Avoiding Perverse Incentives in Affine Congestion Games

Philip N. Brown and Jason R. Marden

***Abstract*— In engineered systems whose performance depends on user behavior, it is often desirable to influence behavior in an effort to achieve performance objectives. However, doing so naively can have unintended consequences; in the worst cases, a poorly-designed behavior-influencing mechanism can create a perverse incentive which encourages adverse user behavior. For example, in transportation networks, marginal-cost tolls have been studied as a means to incentivize low-congestion network routing, but have typically been analyzed under the assumption that all network users value their time equally. If this assumption is relaxed, marginal-cost tolls can create perverse incentives which increase network congestion above un-tolled levels. In this paper, we prove that if some network users are unresponsive to tolls, any taxation mechanism that does not depend on network structure can create perverse incentives. Thus, to systematically avoid perverse incentives, a taxation mechanism *must* be network-aware to some extent. On the other hand, we show that a small amount of additional information can mitigate this negative result; for example, we show that it is relatively easy to avoid perverse incentives on affine-cost parallel-path networks, and we fully characterize the taxation mechanisms that minimize congestion for worst-case user populations on such networks.**

I. Introduction

Many of today's engineered systems are richly connected with their users; economic, social, and technical objectives are often interconnected in complex ways. Examples of this can be found in ridesharing systems [1], transportation networks [2], and power grids [3]. As the interconnections between social and engineered systems increase, the engineer's task increasingly includes influencing the behavior of system users. Accordingly, recent research has focused on developing new analytical tools for influencing social behavior to achieve engineering objectives [4]–[8]. One intrinsic challenge in this setting is that of uncertainty in the social systems to be influenced. It can be difficult to characterize user preferences and decision-making processes in a social system, and any behavior-influencing mechanism must take this uncertainty into account. This has led to recent efforts to develop behavior-influencing mechanisms that are robust to a wide range of possible mischaracterizations [9]–[12].

One particular area of interest is that of influencing drivers' routing choices in transportation networks. It is well-known that if individual drivers choose their own routes through a congestible network to minimize their personal delay, the resulting aggregate network delay can be significantly worse than optimal [13]. Recent research has suggested levying road taxes, thereby modifying agents' costs and incentivizing more-efficient network flows. Many such taxation schemes require the tax-designer to have a perfect characterization of all underlying system variables: network topology, road congestion characteristics, user tax-sensitivities, and overall demand. Given such a perfect characterization, it is known that a system planner can design taxes which incentivize optimal network flows [14]–[16]. Unfortunately, recent results have demonstrated that taxes designed for one problem instance can incentivize inefficient behavior on different (yet closely-related) instances, indicating that these taxes lack robustness to system mischaracterizations.

In [17], the authors propose a semantic framework for robustness as applied to behavior-influencing mechanisms. In this framework, taxes are designed for some nominal routing problem, and the problem is then perturbed in a variety of ways. For each perturbed routing problem, the congestion induced by the original taxes is then compared to two benchmarks: the optimal aggregate delay on the perturbed problem, and the delay of an un-influenced flow on the perturbed problem. The taxes are said to be *strongly robust* to that type of perturbation if the nominal taxes induce optimal flows on all perturbed networks. The taxes are said to be *weakly robust* to the perturbation if on the perturbed networks the flows incentivized by the nominal taxes are never worse than un-influenced flows. That is, if a taxation mechanism is weakly robust, the system planner can be certain that taxing is always at least as good as not taxing; alternatively, these taxes will never create perverse incentives.

A prominent example of a strongly-robust taxation mechanism is that of marginal-cost tolls, which are known to incentivize optimal network flows without requiring *a priori* knowledge of user demand or network topology provided that all users trade off time and money equally [18], [19]. An attractive feature of marginal-cost tolls is that the toll on each network link depends only on that link's flow and congestion properties, and can be computed without any information about overall network topology. This property is known as *network-agnosticity*, and is a desirable characteristic of any taxation mechanism since by construction it confers robustness to variations in network structure.

However, recent research has suggested that this network-agnosticity comes at a price; if network users have unknown price-sensitivities, marginal-cost tolls fail to be strongly-robust to mischaracterizations of user price-sensitivity [17]. Unfortunately, in the most general networks, if users have

This work is supported ONR Grant #N00014-15-1-2762 and NSF Grant #ECCS-1351866.

P. N. Brown is a graduate student researcher with the Department of Electrical and Computer Engineering, University of California, Santa Barbara, CA pnbrown@ece.ucsb.edu. Corresponding author.

J. R. Marden is with the Department of Electrical and Computer Engineering, University of California, Santa Barbara, jrmarden@ece.ucsb.edu.

diverse price-sensitivities, off-the-shelf marginal-cost tolls are not even weakly robust [17]; this is demonstrated in Example 2.1 in this paper. Put plainly, marginal-cost tolls can create perverse incentives if applied naively. Despite this fact, [9], [20] show that for parallel networks subject to a particular utilization constraint, scaled marginal-cost tolls (and affine tolls, a closely-related variant) are weakly robust. This indicates that if the toll-designer has additional information about the routing problem setting (e.g., information about the allowable class of networks), the weak-robustness of marginal-cost tolls may be recovered.

Accordingly, the central goal of this paper is to understand the relationship between robustness and network-agnosticity more fully, and determine specifically under what conditions network-agnostic tolls can be weakly robust. To that end, in our first and most general result we demonstrate that if some users are unresponsive to tolls, the only weakly-robust network-agnostic taxation mechanism is essentially the taxation mechanism which charges zero tolls. That is, to avoid perverse incentives, taxes must depend on some information regarding network structure.

Fortunately, a taxation mechanism need not depend on much additional information: our second result states that for the class of affine-cost parallel-path networks, a weakly-robust network-agnostic taxation mechanism *always* exists; that is, knowledge of the class of allowable networks can render weak robustness possible. We give a full characterization of the space of weakly-robust network-agnostic taxation mechanisms for this setting, and we derive the taxation mechanism that minimizes network congestion in worst case over any distribution of unknown user toll-sensitivities. This is encouraging from the standpoint of a toll-designer, since it suggests that a little additional information about the problem setting can greatly expand the designer's toolbox.

II. Model and Related Work

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 network is called a *parallel-path* network if for all paths $p, p' \in \mathcal{P}$, $p \cap p' = \emptyset$.

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.

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 affine latency function $\ell_e(f_e) = a_e f_e + b_e$, where $a_e \geq 0$ and $b_e \geq 0$ are edge-specific constants. We measure the cost of a flow f by the *total latency*, given by

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

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

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

All optimal flows f^* have the same total latency.

A *routing problem* is given by $G = (V, E, \mathcal{C}, \{\ell_e\})$. The set of all routing problems is written $\mathcal{G}$.

To study the effect of taxes on self-interested behavior, we model the above routing problem as a non-atomic congestion game. We assign each edge $e \in E$ a flow-dependent taxation function $\tau_e : \mathbb{R}^+ \to \mathbb{R}^+$. To characterize users' taxation sensitivities, for each $c \in \mathcal{C}$, let each user $x \in [0, r_c]$ have a taxation sensitivity $s_x^c \in [S_{\rm L}, S_{\rm U}] \subseteq \mathbb{R}^+$, where $S_{\rm L} \geq 0$ and $S_{\rm U} \leq +\infty$ are lower and upper 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

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

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

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

It is well-known that a Nash flow exists for any non-atomic congestion game of the above form [21].

We assume that the sensitivity distribution function s is unknown; for a given routing problem G, we define the set of possible sensitivity distributions as the set of Lebesgue-measurable functions $\mathcal{S}_G = \{s^c : [0, r_c] \to [S_{\rm L}, S_{\rm U}]\}_{c \in \mathcal{C}}$.

B. Taxation Mechanisms and Robustness

In this paper, we consider a particular type of taxation mechanism which we term *network-agnostic*. Here, each edge's taxation function is computed using only locally-available information. That is, $\tau_e(f_e)$ depends only on ℓ_e, not on edge e's location in the network, the overall network topology, the overall traffic rate, or the congestion properties of any other edge. Network agnosticity is an abstraction that allows us to discuss taxation mechanisms that do not depend on network structure. A network-agnostic taxation mechanism τ is essentially a mapping from latency functions to taxation functions, and any edge taxation function is

$$\tau_e(\cdot) = \tau(\ell_e). \tag{5}$$

To evaluate taxation mechanisms, we adopt the robustness framework introduced in [17] in which a distinction is drawn between *strong* and *weak* robustness. In either case, for each network we compare the worst-case Nash congestion induced by a particular taxation mechanism to two benchmarks: first, the optimal total latency on the network; second, the total latency of an un-influenced Nash flow on the network.

Formally, in the context of network-agnostic taxation mechanisms, we write $\mathcal{L}^{\mathrm{nf}}(G, s, \tau)$ to denote the total latency of a Nash flow for routing problem G and population s induced by taxation mechanism τ. We write $\mathcal{L}^{\mathrm{nf}}(G, \emptyset)$ to denote the total latency of an un-influenced Nash flow (note that when there are no tolls, the sensitivity distribution plays no role), and $\mathcal{L}^*(G) = \mathcal{L}(f^*)$ to denote the optimal latency.

Taxation mechanism τ is said to be *strongly robust* if for every network it induces optimal Nash flows for any sensitivity distribution; that is, for all $G \in \mathcal{G}$,

$$\sup_{s \in \mathcal{S}_G} \mathcal{L}^{\mathrm{nf}}(G, s, \tau) = \mathcal{L}^*(G). \tag{6}$$

On the other hand, τ is said to be *weakly robust* if for every network and sensitivity distribution, the total latency induced by τ never exceeds the total latency of an un-influenced Nash flow; i.e., for all $G \in \mathcal{G}$,

$$\sup_{s \in \mathcal{S}_G} \mathcal{L}^{\mathrm{nf}}(G, s, \tau) \leq \mathcal{L}^{\mathrm{nf}}(G, \emptyset). \tag{7}$$

Weak robustness is a guarantee that perverse incentives will never arise; note that at a minimum, it is always trivially true that the zero-toll is weakly robust.

C. Related Work

Among the simplest tolls are *fixed tolls*, which for any $e \in G$, $\tau_e(f_e) = q_e$ for some $q_e \geq 0$. However, if network, traffic-rate, and user sensitivity specifications are not known precisely and fixed tolls are restricted to be network-agnostic, they fail to be weakly-robust [17].

A classic example of strongly-robust network-agnostic tolls is that of the *marginal-cost* or *Pigovian* taxation mechanism τ^{mc}, which assigns taxation functions of

$$\tau^{\mathrm{mc}}(f_e) = f_e \cdot \ell'_e(f_e), \ \forall f_e \geq 0, \tag{8}$$

where ℓ' represents the flow derivative of ℓ. In [18] the authors show that for any $G \in \mathcal{G}$, it is true that $\mathcal{L}^*(G) = \mathcal{L}^{\mathrm{nf}}(G, s, \tau^{\mathrm{mc}})$, provided that all users have a sensitivity equal to 1. Thus, by construction, marginal-cost tolls are strongly-robust to perturbations of network structure and traffic rate, since each taxation function has no dependence on either. Unfortunately, the following example demonstrates that marginal-cost tolls are not even weakly robust in multi-commodity networks for heterogeneous populations.

Example 2.1: Consider the network depicted in Figure 1. There are two source nodes; 0.5 units of traffic from the upper source with sensitivity s_1 share a common destination with 1 unit of traffic from the lower source with sensitivity s_2. Marginal-cost tolls on this network are $\tau_i(f_i) = f_i$ on paths 1 and 2. It is simple to verify that if all traffic trades off time and money equally (i.e., $s_1 = s_2 = 1$), marginal-cost tolls incentivize the optimal flow depicted on the left of the figure (since at this flow, all agents have a cost of 1). However, if the upper-source traffic keeps $s_1 = 1$ but the lower-source traffic has $s_2 = 0$ (i.e., they care only about time), marginal-cost tolls create a perverse incentive, resulting in the configuration depicted on the right in which the high-sensitivity traffic experiences a cost of 2 on paths 1 and 2, and the low-sensitivity traffic experiences a cost of 1 on paths 2 and 3. This second flow exhibits higher total latency than the un-tolled configuration. This proves the lack of weak robustness of marginal-cost tolls in multicommodity heterogeneous networks.

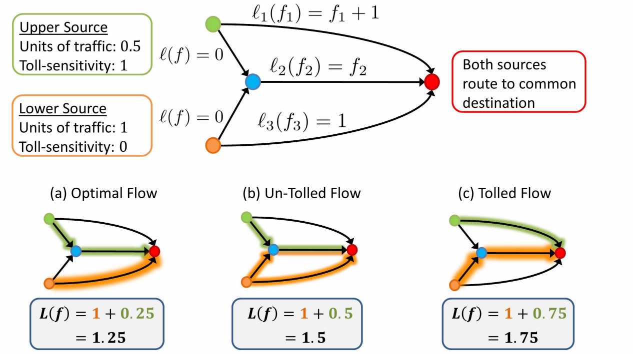

Fig. 1. Example 2.1: A network demonstrating that marginal-cost tolls are not weakly robust to user heterogeneity. This figure depicts a simple two-source network in which 0.5 units of traffic route from the upper source, and 1 unit of traffic routes from the lower source. If traffic from the upper source trades off time and money equally (i.e., $s \equiv 1$), but traffic from the lower source cares only about time (i.e., $s \equiv 0$), marginal-cost tolls create a perverse incentive. The optimal flow here requires all of the traffic from the lower source to the constant-latency link 3. However, only the traffic from the upper source responds to tolls; when marginal-cost tolls are levied (a toll of $\tau_i(f_i) = f_i$ on links 1 and 2), all of the upper-source traffic moves to the inefficient path 1, and the lower-source traffic moves to replace it on path 2, as depicted on the right. This results in a tolled total latency of 1.75 (0.5 units of traffic have a latency of 1.5, and 1 unit of traffic has a latency of 1), greater than the un-tolled value of 1.5 (in which all 1.5 units of traffic have latency 1). Note that this example exhibits perverse incentives even in the case that user sensitivities do not actually take the value 0; all that is required is that the lower-source traffic have low sensitivity.

Moving past marginal-cost tolls, the authors of [20] exhibit a universal taxation mechanism τ^{u} which assigns taxes of

$$\tau_e^{\mathrm{u}}(f_e) = \kappa_{\mathrm{u}} \left(a_e f_e + b_e/2 \right), \tag{9}$$

and show that on any multicommodity network, if $S_{\mathrm{L}} > 0$, it is weakly-robust for large-enough κ_{u}. Here, $S_{\mathrm{L}} > 0$ implies that user sensitivities are bounded away from zero; the κ_{u} required for weak robustness depends on S_{L} in general. However, if the tax designer does not know a lower bound on user sensitivities (i.e., $S_{\mathrm{L}} = 0$), our Theorem 3.1 shows that τ^{u} fails to be weakly robust.

D. Summary of Our Contributions

In this paper, we present several contributions to the theory of robust social influence for congestion games. We first prove that weakly-robust network-agnostic tolls do not exist in the most general settings; subsequently, we examine a special subset of congestion games for which we prove that weakly-robust network-agnostic tolls always exist.

In Theorem 3.1, we show that if $S_{\mathrm{L}} = 0$ (i.e., user sensitivities can take arbitrarily-small values), network-agnosticity and weak robustness are mutually exclusive for general multicommodity networks with heterogeneous populations. That is, given any network-agnostic taxation mechanism that improves congestion on one routing problem, there exists a network and user sensitivity distribution for which that taxation mechanism *increases* congestion above un-tolled levels.

Thus, perverse incentives cannot be systematically avoided unless tolls exhibit some degree of network-dependence. This is in sharp contrast to the universal tolls of [20] (see Equation (9) in this paper), which are weakly-robust for any network provided that $\kappa_{\rm u}$ is large enough and user sensitivities are bounded away from zero.

In light of this negative result, we subsequently ask if network-agnostic tolls can *ever* be weakly robust, and show in Theorem 3.3 that knowledge of the class of allowable networks can mitigate the impossibility result of Theorem 3.1. Specifically, we show that if networks are known to be single-commodity parallel-path, there always exist network-agnostic taxation mechanisms which are weakly robust. Furthermore, for any toll upper-bound, we derive the specific taxation mechanism that minimizes worst-case congestion for any $S_{\rm L}$ and $S_{\rm U}$. This is true without any restrictions on the range of user sensitivities: unlike in Theorem 3.1, here we allow $S_{\rm L} = 0$ and $S_{\rm U} = +\infty$.

III. Our Contributions

A. Impossibility in Multicommodity Networks

Our first result considers the case that a toll-designer wishes to design a network-agnostic taxation mechanism that improves the efficiency of Nash flows on every network. Theorem 3.1 shows that this is impossible on multicommodity networks if user sensitivities are heterogeneous and are not bounded away from 0. That is, if a network-agnostic taxation mechanism improves congestion for one routing problem, it must degrade congestion for another.

Theorem 3.1: For the class of all multicommodity affine-latency networks, if $S_{\rm L} = 0$ and $S_{\rm U} > 0$, a network-agnostic taxation mechanism τ is weakly robust if and only if for every network $G \in \mathcal{G}$ and population s,

$$\mathcal{L}^{\rm nf}(G, s, \tau) = \mathcal{L}^{\rm nf}(G, \emptyset). \tag{10}$$

It is important to note that (10) is not a design criterion, but rather a definition of triviality. We include this formalism because there exist taxes which are non-zero, but exactly mimic the effect of zero tolls. In general, these trivial taxes are of the form $\tau_e(f_e) = \mu \ell_e(f_e)$ for some $\mu \geq 0$; since these merely scale the game's original latency functions, they induce exactly the same Nash flows as un-influenced flows.

As a first step towards proving Theorem 3.1, in Lemma 3.2 we present necessary conditions for the weak robustness of any network-agnostic taxation mechanism; we will subsequently show that any taxation mechanism satisfying these conditions will fail to be weakly robust in certain multicommodity routing problems.

Lemma 3.2: For networks with affine cost functions, if a network-agnostic taxation mechanism is weakly robust, it assigns taxes given by

$$\tau_e(f_e) = \kappa_1 a_e f_e + \kappa_2 b_e, \tag{11}$$

for some $\kappa_1 \geq 0$ and $\kappa_2 \geq 0$, for every $c \in \mathcal{C}$ and every user $x \in [0, r_c]$ coefficients κ_1 and κ_2 satisfy

$$\frac{s_x^c\,(\kappa_1 - \kappa_2)}{1 + s_x^c \kappa_2} \in [0, 1]\,. \tag{12}$$

The proof of Lemma 3.2 appears in the Appendix.

Proof of Theorem 3.1: We prove this by showing on the network of Example 2.1 that if users are heterogeneous and price-sensitivities can take a value of 0, no tolls satisfying the conditions of Lemma 3.2 can be weakly-robust. Consider the network depicted in Figure 1. Commodity c_1 is a high-sensitivity population with sensitivity s_1 and mass $1/2$ is traveling from the upper source; commodity c_2 is a low-sensitivity population with sensitivity s_2 and mass 1 is traveling from the lower source; the populations (commodities) share a common destination. The paths are enumerated $\{1, 2, 3\}$ (from top to bottom, as in Figure 1), so $\mathcal{P}_1 = \{1, 2\}$ and $\mathcal{P}_2 = \{2, 3\}$.

Charge tolls on this network in accordance with Lemma 3.2 (that is, specify κ_1 and κ_2 satisfying (12) for both $s_x^{c_1} = s_1$ and $s_x^{c_2} = s_2$), and define γ_1 and γ_2 as follows:

$$\gamma_1 = \frac{s_1(\kappa_1 - \kappa_2)}{1 + s_1 \kappa_2}, \quad \text{and} \quad \gamma_2 = \frac{s_2(\kappa_1 - \kappa_2)}{1 + s_2 \kappa_2}. \tag{13}$$

The upper population's Nash-flow incentive constraint induced by these tolls is given by

$$(1 + \gamma_1) f_1 + 1 = (1 + \gamma_1) f_2, \tag{14}$$

and the lower population's by

$$(1 + \gamma_2) f_2 = 1. \tag{15}$$

It can be shown that any Nash flow for which $\gamma_2 \leq \gamma_1$ must satisfy (14) and (15); since the system of equations is upper-triangular, f_2 depends only on γ_2:

$$f_1 = \frac{1}{1 + \gamma_2} - \frac{1}{1 + \gamma_1}, \quad \text{and} \quad f_2 = \frac{1}{1 + \gamma_2}. \tag{16}$$

In essence, the low-sensitivity population holds all the power; the flow on path 2 is not affected by the sensitivities of the high-sensitivity population.

Consider the definitions of γ_1 and γ_2; note that for any fixed choice of κ_1 and κ_2, γ_2 can be made arbitrarily-close to 0 by choosing a very low s_2. To model the extreme case, let $s_2 = 0$ so that $\gamma_2 = 0$. Then $f_2 = 1$ and the total latency on the network as a function of $\gamma_1 > 0$ is given by

$$\begin{aligned} \mathcal{L}(\gamma_1) &= 1 + \left(\frac{\gamma_1}{1 + \gamma_1}\right)^2 + \frac{\gamma_1}{1 + \gamma_1} + \frac{1 - \gamma_1}{2(1 + \gamma_1)} \\ &= 1.5 + \left(\frac{\gamma_1}{1 + \gamma_1}\right)^2 > 1.5. \end{aligned}$$

Thus, charging tolls that induce $\gamma_1 > 0$ cause the total latency of a tolled Nash flow to be greater than that of the un-tolled Nash flow. The only tolls that guarantee $\gamma_1 = 0$ have tolling coefficients $\kappa_1 = \kappa_2$; any tolls of this form have $\mathcal{L}^{\rm nf}(G, s, \tau) = \mathcal{L}^{\rm nf}(G, \emptyset)$. ■

B. Weakly Robust Tolls for Parallel-Path Networks

We now ask if knowledge of a small amount of information can mitigate the negative result of the previous section. Indeed, if networks are known to be single-commodity parallel-path networks, Theorem 3.3 shows that weakly-robust network-agnostic taxation mechanisms always exist. This is true even in extreme cases when $S_{\rm L} = 0$ or $S_{\rm U} = \infty$.

The implications for a toll-designer are encouraging: this result demonstrates the intuitive principle that information regarding the possible class of networks can greatly expand the designer's toolbox.

Before stating the theorem, we point out that worst-case performance guarantees provided by a taxation mechanism can often be improved by increasing all edge tolls appropriately (as discussed in [20]). Thus, in order to make meaningful statements about congestion-minimizing tolls, it is useful to parameterize tolls by a stylized upper-bound; the parameter $\kappa_{\max} > 0$ plays this role in the following theorem.

Theorem 3.3: For single-commodity parallel-path networks with affine latency functions, a network-agnostic taxation mechanism is weakly robust if and only if it satisfies the conditions of Lemma 3.2. Furthermore, for any $S_{\rm U} < \infty$ and toll-scalar upper-bound $\kappa_{\max} \geq 0$ the congestion-minimizing taxation mechanism assigns tolling functions

$$\tau_e(f_e) = \kappa_{\max} a_e f_e + b_e \max\left\{0, \frac{\kappa_{\max} S_{\rm U} - 1}{2S_{\rm U}}\right\}. \quad (17)$$

If $S_{\rm U} = +\infty$, then for any $S_{\rm L} \geq 0$, (17) simplifies to

$$\tau_e(f_e) = \kappa_{\max}\left(a_e f_e + b_e/2\right). \quad (18)$$

The proof of Theorem 3.3 appears in the Appendix.

Here, we find that the universal tolls of [20] are in fact weakly robust for parallel-path networks. Another important fact to note is that (17) gives some insight into the robustness of scaled marginal-cost tolls: Recall that for affine-latency congestion games, marginal-cost tolls are given by $\tau_e^{\rm mc}(f_e) = a_e f_e$, and incentivize optimal Nash flows for unit-sensitivity homogeneous populations. If $\kappa_{\max} = 1/S_{\rm U}$, (17) gives the congestion-minimizing scaled marginal-cost toll as $\tau_e(f_e) = a_e f_e / S_{\rm U}$. This can be interpreted as a conservatively-scaled marginal-cost toll; it implies that the best way to avoid perverse incentives with marginal-cost tolls is to charge tolls as though all users have sensitivity equal to $S_{\rm U}$.

IV. Conclusion

In this paper, we have presented initial findings on the weak robustness of network-agnostic taxation mechanisms; we showed that in general routing problems, network-agnosticity carries the risk of perverse incentives if some network users are unresponsive to tolls. On the other hand, we showed that on parallel-path networks, perverse incentives can be systematically avoided, and we characterized the full space of weakly robust network-agnostic taxation mechanisms for this setting. Throughout, we gauged everything from a draconian worst case perspective; relaxing this approach slightly may yield significant robustness gains. For example, if a tax-designer has coarse distributional knowledge of a user population's price sensitivities, it is possible that this knowledge can be exploited to employ more aggressive taxation schemes which remain weakly robust.

References

[1] A. Kleiner, B. Nebel, and V. A. Ziparo, "A mechanism for dynamic ride sharing based on parallel auctions," in *International Joint Conference on Artificial Intelligence*, pp. 266–272, 2011.

[2] I. Arieli, "Transfer Implementation in Congestion Games," *Dynamic Games and Applications*, vol. 5, no. 2, pp. 228–238, 2015.

[3] C. Silva, "Application of Mechanism Design to Electric Power Markets (Republished)," *IEEE Trans. on Power Systems*, vol. 16, no. 4, pp. 862–869, 2001.

[4] P. N. Brown and J. R. Marden, "A Study on Price-Discrimination for Robust Social Coordination," in *American Control Conf.*, (Boston, MA, USA), pp. 1699–1704, 2016.

[5] M. Alizadeh, Y. Xiao, A. Scaglione, and M. van der Schaar, "Dynamic Incentive Design for Participation in Direct Load Scheduling Programs," *IEEE Journal of Selected Topics in Signal Processing*, vol. 8, no. 6, pp. 1111–1126, 2014.

[6] R. Gopalakrishnan, J. R. Marden, and A. Wierman, "Potential games are necessary to ensure pure nash equilibria in cost sharing games," *Proc. 14th ACM Conf. on Electronic Commerce*, pp. 563–564, 2013.

[7] V. Bonifaci, M. Salek, and G. Schäfer, "Efficiency of restricted tolls in non-atomic network routing games," in *Algorithmic Game Theory*, vol. 6982 LNCS, pp. 302–313, 2011.

[8] M. Gairing, "Covering games: Approximation through non-cooperation," in *Internet and Economics*, pp. 184–195, 2009.

[9] P. N. Brown and J. R. Marden, "The Robustness of Marginal-Cost Taxes in Affine Congestion Games," *Transactions on Automatic Control, to appear*, 2016.

[10] G. Christodoulou, K. Mehlhorn, and E. Pyrga, "Improving the price of Anarchy for Selfish Routing via coordination mechanisms," *Algorithmica*, vol. 69, no. 3, pp. 619–640, 2014.

[11] D. Bergemann and K. H. Schlag, "Pricing without Priors," *Journal of the European Economic Association*, vol. 6, no. 2-3, pp. 560–569, 2008.

[12] U. Bhaskar, K. Ligett, L. Schulman, and C. Swamy, "Achieving Target Equilibria in Network Routing Games without Knowing the Latency Functions," in *Proc. IEEE Symp. Foundations of Computer Science*, (Philadelphia, PA, USA), pp. 31–40, 2014.

[13] T. Roughgarden, *Selfish Routing and the Price of Anarchy*. MIT Press, 2005.

[14] L. Fleischer, K. Jain, and M. Mahdian, "Tolls for Heterogeneous Selfish Users in Multicommodity Networks and Generalized Congestion Games," in *Proc. 45th IEEE Symp. on Foundations of Computer Science*, (Rome, Italy), pp. 277–285, 2004.

[15] G. Karakostas and S. Kolliopoulos, "Edge pricing of multicommodity networks for heterogeneous selfish users," in *Proc. 45th IEEE Symp. on Foundations of Computer Science*, (Rome, Italy), pp. 268–276, 2004.

[16] 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, 2003.

[17] P. N. Brown and J. R. Marden, "Studies on Robust Social Influence Mechanisms: Incentives for Efficient Network Routing in Uncertain Settings," *Control Systems Magazine, to appear*, 2016.

[18] M. Beckmann, 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] P. N. Brown and J. R. Marden, "Optimal Mechanisms for Robust Coordination in Congestion Games," in *54th IEEE Conf. Decision and Control*, (Osaka, Japan), pp. 2283–2288, 2015.

[21] 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.

Appendix: Proofs

Proof of Lemma 3.2: Let $\tau^{\rm na}$ be a network-agnostic taxation mechanism; that is, for any affine latency function $\ell_e(f_e) = a_e f_e + b_e$, the taxation function on edge e is given by $\tau_e(f_e) = \tau^{\rm na}(a_e f_e + b_e)$. Consider the network in Figure 2(a); the un-tolled and optimal flow on this network is always $(r/2, r/2)$. Thus, weakly-robust tolls must charge the same total amount on the upper path as they do on the lower path, or $\tau^{\rm na}(\ell_3) = \tau^{\rm na}(\ell_1) + \tau^{\rm na}(\ell_2)$. By replacing ℓ_1, ℓ_2, and ℓ_3 with various combinations of linear and constant functions, it can be shown that this additivity implies that

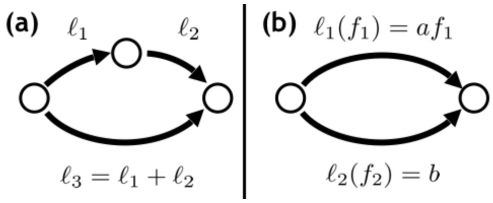

Fig. 2. Networks used to prove necessary conditions for weak robustness of network-agnostic taxation mechanisms.

$\tau^{\rm na}$ is linear, or there exist scalars $\kappa_1 \in \mathbb{R}^+$ and $\kappa_2 \in \mathbb{R}^+$ such that

$$\tau^{\rm na}(af+b) = \kappa_1 af + \kappa_2 b. \tag{19}$$

For user x, tolls of this form induce effective costs of

$$J_x^e(f_e) = a_e f_e + b_e + \gamma_x a_e f_e, \tag{20}$$

where for all x, $\gamma_x \in \mathbb{R}$ and $\gamma_x \triangleq \frac{s_x(\kappa_1-\kappa_2)}{1+s_x\kappa_2}$. In the following, we assume that for all x, $\gamma_x = \gamma$.

If $\gamma < 0$ (i.e., $\kappa_1 < \kappa_2$), it is simple to show that this is a perverse incentive on any network of the form in Figure 2(b) whenever r is sufficiently high, since it incentivizes over-congestion on link 1. Therefore, suppose that $\gamma > 1$. In Figure 2(b) let $a = b = 1$. If $r = 1/2$, $f^{\rm nf} = f^{\rm opt} = (1/2, 0)$ with total latency $L(f^{\rm nf}) = 1/4$. However, for any $\gamma > 1$, some positive mass of traffic routes on the lower edge, increasing the total latency above $1/4$. Thus, affine tolls can only be weakly-robust if for all x, $\gamma_x \in [0,1]$, completing the proof. ∎

A. Proof of Theorem 3.3 and associated Lemmas

To facilitate our arguments, we assign labels of $i \in \{1,\ldots,|\mathcal{P}|\}$ to a network's paths such that if $b_i \triangleq \sum_{e\in p_i} b_e$, for all i we have $b_i < b_{i+1}$. Similarly, we write $f_i \triangleq \sum_{e\in p_i} f_e$ and $a_i \triangleq \sum_{e\in p_i} a_e$.

We complete the proof for the case that every b_i is distinct and at most one path has $a_i = 0$. This is without loss of generality since if paths i,j have $b_i = b_j$, they may be combined into one path; if two paths have $a_i = a_j = 0$, the path with the higher constant coefficient may be ignored.

Lemma 4.1: Any Nash flow $f^{\rm nf}$ induced by tolls according to Lemma 3.2 satisfies the following: if $b_i < b_{i+1}$, a user x is using path i, and a user y is using path $i+1$,

1) $s_x \le s_y$,
2) $2a_i f_i + b_i \ge 2a_{i+1} f_{i+1} + b_{i+1}$.

Proof: Because user cost functions are equivalent to (20) and $\frac{s_x(\kappa_1-\kappa_2)}{1+s_x\kappa_2}$ is monotone-increasing in s_x, we can assume without loss of generality that tolls are of the form $\tau_e(f_e) = a_e f_e$ and that $[S_{\rm L}, S_{\rm U}] \subseteq [0,1]$.

Item 1 is proved in Claim 1.1.2 of [9], which also shows that $a_i f_i^{\rm nf} \le a_{i+1} f_{i+1}^{\rm nf}$. If agent y is using path $i+1$ in a Nash flow, this means that

$$(1+s_y)a_i f_i^{\rm nf} + b_i \ge (1+s_y)a_{i+1} f_{i+1}^{\rm nf} + b_{i+1}. \tag{21}$$

Since $s_y \le 1$ and $a_i f_i^{\rm nf} \le a_{i+1} f_{i+1}^{\rm nf}$, (21) implies item 2. ∎

Proof of Theorem 3.3: It remains to show that reducing a group of users' sensitivities will cause a net transfer of traffic from low marginal-cost paths to high marginal-cost paths (that is, from higher-order paths to lower-order). Given a Nash flow $f^{\rm nf}$ on G of a population s, suppose a small mass ϵ of users have their types changed to $S_{\rm L}$; this perturbs the original Nash flow by some $\delta \in \mathbb{R}^{|\mathcal{P}|}$. For simplicity, assume that all of these users were originally using path p_i in $f^{\rm nf}$. Lemma 4.1 implies that after the sensitivity change, these users will prefer p_1 to any other path, so the primary effect of the sensitivity change is that these users will switch from p_i to p_1.

This transfer of traffic from p_i to p_1 increases the cost of p_1 and may decrease the cost of p_i; this could cause some agents on p_1 to switch to p_2, and some agents on paths p_{i+1} and/or p_{i-1} to switch to p_i. Note that it is *not* possible for the flow on p_1 to decrease or the flow on p_i to increase. Denote by $\mathcal{P}_1$ the sequence of paths on which flow increases due to the transfer from p_i to p_1: these paths have labels $\{1,2,\ldots,j\}$. Likewise, denote by $\mathcal{P}_i$ the sequence of paths on which flow decreases due to the shift: these paths are $\{k,\ldots,i,\ldots,h\}$. Note that the path flow on any path of lower order than j cannot decrease, and the flow on any path of higher order than h cannot increase. That is, there is a net change of some nonnegative mass η of traffic from $\mathcal{P}_i$ to $\mathcal{P}_1$, and no other path flows change. We can thus lower-bound the change in total latency by applying the ordering over marginal-costs from Lemma 4.1:

$$\begin{aligned}\mathcal{L}\left(f^{\rm nf}+\delta\right) - \mathcal{L}\left(f^{\rm nf}\right) &= \sum_{i=1}^{n} a_i\delta_i^2 + \delta_i\left(2a_i f_i^{\rm nf} + b_i\right)\\ &\ge \sum_{i=1}^{n} \delta_i\left(2a_i f_i^{\rm nf} + b_i\right)\\ &\ge \eta\left(2a_j f_j^{\rm nf} + b_j - 2a_k f_k^{\rm nf} - b_k\right)\\ &\ge 0. \end{aligned} \tag{22}$$

Thus, at any Nash flow, changing the sensitivity of a small group of users to $S_{\rm L}$ results in a new Nash flow which has greater total latency than the original Nash flow. This implies for any heterogeneous population and any network $G \in \mathcal{G}_p$, denoting the taxation mechanism of Lemma 3.2 by $\tau^{\rm wr}$ and denoting a homogeneous population with sensitivity $S_{\rm L}$ by $s^{\rm L}$, that

$$\mathcal{L}^{\rm nf}\left(G, s, \tau^{\rm wr}\right) \le \mathcal{L}^{\rm nf}\left(G, s^{\rm L}, \tau^{\rm wr}\right). \tag{23}$$

Since an un-influenced Nash flow can be represented as a Nash flow for homogeneous population with 0 sensitivity, we may simply apply (23) a second time to obtain

$$\mathcal{L}^{\rm nf}\left(G, s, \tau^{\rm wr}\right) \le \mathcal{L}^{\rm nf}\left(G, \emptyset\right), \tag{24}$$

or $\tau^{\rm wr}$ is weakly robust on parallel networks.

Finally, since the worst-case Nash flows are caused by extreme low-sensitivity homogeneous populations; congestion-minimizing tolls should be designed to maximize $\frac{S_{\rm L}(\kappa_1-\kappa_2)}{1+\kappa_2 S_{\rm L}}$ subject to $\frac{S_{\rm U}(\kappa_1-\kappa_2)}{1+\kappa_2 S_{\rm U}} \le 1$. This is done by choosing κ_1 as large as possible and κ_2 such that $\frac{S_{\rm U}(\kappa_1-\kappa_2)}{1+\kappa_2 S_{\rm U}} = 1$. Given this fact, it is easy to compute (17), and (18) follows by taking the limit as $S_{\rm U} \to +\infty$. ∎