Utilizing Information Optimally to Influence Distributed Network Routing

Bryce L. Ferguson, Philip N. Brown, and Jason R. Marden

Abstract—How can a system designer exploit system-level knowledge to derive incentives to optimally influence social behavior? The literature on network routing contains many results studying the application of monetary tolls to influence behavior and improve the efficiency of self-interested network traffic routing. These results typically fall into two categories: (1) optimal tolls which incentivize socially-optimal behavior for a known realization of the network and population, or (2) robust tolls which provably reduce congestion given uncertainty regarding networks and user types, but may fail to optimize routing in general. This paper advances the study of robust influencing, mechanisms asking how a system designer can optimally exploit additional information regarding the network structure and user price sensitivities to design pricing mechanisms which influence behavior. We design optimal scaled marginal-cost pricing mechanisms for a class of parallel-network routing games and derive the tight performance guarantees when the network structure and/or the average user price-sensitivity is known. Our results demonstrate that from the standpoint of the system operator, in general it is more important to know the structure of the network than it is to know distributional information regarding the user population.

I. INTRODUCTION

In systems whose performance is driven by social behavior, it is well known that the self-interested choices of individuals can dramatically degrade system performance. This inefficiency resulting from selfish behavior is commonly characterized by the ratio between the worst-case social welfare resulting from choices of self-interested users and the optimal social welfare; this is typically referred to as the *price of anarchy* [1] and has become a highly studied area in resource allocation [2], [3], distributed control [4], and transportation [5], [6]. A common line of research studies how this inefficiency can be mitigated by using pricing mechanisms and information systems which incentivize users to make decisions more in line with the social optimum [7], [8]. Naturally, an effective implementation of incentives is heavily dependent on how the users respond to the incentives.

In this paper, we study the design of incentive mechanisms to influence social behavior in a class of congestion games; our particular focus here is on the relationship between a designer's ability to effectively influence selfish behavior and the designer's uncertainty regarding various system parameters.

Specifically, we consider a network routing problem in which a unit mass of users need to be routed across a network of two parallel edges with affine latency functions. Finding a flow that minimizes the total latency in the system is straightforward if the system designer has full control over deciding the path of each user. However, a network flow resulting from users' self-interested decisions need not be optimal [9]–[12]. Modeling the routing problem as a non-atomic congestion game, we characterize this self-interested behavior as a *Nash flow*; that is, a network flow in which every user individually chooses the lowest-latency path, given the choices of other users.

Financial incentives in the form of road tolls are a common approach to mitigate the inefficiency due to selfish behavior; these tolls modify users' preferences, inducing new, more-efficient Nash flows. In the case where the system designer is fully aware of the network topology, population size, and the users' price-sensitivity, tolls can be designed that influence users to self-route in a manner that minimizes total latency [13]–[15]. Though these tolls can induce optimal behavior, the required information can be difficult or impossible to obtain, and lack robustness to uncertainty in various system parameters [16].

Alternatively, in the case where the system designer is unaware of the structure of the network or population size, and knows only the possible support of user price-sensitivity, there exist robust taxation mechanisms which improve the efficiency for certain classes of networks. These tolls are not guaranteed to induce perfectly optimal behavior, but they bound the price of anarchy strictly below the nominal un-influenced value of $4/3$ [17]. The study of information value in system design has been considered in areas such as optimal control [18] and multi agent systems [19].

In this paper, we seek to bridge the gap between optimal taxation mechanisms which require detailed information, and robust tolls that require less information but fail to perfectly optimize routing. In [17], the system designer is oblivious to the network structure and knows only the possible support of the distribution of user price-sensitivities; using only this information, the designer selects optimal robust tolls to minimize the price of anarchy. Our contributions augment this information in three ways: one in which the designer additionally knows the network structure, one in which the designer additionally knows the average user-sensitivity, and one in which both structure and average sensitivity are known. For each information environment, we determine the tolling scheme that uses all available information *optimally*, and the resulting efficiency guarantees are compared to previously-known results for robust tolls.

In comparing the performance of these mechanisms, we ask which piece of information is more valuable to the system designer when used optimally: network structure or users' average price sensitivity. Intriguingly, the answer is highly context-dependent. In some cases, we find that the mean-sensitivity is highly informative and is thus a more valuable piece of information than the network structure. However, taken in worst-case, we find that the network structure is more valuable than the mean.

II. MODEL AND PERFORMANCE METRICS

A. Routing Game

Consider a routing problem, in which a unit mass of traffic must be directed through a parallel network from an origin node to a destination node over a set of edges E. A *feasible flow* over the network is an assignment of traffic to each edge $f = \{f_e\}_{e \in E} \in \Delta(E)$ where $f_e \geq 0$ denotes the flow on an edge e and $\Delta(E)$ denotes the standard probability simplex over the set E; that is, $\sum_{e \in E} f_e = 1$. To characterize transit delay, each edge $e \in E$ in the network has a latency function of the form

$$\ell_e(f_e) = a_e f_e + b_e. \tag{1}$$

where $a_e \geq 0$ and $b_e \geq 0$. The latency on an edge is thus a non-decreasing, non-negative function of the flow on that edge. The system cost of a flow f is characterized by the *total latency* in the network, defined as

$$\mathcal{L}(f) = \sum_{e \in E} f_e \cdot \ell_e(f_e), \tag{2}$$

and we denote the flow that minimizes this total latency as $f^{\text{opt}} \in \arg\min_{f \in \Delta E} \mathcal{L}(f)$. We specify a particular network by the tuple $G = (E, \{\ell_e\}_{e \in E})$.

This paper studies taxation mechanisms designed to influence the emergent collective behavior of self-interested, price-sensitive users. We model this routing problem as a congestion game where each edge $e \in E$ is assigned a flow dependent tolling function $\tau_e : [0, 1] \rightarrow \mathbb{R}^+$. A user $x \in [0, 1]$ has a price-sensitivity $s(x) > 0$; this price-sensitivity is subjective for each user and relates the user's cost from being tolled to their cost from experiencing delays and is the reciprocal of the user's value of time. We write $e(x)$ to denote the edge chosen by user x. The cost function for a user x that is on an edge $e(x) \in E$ can be expressed as

$$J_x(f) = \ell_{e(x)}\left(f_{e(x)}\right) + s(x)\tau_{e(x)}\left(f_{e(x)}\right). \tag{3}$$

A flow f is a *Nash flow* if for every user $x \in [0, 1]$

$$J_x(f) \in \arg\min_{e \in E}\{\ell_e(f_e) + s(x)\tau_e(f_e)\}. \tag{4}$$

A game is therefore characterized by a network G, player sensitivity distribution $s : [0, 1] \rightarrow \mathbb{R}^+$, and a set of tolling functions $\{\tau_e\}_{e \in E}$, denoted by the tuple $(G, s, \{\tau_e\}_{e \in E})$. It is shown in [20] that a Nash flow will always exist in a congestion game of this form.

B. Taxation Mechanisms & Performance Metrics

To understand the robustness of a tolling scheme, we consider the performance over a class of networks and users' sensitivities. For a network, G, we identify the latency functions which constitute the network by $L(G)$; further, for a family of networks $\mathcal{G}$, let $L(\mathcal{G}) = \bigcup_{G \in \mathcal{G}} L(G)$ be the set of all latency functions that compose the networks in $\mathcal{G}$.

A *taxation mechanism* T maps latency functions ℓ_e to tolling functions τ_e. For a family of networks $\mathcal{G}$, this mapping is denoted $T : L(\mathcal{G}) \rightarrow \mathcal{T}$, where $\mathcal{T}$ is the set of all linear tolling functions on $[0, 1]$. These tolling functions are termed *scaled marginal-cost tolls* and as discussed in Section III, are an important class of tolling functions which are optimal in many contexts. For a class of games $\mathcal{G}$, a taxation mechanism assigns a tolling function to an edge by the latency function on said edge, independently of what specific network in $\mathcal{G}$ is realized. It is important to note that our formulation fits that of the classic Pigouvian-taxes [10] given by

$$\tau_e(f) = f_e \cdot \frac{d}{df_e}\ell_e(f_e) = a_e f_e, \quad \forall\, e \in E. \tag{5}$$

To formalize the notion of uncertainty in users' response to a taxation mechanism, we consider families of sensitivity distributions that can occur when the system designer is only aware of the lower bound S_{L} and upper bound S_{U} on users price sensitives. We define the set of possible sensitivity distributions as $\mathcal{S} = \{s : [0, 1] \rightarrow [S_{\text{L}}, S_{\text{U}}]\}$. When the average price sensitivity $\overline{s}$ of the users is introduced to the system designer, the set of possible distributions becomes $\mathcal{S}(\overline{s}) = \{s \in \mathcal{S}| \int_0^1 s(x)dx = \overline{s}\}$; it is clear that $\mathcal{S}(\overline{s}) \subset \mathcal{S}$.

To evaluate the performance of a tolling mechanism, let $\mathcal{L}^{\text{nf}}(G, s, T)$ be the total latency on a network G, with price sensitivity distribution s, in the Nash flow f^{nf} when tolls are assigned according to taxation mechanism[1] T, and let $\mathcal{L}^{\text{opt}}(G)$ be the minimum total latency which occurs under the optimal flow f^{opt}. The *price of anarchy* compares the Nash flow on a network with the optimal flow; this characterizes the inefficiency of the network and can be defined as

$$\text{PoA}(G, s, T) = \frac{\mathcal{L}^{\text{nf}}(G, s, T)}{\mathcal{L}^{\text{opt}}(G)} \geq 1. \tag{6}$$

We extend this definition to include families of networks and sensitivity distributions, i.e.,

$$\text{PoA}(\mathcal{G}, \mathcal{S}, T) = \sup_{G \in \mathcal{G}} \sup_{s \in \mathcal{S}} \left\{ \frac{\mathcal{L}^{\text{nf}}(G, s, T)}{\mathcal{L}^{\text{opt}}(G)} \right\}, \tag{7}$$

such that the price of anarchy is now the worst-case inefficiency over possible networks and sensitivity distributions[2]. Note that the same taxation mechanism T is applied to any realized network-sensitivity distribution pair.

---

[1] The taxation mechanism is a mapping from latency functions to tolling functions. A game with taxation mechanism T is therefore denoted $(G, s, \{T(\ell_e)|\ell_e \in G\})$. For brevity, we simply denote this as (G, s, T).

[2] A trivial taxation mechanism of assigning zero tolling functions to each edge (i.e., $\tau_e = 0$ for all $e \in E$) will have a price of anarchy of 4/3 for this class of networks [11].

	Any sensitivity distribution	$\mathcal{S}$ *mean-agnostic*	$\mathcal{S}(\bar{s})$ *mean-aware*	$s(x)$ *population-aware*
$\mathcal{G}$ *network-agnostic*	$\mathrm{PoA}\,(\mathcal{G}) = 1.\overline{33}$ [11]	$\mathrm{PoA}^*\,(\mathcal{G},\mathcal{S}) \approx 1.176$ (A), [17]	$\mathrm{PoA}^*\,(\mathcal{G},\mathcal{S}(\bar{s})) \approx 1.1385$ (B)	
G *network-aware*		$\mathrm{PoA}^*\,(G,\mathcal{S}) \leq 1.09$ (C)	$\mathrm{PoA}^*\,(G,\mathcal{S}(\bar{s})) \leq 1.0494$ (D)	$\mathrm{PoA}\,(G,T) = 1$ [14]

Fig. 1. Price of anarchy bounds over different classes of games with optimal taxation mechanisms. Each class of games consists of parallel networks with affine latency functions. Information available to the system designer is embedded in the class of games (e.g. the network-aware classes only contain games with the same network). Smaller classes denote more certainty in the structure of the composing games, allowing for more effective toll design. These numbers are generated for $S_\mathrm{L} = 1$ and $S_\mathrm{U} = 10$, and a worst case over networks and mean sensitivities.

C. Optimal Tolling & Our Contributions

When designing a taxation mechanism, the goal of the system designer is to minimize worst-case inefficiency while the network and/or users' sensitivities are not known. Thus, we define an optimal tolling mechanism as

$$T^* \in \arg\inf_{T:L(\mathcal{G})\to\mathcal{T}} \mathrm{PoA}(\mathcal{G},\mathcal{S},T),$$

such that it is the scaled marginal-cost taxation mechanism which minimizes the price of anarchy expressed in (7) for a given family of networks $\mathcal{G}$ and sensitivity distributions $\mathcal{S}$. Further, we define the price of anarchy bound under an optimal tolling mechanism as

$$\mathrm{PoA}^*(\mathcal{G},\mathcal{S}) \triangleq \inf_{T:L(\mathcal{G})\to\mathcal{T}} \mathrm{PoA}(\mathcal{G},\mathcal{S},T). \quad (8)$$

In this paper, we consider networks with two parallel links (the family of which we simply denote $\mathcal{G}$), to demonstrate the value of information to a system designer by comparing the price of anarchy bounds when different amounts of information are available. It is not currently known if these bounds extend to richer classes of networks, but these networks often display worst-case inefficiency over larger classes of networks and allow us to analyze the benefit of these partially informed tolls [21].

The information available to the system designer is encoded in these families of networks and sensitivity distributions. When the system designer is *network-aware*, the family of networks for which the taxation mechanism is defined is a singleton (i.e., the system designer knows the network). Otherwise, we say the system designer is *network-agnostic* when only a family of possible networks is known. We look at the following four cases when the system designer is aware of the network, mean sensitivity, both, or neither:

Network-agnostic & mean-agnostic: We first consider the case that the system designer knows only the lower and upper bound on user's sensitivities. This was originally considered in [17] for parallel routing problems under a utilization constraint; here, we generalize this known result for networks with no utilization constraint. The price of anarchy bound under an optimal toll in this scenario is denoted

$$\mathrm{PoA}^*(\mathcal{G},\mathcal{S}) = \inf_{T:L(\mathcal{G})\to\mathcal{T}} \mathrm{PoA}(\mathcal{G},\mathcal{S},T), \quad (9)$$

and a general expression is given in Theorem 1. A realization of this is given in box (A) of Figure 1.

Network-agnostic & mean-aware: Second, we consider the scenario in which the system designer knows the mean of users' sensitivities, as well as the lower and upper bound. Here, the number of possible sensitivity distributions is reduced, and the system designer will be able to make better guarantees on worst-case inefficiency. The price of anarchy bound under an optimal toll in this scenario is denoted

$$\mathrm{PoA}^*(\mathcal{G},\mathcal{S}(\bar{s})) = \inf_{T:L(\mathcal{G})\to\mathcal{T}} \mathrm{PoA}(\mathcal{G},\mathcal{S}(\bar{s}),T), \quad (10)$$

and a general expression is given in Theorem 2. A realization of this is given in box (B) of Figure 1.

Network-aware & mean-agnostic: Next, we consider when the system designer again knows only the lower and upper bound on users' sensitivities, but is also aware of the network structure. The system designer will be able to make better guarantees on worst-case inefficiency when there is no uncertainty in the network that was realized. The price of anarchy under an optimal toll in this scenario is denoted

$$\mathrm{PoA}^*(G,\mathcal{S}) = \inf_{T:L(G)\to\mathcal{T}} \mathrm{PoA}(G,\mathcal{S},T), \quad (11)$$

and a general expression is given in Theorem 3. A realization of this is given in box (C) of Figure 1.

Network-aware & mean-aware: Finally, we consider the case when the system designer knows the lower bound, upper bound, and mean of users' sensitivities and also knows the network structure. Here, the system designer will be able to guarantee the best performance of any of the scenarios we consider. The price of anarchy bound under an optimal toll in this scenario is denoted

$$\mathrm{PoA}^*(G,\mathcal{S}(\bar{s})) = \inf_{T:L(G)\to\mathcal{T}} \mathrm{PoA}(G,\mathcal{S}(\bar{s}),T), \quad (12)$$

and a general expression is given in Theorem 4. A realization of this is given in box (D) of Figure 1.

Though it is clear that a network-aware mechanism will provide better guarantees than a similar network-agnostic mechanism, and that a mean-aware mechanism will provide better guarantees than a similar mean-agnostic mechanism, it is not obvious what will provide a better gain in performance when introduced to an uninformed system designer: network-awareness or mean-awareness.

To bring clarity to the results of this paper, we plot the price of anarchy bounds with respect to $\bar{s}$ for the case $S_\mathrm{L} = 1$ and $S_\mathrm{U} = 10$ in Figure 2. As expected, the
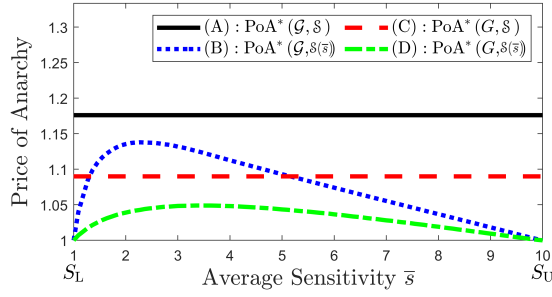
Fig. 2. Price of anarchy with respect to the mean sensitivity. Each plot represents a bound for one of the four introduced tolling mechanisms: (A) network agnostic, mean agnostic toll, (B) network agnostic, mean aware toll, (C) network aware, mean agnostic toll, and (D) network aware, mean aware toll. Each is below the untolled price of anarchy of 4/3. Price sensitivity bounds $S_L = 1$ and $S_U = 10$ are shown; changing these values has a minimal effect on the relation between the lines.

addition of information provides better guarantees on worst-case inefficiency. Notice when $\overline{s} = S_L$ or $\overline{s} = S_U$, the price of anarchy is 1 for the mean-aware tolls, i.e., the toll is optimal for any network. When the mean of user sensitivity is at one of the bounds, then the only feasible distribution is each user sharing that sensitivity. When all users share a common sensitivity, tolls can by applied to induce optimal routing [22]. However, the network-aware, mean-agnostic taxation mechanism performs better than the network-agnostic, mean-aware for means that are far from S_L and S_U. This implies that knowing the mean sensitivity of users is not enough to guarantee better performance than knowing the structure of the network.

III. MAIN RESULTS

For each scenario we consider, we give a theorem reporting the price of anarchy bound, as well as a proposition reporting the scaling factor of the optimal scaled marginal-cost toll. Note that in this paper, we focus our search for optimal tolls to a search for scaled marginal-cost tolls. These taxation mechanisms can be defined by a single scaling factor k, and will be denoted $T(k)$. It is shown in Lemma 2.2 in [17] that the search for an optimal *bounded* toll is equivalent to a search for an optimal linear toll in the network-agnostic case. The form of the optimal taxation mechanism in the network aware case remains an open question.

A. Network-agnostic, mean-agnostic

We first look at the case where the system designer has the least amount of information available to them, and formulate a price of anarchy bound under a taxation mechanism that is optimal in this scenario.

Theorem 1: When only S_L and S_U are known, the price of anarchy under an optimal, scaled marginal-cost tolling mechanism is

$$\text{PoA}^* (\mathcal{G}, \mathcal{S}) = \frac{\left(q - 1 + \sqrt{q^2 + 14q + 1} \right)^2}{8q \left(-q - 1 + \sqrt{q^2 + 14q + 1} \right)}, \quad (13)$$

where $q := S_L/S_U$.

To find this upper bound, we find the optimal scaled marginal-cost toll. We then use this toll to find the worst-case game which maximizes the price of anarchy under this toll.

Proposition 1: When only S_L and S_U are known, the optimal network-agnostic marginal-cost toll scaling factor is

$$k_{\text{agn}} = \frac{-S_L - S_U + \sqrt{S_L^2 + 14 S_L S_U + S_U^2}}{2 S_L S_U}. \quad (14)$$

Proof: The expression (14) can be shown to be a solution to the equation

$$\frac{4}{4(1 + k_{\text{agn}} S_L) - (1 + k_{\text{agn}} S_L)^2} = \frac{(1 + k_{\text{agn}} S_U)^2}{4 k_{\text{agn}} S_U}. \quad (15)$$

It is shown in [23] that when $S_L < S_U$, (15) always has a exactly one solution on the interval $(1/S_U, 1/S_L)$, and that solution is the desired optimal scale factor. We show here that (14) describes this particular solution. Let $p = S_U/S_L$. In this regime, $p > 1$, so we have

$$1 + 14p + p^2 > 1 + 14p + p^2 + 8(1 - p) = (p + 3)^2.$$

Thus,

$$k_{\text{agn}} > \frac{-1 - p + \sqrt{(p+3)^2}}{2 S_U} = \frac{1}{S_U}. \quad (16)$$

Likewise, letting $q = S_L/S_U$ (so that $q < 1$), we have

$$1 + 14q + q^2 < 1 + 14q + q^2 + 8(1 - q) = (q + 3)^2,$$

yielding

$$k_{\text{agn}} < \frac{-1 - q + \sqrt{(q+3)^2}}{2 S_U} = \frac{1}{S_L}. \quad (17)$$

∎

Proof of Theorem 1: Using the optimal scaling factor from Proposition 1, expression (13) can be found by substituting (14) into (15). ∎

B. Network-agnostic, mean-aware

When the mean sensitivity of users in the population is available to the system designer, the set of possible distributions is reduced to the set $\mathcal{S}(\overline{s}) \subset \mathcal{S}$. This additional information allows the system designer to find a new optimal taxation mechanism that improves the inefficiency bound.

Theorem 2: When S_L, S_U and the mean sensitivity $\overline{s}$ are known, the price of anarchy under an optimal, scaled marginal-cost toll is given by

$$\text{PoA}^* (\mathcal{G}, \mathcal{S}(\overline{s})) = \frac{R^2 - \alpha R + \alpha}{\alpha - \alpha^2/4}, \quad (18)$$

where $R := (S_U - \overline{s})/(S_U - S_L)$, and $\alpha = (1 + S_U k_{(\overline{s})}) R$, with $k_{(\overline{s})}$ being the solution to (26).

To find this bound, we perform a series of reductions in the set of distributions and networks we need consider in our search for one which realizes worst-case inefficiency. First, we show that for each network there exist two sensitivity distributions which realize the price of anarchy over all possible

distributions. Then, we demonstrate a method to transform a network to one that has one linear and one constant latency function and higher price of anarchy. Finally, we prove the optimal scaling factor for a scaled marginal-cost toll equates the price of anarchy for two specific networks and show that these realize the given price of anarchy bound.

We first give a list of supporting Lemmas, then use these to complete the proof of the theorem.

We say users x, y have the same *type* if $s(x) = s(y)$. Further let a *bimodal* distribution be one in which there exist exactly two user types; the set of such distributions is denoted $\mathcal{S}^{\mathrm{bi}}(\overline{s}) \subset \mathcal{S}(\overline{s})$. We denote a bimodal distribution with types S_1 and S_2 by (S_1, S_2). Note that for a given $\overline{s}$, S_1 and S_2, the mass of users with each sensitivity is well defined. Additionally, we adopt the convention used elsewhere that the network links are indexed such that $b_1 \leq b_2$.

Lemma 2.1: A Nash flow f for a sensitivity distribution $s \in \mathcal{S}(\overline{s})$, under a scaled marginal cost tax T, is likewise a Nash flow for some distribution $s' \in \mathcal{S}^{\mathrm{bi}}(\overline{s})$ in which one type of user is indifferent between the two edges and all users on each edge are of a single type. This implies the price of anarchy over sensitivities in $\mathcal{S}(\overline{s})$ is equal to the price of anarchy over bimodal distributions in $\mathcal{S}^{\mathrm{bi}}(\overline{s})$, i.e.,

$$\mathrm{PoA}(\mathcal{G}, \mathcal{S}(\overline{s}), T) = \mathrm{PoA}(\mathcal{G}, \mathcal{S}^{\mathrm{bi}}(\overline{s}), T). \qquad (19)$$

Proof: Let $s_1 \in \mathcal{S}(\overline{s})$ be some distribution of users' sensitivities, and let S_{ind} be the sensitivity that has equal cost between the two links in the Nash flow f^{nf}, i.e., solution to

$$(1 + S_{\mathrm{ind}}k)a_1 f_1^{\mathrm{nf}} + b_1 = (1 + S_{\mathrm{ind}}k)a_2 f_2^{\mathrm{nf}} + b_2. \quad (20)$$

Note that in the case where $S_{\mathrm{ind}} > S_{\mathrm{U}}$ or $S_{\mathrm{ind}} < S_{\mathrm{L}}$, any distribution $s \in \mathcal{S}$ will have the same Nash flow with all users choosing the same edge. First, consider the case where $S_{\mathrm{ind}} < \mu(s_1)$, where $\mu(\cdot)$ is the mean of the distribution. From Claim 1.1.2 in [24], if a user has a sensitivity $S < S_{\mathrm{ind}}$, then they strictly prefer the first link; if they have a sensitivity $S > S_{\mathrm{ind}}$ then they strictly prefer the second.

Now, let s_2 be a new distribution where each user who had chosen edge 1 now has sensitivity S_{ind}. The Nash flows from s_1 and s_2 are the same, as the same number of users have a sensitivity $S \leq S_{\mathrm{ind}}$ and thus the same users choose the first edge. It is clear that $\mu(s_2) > \mu(s_1)$ as no user has a lower sensitivity and some have higher.

Now, consider a third distribution s_3, where users who chose edge 2 now have some sensitivity $S' \in (S_{\mathrm{ind}}, S_{\mathrm{U}}]$; these users will now strictly prefer the second edge of the network but the Nash flow will remain unchanged. If we pick $S' = S_{\mathrm{U}}$, the mean has surely increased again; if we pick $S' = S_{\mathrm{ind}}$, because we are in the case $S_{\mathrm{ind}} < \overline{s}$, the mean is lower than $\mu(s_1)$. Because $\mu(s_3)$ is continuous with S', we can select S' so that $\mu(s_3) = \mu(s_1)$. The case of $S_{\mathrm{ind}} > \mu(s_1)$ is similar.

The distribution $s_3 = (S_{\mathrm{ind}}, S')$ induces the same Nash flow as s_1 and now, one set of users is indifferent and users of the same type exist on the same edge only. ∎

Having shown in Lemma 2.1 that the price of anarchy is realized by bimodal distributions, we further refine our

search for worst-case populations to just two specific bimodal distributions with simple characterizations.

Lemma 2.2: For a given network $G \in \mathcal{G}$ and scaled marginal-cost tax T with toll scaling factor k, two distributions $s_{\ell}^{(\overline{s}, G, k)}$ and $s_u^{(\overline{s}, G, k)}$, that maximize and minimize (respectively) the flow on the first edge of the network, realize the price of anarchy over those in $\mathcal{S}^{\mathrm{bi}}(\overline{s})$,

$$\mathrm{PoA}(G, \mathcal{S}^{\mathrm{bi}}(\overline{s}), T) = \mathrm{PoA}(G, \{s_{\ell}^{(\overline{s}, G, k)}, s_u^{(\overline{s}, G, k)}\}, T) \quad (21)$$

Proof: The proof follows from the fact that total latency is quadratic with the flow, thus the largest price of anarchy will come from the flow that is furthest from optimal. From Lemma 2.1, we see that any flow induced by a distribution $s \in \mathcal{S}(\overline{s})$ can be realized by a bimodal distribution that has one set of users observing equal cost between the links and each edge containing only one sensitivity type. We therefore define $s_{\ell}^{(\overline{s}, G, k)}$ as the distribution that maximizes f_1^{nf} and $s_u^{(\overline{s}, G, k)}$ as the distribution which maximizes f_1^{nf}. ∎

We next reduce our search for a worst-case network to a set of graphs that have a linear latency function on the first edge and a constant latency on the second, denoted $\mathcal{G}^{\mathrm{lc}}$.

Lemma 2.3: For any $G \in \mathcal{G}$, there exists a $\hat{G} \in \mathcal{G}^{\mathrm{lc}} \subset \mathcal{G}$ that, under the same scaled marginal cost tolling mechanism $T(k)$, has a higher price of anarchy, implying,

$$\mathrm{PoA}(\mathcal{G}, \mathcal{S}(\overline{s}), T(k)) = \mathrm{PoA}(\mathcal{G}^{\mathrm{lc}}, \mathcal{S}(\overline{s}), T(k)). \quad (22)$$

Proof: Consider a network $G \in \mathcal{G}$ with affine latency functions on each link $\ell_i(f) = a_i f + b_i$. Let $\hat{G}$ have cost functions $\hat{\ell}_i(f) = \hat{a}_i f + \hat{b}_i$ with $\hat{a}_i \geq 0$ and $\hat{b}_i \geq 0$.

We first show that simply removing the constant latency term on the first edge b_1 strictly increases the price of anarchy under any scaled marginal-cost toll.

Using the optimal and Nash flow in (23), if $\hat{b}_2 = b_2 - b_1$ and $\hat{b}_1 = 0$ then G and $\hat{G}$ will have the same optimal flow and Nash flow for a distribution s. From (2), we observe that $\mathcal{L}^{\mathrm{opt}}(G) = \mathcal{L}^{\mathrm{opt}}(\hat{G}) + b_1$ as well as $\mathcal{L}^{\mathrm{nf}}(G, s, k) = \mathcal{L}^{\mathrm{nf}}(\hat{G}, s, k) + b_1$; therefore,

$$\mathrm{PoA}(G, \mathcal{S}(\overline{s}), T(k)) = \frac{\mathcal{L}^{\mathrm{nf}}\left(\hat{G}, s, T(k)\right) + b_1}{\mathcal{L}^{\mathrm{opt}}(\hat{G}) + b_1}$$

$$\leq \frac{\mathcal{L}^{\mathrm{nf}}\left(\hat{G}, s, T(k)\right)}{\mathcal{L}^{\mathrm{opt}}(\hat{G})} = \mathrm{PoA}(\hat{G}, \mathcal{S}(\overline{s}), T(k)).$$

Thus, for any network $G \in \mathcal{G}$, there exists a network $\hat{G}$ with a linear latency function on an edge with higher price of anarchy.

Next, we show a network $G \in \mathcal{G}$ will have the same price of anarchy as a network $\hat{G} \in \mathcal{G}$ under the same scaled marginal cost toll if the latency functions of $\hat{G}$ equal the latency functions of G times a scaling factor c.

Under a distribution $s \in \mathcal{S}$, G and $\hat{G}$ will have the same Nash flow. Using the indifferent sensitivity S_{ind} that is the solution to (20), the Nash flow and optimal flow on the first edge are

$$f_1^{\mathrm{opt}} = \frac{2a_2 + b_2 - b_1}{2(a_1 + a_2)}, \quad f_1^{\mathrm{nf}} = \frac{(1 + S_{\mathrm{ind}}k)a_2 + b_2 - b_1}{(1 + S_{\mathrm{ind}}k)(a_1 + a_2)}. \quad (23)$$

Under the same distribution s, the same sensitivity S_{ind} will satisfy

$$(1 + S_{\text{ind}}k)ca_1 f_1^{\text{nf}} + cb_1 = (1 + S_{\text{ind}}k)ca_2 f_2^{\text{nf}} + cb_2, \quad (24)$$

which are the latency functions for the network $\hat{G}$. It is now clear that G and $\hat{G}$ will have the same Nash and optimal flows. From the definition of total latency in (2), the latency in $\hat{G}$ will be c times the latency in G under the same flow. The price of anarchy, which is the ratio of two total latencies, will be identical in G and G'.

Lastly, we show that by decreasing a_2 in a network, the price of anarchy will increase. In Lemma 2.1, it was shown that any feasible Nash flow can be induced by a bimodal sensitivity distribution in which users are segregated on either link by their sensitivity. The price of anarchy for the network G with a Nash flow caused by s will therefore be,

$$\text{PoA}(G, s, T(k)) = \frac{\ell_1(f_1^{\text{nf}})f_1^{\text{nf}} + \ell_2(f_2^{\text{nf}})f_2^{\text{nf}}}{\ell_1(f_1^{\text{opt}})f_1^{\text{opt}} + \ell_2(f_2^{\text{opt}})f_2^{\text{opt}}}. \quad (25)$$

Let us consider the case where $f_2^{\text{nf}} > f_2^{\text{opt}}$. Now, consider a new network, $\hat{G}$ which replaces latency function $\ell_2(f) = a_2 f + b_2$ in G with $\hat{\ell}_2(f) = a_2 f + \hat{b}_2$ where $\hat{b}_2 = b_2 + \delta$ such that $\delta > 0$. Because the users are segregated on the links, the Nash flow will not change. Note that because $f_2^{\text{nf}} > f_2^{\text{opt}}$

$$\frac{\ell_2(f_2^{\text{nf}})}{\ell_2(f_2^{\text{opt}})} = \frac{a_2 f_2^{\text{nf}} + b_2}{a_2 f_2^{\text{opt}} + b_2} < \frac{f_2^{\text{nf}}}{f_2^{\text{opt}}}.$$

It can now be shown that

$$\frac{\mathcal{L}^{\text{nf}}(G, s, T(k))}{\mathcal{L}^{\text{opt}}(G)} < \frac{\mathcal{L}^{\text{nf}}(G, s, T(k)) + \delta f_2^{\text{nf}}}{\mathcal{L}^{\text{opt}}(G) + \delta f_2^{\text{opt}}}$$
$$= \frac{\mathcal{L}^{\text{nf}}(\hat{G}, s, T(k))}{\mathcal{L}^{\text{opt}}(\hat{G})}.$$

Thus the price of anarchy has increased in the new network $\hat{G}$, under the same sensitivity distribution and scaled marginal cost toll, when b_2 was increased, which has the same effect as decreasing the other terms and holding b_2 constant. A very similar argument can be followed for when $f_2^{\text{nf}} < f_2^{\text{opt}}$ by picking $\hat{a}_2 = a_2 - \delta$, and the price of anarchy again increases. ∎

We further reduce our search for a worst case network to two specific networks. For a given set of distributions $\mathcal{S}(\overline{s})$ and toll scaling factor k, we define two networks:

(1) $G_\beta \in \mathcal{G}^{\text{lc}}$ with latency functions $\ell_1(f_1) = f_1$ and $\ell_2(f_2) = \beta$ and satisfies $s_\ell^{(\overline{s}, G_\beta, k)} = (S_{\text{L}}, S_{\text{U}})$, and

(2) $G_\alpha \in \mathcal{G}^{\text{lc}}$ with latency functions $\ell_1(f_1) = f_1$ and $\ell_2(f_2) = \alpha$ and satisfies $s_u^{(\overline{s}, G_\alpha, k)} = (S_{\text{L}}, S_{\text{U}})$. Due to the discussion in the proof of Lemma 2.3, any network in $\mathcal{G}^{\text{lc}}$ with cost functions satisfying $b_2/a_1 = \beta$ will have the same price of anarchy as G_β, and the same is true for for G_α.

Lemma 2.4: For linear constant networks, under sensitivity distributions in $\mathcal{S}(\overline{s})$ with toll scaling factor k, the network G_α or G_β will realize the upper bound on the price of anarchy, i.e.,

$$\text{PoA}(\mathcal{G}^{\text{lc}}, \mathcal{S}(\overline{s}), T(k)) = \text{PoA}(\{G_\alpha, G_\beta\}, \mathcal{S}(\overline{s}), T(k)).$$

Proof: It can be seen by differentiation of (23), the price of anarchy increases with the value of the indifferent sensitivity when $f_1^{\text{nf}} < f_1^{\text{opt}}$ and decreases when $f_1^{\text{nf}} < f_1^{\text{opt}}$. Recall that $s_\ell^{(\overline{s}, G, k)}$ has $f_{1\ell} > f_1^{\text{opt}}$ and indifferent sensitivity $S_{\ell 1}$; similarly, $s_u^{(\overline{s}, G, k)}$ has $f_{1u} < f_1^{\text{opt}}$ and indifferent sensitivity S_{u2}. It is therefore true that having $S_{\ell 1} = S_{\text{L}}$ or $S_{u2} = S_{\text{U}}$ is a necessary condition for the network which maximizes the price of anarchy.

Further, in bimodal distributions (S_1, S_2) where users are homogeneous on either link, $f_1^{\text{nf}} = (S_2 - \overline{s})/(S_2 - S_1)$. For $s_\ell^{(\overline{s}, G, k)}$ when users with sensitivity $S_{\ell 1} = S_{\text{L}}$ are indifferent, the largest flow that can occur on f_1 occurs when $s_\ell^{(\overline{s}, G, k)} = (S_{\text{L}}, S_{\text{U}})$. Similarly, for $s_u^{(\overline{s}, G, k)}$, when users with sensitivity $S_{u2} = S_{\text{U}}$ are indifferent, the least flow in f_1 has $s_u^{(\overline{s}, G, k)} = (S_{\text{L}}, S_{\text{U}})$. One of these two conditions must be met by a network $G \in \mathcal{G}^{\text{lc}}$ to maximize the price of anarchy. Those networks are the defined G_α and G_β. ∎

Proposition 2: When S_{L}, S_{U} and the mean sensitivity $\overline{s}$ are known, the optimal network-agnostic marginal-cost toll scaling factor $k_{(\overline{s})}$ will be the solution on $(1/S_{\text{U}}, 1/S_{\text{L}})$ to

$$\frac{4(1 + S_{\text{U}}k - S_{\text{U}}kR)}{(4 + R)(1 + S_{\text{U}}k) + (S_{\text{U}}k + S_{\text{U}}^2 k^2)R}$$
$$= \frac{4(1 + S_{\text{L}}k - S_{\text{L}}kR)}{(4 + R)(1 + S_{\text{L}}k) + (S_{\text{L}}k + S_{\text{L}}^2 k^2)R}, \quad (26)$$

Proof: The k that solves (26) equates the price of anarchy for G_β and G_α. To show this is optimal, it is sufficient to show that the price of anarchy for the network G_β is monotone decreasing with k while the price of anarchy for the network G_α is monotone increasing with k. If the networks have this monotonic relation with k, then the k that minimizes the price of anarchy must equalize them.

Consider a network $G \in \mathcal{G}^{\text{lc}}$ characterized by $\gamma = b_2/a_1$. If this satisfies that $s_\ell^{(\overline{s}, G, k)} = (S_{\text{L}}, S_{\text{U}})$ or $s_u^{(\overline{s}, G, k)} = (S_{\text{L}}, S_{\text{U}})$, then the price of anarchy for this network will be

$$\text{PoA}(G, \mathcal{S}(\overline{s}), T) = \frac{R^2 - \gamma R + \gamma}{\gamma - \gamma^2/4}. \quad (27)$$

This expression is locally minimized by $\gamma = 2R$, further, by differentiation, it can be observed that it is monotone decreasing for $0 < \gamma < 2R$ and monotone increasing for $\gamma > 2R$.

For the previously defined network G_β, under the bimodal distribution $(S_{\text{L}}, S_{\text{U}})$,

$$(1 + S_{\text{L}}k)R = \beta, \quad (28)$$

where β is dependent on the scaling factor k. From [23], the optimal scaling factor k will be in $(1/S_{\text{U}}, 1/S_{\text{L}})$. Therefore, for any k, $\beta < 2R$. The price of anarchy for this network is therefore monotone decreasing with β, and from (28), β is clearly increasing with k. The price of anarchy of the network is therefore decreasing with k. Similarly for G_α, under the distribution $(S_{\text{L}}, S_{\text{U}})$,

$$(1 + S_{\text{U}}k)R = \alpha > 2R, \quad (29)$$

and the price of anarchy will be increasing with k. ∎

Proof of Theorem 2: From Lemma 2.4, a network G_α realizes the price of anarchy when the toll scaling factor is chosen optimally as in Proposition 2. The price of anarchy for this network is found by substituting α from (29) into the latency function ratio in (27). ∎

C. Network-aware, mean-agnostic

If the system designer is made aware of the structure of the network, then the optimal network-aware taxation mechanism will out-perform one that is designed without such knowledge. Though it is still an open question as to what tolls are optimal for a network-aware system designer in this scenario, we opt to still consider linear tolling functions to provide a means of comparing with the previously introduced, network-agnostic taxation mechanisms.

Theorem 3: When only S_L and S_U are known, the price of anarchy under an optimal network-aware scaled marginal-cost toll is tightly upper bounded by

$$\text{PoA}^*(G, \mathcal{S}) \leq \frac{4}{3}\left(1 - \frac{\sqrt{q}}{\left(1 + \sqrt{q}\right)^2}\right) \tag{30}$$

where $q := S_L/S_U$.

For the proof of this bound, we implicitly assume that the system operator can determine hypothetical homogeneous low-sensitivity Nash flows associated with each tolling factor; this is reasonable since the Nash flow for a homogeneous population is known to be the solution to a convex optimization problem [11]. In the following, we write $f_i^{\text{nf}}(G, S, k)$ to denote the amount of traffic on link i in a Nash flow on network G with homogeneous sensitivity S and toll scale factor k.

Proposition 3: For any network $G \in \mathcal{G}$ and any $S_U \geq S_L > 0$, let $k^{\text{gm}} = (S_L S_U)^{-1/2}$. The following is an optimal network-aware marginal-cost toll scaling factor:

$$k_{(G)} = \begin{cases} 0 & \text{if } f_2^{\text{nf}}(G, S_L, k^{\text{gm}}) = 0, \\ k^{\text{gm}} & \text{otherwise.} \end{cases} \tag{31}$$

Proof: Consider the following cases, differentiated by the structure of Nash flows resulting from $k = k^{\text{gm}} := (S_L S_U)^{-1/2}$:

1) $f_2^{\text{nf}}(G, S_L, k^{\text{gm}}) > 0$, and
2) $f_2^{\text{nf}}(G, S_L, k^{\text{gm}}) = 0$.

It is shown in [24] that in Case (1), it must be true that $\mathcal{L}^{\text{nf}}(G, S_L, k) = \mathcal{L}^{\text{nf}}(G, S_U, k)$ and that this choice of k is uniquely optimal, resulting in the price of anarchy given in (30).

Consider Case (2). Here, the extreme low-sensitivity population with $s = S_L$ strictly prefers link 1 when $k = k^{\text{gm}}$, effectively stripping the designer of her influence over the price of anarchy. It can easily be shown (using, e.g., tools from [24]) that

$$k' \leq k^{\text{gm}} \implies \mathcal{L}^{\text{nf}}(G, S_L, k) = \mathcal{L}^{\text{nf}}(G, \emptyset), \tag{32}$$

but that

$$k^\dagger > k^{\text{gm}} \implies \mathcal{L}^{\text{nf}}(G, S_U, k^\dagger) > \mathcal{L}^{\text{nf}}(G, \emptyset). \tag{33}$$

That is, in this regime, the designer cannot change the behavior of $s = S_L$ without increasing tolls, but cannot increase tolls because this would cause the high-sensitivity population with $s = S_U$ to route more inefficiently. That is, $k = 0$ is an optimal tolling coefficient in this case.[3] ∎

Proof of Theorem 3 It follows easily from the results in [24] that in Case (2) when $k \leq k^{\text{gm}}$, it is true for any s that $\mathcal{L}^{\text{nf}}(G, s, k) \leq \mathcal{L}^{\text{nf}}(G, S_U, k^{\text{gm}})$; the price of anarchy bound for this scenario is thus precisely that in [24], where now we include games which need not have flow on every edge in an untolled Nash flow. ∎

D. Network-aware, mean-aware

Finally, we consider the network-aware, mean-aware scaled marginal-cost toll to illustrate the gain in performance when both network and population information is available.

Theorem 4: When S_L, S_U and the mean sensitivity $\bar{s}$ are known, under an optimal network-aware scaled marginal-cost tolling mechanism, the price of anarchy is tightly upper bounded by

$$\text{PoA}^*(G, \mathcal{S}(\bar{s})) \leq \frac{R^2 - \beta R + \beta}{\beta - \beta^2/4}, \tag{34}$$

where $R = (S_U - \bar{s})/(S_U - S_L)$ and β is the unique solution on the interval $[0, 2]$ to

$$\beta = R\left(1 + \sqrt{\frac{1 + R - \beta}{\bar{s}/S_L + R - \beta}}\right). \tag{35}$$

We prove the theorem by making several of the same reductions as in Theorem 2. However, we now derive an optimal network-aware toll scaling factor $k_{(\bar{s}, G)}$ and show a unique network will realize the upper bound on inefficiency in this setting.

Proposition 4: For a network $G \in \mathcal{G}$ with price sensitivity distributions $s \in \mathcal{S}(\bar{s})$ with extreme sensitivity distributions $s_\ell^{(\bar{s}, G, k)} = (S_{\ell 1}, S_{\ell 2})$ and $s_u^{(\bar{s}, G, k)} = (S_{u1}, S_{u2})$, the optimal toll scaling factor for a scaled marginal cost toll will take the form,

$$k_{(\bar{s}, G)} = \frac{1}{\sqrt{S_{\ell 1} S_{u2}}}. \tag{36}$$

Proof of Proposition 4: From Lemma 2.1, under the same tolling mechanism, the set of Nash flows caused by $\mathcal{S}(\bar{s})$ is equivalent to those caused by distributions with bounds $[S_{\ell 1}, S_{u2}]$ and no mean constraint. The optimal scaling factor will therefore minimize the price of anarchy over this set of distributions. From [24], the optimal scaling factor for a scaled marginal-cost toll will take this form. ∎

In Lemma 2.3, it was shown that a transformation from a network $G \in \mathcal{G}$ to a network $\hat{G} \in \mathcal{G}^{\text{lc}}$ will increase the price of anarchy; we also note that this transformation had no dependence on the toll scaling factor and we can thus choose a k that is optimal for the resulting network.

---

[3]In this case, the set of price-of-anarchy-minimizing tolling coefficients is not a singleton in general: any coefficient satisfying $\mathcal{L}^{\text{nf}}(G, S_L, k) \geq \mathcal{L}^{\text{nf}}(G, S_U, k)$ is optimal. Implication (32) means that this set always contains $k = 0$.

Corollary 1: When making a reduction from $G \in \mathcal{G}$ to $\hat{G} \in \mathcal{G}^{\mathrm{lc}}$, the price of anarchy increases regardless of the toll scaling factor k, including when $k = k_{(\overline{s},G)}$ for each network before and after the reduction.

Proof: In the proof of Lemma 2.3, the relation between k and the price of anarchy was not used; instead, it was shown that the price of anarchy increases as the network is transformed from any two link network, to one that was in $\mathcal{G}^{\mathrm{lc}}$. Consider having network G with the non-optimal toll scaling factor $k_{(\overline{s},\hat{G})}$. When the reduction from G to $\hat{G}$ is done, by Lemma 2.3 we have

$$\mathrm{PoA}(G, \mathcal{S}(\overline{s}), T(k_{(\overline{s},G)})) \leq \mathrm{PoA}(G, \mathcal{S}(\overline{s}), T(k_{(\overline{s},\hat{G})}))$$
$$\leq \mathrm{PoA}(\hat{G}, \mathcal{S}(\overline{s}), T(k_{(\overline{s},\hat{G})})). \quad \blacksquare$$

Proof of Theorem 4: It is shown in Lemma 2.4 that a set of two networks realizes the price of anarchy. The price of anarchy for the network G_β is found by (27). Now, let G', defined by β', be a network that has the same price of anarchy when the flow R is on the first link, One solution is clearly $\beta = \beta'$, the other is $\beta' = \frac{(4-\beta)R}{R^2 - \beta R + \beta}$. Using the cost function of network G, we have $\beta = (1 + S_{\mathrm{L}} k_{(\overline{s},G)})R$. Thus, if β' satisfied $s_u^{(\overline{s},G,k)} = (S_{\mathrm{L}}, S_{\mathrm{U}})$ for the same mean sensitivity, then

$$\beta' = \frac{(4-\beta)R}{R^2 - \beta R + \beta} = (1 + S_{\mathrm{U}} k_{(\overline{s},G')})R. \quad (37)$$

However, it can be shown that the right hand side of (28) is strictly less than the left hand side. This imposes that the flow $f_1 = R$ can not be a Nash flow in G' under distributions in $\mathcal{S}(\overline{s})$ and therefore not achieve the same price of anarchy as G. This implies that the price of anarchy for G_β is greater than that of G_α when both are tolled optimally with respect to Proposition 4.

As this network is optimally tolled, from Proposition 4, it will be the case that

$$S_{u2} = \frac{\overline{s} - S_{\mathrm{L}}}{1 + R - \beta} + S_{\mathrm{L}}. \quad (38)$$

Now, in sensitivity distribution $s_f^{(\overline{s},G,k)} = (S_{\mathrm{L}}, S_{\mathrm{U}})$, users with sensitivity S_{L} are indifferent with optimally scaled toll $k_{(\overline{s},G)}$. Using β from (28) and substituting the optimal scaling factor with extreme sensitivity from (38) leads to the characterization of β in the theorem statement, and the price of anarchy is found by substituting this into (27). $\quad \blacksquare$

IV. Conclusion

This paper represents a study of the challenges and opportunities afforded to a system designer when faced with information. While it is clear that gathering additional information about a problem can help a designer to influence behavior more effectively, the question of how to use the information optimally is not trivial, even for simple classes of problems. Ongoing work focuses on extending this paper's results to more general classes of networks and determining when precisely scaled marginal-cost taxes are optimal over all taxation mechanisms.

References

[1] C. Papadimitriou, "Algorithms, games, and the internet," in *STOC '01: Proceedings of the thirty-third annual ACM symposium on Theory of computing*, 2001, pp. 749–753.

[2] R. Johari and J. N. Tsitsiklis, "Efficiency Loss in a Network Resource Allocation Game," *Mathematics of Operations Research*, vol. 29, no. 3, pp. 407–435, aug 2004.

[3] D. Paccagnan, R. Chandan, and J. R. Marden, "Distributed resource allocation through utility design - part I: optimizing the performance certificates via the price of anarchy," *CoRR*, vol. abs/1807.01333, 2018.

[4] J. R. Marden and J. S. Shamma, "Game Theory and Distributed Control," in *Handbook of Game Theory Vol. 4*, H. Young and S. Zamir, Eds. Elsevier Science, 2014.

[5] G. Piliouras, E. Nikolova, and J. S. Shamma, "Risk sensitivity of price of anarchy under uncertainty," in *Proceedings of the 14th ACM Conference on Electronic Commerce*, vol. 9, no. 4, New York, New York, USA, 2013, pp. 715–732.

[6] A. Moradipari and M. Alizadeh, "Electric vehicle charging station network equilibrium models and pricing schemes," *CoRR*, vol. abs/1811.08582, 2018.

[7] L. J. Ratliff, R. Dong, S. Sekar, and T. Fiez, "A Perspective on Incentive Design: Challenges and Opportunities," *Annual Review of Control, Robotics, and Autonomous Systems*, vol. 2, no. 1, pp. 1–34, 2018.

[8] D. Acemoglu, A. Makhdoumi, A. Malekian, and A. Ozdaglar, "Informational Braess' Paradox: The Effect of Information on Traffic Congestion," *Operations Research*, vol. 66, no. 4, 2018.

[9] J. Wardrop, "Some Theoretical Aspects of Road Traffic Research," in *Proc. of the Institution of Civil Engineers, Part II, Vol. 1, No. 36*, 1952, pp. 352–362.

[10] A. C. Pigou, *The Economics of Welfare*. Macmillan, New York, 1920.

[11] T. Roughgarden, "How Bad Is Selfish Routing?" *J. ACM*, vol. 49, no. 2, pp. 236–259, 2002.

[12] R. Cole, T. Lianeas, and E. Nikolova, "When does diversity of agent preferences improve outcomes in selfish routing?" in *IJCAI International Joint Conference on Artificial Intelligence*, 2018.

[13] 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, 2003, pp. 521–530.

[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, 2004, pp. 277–285.

[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, 2004, pp. 268–276.

[16] P. N. Brown and J. R. Marden, "Studies on Robust Social Influence Mechanisms: Incentives for Efficient Network Routing in Uncertain Settings," *IEEE Control Systems Magazine*, vol. 37, no. 1, pp. 98–115, 2017.

[17] ——, "Optimal Mechanisms for Robust Coordination in Congestion Games," *IEEE Transactions on Automatic Control*, vol. 63, no. 8, pp. 2437–2448, 2018.

[18] F. Farokhi and K. H. Johansson, "Optimal control design under limited model information for discrete-time linear systems with stochastically-varying parameters," *IEEE Transactions on Automatic Control*, vol. 60, no. 3, pp. 684–699, March 2015.

[19] D. Paccagnan and J. R. Marden, "The importance of system-level information in multiagent systems design: Cardinality and covering problems," *IEEE Transactions on Automatic Control*, pp. 1–1, 2018.

[20] A. Mas-Colell, "On a Theorem of Schmeidler," *Mathematical Economics*, vol. 13, pp. 201–206, 1984.

[21] T. Roughgarden, "The price of anarchy is independent of the network topology," *J. Computer and System Sciences*, vol. 67, no. 2, pp. 341–364, sep 2003.

[22] M. J. Beckmann, C. B. McGuire, and C. B. Winsten, "Studies in the Economics of Transportation," p. 359, 1955.

[23] P. N. Brown and J. R. Marden, "The Benefit of Perversity in Distributed Network Routing," in *56th IEEE Conf. on Decision and Control*, Melbourne, Australia, 2017, pp. 6229–6234.

[24] ——, "The Robustness of Marginal-Cost Taxes in Affine Congestion Games," *Transactions on Automatic Control*, vol. 62, no. 8, pp. 3999–4004, 2017.