1 + \kappa_1^* S_U)f_H(c) - 1] \cdot \frac{1}{S_U} \\ &= \kappa_2^* c S_U = \tau_2^A(f). \end{aligned} \quad (44)$$

This implies that $\tau_2^*(f) > \kappa_2^* c$ for all $f > 0$, or that

$$\tau^*(f; 0, b) > \tau^A(f; 0, b) \quad (45)$$

for all positive f and b .

Finally, note that the additivity assumption of Definition 2 implies that $\tau^*(f; a, b)$ is additive in its second and third arguments. That is, we may add inequalities (43) and (45) to conclude that for all nonnegative f , a , and b , that

$$\tau^*(f; a, b) > \kappa_1^* af + \kappa_2^* b, \quad (46)$$

or that a necessary condition for $\text{PoA}(G, \tau^*) < \text{PoA}(G, \tau^A)$ is that τ^* must charge higher tolls on *every* edge in *every* network. ■