This article has been accepted for publication in a future issue of this journal, but has not been fully edited. Content may change prior to final publication. Citation information: DOI 10.1109/TAC.2019.2962102, IEEE Transactions on Automatic Control

Uncertainty in Multi-Commodity Routing Networks: When does it help?

Shreyas Sekar, Liyuan Zheng, Lillian J. Ratliff, and Baosen Zhang

***Abstract*—We study the equilibrium behavior in a multi-commodity selfish routing game with uncertain users where each user over- or under-estimates their congestion costs by a multiplicative factor. Surprisingly, we find that uncertainties in different directions have qualitatively distinct impacts on equilibria. Namely, contrary to the usual notion that uncertainty increases inefficiencies, network congestion decreases when users over-estimate their costs. On the other hand, under-estimation of costs leads to increased congestion.**

We apply these results to urban transportation networks, where drivers have different estimates about the cost of congestion. In light of the dynamic pricing policies aimed at tackling congestion, our results indicate that users' perception of these prices can significantly impact the policy's efficacy, and "caution in the face of uncertainty" leads to favorable network conditions.

***Index Terms*—Network Routing, Uncertainty, Nash Equilibrium, Transportation**

I. Introduction

Multi-commodity routing networks that allocate resources to self-interested users lie at the heart of many systems such as communication, transportation, and power networks (see, e.g., [1] for an overview). In all of these systems, users *inherently face uncertainty and are heterogeneous*. These users rarely have perfect information about the state of the system, and each have their own idiosyncratic objectives and trade-offs between time, money, and risk [2]–[4]. Naturally, users' personalized beliefs or preferences regarding system costs and delays influence their decision and, in turn, the welfare of the overall system. In this paper, we provide an understanding of the effects of certain classes of uncertainties and limited user heterogeneity with respect to such uncertainties on network performance—i.e. we establish conditions on when they are helpful and harmful to the overall social welfare.

A motivating example of a routing network that we use throughout this paper is the *urban transportation network*. Commuters in road networks simultaneously trade-off between diverse objectives such as total travel time, road taxes, parking costs, waiting delays, walking distance and environmental impact. At the same time, these users tend to possess varying levels of information, and there is evidence [5], [6] to suggest that the routes adopted depend not on the true costs but on how they are perceived by the users. For instance, users prefer more consistent routes over those with high variance [7], seek to minimize travel time over parking costs [8], and react adversely to per-mile road taxes [9].

S. Sekar is affiliated with Harvard University, Cambridge, MA 02138. L. Zheng, L. Ratliff, and B. Zhang are with the Department of Electrical Engineering, University of Washington, Seattle, WA 98195. (email: ssekar@hbs.edu, {liyuanz8, ratliffl, zhangbao}@uw.edu).

This work was supported by NSF awards CNS-1634136 and CNS-1646912.

Furthermore, the technological and economic incentives employed by planners interact with user beliefs in a complex manner [10]. For example, to limit the economic losses arising from urban congestion, cities across the world have introduced a number of solutions including road taxes, time-of-day-pricing, road-side message signs and route recommendations [11]–[13]. However, the dynamic nature of these incentives (e.g., frequent price updates) and the limited availability of information dispersal mechanisms may add to users' uncertainties and asymmetries in beliefs. The effect of uncertainties on network equilibria has been examined in recent work [10], [14]–[16] where each user perceives the network condition to be different than the true conditions. The current results have mostly focused on simple network topologies (e.g., parallel links) or networks where a fixed percentage of the population is endowed with a specific level of uncertainty. Given the complexity of most practical networks, it is natural to ask how uncertainty (and user beliefs on network costs) affects equilibria in scenarios with at least two types of users, whose perceptions vary according to the user type. Specifically, in this work we are motivated by the following two questions: (i) *how do equilibria depend on the type and level of uncertainty among network users*, and (ii) *when does uncertainty improve or degrade equilibrium quality?*

To address these questions, we turn to a *multi-commodity selfish routing* framework commonly employed by many disciplines. In our model, users seek to route some flow from a source to a destination across a network and they face congestion costs on each link. Crucially, users' perception of these congestion costs may differ from the true cost or travel time. It is well-known that even in the presence of perfect information (every user knows the exact true cost), strategic behavior by the users can result in considerably worse congestion at equilibrium when compared to a centrally optimized routing solution [17]. Against this backdrop, we analyze what happens when users have imperfect knowledge of the congestion costs. A surprising outcome arises: in the presence of uncertainty, if users select routes based on perceived costs that over-estimate the true cost, *the equilibrium quality is better compared to perfect information case*. Conversely, if the users are not cautious and under-estimate the costs, the equilibrium quality becomes worse.

A. Contributions

We introduce the notion of *type-dependent uncertainty* in multi-commodity routing networks, where the uncertainty of users belonging to type θ is captured by a single parameter $r_\theta > 0$. Specifically, for each user of type θ, if their true cost on edge e is given by $C_e(x) = a_e x + b_e$, where x is the total population of users on this edge, then their *perceived cost* is $r_\theta a_e x + b_e$. For the majority of this work, we focus on *cautious behavior*, where users over-estimate the costs ($r_\theta \geq 1$), for all types θ. Some of our results will also hold for the case where users under-estimate the costs ($r_\theta \leq 1$), for all types θ. The central message of this paper is that when users exhibit "caution in the face of uncertainty", the social cost at the equilibrium is smaller compared to the case where users have perfect information (i.e. know the true congestion costs).

The following results are *independent of network topology*:

(a) The social cost—i.e. $\mathrm{C}(\mathbf{x}) = \sum_{e\in\mathcal{E}} x_e C_e(x_e)$ where x_e is the total population mass, summed over all user types, flowing on edge e—of the equilibrium solution(s) where all users have the same level of uncertainty ($r_\theta = r$ for all θ) is *always smaller than or equal to* the cost of the equilibrium solution without uncertainty when $r \in [1, 2]$ and vice-versa when $r < 1$.

(b) The worst-case ratio of the social cost of the equilibrium to that of the socially optimal solution (i.e. the price of anarchy [17]) is $4/(4r_{\max}\gamma - r_{\max}^2)$, where $r_{\max} = \max_\theta r_\theta$ and $\gamma \leq 1$ is the ratio of the minimum to the maximum uncertainty over user types.

Constraining network topology, we show the following:

(a) The social cost of the equilibrium where a fraction of the users exhibit an uncertainty of $r \in [1, 2]$ and the rest have no uncertainty *is always smaller* than or equal to the social cost of an identical system without uncertainty, as long as the network has the serially linearly independent topology [18].

(b) In systems having users with and without uncertainty, the routing choices adopted by the uncertain users always results in an improvement in the costs experienced by users without uncertainty, as long as the graph has a series-parallel topology [18].

Finally, we prove that all of our results generalize gracefully to a class of well-motivated polynomial functions known as *shifted monomials*, where $C_e(x) = a_e x^d + b_e$ for $d \geq 1$. In fact, for these general functions, we show that uncertainty is typically beneficial over a larger range of the parameter r, i.e., when $r \in [1, d+1]$. Our results provide a complete characterization for routing games with two user types or uncertainty levels and a worst-case price of anarchy bound for instances having more than two types. Moreover, we show that our characterizations are tight by means of illustrative examples where uncertainty leads to an increase in the social cost when our characterization conditions are violated.

To validate the theoretical results, we present illustrative simulation results. We focus specifically on the application of parking in urban transportation networks and consider realistic urban network topologies with two types of users: through traffic and parking users. Given a parking population with uncertainties, we show that *cautious behavior improves equilibrium quality while lack of caution degrades it* even when uncertainty is asymmetric across user types and when the same user faces different levels of uncertainty on different parts of the network.

A preliminary version of this paper [19] appeared in the 2018 American Control Conference. The current work is a significant generalization of that version including (i) several new results (Theorems 3, 4, 5, 6); (ii) a new model and simulations and a detailed discussion of the modeling choices (Section II-B); (iii) expanded related work and connections between our results and those in the literature on tolling.

B. Related Work

There is an extensive literature on congestion games and more general potential games [20] which has focused on quantifying the equilibrium cost as a function of system parameters such as the network topology [21] and the degree of cost functions [17]. Although we study the same research questions as this literature, we look at settings where the users' perceived cost on a route may not be equal to the true expected cost. In contrast, much of the traditional work in this domain look at models where the users are aware of the precise costs which leads to considerably different results.

A notable exception is the body of work on routing games with player specific costs which may or may not align with the true costs [22]–[25]. First, a majority of these works look at games with atomic users, which are known to have qualitatively different results than the non-atomic game that we consider. Second, they primarily analyze the existence of and convergence to equilibrium solutions, whereas we study more quantitative questions such as the price of anarchy, and comparing the social cost of the equilibrium solutions with and without cost misalignment.

Our work is closely related to the extensive body of work on risk-averse selfish routing [26]–[28] and pricing tolls in congestion networks [2], [29], [30]. The former line of research focuses on the well known *mean-variance* model where each self-interested user selects a path that minimizes a linear combination of their expected travel time and standard deviation. While such an objective is desirable from a central planner's perspective, experimental studies suggest that individuals tend to employ simpler heuristics when faced with uncertainty [31]. Motivated by this, we adopt a multiplicative model of uncertainty similar to [24], [32], [33].

The literature on computing tolls for network users is driven by the need to *implement the optimum routing* by adjusting the toll amount, on each edge and was originally pioneered by Vickrey [29] and Walters [34]. It is possible to draw parallels between our model where users over-estimate costs and tolling; specifically, tolls leverage users' time-money trade-offs to alter their perceptions of the cost of each edge. While tolls can be (within reason) arbitrarily decided, the system planner has little influence over the level of uncertainty among the users. Bearing this in mind, we strive for a more subtle understanding of how equilibrium congestion depends on the level of uncertainty. Moreover, different than the existing literature, we also study the effect of cost under-estimation. It is also is worth noting that the classical works on network pricing have focused on edge-specific tolls applicable to all users on a link, yet recent works such as [35] have sought out more sophisticated congestion pricing techniques including trip-dependent and time-sensitive pricing.

While tolls can be viewed as monetary disincentives that alter the behavior of self-interested users in order to reduce congestion, there has been a growing interest in using *information as an incentive* to achieve the same effect. This line of work [16], [36] looks at principal-agent settings, where the principal can mis-represent or limit the information provided to agents. Similar to the present work, these papers reach the conclusion that when users' perceived costs do not align with the true costs, it is possible to achieve outcomes with lower social costs. That said, as noted previously, existing work only focuses on the problem of identifying the single mis-representation that optimizes the social cost for largely simple settings (e.g., parallel links). On the other hand, we seek to comprehensively characterize the landscape of user perceptions and how they impact congestion even for complex network structures.

We remark that our work is thematically similar to the recent paper on the *informational Braess paradox* [37] whose framework can be viewed as an extreme case of our model where the uncertainty parameter $r_\theta \to \infty$ on some edges. Our model is more continuous, as user attitudes are parameterized by a finite value of r_θ, which allows for a more descriptive depiction of the trade-offs faced by users who must balance travel time, congestion, and uncertainty. Finally, some recent works also study the effect of cost misperceptions in non-atomic routing games [30], [38], [39]. The present work partially extends a few of the results in these papers—e.g., from parallel links to more general structures [30], and from one user type to multiple types [38], [39]. Moreover, while [38], [39] focus on price of anarchy bounds, we also provide an instance-wise characterization (Theorems 1,3).

C. Organization

The rest of the paper is structured as follows. In Section II, we formally introduce our model followed by our main results in Sections III, IV, and V. Section VI presents our simulation results on urban transportation networks with parking and routing users who face different levels of uncertainty. Finally, we conclude with discussion in Section VII.

II. Model and Preliminaries

We consider a non-atomic, multi-commodity selfish routing game with multiple types of users. Specifically, we consider a network represented as $G = (V, \mathcal{E})$ where V is the set of nodes and $\mathcal{E}$ is the set of edges. For each edge $e \in \mathcal{E}$, we define a linear cost function

$$C_e(x_e) = a_e x_e + b_e, \tag{1}$$

where $x_e \geq 0$ is the total population (or flow) of users on that edge and $a_e, b_e \geq 0$. One can interpret $C_e(\cdot)$ as the true cost or expected congestion felt by the users on this edge. However, due to uncertainty, users may perceive the cost on each edge $e \in \mathcal{E}$ to be different from its true cost.

To capture that users may have different perceived uncertainties, we introduce the notion of *type*. Specifically, we consider a finite set of user types $\mathcal{T}$, where each type $\theta \in \mathcal{T}$ is uniquely defined by the following tuple $(s_\theta, t_\theta, \mu_\theta, r_\theta)$. The parameter $\mu_\theta > 0$ denotes the total population of users belonging to type θ such that each of these infinitesimal users seeks to route *some flow* from its source node $s_\theta \in V$ to the destination node $t_\theta \in V$. Moreover, given parameter $r_\theta > 0$, users of type θ perceive the cost of edge $e \in \mathcal{E}$ to be

$$\hat{C}^\theta_e(x_e) = r_\theta a_e x_e + b_e. \tag{2}$$

The uncertainty parameter r_θ denotes the personalized beliefs of the non-atomic population of type θ and when considered across types, captures the heterogeneity in preferences. As we articulate in Section II-B, this term can be viewed as a user's belief or preference stemming from a lack of information regarding the true costs or aversion to congestion or wait times. For illustration, consider an urban transportation network. Then b_e may represent the constant travel time on a link (in the absence of other vehicles) and $a_e x_e$, the congestion-dependent component of the travel time. A multiplicative uncertainty of $r_\theta > 1$ indicates that users of type θ adversely view costs arising due to congestion (e.g., waiting in traffic) when compared to other costs .

A path $p \in \mathcal{P}_\theta$ is a sequence of edges connecting s_θ to t_θ. Define $\mathcal{P}_\theta$ to be the set of all s_θ–t_θ paths in G. Let $x^\theta_p \in \mathbb{R}$ be the total flow routed by users of type θ on path $p \in \mathcal{P}_\theta$. We use the notation $\mathbf{x} = (x^\theta_p)_{\theta \in \mathcal{T}, p \in \mathcal{P}_\theta} \in \mathbb{R}^{|\mathcal{T}|\cdot|\mathcal{P}_\theta|}$ for a network flow and $\mathbf{x}^\theta = (x^\theta_p)_{p \in \mathcal{P}_\theta}$ to denote the network flow of type $\theta \in \mathcal{T}$. Then, for each type $\theta \in \mathcal{T}$, define the set of feasible flows to be

$$\mathcal{X}_\theta = \{\mathbf{x}^\theta |\ \textstyle\sum_{p \in \mathcal{P}_\theta} x^\theta_p = \mu_\theta,\ x^\theta_p \geq 0,\ \forall p \in \mathcal{P}_\theta\}. \tag{3}$$

The action space of users of type θ is $\mathcal{X}_\theta$—that is, users of type θ choose a feasible flow $\mathbf{x}^\theta \in \mathcal{X}_\theta$. Further, define the joint action space $\mathcal{X} = (\mathcal{X}_\theta)_{\theta \in \mathcal{T}}$—i.e. the space of feasible flows for all user types.

Path flows induce edge flows. Let $x^\theta_e \in \mathbb{R}$ be the flow on edge e due to users of type θ. The edge and path flow for users of type θ are related by

$$x^\theta_e = \textstyle\sum_{p \in \mathcal{P}_\theta, e \in p} x^\theta_p.$$

Define the total flow on edge e to be $x_e = \sum_{\theta \in \mathcal{T}} x^\theta_e$. Then, using this notation, we write the path cost in terms of edge flow. For any path p,

$$C_p(\mathbf{x}) = \textstyle\sum_{e \in p} C_e(x_e) = \sum_{e \in p}(a_e x_e + b_e). \tag{4}$$

Similarly, the perceived path costs are given by

$$\hat{C}^\theta_p(\mathbf{x}) = \textstyle\sum_{e \in p} \hat{C}^\theta_e(x_e) = \sum_{e \in p} r_\theta a_e x_e + b_e. \tag{5}$$

The following definition of a game instance $\mathcal{G}$ captures all of the relevant information about the multi-commodity routing game including the notion of type-based uncertainty we are interested in studying.

Definition 1 (Instance): An instance of the multi-commodity routing game is a tuple $\mathcal{G} = \{(V, \mathcal{E}), \mathcal{T}, \mathcal{X}, (\mathcal{P}_\theta)_{\theta \in \mathcal{T}}, (s_\theta, t_\theta, \mu_\theta, r_\theta)_{\theta \in \mathcal{T}}, (C_e)_{e \in \mathcal{E}}\}$, where

$(V, \mathcal{E})$ denotes the network, $(C_e)_{e\in\mathcal{E}}$ is the set of cost functions on each link, $\mathcal{T}$ is the set of types and each type $\theta \in \mathcal{T}$ is specified by the entities $\mathcal{P}_\theta, (s_\theta, t_\theta, \mu_\theta, r_\theta)_{\theta\in\mathcal{T}}$ that denotes the set of paths, origin-destination pair, total population mass and uncertainty level, respectively.

A. Nash Equilibrium Concept

We assume that the users in the system are self-interested with the goal of minimizing their individual cost. Therefore, the solution concept of interest in such a setting is a Nash equilibrium, where each user routes their flow on minimum cost paths with respect to their perceived cost functions and the actions of the other users.

Definition 2 (Nash Equilibrium): Given a game instance $\mathcal{G}$, a feasible flow $\mathbf{x} \in \mathcal{X}$ is said to be a *Nash equilibrium* if for every $\theta \in \mathcal{T}$, for all $p \in \mathcal{P}_\theta$ with positive flow, $x_p^\theta > 0$,

$$\hat{C}_p^\theta(\mathbf{x}) \le \hat{C}_{p'}^\theta(\mathbf{x}), \quad \forall\ p' \in \mathcal{P}_\theta \tag{6}$$

In the transportation literature, this solution concept is referred to as a *Wardrop equilibrium*. For the sake of consistency with the body of work pertaining to price of anarchy [20], we will continue using the term Nash equilibrium as the two concepts lead to completely equivalent solutions. For the rest of this work, we also will assume that all the flows considered are feasible. Finally, it is useful to point out that our model and solution allow for users of the same type to select different paths that connect the source and destination. Given a feasible flow $\mathbf{x} \in \mathcal{X}$, it is possible that $x_p^\theta, x_{p'}^\theta > 0$ for $p, p' \in \mathcal{P}_\theta$; however, if $\mathbf{x}$ is a Nash equilibrium then, it must be the case that $\hat{C}_p^\theta(\mathbf{x}) = \hat{C}_{p'}^\theta(\mathbf{x})$—i.e., users do not choose sub-optimal paths with respect to their perceived costs at equilibrium.

B. Model Discussion

Although users' perceived cost is deterministic given the type—i.e., $\hat{C}_e^\theta(x) = r_\theta a_e x + b_e$—this modeling choice allows us to capture decision-making in the presence of uncertainty in costs. There is extensive experimental evidence showing that users do not always minimize expected costs under uncertainty and instead adopt simpler heuristics [6], [40], [41]. Consider an example where the true cost on edge $e \in \mathcal{E}$ is $\beta a_e x + b_e$, and $\beta \sim \mathcal{F}$ is a random variable in the range $[\beta^{\min}, \beta^{\max}]$ with $E[\beta] = 1$. By selecting the parameter r_θ appropriately, we can model several user heuristics studied in the literature:

1) If the cost distribution ($\mathcal{F}$) is unknown, a risk-averse user (of type θ) typically pads their estimated cost on edge e with a parameter $\delta_e^\theta > 0$ that represents the 'margin of safety' [42]–[44]. For instance, the perceived cost is given by $a_e x + b_e + \delta_e^\theta$, where $\delta_e^\theta = (r_\theta - 1)a_e x$. Then a risk-seeking user may have $\delta_e < 0$ due to their optimism.
2) A body of evidence suggests that they prefer to optimize convenient heuristic functions as opposed to more involved probabilistic reasoning even if knowledge of the distribution is available [41].A well-studied example is the *Hurwicz criterion* [44], [45] that represents a compromise between the best-case and worst-case travel costs. Under this model, $\hat{C}_e^\theta(x) = \alpha^\theta(\beta^{\min} a_e x + b_e) + (1-\alpha^\theta)(\beta^{\max} a_e x + b_e)$, where $r_\theta = \alpha^\theta \beta^{\min} + (1-\alpha^\theta)\beta^{\max}$ and $0 \le \alpha^\theta \le 1$ is a parameter that denotes the level of caution.
3) In transportation, drivers exhibit *delay aversion* by preferring routes with lower congestion even under larger travel times [46], [47]. In our model, the factor r_θ on the a_e term represents this delay aversion.

Note that although our model only induces multiplicative uncertainty on the congestion-dependent component and not the b_e term, this is without loss of generality. If the users' perceived cost on edge e is $r_\theta^1 a_e x + r_\theta^2 b_e$, it can be transformed to an instance of *our routing game*, where $r_\theta = \frac{r_\theta^1}{r_\theta^2}$.

Furthermore, the equilibrium concept that we consider is consistent with previous studies of non-atomic routing games with heterogeneous cost functions users have heterogeneous perceptions of the costs [22], [39], [48]. It is interesting to consider how a Nash equilibrium is reached and the ties between our solution concept and the *Bayes Nash equilibrium* (BNE) [49]. Given that we study a non-atomic routing game with infinitely many users, in order to reason about the Nash equilibrium, it is sufficient to know the probability distribution on the users' types. Formally, suppose that each user's (private) type θ is drawn i.i.d from a publicly known distribution $\mathcal{F}_\mathcal{T}$—i.e., user has type θ with probability $\frac{\mu_\theta}{\sum_{\theta'} \mu_{\theta'}}$, where μ_θ is as given in Definition 1. Then, the equilibrium solution that is achieved when each infinitesimal user makes a routing decision with respect to its own parameter r_θ and belief $F_\mathcal{T}$ is equivalent to that given in Definition 2 *almost surely*. This is because the cost of each edge at equilibrium depends only on the aggregate flow on that edge.

Moreover, since the equilibrium in non-atomic games only depends on the distribution of user types, it is analogous to that a Bayes Nash equilibrium (see Section VII for a discussion of other models). Owing to this similarity, the BNE in network routing is more commonly studied in the context of *atomic* games with a finite number of agents [50], [51].

When users do not possess knowledge of others, our equilibrium notion occurs when the users are able to learn. In routing games such as ours which admit a potential function (see Proposition 1), a number of myopic user policies such as best-response [52] and no-regret learning [53], [54] are known to converge to equilibrium under iterative play. Under these policies, each user's route choice in a given round depends only the observed costs on various paths in prior rounds. We reiterate that although users' perceived costs could stem from information asymmetry, our setting still gives rise to a deterministic game since each user type routes their flow according to a known cost function.

C. Social Cost and Price of Anarchy

One of the central goals in this work is to compare the quality of the equilibrium solution in the presence of uncertainty to the socially optimal flow as, e.g., computed by a centralized planner with the goal of minimizing the aggregate cost in the system. Specifically, the social cost of a flow $\mathbf{x}$ is given by

$$\mathrm{C}(\mathbf{x}) = \textstyle\sum_{e\in\mathcal{E}} C_e(x_e)x_e. \tag{7}$$

This article has been accepted for publication in a future issue of this journal, but has not been fully edited. Content may change prior to final publication. Citation information: DOI 10.1109/TAC.2019.2962102, IEEE Transactions on Automatic Control

Note that the social cost is only measured with respect to the true congestion costs and thus does not reflect users' beliefs or uncertainties.

To capture inefficiencies, we leverage the well-studied notion of the *price of anarchy* which is the ratio of the social cost of the *worst-case Nash equilibrium* to that of the socially optimal solution. Specifically, we are interested in bounding price of anarchy as a function of the maximum and minimum uncertainty levels in the system, i.e., $r_{\max} = \max_\theta r_\theta$ and $r_{\min} = \min_\theta r_\theta$. Formally, given parameters $r_{\max}, r_{\min}$, we use $\mathcal{C}(r_{\max}, r_{\min})$ to refer to a set or class of instances (as per Definition 1), where the maximum uncertainty level of any type θ is $r_{\max}$ and the minimum uncertainty level is $r_{\min}$. Given an instance $\mathcal{G} \in \mathcal{C}(r_{\max}, r_{\min})$, we suppose that $\mathbf{x}^*$ is the flow that minimizes the social cost $\mathrm{C}(\mathbf{x})$ and that $\tilde{\mathbf{x}}$ is the Nash equilibrium for the given instance.

Definition 3 (Price of Anarchy): Given a class of instances $\mathcal{C}(r_{\max}, r_{\min})$, the *price of anarchy* for this class is

$$\max_{\mathcal{G}\in\mathcal{C}(r_{\max}, r_{\min})} \mathrm{C}(\tilde{\mathbf{x}})/\mathrm{C}(\mathbf{x}^*). \tag{8}$$

Since, we study a cost-minimizing game, the price of anarchy is always greater than or equal to one.

III. Main Results

To support the main theoretical results, we first show that our game is a *weighted potential game*. Routing games that fall into the general class of potential games have a number of nice properties in terms of existence, uniqueness, and computability [52]. In our case, the existence of a weighted potential function indicates that a Nash equilibrium always exists and moreover, best-response behavior by the users converges to such an equilibrium. General multi-commodity, selfish routing games with heterogeneous users, however, do not belong to the class of potential games unless certain assumptions on the edge cost structure are met [22].

The following proposition states that the game instances of the form we consider admit a (weighted) potential function and hence, there always exists a Nash equilibrium [52].

Proposition 1: A feasible flow $\mathbf{x}$ is a Nash equilibrium for a given instance $\mathcal{G}$ of a multi-commodity routing game with uncertainty vector $(r_\theta)_{\theta\in\mathcal{T}}$ if and only if it minimizes the following potential function:

$$\Phi_r(\mathbf{x}) = \sum_{e\in\mathcal{E}} \left(\tfrac{1}{2} a_e x_e^2 + b_e \sum_{\theta\in\mathcal{T}} \tfrac{1}{r_\theta} x_e^\theta\right) \tag{9}$$

Moreover, for any two minimizers $\mathbf{x}, \mathbf{x}'$, $C_e(x_e) = C_e(x'_e)$ for every edge $e \in \mathcal{E}$.

Note that although users perceive the multiplicative uncertainty r_θ on the a_e term (see Equation (2)), the parameter appears in the denominator of the b_e term in the potential function above. Conceptually, these have a similar effect: dividing Equation (6) by r_θ on both sides, one can obtain equivalent equilibrium conditions where the r_θ term is present in the denominator of the constant b_e.

Proof: By definition, a feasible flow $\mathbf{x} \in \mathcal{X}$ is a Nash equilibrium if the following condition is satisfied for all $\theta \in \mathcal{T}$ and for all $p, p' \in \mathcal{P}_\theta$ with $x_p^\theta > 0$:

$$\sum_{e\in p}(r_\theta a_e x_e + b_e) \le \sum_{e\in p'}(r_\theta a_e x_e + b_e).$$

Since $r_\theta > 0$, this is equivalent to $\sum_{e\in p}(a_e x_e + \frac{b_e}{r_\theta}) \le \sum_{e\in p'}(a_e x_e + \frac{b_e}{r_\theta})$. The remainder of proof trivially follows from standard arguments pertaining to the minimizer of a convex function. See [52] for more detail. ■

The second part of the proposition indicates that the equilibria are essentially unique as the cost on every edge is the same across solutions.

A. Effect of Uncertainty on Equilibrium Quality

Our first main result identifies a special case of the general multi-commodity game for which uncertainty helps improve equilibrium quality—i.e. decreases the social cost—whenever users over-estimate their costs by a small factor and vice-versa when they under-estimate costs. To show this result, we need the following technical lemma.

Lemma 1: Given an instance $\mathcal{G}$ of a multi-commodity selfish routing game with Nash equilibrium $\tilde{\mathbf{x}} = (\tilde{x}_e)_{e\in\mathcal{E}}$, we have that for any feasible flow $\mathbf{x}$,

$$\mathrm{C}(\tilde{\mathbf{x}}) - \mathrm{C}(\mathbf{x}) \le -\sum_{\theta\in\mathcal{T}}(\tfrac{2}{r_\theta} - 1)\sum_{e\in\mathcal{E}} b_e \Delta x_e^\theta, \tag{10}$$

where $\Delta x_e^\theta = \tilde{x}_e^\theta - x_e^\theta$.
The proof of the above lemma is provided in Appendix A.

Given an instance $\mathcal{G}$ of the multi-commodity routing game, we define $\mathcal{G}^1$ to be the corresponding game instance with no uncertainty—that is, $\mathcal{G}^1$ has the same graph, cost functions, and user types as $\mathcal{G}$, yet $r_\theta = 1$ for all $\theta \in \mathcal{T}$.

Theorem 1: Consider any given instance $\mathcal{G}$ of the multi-commodity routing game with Nash equilibrium $\tilde{\mathbf{x}}$ and corresponding game instance $\mathcal{G}^1$, having no uncertainty, with Nash equilibrium is $\mathbf{x}^1$. Suppose $r_\theta = r$ for all $\theta \in \mathcal{T}$. Then, $\mathrm{C}(\tilde{\mathbf{x}}) \le \mathrm{C}(\mathbf{x}^1)$ if $1 \le r \le 2$ and $\mathrm{C}(\tilde{\mathbf{x}}) \ge \mathrm{C}(\mathbf{x}^1)$ if $0 \le r \le 1$.

Remark: What happens when the users are highly cautious, i.e., $r_\theta > 2$ for all θ? Due to the presence of a few negative examples where the social cost increases in the presence of uncertainty, we cannot conclusively state that uncertainty helps or hurts for all instances. However, these negative instances appear to be isolated—both our price of anarchy result (Theorem 2) and our experiments (Section VI) validate our claim that caution in the face of uncertainty helps the users by lowering equilibrium social costs even when $r_\theta > 2$—i.e., uncertainty is favorable when the users are very cautious. It is however, interesting to note that although under-estimation always leads to a worse equilibrium, over-estimation may lead to better or worse equilibria.

Proof of Theorem 1: Let $\Phi_r(\mathbf{x})$ denote the potential function for the instance $\mathcal{G}$ and $\Phi_1(\mathbf{x})$ denote the potential function for $\mathcal{G}^1$ where Φ_1 is given in (9) with $r = 1$. By definition of the potential function, we know that $\Phi_r(\tilde{\mathbf{x}}) - \Phi_r(\mathbf{x}^1) \le 0$ and $\Phi_1(\mathbf{x}^1) - \Phi_1(\tilde{\mathbf{x}}) \le 0$. Expanding them gives

$$\sum_{e\in\mathcal{E}} \left(\tfrac{a_e \tilde{x}_e^2}{2} + b_e \sum_{\theta\in\mathcal{T}} \tfrac{\tilde{x}_e^\theta}{r} - \tfrac{a_e (x_e^1)^2}{2} - b_e \sum_{\theta\in\mathcal{T}} \tfrac{x_e^{1,\theta}}{r}\right) \le 0,$$

and

$$\sum_{e\in\mathcal{E}} \left(\tfrac{a_e (x_e^1)^2}{2} + b_e \sum_{\theta\in\mathcal{T}} x_e^{1,\theta} - \tfrac{a_e \tilde{x}_e^2}{2} - b_e \sum_{\theta\in\mathcal{T}} \tilde{x}_e^\theta\right) \le 0,$$

where $x_e^{1,\theta}$ denotes the total flow on edge e by users of type θ in the solution $\mathbf{x}^1$. Let us define $\Delta x_e^\theta = \tilde{x}_e^\theta - x_e^{1,\theta}$, and

$\Delta x_e = \tilde{x}_e - x_e^1$. By summing the above inequalities, the a_e terms cancel out and

$$\sum_{e\in\mathcal{E}} b_e \sum_{\theta\in\mathcal{T}} (\frac{\Delta x_e^\theta}{r} - \Delta x_e^\theta) = \sum_{e\in\mathcal{E}} b_e \sum_{\theta\in\mathcal{T}} (\frac{1}{r} - 1)\Delta x_e^\theta \leq 0.$$

Using the fact that $\sum_\theta \Delta x_e^\theta = \Delta x_e$ for all edges, we get that

$$(\tfrac{1}{r} - 1)\sum_{e\in\mathcal{E}} b_e \Delta x_e \leq 0. \quad (11)$$

Hence, $\sum_{e\in\mathcal{E}} b_e \Delta x_e \geq 0$ when $r > 1$ and $\sum_{e\in\mathcal{E}} b_e \Delta x_e \leq 0$ when $r < 1$.

We finish the proof by considering two separate cases: (case 1) $1 \leq r \leq 2$ and (case 2) $r < 1$. In case 1, applying Lemma 1 to the instance $\mathcal{G}$ with $\mathbf{x} = \mathbf{x}^1$, we obtain that

$$\mathrm{C}(\tilde{\mathbf{x}}) - \mathrm{C}(\mathbf{x}^1) \leq (1 - \tfrac{2}{r})\sum_{e\in\mathcal{E}} b_e \Delta x_e. \quad (12)$$

We claim that when $r \in (1, 2]$, the right-hand side of (12) is lesser than or equal to zero. This is not particularly hard to deduce owing to the fact that $(1 - \frac{2}{r}) < 0$ in the given range and that $\sum_{e\in\mathcal{E}} b_e \Delta x_e \geq 0$ as deduced from (11). Therefore, $\mathrm{C}(\tilde{\mathbf{x}}) - \mathrm{C}(\mathbf{x}^1) \leq 0$, which proves the claim that uncertainty with a limited amount of caution helps lower equilibrium costs.

Now, let us consider (case 2) where $r < 1$. Applying Lemma 1 to the instance $\mathcal{G}^1$ with $\mathbf{x} = \tilde{\mathbf{x}}$, and using the fact that $\Delta x_e = \tilde{x}_e - x_e^1$, we have that $\mathrm{C}(\mathbf{x}^1) - \mathrm{C}(\tilde{\mathbf{x}}) \leq (\frac{2}{r} - 1)\sum_{e\in\mathcal{E}} b_e \Delta x_e$. Once again when $r < 1$, we know that $\frac{2}{r} - 1 > 0$ and (6) gives $\sum_{e\in\mathcal{E}} b_e \Delta x_e \leq 0$ in the given range. ∎

The following corollary identifies a specific level of uncertainty at which the equilibrium solution is actually optimal.

Corollary 1: Given an instance $\mathcal{G}$ of the multi-commodity routing game, let $\tilde{\mathbf{x}}$ denote its Nash equilibrium and $\mathbf{x}^*$ denote the socially optimal flow. If $r_\theta = 2$ for all $\theta \in \mathcal{T}$, then $\mathrm{C}(\tilde{\mathbf{x}}) = \mathrm{C}(\mathbf{x}^*)$—i.e. the equilibrium is socially optimal.

Proof: Suppose $r_\theta = r = 2$. Applying Lemma 1, we have that $\mathrm{C}(\tilde{\mathbf{x}}) - \mathrm{C}(\mathbf{x}^*) \leq -(\frac{2}{r} - 1)\sum_{e\in\mathcal{E}} b_e - \Delta x_e = 0$. ∎

Corollary 1 has a natural interpretation in terms of the theory of *computing optimal tolls*. When $r_\theta = 2$ for all $\theta \in \mathcal{T}$, each user's cost perceived function on a given edge $e \in \mathcal{E}$ becomes $\hat{C}_e^\theta(x_e) = (a_e x_e + b_e) + (a_e x_e) = C_e(x_e) + a_e x_e$, i.e., the true cost plus 'an additional term'. If we view the term $a_e x_e$ as a congestion-dependent toll paid by the user for routing flow on link e, then Corollary 1 is equivalent to a classical result in tolling theory [20], [29]. In particular, this result states that *if the toll on each link is set to the marginal cost of adding an extra user, then the ensuing Nash equilibria minimize the social cost*. Against this backdrop, Corollary 1 establishes an intuitive connection between cost mis-perceptions under uncertainty and monetary tolls. That said, much of the work in the tolling literature focuses on optimal tolls and thus, Theorem 1 can be viewed as a novel result on the properties of price-bounded and even negative tolls.

B. Price of Anarchy Under Uncertainty

In Theorem 1, we showed that the equilibrium cost under uncertainty decreases (resp. increases) when users are mildly cautious (resp. not cautious) and all user types have the same level of uncertainty. This naturally raises the question of quantifying the improvement (or degradation) in equilibrium quality and whether uncertainty helps when the uncertainty parameter can differ between user types. In the following theorem, we address both of these questions by providing price of anarchy bounds as a function of the maximum uncertainty in the system and γ, which is the ratio between the minimum and maximum uncertainty among user types.

Theorem 2: (Price of Anarchy) For any multi-commodity routing game $\mathcal{G}$, the ratio between the social cost of the Nash equilibrium to that of the socially optimal solution is at most

$$PoA(\mathcal{G}) \leq 4/(4\gamma r_{\max} - r_{\max}^2), \quad (13)$$

where $r_{\max} = \max_\theta r_\theta$, and $\gamma = \frac{\min_\theta r_\theta}{\max_\theta r_\theta}$ if $r_{\max} < 4\gamma$.

Discussion: The price of anarchy bound in (13) is tight when $\gamma = 1$ and $r \in [0, 2]$. For any given $r \leq 2$, this can be confirmed by constructing a Pigouvian network similar to the one in Fig. 3 where the costs are $C_{e_1}(x) = x$ and $C_{e_2}(x) = r$. Further, we note that our result generalizes the price of anarchy bounds in [38], [39] towards instances with heterogeneous uncertainties ($\gamma < 1$). For example, Corollary 6.8 in [38] gives the same expression as (13) for instances where $\gamma = 1$ and $r \in [1, 2]$. Similarly, we note that for the special case when $\gamma = 1$ and $r > 2$, [39] provides a tighter bound for the price of anarchy compared to our result in Theorem 2.

Although the price of anarchy expression in (13) is not tight when $\gamma < 1$, it still serves as a useful upper bound to understand the conditions under which there is a strict improvement in equilibrium quality compared to the no-uncertainty scenario. For example, the price of anarchy in (13) is plotted in Fig. 1 as a function of $r_{\max}$ for three different values of γ. Based on these plots, one can gather that uncertainty lowers equilibrium cost when users over-estimate their latencies by a small margin. This is true for $r \in [1, 3]$ when $\gamma = 1$, and for smaller ranges when $\gamma < 1$. Note that in all these cases, we compare to the price of anarchy of $\frac{4}{3}$ for multi-commodity routing games with linear cost and no uncertainty [17]. Another interesting observation here is that in contrast to Theorem 1, we observe an improvement in equilibrium quality for a small range beyond $r = 2$, e.g., when $\gamma = 1$, this occurs up to $r = 3$. However, we know from [39] when the users are over-pessimistic, and $r_{\max}$ is large, the equilibrium quality degrades.

The price of anarchy bounds for $\gamma < 1$ are also quite revealing. As with the $\gamma = 1$ case, we notice that equilibrium cost can improve even when $r > 2$—albeit for a small range—e.g., for $r_{\max} \in [1.32, 2.28]$ when $\gamma = 0.9$. On the other hand, uncertainty can hurt equilibrium even if $r_{\max} > 1$; this occurs if $r_{\min} = \min_{\theta\in\mathcal{T}} r_\theta < 1$ and a large fraction of users have $r_\theta = r_{\min}$. One can construct simple two network examples to verify this lower bound. Finally, the price of anarchy result leads to a surprising observation: as long as $r_{\max} > 1$ and γ is not too small, for any given instance $\mathcal{G}$ of the multi-commodity routing game, either the equilibrium quality is already good or uncertainty helps lower congestion by a significant amount. Therefore, uncertainty rarely hurts the quality of the equilibrium and often helps.

In order to prove Theorem 2, we need the following technical lemma whose proof is provided in Appendix B.

This article has been accepted for publication in a future issue of this journal, but has not been fully edited. Content may change prior to final publication. Citation information: DOI 10.1109/TAC.2019.2962102, IEEE Transactions on Automatic Control

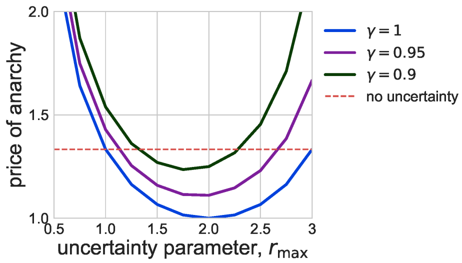

Fig. 1. Price of anarchy as a function of $r_{\max} = \max_\theta r_\theta$ for three different values of $\gamma = \frac{\min_\theta r_\theta}{\max_\theta r_\theta}$. In general, when $r_{\max} > 1$, the price of anarchy under uncertainty is smaller than that without uncertainty and vice-versa for $r_{\max} < 1$. Yet, too much caution or large asymmetries in the uncertainty level across user types can also lead to poor equilibrium quality.

Lemma 2: For any two non-negative vectors of equal length, $(x_1, x_2, \ldots, x_n)$ and $(x'_1, x'_2, \ldots, x'_n)$, let $x = \sum_{i=1}^n x_i$ and $x' = \sum_{i=1}^n x'_i$. Moreover, let $\mathbf{r} = (r_1, r_2, \ldots, r_n)$ be another vector of length n whose entries are strictly positive. Then, if we let $r_* = \min_i r_i$ and $r^* = \max_i r_i$, for any given function $f(y) = ay + b$ with $a, b \geq 0$, we have that

$$\frac{f(x)x}{f(x')x' + \sum_{i=1}^n (x_i - x'_i) f(r_i x)} \leq 4\left(4r_* - (r^*)^2\right)^{-1}, \tag{14}$$

Proof of Theorem 2: Consider some instance $\mathcal{G}$ of the multi-commodity routing game with equilibrium $\tilde{\mathbf{x}}$. Then, adopting the variational inequality conditions for a Nash equilibrium [14], [17], for any other solution $\mathbf{x}$, and every user type $\theta \in \mathcal{T}$, $\sum_{e\in\mathcal{E}} (r_\theta a_e \tilde{x}_e + b_e)\left(\tilde{x}_e^\theta - x_e^\theta\right) = \sum_{e\in\mathcal{E}} C_e(r_\theta \tilde{x}_e)(\tilde{x}_e^\theta - x_e^\theta) \leq 0$. Fix any edge $e \in \mathcal{E}$. Let $(x_1, \ldots, x_n) = (\tilde{x}_e^\theta)_{\theta\in\mathcal{T}}$, $x = \sum_{\theta\in\mathcal{T}} \tilde{x}_e^\theta$, $(x'_1, \ldots, x'_n) = (x_e^{*\theta})_{\theta\in\mathcal{T}}$, and $x' = \sum_{\theta\in\mathcal{T}} x_e^{*\theta}$. Applying Lemma 2 with $f(y) = C_e(y)$ and $\mathbf{r} = (r_\theta)_{\theta\in\mathcal{T}}$, we have that $C_e(\tilde{x}_e)\tilde{x}_e \leq \zeta(C_e(x_e^*)x_e^* + \sum_{\theta\in\mathcal{T}}(\tilde{x}_e^\theta - x_e^{*\theta})C_e(x_e r_\theta))$ where $\zeta = 4/(4r_{\min} - r_{\max}^2) = 4/(4\gamma r_{\max} - r_{\max}^2)$, and where $r_{\min} = \min_{\theta\in\mathcal{T}} r_\theta$ and $r_{\max} = \max_{\theta\in\mathcal{T}} r_\theta$.

We claim that

$$\begin{aligned}\sum_{e\in\mathcal{E}} C_e(\tilde{x}_e)\tilde{x}_e &\leq \zeta\big(\sum_{e\in\mathcal{E}}\sum_{\theta\in\mathcal{T}}(\tilde{x}_e^\theta - x_e^{*\theta})\left(r_\theta a_e \tilde{x}_e + b_e\right) \\ &\quad + \sum_{e\in\mathcal{E}} C_e(x_e^*)x_e^*\big) \\ &\leq \zeta\sum_{e\in\mathcal{E}} C_e(x_e^*)x_e^*. \end{aligned} \tag{15}$$

Because for any type $\theta \in \mathcal{T}$, $\sum_{e\in\mathcal{E}} (r_\theta a_e \tilde{x}_e + b_e)\left(\tilde{x}_e^\theta - x_e^\theta\right) \leq 0$. Since the price of anarchy is defined as $\frac{\sum_{e\in\mathcal{E}} C_e(\tilde{x}_e)\tilde{x}_e}{\sum_{e\in E} C_e(x_e^*)x_e^*}$, a worst case bound of ζ follows from (15), giving us the theorem statement. ■

IV. The Effect of Heterogeneity on Congestion

In this section, we consider a more nuanced setting where different users have different levels of uncertainty. Our goal is to understand the effect of this *heterogeneity in uncertainty* on the equilibrium congestion by studying the following two questions: (i) how do the routing choices adopted by the uncertain users impact the cost of the users without uncertainty, and (ii) when is the equilibrium social cost of a system with heterogeneous uncertainties smaller than the social cost incurred when all users have no uncertainty?

A. Notation and Graph Topologies

To isolate the effect of heterogeneity on equilibrium congestion, we study a selfish routing game on an undirected network. There are two user types, $\mathcal{T} = \{\theta_1, \theta_2\}$. Users belonging to both these types seek to route their flow between source node s and destination node t. The uncertainty levels for the two user types are specified as $r_{\theta_1} = 1$ and $r_{\theta_2} = r > 0$, and so, users of type θ_1 are without uncertainty. We refer to this as the *two-commodity game with and without uncertainty*. We slightly abuse notation and use $\mathcal{G} = \{G, \mathcal{T}, (s,t), (\mu_{\theta_1}, \mu_{\theta_2}), (1, r), (C_e)_{e\in\mathcal{E}}\}$ to refer to a game instance, and $\mathcal{P}$ to represent the set of s-t paths in G.

Unlike the previous sections, where we made no assumptions on the graph structure, network topology matters here. Specifically, we will consider two well-studied topologies: *series-parallel* and *linearly independent* graphs:

Definition 4: (*Series-Parallel* [18]) An undirected graph G with a single source s and destination t is said to be a series-parallel graph if no two s-t paths pass through an edge in opposite directions.

There are a number of other equivalent definitions for this class of graphs; e.g., a graph is said to be series-parallel if it does not contain an embedded *Wheatstone network* [18]. Series-parallel graphs are an extremely well-studied topology that naturally arise in a number of applications pertaining to network routing. We refer the reader to [18] for more details.

Definition 5: (*Linearly Independent* [18]) An undirected graph G with a single source s and destination t is said to be linearly independent if every s-t path contains at least one edge that does not belong to any other s-t path.

Our final definition involves a simple extension of the above topology to include linearly independent graphs connected serially. Formally, a graph $G = (V, \mathcal{E})$ is said to consist of two sub-graphs $G_1 = (V_1, \mathcal{E}_1)$ with source-destination pair (s_1, t_1) and $G_2 = (V_2, \mathcal{E}_2)$ with source-destination pair (s_2, t_2) connected in serial if $V = V_1 \cup V_1$ with $t_1 = s_2$ and $\mathcal{E} = \mathcal{E}_1 \cup \mathcal{E}_2$.

Definition 6: (*Serially Linearly Independent (SLI)*) An undirected graph G with a single source s and destination t is said to belong to the serially linearly independent class if (i) G is linearly independent or (ii) G consists of two linearly independent graphs connected in serial.

This extended topology was first introduced in [37]. Every linearly independent graph belongs to the serially linearly independent class, and every serially linearly independent graph belongs to the series-parallel class [18].

B. Impact of Uncertain Users on Users without Uncertainty

We begin by studying what happens to the congestion cost faced by the users without uncertainty as the uncertainty level increases for the other users. This question is of considerable interest in a number of settings. For example, in urban transportation networks, the uncertainty about where a driver can find available street parking can often cascade into increased congestion for other drivers leading to a detrimental effect on the overall congestion cost [55], [56].

This article has been accepted for publication in a future issue of this journal, but has not been fully edited. Content may change prior to final publication. Citation information: DOI 10.1109/TAC.2019.2962102, IEEE Transactions on Automatic Control

The following proposition shows a somewhat surprising result. As long as the network topology is series-parallel, the aggregate cost felt by users of type θ_1 (users without any uncertainty) always reduces when users of θ_2 are uncertain about the costs. In other words, the behavior under uncertainty by one type of users always decreases the congestion costs of other types of users who do not face any uncertainty.

Proposition 2: Given an instance $\mathcal{G}$ of the two-commodity game with and without uncertainty such that the graph G is series-parallel, let $\mathcal{G}^1$ denote a modified version of this instance with no uncertainty (i.e. $r_{\theta_1} = r_{\theta_2} = 1$). Let $\tilde{\mathbf{x}}$ and $\mathbf{x}^1$ denote the Nash equilibrium for the two instances, respectively. Then, $\mathrm{C}^{\theta_1}(\tilde{\mathbf{x}}) \leq \mathrm{C}^{\theta_1}(\mathbf{x}^1)$, where $\mathrm{C}^{\theta_1}(\mathbf{x}) = \sum_{e\in\mathcal{E}} C_e(x_e)x_e^{\theta_1}$ is the aggregate cost of users of type θ_1.

Proof: It is well-known that [18, Lemma 3] for a series-parallel graph G and any two feasible flows $\mathbf{x}, \mathbf{x}'$, there exists a s–t path p with $x_p > 0$, such that for every edge $e \in p$, $x'_e \leq x_e$. Now, consider flows $\mathbf{x}^1$ and $\tilde{\mathbf{x}}$. Applying the previous property, we get that, there exists a path p with $x^1_p > 0$ such that for all $e \in p$, $\tilde{x}_e \leq x^1_e$.

We now bound both $\mathrm{C}^{\theta_1}(\tilde{\mathbf{x}})$ and $\mathrm{C}^{\theta_1}(\mathbf{x}^1)$ in terms of the cost of the path p. Specifically, note that in the solution $\mathbf{x^1}$, the path p has non-zero flow on it so that $\mathrm{C}^{\theta_1}(\mathbf{x}^1) = \mu_{\theta_1}\sum_{e\in p} C_e(x^1_e)$. However, in the solution $\tilde{\mathbf{x}}$, we know that every user of type θ_1 is using a minimum cost path with respect to the true costs and therefore, the cost of any path used by type θ_1 is at least that of the path p. Formally, $\mathrm{C}^{\theta_1}(\tilde{\mathbf{x}}) \leq \mu_{\theta_1}\sum_{e\in p} C_e(\tilde{x}_e) \leq \mu_{\theta_1}\sum_{e\in p} C_e(x^1_e)$. The final inequality follows from the monotonicity of the cost functions and the fact that $\tilde{x}_e \leq x^1_e$ for all $e \in p$. Therefore, we conclude that $\mathrm{C}^{\theta_1}(\tilde{\mathbf{x}}) \leq \mathrm{C}^{\theta_1}(\mathbf{x}^1)$. ■

C. Characterization of Instances where Heterogeneity Helps

We now consider the impact of heterogeneity on the system performance as a whole and present a simple characterization based on the network topology and the level of uncertainty, where the presence of uncertainty (among a fraction of the user population) results in a decrease in the equilibrium social cost. Specifically, we show that for SLI networks, as long as the uncertainty level of users belonging to type θ_2 is at most two (i.e., $1 \leq r = r_{\theta_2} \leq 2$), the social cost of the equilibrium solution is always smaller than or equal to that of the equilibrium when there is no uncertainty.

Before showing our theorem, we state the following technical lemma whose proof is deferred to Appendix C.

Lemma 3: Given any instance $\mathcal{G}$ of the two-commodity routing game with and without uncertainty where the graph G is linearly independent, let $\mathbf{y}$ and $\mathbf{y}^1$ denote the Nash equilibria of instances $\mathcal{G}$ and $\mathcal{G}^1$ respectively. Then $y^{\theta_1}_p \leq y^1_p \quad \forall p \in \mathcal{P}$. Informally, the above lemma states that given equilibrium flows $\mathbf{y}$ and $\mathbf{y}^1$ for any arbitrary instance $\mathcal{G}$ and its uncertainty-free variant $\mathcal{G}^1$, the equilibrium solutions must satisfy the property that for any path p, the flow on this path in the absence of uncertainty (instance $\mathcal{G}^1$) must be greater than or equal to its magnitude due to the uncertainty-free users in $\mathcal{G}$.

Theorem 3: Consider any given instance $\mathcal{G}$ of the two-commodity game with and without uncertainty. Let $\tilde{\mathbf{x}}$ denote the Nash equilibrium of this game and the corresponding game instance $\mathcal{G}^1$, having no uncertainty has Nash equilibrium $\mathbf{x}^1$. Then, as long as G belongs to the serially linearly independent class and $1 \leq r \leq 2$, $\mathrm{C}(\tilde{\mathbf{x}}) \leq \mathrm{C}(\mathbf{x}^1)$. Moreover, there exist instances of the two-commodity routing game with and without uncertainty on series-parallel networks where the social cost of the equilibrium without uncertainty is strictly smaller than the cost with uncertainty $r \in [1, 2]$.

Proof of Theorem 3: Each SLI network can be broken down into a sequence of linearly independent networks connected in series. Applying Definition 6 recursively, we get a sequence of linearly independent sub-graphs $G(1) = (V(1), \mathcal{E}(1)), G(2) = (V(2), \mathcal{E}(2)), \ldots, G(\ell) = (V(\ell), \mathcal{E}(\ell))$ with source-destination pairs $(t_0, t_1), (t_1, t_2), \ldots, (t_{\ell-1}, t_\ell)$ respectively (note that $t_0 = s, t_\ell = t$), that are connected in series—i.e., G_1 is connected in series with G_2 such that the destination t_1 for G_1 acts as the origin for G_2. By definition, the set of edges in these subgraphs are mutually disjoint.

Secondly, given the equilibrium flow $\tilde{\mathbf{x}}$ on G for instance $\mathcal{G}$, we can divide this flow into components $(\tilde{\mathbf{x}}(1), \tilde{\mathbf{x}}(2), \ldots, \tilde{\mathbf{x}}(\ell))$ such that for every $1 \leq i \leq \ell$, $\tilde{\mathbf{x}}(i)$ is the sub-flow of $\tilde{\mathbf{x}}$ on the graph $G(i)$, and for every $e \in \mathcal{E}(i)$, $\tilde{x}(i)_e = \tilde{x}_e$. Finally, it is not hard to see that $\tilde{\mathbf{x}}(i)$ must be an equilibrium of the sub-instance of $\mathcal{G}$ restricted to $G(i)$.

Given this *decomposition*, we apply Lemma 3 to each $G(i)$. Consider any index i: since the graph $G(i)$ is linearly independent, we can apply Lemma 3 and get that for any t_{i-1}-t_i path p in $G(i)$, $\tilde{x}(i)^{\theta_1}_p \leq x^1(i)_p$.

Suppose that $\mathcal{P}(i)$ denotes the set of t_{i-1}-t_i paths in $G(i)$. Then by standard convex analysis,

$$\begin{aligned} C(\tilde{\mathbf{x}}) &= \textstyle\sum_{e\in\mathcal{E}} C_e(\tilde{x}_e)\tilde{x}_e = \sum_{e\in\mathcal{E}}(a_e\tilde{x}_e + b_e)\tilde{x}_e \\ &\leq \textstyle\sum_{e\in\mathcal{E}}\left(C_e(x^1_e)x^1_e + (2a_e\tilde{x}_e + b_e)(\tilde{x}_e - x^1_e)\right) \\ &\leq C(\mathbf{x}^1) + \textstyle\sum_{e\in\mathcal{E}}(2a_e\tilde{x}_e + b_e)(\tilde{x}_e - x^1_e) \end{aligned} \tag{16}$$

so that, using the above decomposition,

$$\begin{aligned} C(\tilde{\mathbf{x}}) &\leq C(\mathbf{x}^1) + \textstyle\sum_{i=1}^{\ell}\sum_{e\in\mathcal{E}(i)}(2a_e\tilde{x}(i)_e + b_e) \\ &\qquad \cdot (\tilde{x}(i)_e - x^1(i)_e) \\ &= C(\mathbf{x}^1) + \textstyle\sum_{i=1}^{\ell}\sum_{p\in\mathcal{P}(i)}\left(\sum_{e\in p}(2a_e\tilde{x}(i)_e + b_e)\right) \\ &\qquad \cdot (\tilde{x}(i)_p - x^1(i)_p) \end{aligned}$$

To complete the proof we show $\sum_{p\in\mathcal{P}(i)}(\sum_{e\in p}(2a_e\tilde{x}(i)_e + b_e))(\tilde{x}(i)_p - x^1(i)_p) \leq 0$ for all $1 \leq i \leq \ell$. Indeed, fix an arbitrary index i and consider the corresponding graph $G(i)$ and flows $\tilde{\mathbf{x}}(i)$ and $\mathbf{x}^1(i)$. Recall that $\mathbf{x}^1(i)$ is the equilibrium solution in the absence of uncertainty and therefore, minimizes the following potential function $\Phi_{\mathbf{1}}(\mathbf{x}) = \sum_{e\in\mathcal{E}(i)}(\frac{1}{2}a_e(x_e)^2 + b_e x_e)$. Hence, by convexity

$$\begin{aligned} 0 &\leq \Phi_{\mathbf{1}}(\tilde{\mathbf{x}}(i)) - \Phi_{\mathbf{1}}(\mathbf{x}^1(i)) \\ &= \textstyle\sum_{e\in\mathcal{E}(i)}\left(\frac{a_e}{2}(\tilde{x}(i)_e)^2 + b_e\tilde{x}(i)_e - \frac{a_e}{2}(x^1(i)_e)^2 - b_e x^1(i)_e\right) \\ &\leq \textstyle\sum_{e\in\mathcal{E}(i)}(a_e\tilde{x}(i)_e + b_e)(\tilde{x}(i)_e - x^1(i)_e). \end{aligned} \tag{17}$$

Next, we claim that $\sum_{e\in\mathcal{E}(i)}(ra_e\tilde{x}(i)_e + b_e)(\tilde{x}(i)_e - x^1(i)_e) \leq 0$. Indeed, for convenience define $D_r = \sum_{e\in\mathcal{E}(i)}(ra_e\tilde{x}(i)_e +$

$b_e)(\tilde{x}(i)_e - x^1(i)_e)$ so that

$$\begin{aligned} D_r &= \textstyle\sum_{p\in\mathcal{P}(i)} \sum_{e\in p}(ra_e\tilde{x}(i)_e + b_e)(\tilde{x}(i)_p - x^1(i)_p) \\ &= \textstyle\sum_{p\in\mathcal{P}(i)^-} \sum_{e\in p}(ra_e\tilde{x}(i)_e + b_e)(\Delta x_p) \\ &\quad + \textstyle\sum_{p\in\mathcal{P}(i)^+} \sum_{e\in p}(ra_e\tilde{x}(i)_e + b_e)(\Delta x_p), \end{aligned} \tag{18}$$

where $\Delta x_p = \tilde{x}(i)_p - x^1(i)_p$ and $\mathcal{P}(i)^-, \mathcal{P}(i)^+$ refer to the set of the paths where $\Delta x_p \le 0$ and $\Delta x_p > 0$, respectively. From Lemma 3, we know that $\tilde{x}(i)_p^{\theta_1} \le x^1(i)_p$ for all $p \in \mathcal{P}(i)$. Therefore, if $\Delta x_p = \tilde{x}(i)_p^{\theta_1} + \tilde{x}(i)_p^{\theta_2} - x^1(i)_p > 0$ for any $p \in \mathcal{P}(i)$, it must be the case that $\tilde{x}(i)_p^{\theta_2} > 0$ for that path. In other words, we have that for every $p \in \mathcal{P}(i)^+$, $\tilde{x}(i)_p^{\theta_2} > 0$. Recall that $\tilde{x}(i)_p^{\theta_2}$ denotes the total flow due to the users with uncertainty on path p.

We make the following two observations that help us simplify (18). First, by flow conservation, we know that $\sum_{p\in\mathcal{P}(i)} \Delta x_p = 0$. Second, for every $p \in \mathcal{P}(i)^+$, we have that $\tilde{x}(i)_p^{\theta_2} > 0$, and so p is a min-cost path for the users with uncertainty in the flow $\tilde{\mathbf{x}}(i)$. This in turn implies that for any $p' \in \mathcal{P}(i)$ and $p \in \mathcal{P}(i)^+$, the following inequality is valid $\sum_{e\in p}(ra_e\tilde{x}(i)_e + b_e) \le \sum_{e\in p'}(ra_e\tilde{x}(i)_e + b_e)$. Define $c_p^r := \min_{p\in\mathcal{P}(i)} \sum_{e\in p}(ra_e\tilde{x}(i)_e + b_e)$. Then, using the fact $\Delta x_p \le 0$ for all $p \in \mathcal{P}(i)^-$ and that $\sum_{e\in p}(ra_e\tilde{x}(i)_e + b_e) \ge c_p^r$ for all $p \in \mathcal{P}(i)^-$, (18) can be rewritten as

$$\begin{aligned} D_r &= \textstyle\sum_{p\in\mathcal{P}(i)^-} \sum_{e\in p}(ra_e\tilde{x}(i)_e + b_e)(\Delta x_p) \\ &\quad + \textstyle\sum_{p\in\mathcal{P}(i)^+} \sum_{e\in p} c_p^r(\Delta x_p) \\ &\le \textstyle\sum_{p\in\mathcal{P}(i)^-} \sum_{e\in p} c_p^r(\Delta x_p) + \sum_{p\in\mathcal{P}(i)^+} \sum_{e\in p} c_p^r(\Delta x_p) \\ &= c_p^r \textstyle\sum_{p\in\mathcal{P}(i)} \Delta x_p = 0 \end{aligned} \tag{19}$$

so that

$$\textstyle\sum_{e\in\mathcal{E}(i)}(ra_e\tilde{x}(i)_e + b_e)(\tilde{x}(i)_e - x^1(i)_e) \le 0. \tag{20}$$

Summing this inequality with (17), we see that the b_e terms cancel out, leaving us with $(r-1)\sum_{e\in E(i)} a_e\tilde{x}(i)_e(\tilde{x}(i)_e - x^1(i)_e) \le 0$ which implies that $\sum_{e\in\mathcal{E}(i)} a_e\tilde{x}(i)_e(\tilde{x}(i)_e - x^1(i)_e) \le 0$ since $r > 1$. This, in turn, implies that

$$(2-r)\textstyle\sum_{e\in\mathcal{E}(i)} a_e\tilde{x}(i)_e(\tilde{x}(i)_e - x^1(i)_e) \le 0 \tag{21}$$

since $r \le 2$. Adding (21) to (20), we get that $\sum_{e\in E(i)}(2a_e\tilde{x}(i)_e + b_e)(\tilde{x}(i)_e - x^1(i)_e) \le 0$, which in combination with (16) yields the first part of the theorem. The negative example for networks that are series-parallel but not SLI is provided in Section IV-D1. ■

Although we only provide an instance-wise characterization for routing games with two types of users, we conjecture that uncertainty can help lower equilibrium cost with three or more types as long as $r \in [1, 2]$ for all types and the network is SLI. It may be possible to leverage the proof of Theorem 3 in an inductive fashion by gradually adding more types to prove such a result. More precisely, consider an arbitrary instance where the types are numbered such that $r_{\theta_1} < r_{\theta_2} \le \ldots \le r_{\theta_{|\mathcal{T}|}}$. Proceeding inductively and supposing that the claim holds for all instances with $j-1$ types, one could then add the j–th uncertainty level r_{θ_j}. In this case, comparing the equilibrium cost for the instance with $j-1$ and j types could be analogous to the proof of Theorem 3 with only two types since the population of users having all uncertainties except $r_{\theta_{j-1}}$ and r_{θ_j} remains static across the two instances. We leave this analysis for future work.

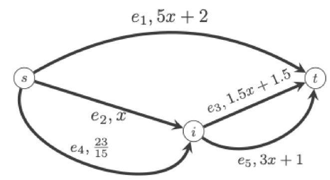

Fig. 2. A series-parallel network that is not SLI, where all traffic originates at node s and terminates at t. The label on each edge represents its identity and cost. For the two-commodity game with and without uncertainty, the total population is 1 and a population of $\epsilon = 0.05$ users have an uncertainty factor of $r = 2$. The social cost for this example in the presence of uncertainty is strictly larger than the social cost in the absence of uncertainty.

D. Tightness of Results: Negative Examples

Our central result in this section (Theorem 3) shows that for networks having the SLI topology with a limited amount of uncertainty ($r \in [1, 2]$), the social cost at equilibrium in the presence of uncertainty is always smaller than or equal to the social cost without any uncertainty. Naturally, this raises the question of whether the result is tight—i.e. what happens when the graph does not belong to the SLI class or if $r > 2$. In this section, we show the tightness of our results by illustrating examples where uncertainty leads to an increase in congestion costs when either one of the above requirements fail.

1) Series-Parallel Networks where Uncertainty Hurts: Fig. 2 depicts an example of a network that is series-parallel but not SLI, where the equilibrium cost when a small fraction of users have an uncertainty of $r = 2$ becomes strictly larger than the cost when all users have no uncertainty. The details of the equilibrium solutions for this example are listed in Table I. Contrasting our result in Theorem 3, this example indicates that for networks that violate the SLI topology, the helpful effects of uncertainty are not guaranteed even when $r \in [1, 2]$.

TABLE I
FLOWS ON EACH PATH AND SOCIAL COST IN THE ABSENCE AND PRESENCE OF UNCERTAINTY: IN THE LATTER CASE, A SMALL POPULATION OF $\epsilon = 0.05$ USERS HAVE AN UNCERTAINTY OF $r = 2$. THE EQUILIBRIUM SOCIAL COST IN THE ABSENCE OF UNCERTAINTY IS SMALLER.

Instance	Equilibrium Flow	Cost
Certain	$x_{p_1} = \frac{4}{21}$, $x_{p_2} = \frac{3}{7}$, $x_{p_3} = \frac{8}{21}$	62/21
Uncertain	$x_{p_1} = \frac{11}{60}$, $x_{p_2} = \frac{23}{60}$, $x_{p_3} = \frac{23}{60}$, $x_{p_4} = \frac{1}{2}$	591/200

2) Networks with Large Uncertainty, i.e., $r > 2$: We now provide an example of a simple two-link Pigou network where the equilibrium social cost with uncertainty $r = 3$ is strictly larger than the equilibrium cost without any uncertainty. This example illustrates that even for networks that fall within the SLI class (parallel links are the most trivial class of networks), uncertainty can be detrimental to social cost when $r > 2$.

Consider the instance depicted in Fig. 3. In the absence of uncertainty, the entire population uses e_2, resulting in a social cost of $C_{e_2}(1) = 1$. Next, suppose that a population of $\epsilon \le 1$ users have an uncertainty level of $r = 3$. If $\epsilon \le$

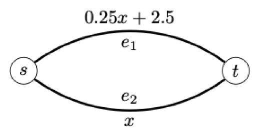

Fig. 3. A two link Pigou network. Consider a two-commodity game with and without uncertainty where a fraction $0 < \epsilon \leq 1$ of users have an uncertainty parameter of $r = 3$. For every $\epsilon > 0$, the equilibrium social cost without uncertainty is strictly smaller than the social cost with uncertainty.

$2/15$, then all ϵ users route their flow on e_1, and the social cost of the equilibrium solution with uncertainty is given by $C_{e_1}(\epsilon)\epsilon + C_2(1-\epsilon)1-\epsilon = 1+0.5\epsilon+1.25\epsilon^2$ which is strictly greater than $C_{e_2}(1)$, the social cost without uncertainty. When $\epsilon > 2/15$, only $2/15$–th of the users prefer e_1, and the increase in social cost follows from the $\epsilon = 2/15$ case.

V. Generalizations

In this section, we generalize our model and consider multi-commodity routing games where the edge costs belong to a class of well-motivated polynomial functions. Specifically, we focus on *shifted monomial* [57] cost functions, where $C_e(x) = a_e x^{d_e} + b_e$ and $d_e \geq 1$ for all $e \in \mathcal{E}$. These functions have been the subject of a considerable literature (e.g., see [58]) across disciplines owing to their real-world applications—it is known that actual road latency functions can be modeled as shifted monomials of degree four [59]. In what follows, we show that all of the results from Sections III and IV generalize smoothly to this new class of functions.

We begin by a identifying a potential function $\Phi_{d,r}(\mathbf{x})$ that the Nash equilibrium solution minimizes for a generalization of the multi-commodity routing game given by Definition 1 where every edge $e \in \mathcal{E}$ has the true cost function $C_e(x) = a_e x^{d_e} + b_e$, which users of type $\theta \in \mathcal{T}$ perceive as $\hat{C}_e^{\theta}(x) = r_\theta a_e x^{d_e} + b_e$. Therefore, equilibrium existence is guaranteed for this more general class of functions.

$$\Phi_{d,r}(\mathbf{x}) = \sum_{e\in\mathcal{E}} \left(\tfrac{1}{d_e+1} a_e x_e^{d_e+1} + b_e \sum_{\theta\in\mathcal{T}} \tfrac{1}{r_\theta} x_e^{\theta}\right) \quad (22)$$

For convenience of exposition, we assume that $d_e = d \geq 1$ for all $e \in \mathcal{E}$ for the rest of this section. The following result generalizes Theorem 1 to games with shifted monomial costs.

Theorem 4: Consider an instance $\mathcal{G}$ of the multi-commodity routing game with shifted monomial costs of degree d having Nash equilibrium $\tilde{\mathbf{x}}$ and the corresponding game instance $\mathcal{G}^1$, having no uncertainty, with equilibrium $\mathbf{x}^1$. Suppose $r_\theta = r$ for all $\theta \in \mathcal{T}$. Then, the following hold: (i) $\mathrm{C}(\tilde{\mathbf{x}}) \leq \mathrm{C}(\mathbf{x}^1)$ if $1 \leq r \leq d+1$; (ii) $\mathrm{C}(\tilde{\mathbf{x}}) \geq \mathrm{C}(\mathbf{x}^1)$ if $0 \leq r \leq 1$.

Unlike Theorem 1 where uncertainty leads to decrease in social cost only when $r \in [1, 2]$, we observe (surprisingly) that as d grows, so does the range of r under which uncertainty is helpful. The proof of Theorem 4 is very similar to that the original result for linear costs (Theorem 1); in fact, this is true for all of the results in this section. To avoid repetition, we only sketch the key differences and defer the details to the extended version on arXiv [60]. The central idea involved in the proof of Theorem 4 is a strict generalization of Lemma 1 via a potential function argument, to obtain the following difference in costs: $\mathrm{C}(\tilde{\mathbf{x}}) - \mathrm{C}(\mathbf{x}) \leq -\sum_{\theta\in\mathcal{T}}(\frac{d+1}{r_\theta} - 1)\sum_{e\in\mathcal{E}} b_e(\tilde{x}_e^{\theta} - x_e^{\theta})$, where $\mathbf{x}$ is an arbitrary feasible flow. The rest of the proof is analogous to that of Theorem 1. Next, we generalize the price of anarchy bounds from Theorem 2.

Theorem 5: For any multi-commodity routing game $\mathcal{G}$ with shifted monomial cost functions of degree d, the ratio between the social cost of the Nash equilibrium to that of the socially optimal solution is at most

$$PoA(\mathcal{G}) \leq \frac{(d+1)^{(d+1)/d}}{\gamma r_{\max}(d+1)^{(d+1)/d} - d r_{\max}^{(d+1)/d}}, \quad (23)$$

where $r_{\max} = \max_{\theta\in\mathcal{T}} r_\theta$, and $\gamma = \frac{\min_\theta r_\theta}{\max_\theta r_\theta}$ as long as $r_{\max} < (\frac{\gamma}{d})^d (d+1)^{(d+1)}$.

Substituting $d = 1$ in the above equation, we obtain Theorem 2 as a special case. In the absence of uncertainty, it is known that routing games with monomial functions having large degree d, tend to exhibit poor price of anarchy, i.e., the PoA bound grows as $O(\frac{d}{\log(d)})$ [20]. Theorem 5 makes a strong case that increasing uncertainty can counteract the poor price of anarchy, e.g., if we assume $\gamma = 1$, the PoA bound from Equation (23) grows as $\Theta(\frac{d}{r_{\max}\log(d)})$ when $r_{\max} = O(\frac{d}{\log(d)})$.

We now generalize the results from Section IV to shifted monomial costs. First, it is not hard to see that Proposition 2 can be immediately extended to arbitrary cost functions including monomials as the proof does not depend on any specific cost function. Hence, we have the following generalization of Theorem 3 in the context of shifted monomial functions.

Theorem 6: Consider an instance $\mathcal{G}$ of the two-commodity game with and without uncertainty and with shifted monomial cost functions of degree $d \geq 1$. Let $\tilde{\mathbf{x}}$ denote the Nash equilibrium of this game and the corresponding game instance $\mathcal{G}^1$, having no uncertainty has equilibrium $\mathbf{x}^1$. Then, as long as G belongs to the SLI class and $1 \leq r \leq d+1$: $\mathrm{C}(\tilde{\mathbf{x}}) \leq \mathrm{C}(\mathbf{x}^1)$.

Discussion: We remark that Theorem 6 (and Theorem 3) partially extends a similar result from [30]. In particular, in Lemma 5.2 in [30] it is shown that for parallel link networks, the social cost of the equilibrium with no uncertainty is always larger than or equal to that of the equilibrium solution when the users have (possibly heterogeneous) uncertainty levels in the range $r \in [1, 2]$. We generalize this result from parallel links to SLI networks but for a more specialized model with two types of users and a more restricted class of cost functions.

Proof sketch: We begin by observing that Lemma 3 holds for shifted monomials as the proof does not explicitly refer to the linear functional form. Leveraging the SLI decomposition, convexity leads to $\sum_{e\in\mathcal{E}} C_e(\tilde{x}_e)(\tilde{x}_e - x_e^1) \geq 0$. Similarly, we can decompose the paths into those that gain and lose flow respectively to prove that $\sum_{e\in\mathcal{E}}(r a_e \tilde{x}_e^d + b_e)(\tilde{x}_e - x_e^1) \leq 0$. The rest of the proof follows from simple algebra. ■

VI. Case Studies

In this section, we present our main simulation results on both stylistic as well as realistic urban network topologies comprising of two types of users—i.e. *through traffic, parking users* (types θ_1, θ_2, respectively)—each associated with a single commodity. We consider a more general edge-dependent uncertainty model and assume that the parking users have different uncertainty levels on different parts of the network and the through traffic does not suffer from uncertainty at all.

This article has been accepted for publication in a future issue of this journal, but has not been fully edited. Content may change prior to final publication. Citation information: DOI 10.1109/TAC.2019.2962102, IEEE Transactions on Automatic Control

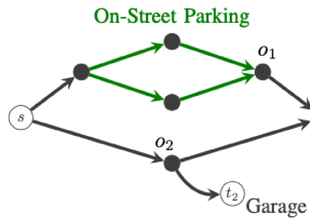

Fig. 4. A multi-commodity network with two types of users: parking users and through traffic. All traffic originates at the source node s. Users belonging to the through traffic population simply select a (minimum-cost) path from s to t and incur the latencies on each link. The parking users select between one of two parking structures: on-street parking (indicated in green) with additional circling costs and off-street parking (e.g., parking garage).

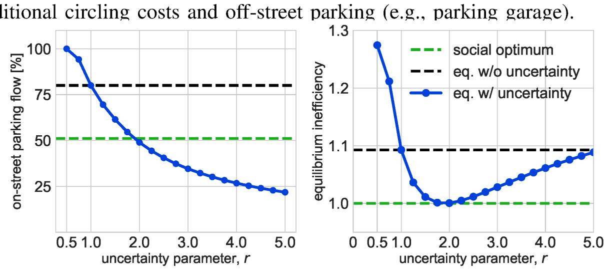

Fig. 5. (Left) The on-street parking population flow (mass) under the social optimum and the Nash equilibrium with and without uncertainty as the level of uncertainty r varies. (Right) The equilibrium quality as measured by the ratio of the social cost of the Nash equilibrium to that of the social optimum. As r increases, we observe from the left plot that more users move away from on-street parking, which in turn affects the social cost as seen from the right plot. In particular, when $r > 1$, the users over-estimate both on-street and garage wait times leading to a decrease in social cost as more users choose the garage option. When $r < 1$, users view the garage option adversely, which leads to more congestion.

We vary the level of uncertainty faced by the parking users, and observe its effect on the social cost at equilibrium.

Despite the generality of the model considered here—different user types have different beliefs and the level of uncertainty depends on the edge under consideration—our simulations validate the theoretical results presented in the previous section. In particular, when the parking users are cautious (over-estimate the congestion cost due to cruising and waiting for a parking spot), we observe that the social cost of the equilibrium decreases in both our experiments. On the other hand, the behavior of parking users who under-estimate the cost incurred due to searching for a parking spot leads to a worse-cost equilibrium.

A. *Effect of Uncertainty on On-Street vs Garage Parking*

Inspired by the work in [55] which provides a framework for integrating parking into a classical routing game that abstracts route choices in urban networks, we begin with a somewhat stylized example of an urban network, depicted and described in Fig. 4. The users looking for a parking spot are faced with two options: (i) *on-street parking* which is cheaper but leads to larger wait times; (ii) an *off-street or a private garage option* that has lower wait times at the expense of a higher price.

To understand the costs faced by the parking users (type θ_2), let $\mathcal{E}_{os}$ be the set of edges in the on-street parking structure (the green edges in Fig. 4). For parking users that select the on-street option, the cost on edges $e \in \mathcal{E}_{os}$ are of the form $C_e^{\theta_2}(x_e) = C_{e,\ell}^{\theta_2}(x_e) + C_{e,os}^{\theta_2}(x_e)$, where $C_{e,\ell}^{\theta_2}(x_e) = a_e x_e + b_e$ is the travel latency part of the cost and $C_{e,os}^{\theta_2}(x_e) = a_{os}x_e + b_{os}$ is the parking part of the cost. Fig. 4 is easily converted into a two-commodity network by creating a *fake* edge $\tilde{e}$ from node o_1 to t_2 that has the accumulated parking costs from the edges in $\mathcal{E}_{os}$—i.e. $C_{\tilde{e}}^{\theta_2}(x_{\tilde{e}}) = \sum_{e\in\mathcal{E}_{os}} C_{e,os}^{\theta_2}(x_{\tilde{e}}) = \bar{a}_{os}x_{\tilde{e}} + \bar{b}_{os}$. Then, the costs on edges in $\mathcal{E}_{os}$ are re-defined to only contain the travel latency component of the cost, and this is the same for both types of users: for $e \in \mathcal{E}_{os}$, $C_e^{\theta_1}(x_e) = C_e^{\theta_2}(x_e) = a_e x_e + b_e$. For the off-street parking structure, the edge, say e', from o_2 to t_2 has cost $C_{e'}^{\theta_2}(x_{e'}) = a_{pg}x_{e'} + b_{pg}$ and all other edges have costs $C_e^{\theta_1}(x_e) = C_e^{\theta_2}(x_e) = a_e x_e + b_e$.

The price on-street is generally lower than that of a private garage, whereas the inequality is reversed for wait times. Parking garages typically comprise of a larger capacity than on-street parking options and one expects this to be reflected in the $\bar{a}_{os}$ and a_{pg} terms.

The uncertainty is faced by users of type θ_2 only on the costs pertaining to parking such that the congestion-dependent component of their parking cost is multiplied by a parameter $r > 0$. This captures the notion that the users may face uncertainty regarding the number of other users competing for the same parking spots(s) or the capacity of each structure. In the converted two commodity network, this translates to $\hat{C}_{\tilde{e}}^{\theta_2}(x_{\tilde{e}}) = r\bar{a}_{os}x_{\tilde{e}} + \bar{b}_{os}$ where $\tilde{e}$ is the fake edge from o_1 to t_2 and $\hat{C}_{e'}^{\theta_2}(x_{e'}) = ra_{pg}x_{e'} + b_{pg}$ where e' is the edge from o_2 to t_2. As mentioned previously, the through traffic population (type θ_1) does not suffer from uncertainty on any of its edges and therefore, $r_{\theta_1}(e) = 1$ with respect to every edge e.

For the simulations, we assume that there is a total population mass of 2 originating at the source node s, comprising of an equal number of parking and through traffic users. The edge congestion functions are selected randomly from a suitable range. The costs on the parking structures are set as:

1) *On-Street Parking*: On-street parking contains a fixed number of parking spots, and the parameter $\bar{a}_{os}$ is chosen to be inversely proportional to this quantity. The constant term $\bar{b}_{os}$ captures the price paid by drivers for on-street parking and is selected based on parking prices multiplied by a constant that captures how users tend to trade-off between time (congestion) and money.
2) *Off-street Parking*: Off-street parking is assumed to have a large number of available parking spots; thus, we set $a_{pg} = 0$. The parameter b_{pg} is set to be the price of off-street parking (e.g., a garage) multiplied by the same trade-off parameter as above.

Fig. 5 shows how the parking users divide themselves among the on-street and garage option (left plot) and how this affects equilibrium quality as r varies (right plot). From the left plot, we observe that at the social optimum, about half of the parking population prefers on-street parking. With no uncertainty (i.e. $r = 1$), at the Nash equilibrium more parking users (80%) gravitate towards the cheaper on-street option leading to higher congestion and inefficiency. So even without uncertainty, the system is inefficient. This is reminiscent of the classic Pigou example [20] in traffic networks.

As r increases—that is, as users become more cautious in their beliefs about parking congestion—more users start flocking to the off-street option due to the fact that they

This article has been accepted for publication in a future issue of this journal, but has not been fully edited. Content may change prior to final publication. Citation information: DOI 10.1109/TAC.2019.2962102, IEEE Transactions on Automatic Control

perceive a multiplicative increase in the congestion-dependent term. This results in an improvement to efficiency. On the contrary, for users who tend to under-estimate parking costs ($r < 1$), the appeal of parking in off-street options decreases and more users flock towards on-street parking leading to increased congestion and poor equilibrium quality.

The effect of the above behavior on equilibrium quality can be seen in the right-hand plot in Fig. 5. The graph indicates that a cautious approach under uncertainty (users over-estimating costs) helps improve equilibrium quality, whereas lack of caution results in enhanced congestion in the network. In fact, we notice when $r < 1$, the price of anarchy increases in a rather steep fashion as r decreases. Finally, we observe that even though our theoretical results guarantee an improvement in equilibrium congestion only in the range of $1 \leq r \leq 2$, our simulation results lend credence to the claim that even highly cautious behavior ($r > 2$) can lead to a decrease in the social cost. This can occur due to two reasons. First, according to Fig. 5 (left), the fraction of users parking on-street decreases in a convex fashion as r increases. Therefore, the rate at which users gravitate from on-street parking to garage parking is rather slow as r increases beyond two and consequently, the change in social cost is also bounded. Second, since the garage option has $a_{pg} = 0$, there are no congestion costs and so an increase in the population parking off-street at $r > 2$ leads to a limited increase in cost. In conclusion, we see that when $r > 2$ but is not too large, we still obtain reasonable efficiency because the flow is still comparable to the $r = 2$ case and changes slowly when r increases.

B. Parking vs. Through Traffic in Downtown Seattle

In a similar manner to the toy example, we take a real-world urban traffic network (depicted in Fig. 6) that captures a slice of a highly congested area in downtown Seattle. The network contains both on-street and off-street parking options and the parking population experiences both a travel latency cost and a parking cost, both of which we model as affine functions. Once again, this can be converted to a standard two-commodity instance by adding a new (fake) destination node t_{θ_2} and including *fake edges* from (i) the top left and bottom right nodes in the on-street parking zone (boundary nodes on the blue colored dotted area in Fig. 6a) to t_{θ_2}; (ii) the node containing the parking garage (marked with a 'P' symbol in Fig. 6a) to t_{θ_2}. By adding fake edges from only specific boundary nodes in the on-street parking area, we are able to capture the added congestion due to the cruising and circling behavior exhibited by users searching for parking spots. As with our previous example, we assume that the through traffic faces no uncertainty ($r_{\theta_1}(e) = 1$ for all edges) and the parking users face an uncertainty parameter of $r > 0$ only on the fake edges, which affects their perception of the parking costs.

For the simulations, we assume that the parking traffic originates at a few select nodes in the network (indicated in magenta in Fig. 6) and wishes to route their flow to either an on-street parking slot or a garage. On the other hand, the through traffic originates at every node in the network and has a single destination, which represents drivers seeking to leave the downtown area via state highway 99.

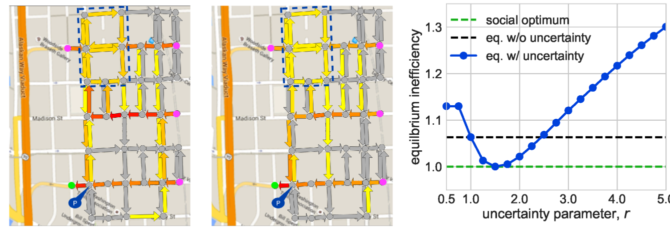

(a) $r = 1$ (b) $r = 1.5$ (c) Inefficiency

Fig. 6. Setup and results for the parking and through traffic example modeled on a region of downtown Seattle. In (a) and (b), the network superimposed on the corresponding area in downtown Seattle is shown. The parking population originates at the magenta nodes, whereas the through traffic begins at every grey node and terminates at the green node. The blue dotted box represents the on-street parking zone and the parking symbol ('P') is the location of the off-street parking garage. The color on each edge depicts the intensity of flow on that edge: the intensity increases as the color transitions from gray to red.

The parameters for our simulations were chosen similarly to the example in Section VI-A. Specifically, the cost functions for the two parking structures were based on an estimation of the number of available parking slots, and the hourly price for parking for both on-street and garage parking in Seattle. Furthermore, owing to the uniformity of the downtown roads, we assumed that all edges in the network have the same congestion cost function, which were sufficiently scaled in order to ensure that they are comparable to the parking costs.

As seen in Fig. 6a, in the downtown Seattle network, the equilibrium without uncertainty is sub-optimal as more parking users select the cheaper on-street option. This leads to heavy congestion in the network (indicated by the red edges) as parking users distributed across the network approach the on-street parking area. That being said, the equilibrium without uncertainty is not *as sub-optimal* as the simple network in Fig. 4 as the cost functions are more symmetric and parking users who originate closer to the parking garage prefer using that option despite the higher price.

In the presence of uncertainty, we observe an interesting phenomenon. When the parking users are cautious, the garage option becomes more preferable to users who are equidistant from both parking locations. The parking users distribute themselves more evenly across the two options, which in turn leads to lesser congestion in the middle of the network.

It is well-known that sub-optimal behavior by the parking leads to increased congestion for through traffic [61]. Our simulation results indicate that the routes adopted by that uncertainty helps alleviate some of this congestion.

Fig. 6c shows the inefficiency of each of the equilibria as a function of the uncertainty parameter value of the parking users. Specifically, at $r = 1.5$, the equilibrium solution is close to optimal (inefficiency ≈ 1.02) and at $r = 2$, it coincides with the socially optimal flow. On the other hand, for $r < 1$, the social cost of the equilibrium solution increases because more users select the on-street parking option. This in turn leads to heavier congestion in the rest of the network. Once again, we observe that over-estimation can lead to a decrease in social cost even when $r > 2$ (up to $r \approx 2.5$ in this case). However, unlike our results in Fig. 5, significant cost over-estimation leads to large inefficiency— this is due to the specific placement of the parking garage in the downtown

This article has been accepted for publication in a future issue of this journal, but has not been fully edited. Content may change prior to final publication. Citation information: DOI 10.1109/TAC.2019.2962102, IEEE Transactions on Automatic Control

example. More precisely as r increases, a large fraction of the population favors the garage, whose location intersects with the route adopted by the through traffic.

We conclude by remarking that even though cautious behavior results in only a small improvement in the price of anarchy (see Fig. 6c), even a small improvement in daily congestion in downtown areas could result in economic gains.

VII. Conclusions and Future Work

In this work, we consider a multi-commodity selfish routing game where different types of users face different levels of uncertainty quantified by a multiplicative parameter r_θ. Broadly classifying the user attitudes as cautious and not-cautious, we provide several theoretical results highlighting the effect that when users over-estimate their network costs, equilibrium quality tends to improve and vice-versa when users under-estimate the costs. Although our primary focus is on linear costs, we also show that our results generalize gracefully to more general polynomial functions that are well-motivated by real-world latencies.

Our work presents a number of novel contributions and new insights on 'how much uncertainty' is beneficial for a system. For example, a little pessimism in the face of uncertainty can result in considerable improvements in equilibrium quality whereas too much optimism may lead to increased congestion. While most of the literature looks at worst-case price of anarchy bounds, we show that uncertainty is helpful for every possible instance (Theorems $1, 3, 4, 6$). This is important since the average instance encountered in practice may not resemble the worst-case one. More specifically, we are among the first to provide an instance-wise characterization of social cost for instances containing two types of users and in networks more general than parallel links (e.g., SLI). In terms of price of anarchy, we extend many of the existing results for homogeneous users [38], [39] and provide new bounds for networks with heterogeneous user uncertainties and two or more types. Finally, the techniques that we leverage to prove our results may be of value to future work. In particular, we believe that the proof of Theorem 3 where we combine two disparate approaches: (i) using a potential function to bound social cost, and (ii) topological characterization of SLI networks is applicable broadly.

The results also expose a number of new avenues for future research. Perhaps, the most important direction is to consider more realistic models of cost uncertainty (e.g., the mean-variance approach [27]) and equilibrium notions such as the Bayes Nash equilibrium [49] (BNE). For instance, one could study a BNE by assuming that users of type θ have incomplete information on the uncertainty corresponding to every other type θ'. Although this general model is appealing at first glance, it is fraught with computational challenges. Under a BNE, each individual user's cost could depend on taking the expectation over an exponential number of realized states, and crucially, layered beliefs on the uncertainty parameter—e.g., user i's belief on user j's belief on user i's cost and so on. Further research is required in order to reconcile the differences between the desirable properties of the BNE, its computational difficulties, and users' preferences for simpler decision-making heuristics as advocated by behavioral economics [31].

Appendix

A. Proof of Lemma 1

Recall from Proposition 1 that the equilibrium solution $\tilde{\mathbf{x}}$ minimizes the corresponding potential function $\Phi_r(\mathbf{x})$. So for some $\mathbf{x}'$, we have that $\Phi_r(\tilde{\mathbf{x}}) - \Phi_r(\mathbf{x}') \leq 0$ and

$$\begin{aligned}
\mathrm{C}(\mathbf{x}') &= \textstyle\sum_{e\in\mathcal{E}}(a_e x'_e + b_e)x'_e \\
&= \textstyle\sum_{e\in\mathcal{E}}\left(a_e(x'_e)^2 + b_e\sum_{\theta\in\mathcal{T}} 2\frac{x'^{\theta}_e}{r_\theta} + b_e\sum_{\theta\in\mathcal{T}}(1-\frac{2}{r_\theta})x'^{\theta}_e\right) \\
&= 2\Phi_r(\mathbf{x}') - \textstyle\sum_{e\in\mathcal{E}} b_e\sum_{\theta\in\mathcal{T}}(\frac{2}{r_\theta}-1)x'^{\theta}_e. \qquad (24)
\end{aligned}$$

Applying (24) to the solutions $\tilde{\mathbf{x}}$ and $\mathbf{x}$, we get that

$$\begin{aligned}
\mathrm{C}(\tilde{\mathbf{x}}) - \mathrm{C}(\mathbf{x}) &= 2\Phi_r(\tilde{\mathbf{x}}) - 2\Phi_r(\mathbf{x}) \\
&\quad - \textstyle\sum_{e\in\mathcal{E}} b_e\sum_{\theta\in\mathcal{T}}(\frac{2}{r_\theta}-1)(\tilde{x}^{\theta}_e - x^{\theta}_e) \qquad (25)\\
&\leq -\textstyle\sum_{\theta\in\mathcal{T}}(\frac{2}{r_\theta}-1)\sum_{e\in\mathcal{E}} b_e\Delta x^{\theta}_e. \qquad (26)
\end{aligned}$$

where (26) follows from $\Phi_r(\tilde{\mathbf{x}}) - \Phi_r(\mathbf{x}) \leq 0$. ■

B. Proof of Lemma 2

Fixing the left-hand side of

$$\frac{f(x)x}{f(x')x' + \sum_{i=1}^n (x_i - x'_i)f(r_i x)} \leq 4\left(4r_* - (r^*)^2\right)^{-1}, \qquad (27)$$

we first derive a lower bound on the denominator (i.e. identify when the denominator is minimized over the space of all valid instantiations of the parameter set). Let us begin with the second term in the denominator:

$$\begin{aligned}
\textstyle\sum_{i=1}^n (x_i - x'_i)f(r_i x) &= \textstyle\sum_{i=1}^n (x_i - x'_i)(r_i a x + b) \\
&= \textstyle\sum_{i=1}^n x_i r_i a x - \sum_{i=1}^n x'_i r_i a x + b(x - x') \\
&\geq \textstyle\sum_{i=1}^n x_i r_* a x - \sum_{i=1}^n x'_i r^* a x + b(x - x') \\
&= x r_* a x - x' r^* a x + b(x - x')
\end{aligned}$$

Using this upper bound on the rest of the terms in the denominator, we get that

$$f(x')x' + \textstyle\sum_{i=1}^n (x_i - x'_i)f(r_i x) \geq a(x')^2 + bx + x r_* a x - x' r^* a x.$$

For any given fixed value of x, consider the function $a(x')^2 - x' r^* a x$: by basic calculus, its minimum value is attained if $x' = \frac{r^* x}{2}$. In other words, for any x, x', we can conclude that $a(x')^2 - x' r^* a x \geq (r^*)^2\frac{ax^2}{4} - (r^*)^2\frac{ax^2}{2} = -(r^*)^2\frac{ax^2}{4}$. Substituting this bound into the above equation, we have

$$f(x') + \textstyle\sum_{i=1}^n (x_i - x'_i)f(r_i x) \geq -(r^*)^2\frac{ax^2}{4} + bx + r_* a x^2.$$

Now that we have removed the dependence on x', we can substitute this back into (14) to get that

$$\begin{aligned}
\frac{f(x)x}{f(x')x' + \sum_{i=1}^n (x_i - x'_i)f(r_i x)} &\leq \frac{ax^2 + bx}{-\frac{(r^*)^2 ax^2}{4} + bx + r_* a x^2} \\
&\leq \frac{ax^2}{-\frac{(r^*)^2 ax^2}{4} + r_* a x^2} = \frac{4}{4r_* - (r^*)^2}
\end{aligned}$$

The penultimate inequality is obtained by removing the dependence on b in the denominator. ■

This article has been accepted for publication in a future issue of this journal, but has not been fully edited. Content may change prior to final publication. Citation information: DOI 10.1109/TAC.2019.2962102, IEEE Transactions on Automatic Control

C. Proof of Lemma 3

We prove Lemma 3 via induction.

Claim 1: Given instances $\mathcal{I} = \{G, \mathcal{T}, (s,t), (\mu_{\theta_1}, \mu_{\theta_2}), (1,r), (C_e)_{e\in\mathcal{E}}\}$ and $\mathcal{I}' = \{G, \mathcal{T}, (s,t), (\mu'_{\theta_1}, \mu'_{\theta_2}), (1,1), (C_e)_{e\in\mathcal{E}}\}$ where the graph G is linearly independent, let $\mathbf{y}$ and $\mathbf{y}^1$ denote the Nash equilibria of the two instances, respectively. Then, if

$$\sum_{e\in p} C_e(y_e) \leq \sum_{e\in p} C_e(y^1_e), \ \forall\ p \in \mathcal{P} \text{ s.t. } y^{\theta_1}_p > 0, \quad (28)$$

it must be the case that $y^{\theta_1}_p \leq y^1_p$ for all $p \in \mathcal{P}$.

To prove Lemma 3, it is sufficient to prove the above claim. Indeed, setting $\mathcal{I} = \mathcal{G}$ and $\mathcal{I}' = \mathcal{G}^1$, we can directly apply the claim to prove the lemma statement. In order to verify that the condition from (28) is satisfied, first observe that linearly independent networks are a special case of series-parallel networks. Now, applying [18, Lemma 3] to the flows $\mathbf{y}$ and $\mathbf{y}^1$ as defined in the statement of Lemma 3, we get that there exists a path $p^* \in \mathcal{P}$ with $y^1_{p^*} > 0$ such that $\sum_{e\in p^*} C_e(y_e) \leq \sum_{e\in p^*} C_e(y^1_e)$. However, p^* is a min-cost path in $\mathbf{y}^1$ and its cost in $\mathbf{y}$ is an upper-bound on that of any min-cost path used by users of type θ_1. This immediately implies that for any p with $y^{\theta_1}_p > 0$ its cost in $\mathbf{y}$ (where it is a min-cost path for users of type θ_1) must be smaller than or equal to its cost in $\mathbf{y}^1$ (where all users have no uncertainty).

The proof of Claim 1 proceeds by induction on the number of edges in the graph—that is, given a graph $G = (V, \mathcal{E})$, we assume that the claim holds for all graph $G' = (V', \mathcal{E}')$ such that $|\mathcal{E}'| < |\mathcal{E}|$ and proceed from there. A property of linearly-independent graphs, which will be useful for our proof.

Proposition 3 ([18, Proposition 5]): A two terminal network G is linearly-independent if and only if: (i) it consists of a single edge, or (ii) it is the result of connecting two linearly-independent networks in parallel, or (iii) it is the result of connecting in series a linearly independent network and a network with a single edge.

Proof of Claim 1: Let us start with the base case. The inductive hypothesis is trivially true when the graph G consists of a single edge between the source and sink (i.e. $\mathcal{E} = \{e\}$). All of the users must route their flow on this edge only, regardless of their uncertainty level. Since we are given (by assumption) that the instance satisfies the condition that $C_e(y_e) \leq C_e(y^1_e)$, by monotonicity of the cost functions, we get that $y^{\theta_1}_e \leq y^1_e$ and the inductive claim thus holds for the base case.

Now, let us consider the inductive case. Consider the inductive claim with respect to an arbitrary linearly independent graph $G = (V, \mathcal{E})$ consisting of two or more edges, and assume that the inductive hypothesis holds for all instances defined on linearly independent graphs $G' = (V', \mathcal{E}')$ such that $|\mathcal{E}'| < |\mathcal{E}|$. From Proposition 3, we know that G consists of two subgraphs $G_1 = (V_1, \mathcal{E}_1)$ and $G_2 = (V_2, \mathcal{E}_2)$ connected either in parallel or in series with $|\mathcal{E}_2| = 1$ (i.e., G_2 is simply a single edge). We consider both cases, apply the inductive claim recursively to both sub-graphs, and merge the resulting flows to prove the inductive claim for the original graph G.

Let (s_1, t_1), (s_2, t_2) denote the origin-destination pairs for G_1 and G_2 respectively. We use $\mathcal{P}_1$ to denote the set of s_1-t_1 paths in G_1 and $\mathcal{P}_2$ to denote the set of s_2-t_2 paths in G_2. When G_1 and G_2 are connected in series (case 1), we have that $s_2 = t_1$. When they are connected in parallel (case 2), $s_1 = s_2, t_1 = t_2$. We treat each of these cases separately.

(case 1) *Subgraphs connected in series*. We first introduce some additional notation required for this part of the proof. As we did in the proof of Theorem 3, we use $\mathbf{y}(1)$ and $\mathbf{y}(2)$ to denote the sub-flows of $\mathbf{y}$ with respect to graphs G_1 and G_2. Similarly, we define $\mathbf{y}^1(1)$ and $\mathbf{y}^1(2)$ to be the sub-flows of $\mathbf{y}^1$. Recall that G_2 consists simply of a single edge.

Since $\mathbf{y}$ is an equilibrium for $\mathcal{I}$, it must be the case that $\mathbf{y}(1)$ and $\mathbf{y}(2)$ are equilibria for G_1, G_2 respectively for suitably defined sub-instances, $\mathcal{I}_1 = \{G_1, \mathcal{T}, (s_1, t_1), (\mu_{\theta_1}, \mu_{\theta_2}), (1, r), (C_e)_{e\in\mathcal{E}_1}\}$ and $\mathcal{I}_2 = \{G_2, \mathcal{T}, (s_2, t_2), (\mu_{\theta_1}, \mu_{\theta_2}), (1, r), (C_e)_{e\in\mathcal{E}_2}\}$, respectively. In a similar manner, $\mathbf{y}^1(1)$ and $\mathbf{y}^1(2)$ represent equilibria for the instances $\mathcal{I}'_1 = \{G_1, \mathcal{T}, (s_1, t_1), (\mu'_{\theta_1}, \mu'_{\theta_2}), (1,1), (C_e)_{e\in\mathcal{E}_1}\}$ and $\mathcal{I}'_2 = \{G_2, \mathcal{T}, (s_2, t_2), (\mu'_{\theta_1}, \mu'_{\theta_2}), (1,1), (C_e)_{e\in\mathcal{E}_2}\}$, respectively. These are formally summarized in the notation table below.

TABLE II
INSTANCE DEFINITIONS AND NASH EQUILIBRIA WHEN G_1 AND G_2 ARE CONNECTED IN SERIES.

	Game Instance Definition	N.E.
$\mathcal{I}$	$\{G, \mathcal{T}, (s,t), (\mu_{\theta_1}, \mu_{\theta_2}), (1,r), (C_e)_{e\in\mathcal{E}}\}$	$\mathbf{y}$
$\mathcal{I}_1$	$\{G_1, \mathcal{T}, (s_1,t_1), (\mu_{\theta_1}, \mu_{\theta_2}), (1,r), (C_e)_{e\in\mathcal{E}_1}\}$	$\mathbf{y}(1)$
$\mathcal{I}_2$	$\{G_2, \mathcal{T}, (s_2,t_2), (\mu_{\theta_1}, \mu_{\theta_2}), (1,r), (C_e)_{e\in\mathcal{E}_2}\}$	$\mathbf{y}(2)$
$\mathcal{I}'$	$\{G, \mathcal{T}, (s,t), (\mu'_{\theta_1}, \mu'_{\theta_2}), (1,1), (C_e)_{e\in\mathcal{E}}\}$	$\mathbf{y}^1$
$\mathcal{I}'_1$	$\{G_1, \mathcal{T}, (s_1,t_1), (\mu'_{\theta_1}, \mu'_{\theta_2}), (1,1), (C_e)_{e\in\mathcal{E}_1}\}$	$\mathbf{y}^1(1)$
$\mathcal{I}'_2$	$\{G_2, \mathcal{T}, (s_2,t_2), (\mu'_{\theta_1}, \mu'_{\theta_2}), (1,1), (C_e)_{e\in\mathcal{E}_2}\}$	$\mathbf{y}^1(2)$

From (28) that for any path $p \in \mathcal{P}$ with $y^{\theta_1}_p > 0$, $\sum_{e\in p} C_e(y_e) \leq \sum_{e\in p} C_e(y^1_e)$. Dividing the path p into sub-paths p_1 and p_2 representing its intersection with G_1, G_2 respectively, we get that $\sum_{e\in p_1} C_e(y_e) + \sum_{e\in p_2} C_e(y_e) \leq \sum_{e\in p_1} C_e(y^1_e) + \sum_{e\in p_2} C_e(y^1_e)$.

Based on the above inequality, either $\sum_{e\in p_1} C_e(y_e) \leq \sum_{e\in p_1} C_e(y^1_e)$ or $\sum_{e\in p_2} C_e(y_e) \leq \sum_{e\in p_2} C_e(y^1_e)$.

Suppose that $\sum_{e\in p_1} C_e(y_e) \leq \sum_{e\in p_1} C_e(y^1_e)$. Since the graph G_1 is linearly independent, we can apply the inductive claim to instances $\mathcal{I}_1$ and $\mathcal{I}'_1$ to get that $y(1)^{\theta_1}_p \leq y^1(1)_p$ for every path $p \in \mathcal{P}_1$. However, recall that G_2 consists of only one edge (call it e_2). This implies that there is a one-to-one correspondence between every path in $\mathcal{P}$ and path in $\mathcal{P}_1$. Specifically, for every path $p \in \mathcal{P}$, there exists a unique path $p_1 \in \mathcal{P}_1$ such that $p = p_1 \cup \{e_2\}$. Since the mapping is a bijection, this implies that for every $p = p_1 \cup \{e_2\}$, it must be the case that $y^{\theta_1}_p = y(1)^{\theta_1}_{p_1}$ and $y^1_p = y^1(1)_{p_1}$. Therefore, for every path $p \in \mathcal{P}$, $y^{\theta_1}_p = y(1)^{\theta_1}_{p_1} \leq y^1(1)_{p_1} = y^1_p$. This proves the inductive claim for this case.

Next, suppose that $\sum_{e\in p_2} C_e(y_e) \leq \sum_{e\in p_2} C_e(y^1_e)$. Once again, since G_2 consists of a single edge e_2, the condition implies that $y_{e_2} = y(2)_{e_2} \leq y^1_{e_2} = y^1(2)_{e_2}$. Thus, we infer that conditional upon $\sum_{e\in p_2} C_e(y_e) \leq \sum_{e\in p_2} C_e(y^1_e)$, $\mu_{\theta_1} + \mu_{\theta_2} \leq \mu'_{\theta_1} + \mu'_{\theta_2}$. That is, the total population for instance $\mathcal{I}$ is smaller than or equal to the total population for instance $\mathcal{I}'$.

Since linearly independent networks are a special case of series-parallel, we can now apply [18, Lemma 3] for flows $\mathbf{y}(1)$ and $\mathbf{y}^1(1)$ on graph G_1. Since $\mu_{\theta_1} + \mu_{\theta_2} \leq \mu'_{\theta_1} + \mu'_{\theta_2}$, this of course, implies that there exists at least one origin-destination path p^* in G_1 with $y^1(1)_{p^*} > 0$ satisfying $\sum_{e \in p^*} C_e(y(1)_e) \leq \sum_{e \in p^*} C_e(y^1(1)_e)$. Therefore, we can conclude that for any path $p \in \mathcal{P}_1$ with $y_p^{\theta_1} > 0$, we have, $\sum_{e \in p} C_e(y(1)_e) \leq \sum_{e \in p} C_e(y^1(1)_e)$.

The above equation symbolizes the induction condition as described in (28). Applying the inductive hypothesis recursively to instances $\mathcal{I}_1, \mathcal{I}'_1$, and using the same reasoning as before, we conclude that for every s-t path $p \in \mathcal{P}$, $y_p^{\theta_1} \leq y_p^1$. This concludes the proof for the series case.

(case 2) *Subgraphs G_1, G_2 connected in parallel to obtain G.* The proof for this case proceeds similarly to the case where the networks are connected in series. Once again, we divide the flows $\mathbf{y}$ and $\mathbf{y}^1$ into sub-flows $\mathbf{y}(1), \mathbf{y}(2)$ and $\mathbf{y}^1(1)$ and $\mathbf{y}^1(2)$. Moreover, we use $\mu_{\theta_1}(1), \mu_{\theta_2}(1)$ to denote the total flow without and with uncertainty in sub-graph G_1 and similarly so, for all the other subgraphs and sub-flows. A comprehensive notation table is listed below.

TABLE III
INSTANCE DEFINITIONS AND NASH EQUILIBRIA WHEN G_1 AND G_2 ARE CONNECTED IN PARALLEL.

	Game Instance Definition	N.E.
$\mathcal{I}$	$\{G, \mathcal{T}, (s,t), (\mu_{\theta_1}, \mu_{\theta_2}), (1,r), (C_e)_{e \in \mathcal{E}}\}$	$\mathbf{y}$
$\mathcal{I}_1$	$\{G_1, \mathcal{T}, (s,t), (\mu_{\theta_1}(1), \mu_{\theta_2}(1), (1,r), (C_e)_{e \in \mathcal{E}_1}\}$	$\mathbf{y}(1)$
$\mathcal{I}_2$	$\{G_2, \mathcal{T}, (s,t), (\mu_{\theta_1}(2), \mu_{\theta_2}(2)), (1,r), (C_e)_{e \in \mathcal{E}_2}\}$	$\mathbf{y}(2)$
$\mathcal{I}'$	$\{G, \mathcal{T}, (s,t), (\mu'_{\theta_1}, \mu'_{\theta_2}), (1,1), (C_e)_{e \in \mathcal{E}}\}$	$\mathbf{y}^1$
$\mathcal{I}'_1$	$\{G_1, \mathcal{T}, (s,t), (\mu'_{\theta_1}(1), \mu'_{\theta_2}(1), (1,1), (C_e)_{e \in \mathcal{E}_1}\}$	$\mathbf{y}^1(1)$
$\mathcal{I}'_2$	$\{G_2, \mathcal{T}, (s,t), (\mu'_{\theta_1}(2), \mu'_{\theta_2}(2), (1,1), (C_e)_{e \in \mathcal{E}_2}\}$	$\mathbf{y}^1(2)$

Before applying the inductive claim recursively to the two subgraphs, we need to verify that the condition specified in (28) is satisfied for the subgraphs. We know that the instances $\mathcal{I}, \mathcal{I}'$ satisfy the condition in (28) (this is part of our assumption). That is, for any $p \in \mathcal{P}$ such that $y_p^{\theta_1} > 0$

$$\sum_{e \in p} C_e(y_e) \leq \sum_{e \in p} C_e(y_e^1). \tag{29}$$

Since the graphs G_1, G_2 are connected in parallel, the set of edges and paths in each graph are disjoint. In other words, $\mathcal{P} = \mathcal{P}_1 \cup \mathcal{P}_2$. Moreover, the path-flows in $\mathbf{y}$ coincide with those in $\mathbf{y}(1)$ or $\mathbf{y}(2)$—e.g., for any $p_1 \in \mathcal{P}_1$, we have that $y_{p_1} = y(1)_{p_1}$ and so on. This combined with (29) implies that for any path $p_1 \in \mathcal{P}_1$ with $y(1)_{p_1}^{\theta_1} > 0$, $\sum_{e \in p_1} C_e(y(1)_e) \leq \sum_{e \in p_1} C_e(y^1(1)_e)$. Therefore, the flows $\mathbf{y}(1)$ and $\mathbf{y}^1(1)$ satisfy the inductive condition. Similarly, we can show that the flows $\mathbf{y}(2)$ and $\mathbf{y}^1(2)$ also satisfy (28).

Applying the inductive claim to the instances $\mathcal{I}_1$ and $\mathcal{I}'_1$, we get that $y(1)_{p_1}^{\theta_1} \leq y^1(1)_{p_1} \quad \forall p_1 \in \mathcal{P}_1$. Similarly, we can apply the inductive claim to instance pair $\mathcal{I}_2, \mathcal{I}'_2$ to get that $y(2)_{p_2}^{\theta_1} \leq y^1(2)_{p_2} \quad \forall p_2 \in \mathcal{P}_2$. Since $\mathbf{y} = \mathbf{y}(1) \cup \mathbf{y}(2)$ and $\mathbf{y}^1 = \mathbf{y}^1(1) \cup \mathbf{y}^1(2)$, the claim follows immediately. ■

REFERENCES

[1] A. Parekh and J. Walrand, *Sharing Network Resources*, R. Srikant, Ed. Morgan and Claypool, 2014.

[2] S. Stidham Jr., "Pricing and congestion management in a network with heterogeneous users," *IEEE Trans. Automat. Contr.*, vol. 49, no. 6, pp. 976–981, 2004.

[3] I. Stupia, L. Sanguinetti, G. Bacci, and L. Vandendorpe, "Power control in networks with heterogeneous users: A quasi-variational inequality approach," *IEEE Trans. Signal Processing*, vol. 63, no. 21, pp. 5691–5705, 2015.

[4] R. E. Azouzi and E. Altman, "Constrained traffic equilibrium in routing," *IEEE Trans. Automat. Contr.*, vol. 48, no. 9, pp. 1656–1660, 2003.

[5] G. Devetag and M. Warglien, "Playing the wrong game: An experimental analysis of relational complexity and strategic misrepresentation," *Games and Economic Behavior*, vol. 62, no. 2, pp. 364–382, 2008.

[6] E. Avineri and J. Prashker, "Violations of expected utility theory in route-choice stated preferences: Certainty effect and inflation of small probabilities," *Transportation Research Record: Journal of the Transportation Research Board*, no. 1894, pp. 222–229, 2004.

[7] C. Chen, A. Skabardonis, and P. Varaiya, "Travel-time reliability as a measure of service," *Transportation Research Record: J. Transportation Research Board*, no. 1855, pp. 74–79, 2003.

[8] G. Pierce and D. Shoup, "Getting the prices right: an evaluation of pricing parking by demand in san francisco," *J. Amer. Planning Association*, vol. 79, no. 1, pp. 67–81, 2013.

[9] B. Schaller, "New york city's congestion pricing experience and implications for road pricing acceptance in the united states," *Transport Policy*, vol. 17, no. 4, pp. 266–273, 2010.

[10] P. N. Brown and J. R. Marden, "The robustness of marginal-cost taxes in affine congestion games," *IEEE Trans. Automat. Contr.*, vol. 62, no. 8, pp. 3999–4004, 2017.

[11] D. Schrank, B. Eisele, T. Lomax, and J. Bak, "2015 urban mobility scorecard," Texas A&M Transportation Institute, 2015.

[12] A. Mortazavi, X. Pan, E. Jin, M. Odioso, and Z. Sun, "Travel Times on Changeable Message Signs Volume II–Evaluation of Transit Signs," *University of California, Berkeley*, 2009.

[13] W. Krichene, J. Reilly, S. Amin, and A. M. Bayen, "Stackelberg routing on parallel networks with horizontal queues," *IEEE Trans. Automat. Contr.*, vol. 59, no. 3, pp. 714–727, 2014.

[14] J. Thai, N. Laurent-Brouty, and A. M. Bayen, "Negative externalities of gps-enabled routing applications: A game theoretical approach," in *Proc. 19th IEEE Inter. Conf. Intelligent Transportation Systems*, 2016, pp. 595–601.

[15] M. Wu, J. Liu, and S. Amin, "Informational aspects in a class of bayesian congestion games," in *Proc. Amer. Control Conf.*, 2017, pp. 3650–3657.

[16] H. Tavafoghi and D. Teneketzis, "Informational incetives in congestion games," in *Proc. 55th Ann. Allerton Conf. Communication, Control, and Computing*, 2017.

[17] T. Roughgarden and É. Tardos, "How bad is selfish routing?" *J. ACM*, vol. 49, no. 2, pp. 236–259, 2002.

[18] I. Milchtaich, "Network topology and the efficiency of equilibrium," *Games and Economic Behavior*, vol. 57, no. 2, pp. 321–346, 2006.

[19] S. Sekar, L. Zheng, L. J. Ratliff, and B. Zhang, "Uncertainty in multi-commodity routing networks: When does it help?" in *2018 Annual American Control Conference (ACC)*. IEEE, 2018, pp. 6553–6558.

[20] T. Roughgarden, "Routing games," in *Algorithmic game theory*. Cambridge University Press New York, 2007, ch. 18, pp. 459–484.

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

[22] F. Farokhi, W. Krichene, A. M. Bayen, and K. H. Johansson, "A heterogeneous routing game," in *Proc. 51st Ann. Allerton Conf. Communication, Control, and Computing*, 2013, pp. 448–455.

[23] M. Gairing, B. Monien, and K. Tiemann, "Routing (un-) splittable flow in games with player-specific linear latency functions," in *Inter. Colloq. Automata, Languages, and Programming*. Springer, 2006, pp. 501–512.

[24] M. Mavronicolas, I. Milchtaich, B. Monien, and K. Tiemann, "Congestion games with player-specific constants," in *Proc. 32nd Intl. Symp. Math. Found. Comput. Sci.*, 2007, pp. 633–644.

[25] I. Milchtaich, "Congestion games with player-specific payoff functions," *Games and economic behavior*, vol. 13, no. 1, pp. 111–124, 1996.

[26] D. Fotakis, D. Kalimeris, and T. Lianeas, "Improving selfish routing for risk-averse players," in *Proc. 11th Inter. Conf. Web and Internet Economics*, 2015, pp. 328–342.

[27] E. Nikolova and N. E. S. Moses, "The burden of risk aversion in mean-risk selfish routing," in *Proc. 16th ACM Conf. Economics and Computation*, 2015, pp. 489–506.

[28] E. Altman, T. Basar, and R. Srikant, "Nash equilibria for combined flow control and routing in networks: asymptotic behavior for a large number of users," *IEEE Trans. Automat. Contr.*, vol. 47, no. 6, pp. 917–930, 2002.

[29] W. S. Vickrey, "Pricing in urban and suburban transport," *The Amer. Economic Review*, vol. 53, no. 2, pp. 452–465, 1963.

[30] P. N. Brown and J. R. Marden, "Fundamental limits of locally-computed incentives in network routing," in *2017 American Control Conference (ACC)*. IEEE, 2017, pp. 5263–5268.

[31] A. Tversky and D. Kahneman, "Judgment under uncertainty: Heuristics and biases," in *Utility, probability, and human decision making*. Springer, 1975, pp. 141–162.

[32] R. Meir and D. C. Parkes, "Congestion games with distance-based strict uncertainty," in *Proc. 29th AAAI Conf. Artificial Intelligence*, 2015, pp. 986–992.

[33] T. Qi, J. Chen, W. Su, and M. Fu, "Control under stochastic multiplicative uncertainties: Part i, fundamental conditions of stabilizability," *IEEE Trans. Automat. Contr.*, vol. 62, no. 3, pp. 1269–1284, 2017.

[34] A. A. Walters, "The theory and measurement of private and social cost of highway congestion," *Econometrica: J. Econometric Society*, pp. 676–699, 1961.

[35] M. Ostrovsky and M. Schwarz, "Carpooling and the economics of self-driving cars," National Bureau of Economic Research, Tech. Rep., 2018.

[36] R. Guers, C. Langbort, and D. Work, "On informational nudging and control of payoff-based learning," *IFAC Proc. Vol.*, vol. 46, no. 27, pp. 69–74, 2013.

[37] 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, pp. 893–917, 2018.

[38] P.-A. Chen, B. D. Keijzer, D. Kempe, and G. Schäfer, "Altruism and its impact on the price of anarchy," *ACM Transactions on Economics and Computation*, vol. 2, no. 4, p. 17, 2014.

[39] R. Meir and D. C. Parkes, "Playing the wrong game: Smoothness bounds for congestion games with behavioral biases," *SIGMETRICS Performance Evaluation Review*, vol. 43, no. 3, pp. 67–70, 2015.

[40] S. Zhu and D. Levinson, "Do people use the shortest path? an empirical test of wardrop's first principle," *PloS one*, vol. 10, no. 8, p. e0134322, 2015.

[41] G. Gigerenzer, "Fast and frugal heuristics: The tools of bounded rationality," *Blackwell handbook of judgment and decision making*, vol. 62, p. 88, 2004.

[42] F. Ordóñez and N. E. Stier-Moses, "Wardrop equilibria with risk-averse users," *Transportation Science*, vol. 44, no. 1, pp. 63–86, 2010.

[43] C. Systematics *et al.*, "Traffic congestion and reliability: Trends and advanced strategies for congestion mitigation," *Final Report, Texas Transportation Institute. http://ops. fhwa. dot. gov/congestion_report_04/index. htm*, 2005.

[44] D. A. Seale, A. Rapoport, and D. V. Budescu, "Decision making under strict uncertainty: An experimental test of competitive criteria," *Organizational Behavior and Human Decision Processes*, vol. 64, no. 1, pp. 65–75, 1995.

[45] J. Qi, M. Sim, D. Sun, and X. Yuan, "Preferences for travel time under risk and ambiguity: Implications in path selection and network equilibrium," *Transportation Research Part B: Methodological*, vol. 94, pp. 264–284, 2016.

[46] M. Wittmann, *Felt time: The psychology of how we perceive time*. Mit Press, 2016.

[47] X. Wu, D. Levinson, and H. Liu, "Perception of waiting time at signalized intersections," *Transportation Research Record: J. Transportation Research Board*, no. 2135, pp. 52–59, 2009.

[48] R. Cole, T. Lianeas, and E. Nikolova, "When does diversity of user preferences improve outcomes in selfish routing?" *arXiv:1702.07806*, 2017.

[49] J. C. Harsanyi, "Games with incomplete information played by "bayesian" players, i–iii part i. the basic model," *Management science*, vol. 14, no. 3, pp. 159–182, 1967.

[50] M. Gairing, B. Monien, and K. Tiemann, "Selfish routing with incomplete information," *Theory Comput. Syst.*, vol. 42, no. 1, pp. 91–130, 2008.

[51] T. Roughgarden, "The price of anarchy in games of incomplete information," in *Proc. 13th ACM Conf. Electronic Commerce*. ACM, 2012, pp. 862–879.

[52] D. Monderer and L. S. Shapley, "Potential games," *Games and Economic Behavior*, vol. 14, no. 1, pp. 124–143, 1996.

[53] W. Krichene, B. Drighès, and A. M. Bayen, "Online learning of nash equilibria in congestion games," *SIAM J. Control and Optimization*, vol. 53, no. 2, pp. 1056–1081, 2015.

[54] A. Héliou, J. Cohen, and P. Mertikopoulos, "Learning with bandit feedback in potential games," in *Advances in Neural Information Processing Systems 30*, 2017, pp. 6372–6381.

[55] D. Calderone, E. Mazumdar, L. J. Ratliff, and S. S. Sastry, "Understanding the impact of parking on urban mobility via routing games on queue-flow networks," in *Proc. IEEE 55th Conf. Decision and Control*, 2016, pp. 7605–7610.

[56] L. J. Ratliff, C. Dowling, E. Mazumdar, and B. Zhang, "To observe or not to observe: Queuing game framework for urban parking," in *Proc. 55th IEEE Conf. Decision and Control*, 2016, pp. 5286–5291.

[57] R. Colini-Baldeschi, M. Klimm, and M. Scarsini, "Demand-independent optimal tolls," in *Proc. 45th Inter. Colloquium Automata, Languages, and Programming*, 2018, pp. 151:1–151:14.

[58] E. Altman, T. Basar, T. Jiménez, and N. Shimkin, "Competitive routing in networks with polynomial costs," *IEEE Trans. Automat. Contr.*, vol. 47, no. 1, pp. 92–96, 2002.

[59] H. C. Manual, "Special report 209," *Transportation Research Board, Washington, DC*, vol. 1, p. 985, 1985.

[60] S. Sekar, L. Zheng, L. J. Ratliff, and B. Zhang, "Uncertainty in multicommodity routing networks: When does it help?" *IEEE Transactions on Automatic Control (under review, available: arxiv:1709.08441)*, 2017.

[61] D. Shoup and H. Campbell, "Gone parkin'," *The New York Times*, vol. 29, 2007.

Shreyas Sekar is a Postdoctoral Fellow at Harvard University. He received his Ph.D. in Computer Science from Rensselaer Polytechnic Institute in 2017 and a B.Tech in Electronics and Communications Engineering from the Indian Institute of Technology, Roorkee in 2012. His research interests are in operations research, mechanism design, and game theory. He is a recipient of the 2017 Robert McNaughton prize from the Rensselaer Polytechnic Institute.

Liyuan Zheng is a Ph.D. student in Electrical Engineering at the University of Washington, Seattle, prior to which he obtained a B.S. in Applied Physics from University of Science and Technology of China. He also obtained a M.S. in Physics at University of Washington. His research interests are in Learning, Optimization and Game Theory.

Lillian J. Ratliff (S'08–M'15) is an Assistant Professor in EE at the University of Washington, Seattle, prior to which she obtained her Ph.D. (2015) in EEECS at the University of California, Berkeley. She obtained a B.S. in Mathematics (2008) and a B.S. (2008) and M.S. (2010) in EE from the University of Nevada, Las Vegas. Her research interests lie at the intersection of game theory, optimization, and learning. She is the recipient of the NSF Graduate Research Fellowship (2009) and the CISE Research Initiation Initiative Award (2017).

Baosen Zhang received his Bachelor of Applied Science in Engineering Science degree from the University of Toronto in 2008; and his PhD degree in Electrical Engineering and Computer Sciences from University of California, Berkeley in 2013. He was a Postdoctoral Scholar at Stanford University. He is currently an Assistant Professor in Electrical and Computer Engineering at the University of Washington, Seattle, WA. His research interests are in power systems and cyberphysical systems.