

# 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 quality in a multi-commodity selfish routing game with many types of users, where each user type experiences a different level of uncertainty. We consider a new model of uncertainty where each user-type over- or under-estimates their congestion costs by a multiplicative constant. We present a variety of theoretical results showing that when users under-estimate their costs, the network congestion decreases at equilibrium, whereas over-estimation of costs leads to increased equilibrium congestion. Motivated by applications in urban transportation networks, we perform simulations consisting of parking users and through traffic on synthetic network topologies. Our results strongly indicate that while users' perception of network costs can significantly impact the equilibrium quality, optimism in the face of uncertainty leads to favorable network conditions.

## 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 [1]. In all of these systems, users are *inherently heterogeneous* not only in their demands and objectives, but also in their belief about the state of the system and how they trade-off between time, money, and risk [2], [3]. Naturally, these private beliefs influence each user's decisions and as a consequence, the total welfare of the overall system. Therefore, understanding the effects of these heterogeneities is fundamental to characterizing network state and performance.

A motivating example of a routing network, which we use throughout this paper, is the *urban transportation network*. Travelers in road networks simultaneously trade-off between 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 heterogeneous attitudes, and there is evidence to suggest that the routes adopted depend not on the true costs but on how they are perceived by users. For instance, users prefer safer routes over those with high variance [4], seek to minimize travel time over parking costs [5], and react adversely to per-mile road taxes [6].

Furthermore, the technological and economic incentives employed by network operators to tackle congestion may compound these effects by interacting with user beliefs in a ‘perverse manner’ [7]. For example, to limit the economic loss arising from urban congestion [8], cities across the world have introduced a number of solutions including road taxes,

All of the authors are associated with the Department of Electrical Engineering, University of Washington, Seattle, WA, USA. {sekarsh, liyuanz8, ratliff, zhangbao}@uw.edu.

This work is partially supported by NSF Awards CNS-1736582, CNS-1646912, and CNS-1634136. A full version of this paper can be found on arXiv [19].

time-of-day-pricing, and road-side message signs. 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. Therefore, to truly evaluate the efficacy of such solutions, it is crucial to understand how these changes impact the congestion level of the system.

The effect of uncertainties on network equilibrium was examined in a recent body of work [7], [9]. These *uncertainties are reflected through the beliefs of users*, where each user may perceive the network condition to be different than the ‘true’ conditions. Current results have largely focused on simple networks (e.g., parallel links) where a fixed percentage of the population is endowed with the same level of uncertainty. Given the complexity in most practical networks, it is natural to ask how uncertainty (i.e., user beliefs) affects equilibria when there are many types of users, who are heterogeneous in their perceptions. Specifically, in this work we answer the following two questions: (i) *how do equilibria depend on the type and level of uncertainty in networks with a multitude of users*, and (ii) *when does uncertainty lead to an improvement or degradation in equilibrium quality*?

To address these questions, we turn to a *multi-commodity selfish routing* framework commonly employed by many disciplines (see, e.g., [10]–[12]). In our model, each user seeks to route some flow along links connecting two nodes in a network and faces congestion costs on each link. These congestion costs are perceived differently by each user in the network, representing the uncertainties in their beliefs. 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 an optimum routing solution [13]. Against this backdrop, we analyze what happens when users have imprecise views about the congestion costs. A surprising outcome arises: in the presence of uncertainty, if users under-estimate the costs and select routes based on these perceived costs, *the equilibrium quality is better compared to the full information case*. Conversely, if the users are overly conservative and over-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  $\theta$  is captured by a single parameter  $r_\theta > 0$ . Specifically, for each user of type  $\theta$ , if the 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  $a_e x + r_\theta b_e$ .

We are interested in studying the quality of the equilibrium routing as well as the socially optimal routing solution, using *social cost* as a metric. This quantity denotes the aggregate cost incurred by all of the users in the network (see (8) for a formal definition).

We consider two types of uncertainties: *pessimism* where users over-estimate the costs ( $r_\theta \geq 1$ ) and *optimism* where users under-estimate the costs ( $r_\theta \leq 1$ ), for all types  $\theta$ .

Under this model of type-based uncertainty in congestion costs, we have the following contributions:

- (a) The social cost of the equilibrium solution where all users have the same level of uncertainty ( $r_\theta = r$  for all  $\theta$ ) is *always smaller than or equal to* the cost of the equilibrium solution without uncertainty when  $r \in [0.5, 1]$  and vice-versa when  $r \geq 1$ .
- (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 [13]) is  $\frac{4r_{\min}^2}{4r_{\min}\gamma-1}$ , where  $r_{\min} = \min_\theta r_\theta$  and  $\gamma \leq 1$  is the ratio of the minimum to the maximum uncertainty over user types.
- (c) 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.

To validate the theoretical results, we present a number of simulation results. We focus specifically on the application of parking in urban transportation networks (see, e.g., Fig. 1) and consider simple networks with two types of users: through traffic and parking users. We demonstrate the effects of an uncertain parking population on equilibrium quality. We show via simulations that *optimism improves equilibrium quality while pessimism degrades it* both when uncertainty is asymmetric across user types and when users face different levels of uncertainty on different network edges. In the full version of this paper [19], we show that this behavior persists even in more realistic urban topologies.

### B. Comparison with Other Models of Uncertainties

Our work is closely related to the extensive body of work on risk-averse selfish routing [14], [15] and pricing tolls in congestion networks [2], [11]. The former line of research focuses on the well known *mean-standard deviation* model where each individual 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 [16]. Motivated by this, we adopt a multiplicative model of uncertainty similar to [17], [18].

In regards to the latter line of work, the literature on computing tolls for heterogeneous users is driven by the need to *implement the optimum routing* by adjusting the toll amount, which is often interpreted as the *time–money* tradeoff, on each edge as a function of the congestion. We, on the other hand, assume that the user beliefs are independent

of the congestion in the network, and aim for a more nuanced understanding of the dependence of network costs on the level of uncertainty.

### 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 Section III. Section IV presents our simulation results on transportation networks with parking and routing users who face different levels of uncertainty. Finally, we conclude with some discussion and comments on future directions in Section V.

## 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

$$C_e(x_e) = a_e x_e + b_e, \quad (1)$$

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  $(s_\theta, t_\theta, \mu_\theta, r_\theta)$ . We assume that  $\mu_\theta > 0$  denotes the total population of users belonging to type  $\theta$  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, the parameter  $r_\theta > 0$  captures the beliefs or uncertainties associated with users of type  $\theta$  and affects the edge cost in the following way: users of type  $\theta$  perceive the cost of edge  $e \in \mathcal{E}$  to be

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

If we interpret  $b_e$  as a price or a tax, then  $r_\theta < 1$  denotes the case where users of type  $\theta$  under-estimate prices compared to their actual value and  $r_\theta > 1$  captures situations where users over-estimate prices or view them adversely compared to their other costs.

Let  $\mathcal{P}_\theta$  denote the set of all  $s_\theta-t_\theta$  paths in  $G$  where  $s_\theta$  is the source and  $t_\theta$  is the destination. Formally, let  $x_p^\theta \in \mathbb{R}$  be the total flow routed by users of type  $\theta$  on path  $p \in \mathcal{P}_\theta$ . We define a *feasible flow* to be a vector  $\mathbf{x} = (x_p^\theta)_{\theta \in \mathcal{T}, p \in \mathcal{P}_\theta}$  such that for all  $\theta \in \mathcal{T}$ ,  $\sum_{p \in \mathcal{P}_\theta} x_p^\theta = \mu_\theta$ , and  $x_p^\theta \geq 0$  for all paths  $p \in \mathcal{P}_\theta$ . Path flows are related to edge flows. Indeed, let  $x_e^\theta \in \mathbb{R}$  be the flow on edge  $e$  of users of type  $\theta$ . The edge and path flow for users of type  $\theta$  are related by

$$x_e^\theta = \sum_{p \in \mathcal{P}_\theta, p \ni e} x_p^\theta \quad (3)$$

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

flow; indeed for any path  $p$  in the graph,

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

Similarly, the perceived path costs are given by

$$\hat{C}_p^\theta(\mathbf{x}) = \sum_{e \in p} (\hat{C}_e^\theta(x_e) = \sum_{e \in p} a_e x_e + r_\theta b_e). \quad (5)$$

We define a game instance as the tuple

$$\mathcal{G} = \{(V, \mathcal{E}), \mathcal{T}, (s_\theta, t_\theta, \mu_\theta, r_\theta)_{\theta \in \mathcal{T}}, (C_e)_{e \in \mathcal{E}}\}. \quad (6)$$

A simple example of a multi-commodity routing network that arises in real-world transportation networks is depicted in Fig. 1. In Section IV, we perform extensive simulations on this stylized network and highlight the effect of uncertainty on network conditions in a more general model where users experience different uncertainties on each of the edges.

#### A. Nash Equilibrium Concept and its Efficiency

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

*Definition 2.1 (Nash Equilibrium):* Given a game instance  $\mathcal{G}$ , a feasible flow  $\mathbf{x}$  is said to be a *Nash equilibrium* if for every  $\theta \in \mathcal{T}$ , for all  $p \in \mathcal{P}_\theta$  with positive flow,  $x_p^\theta > 0$ ,

$$\hat{C}_p^\theta(\mathbf{x}) \leq \hat{C}_{p'}^\theta(\mathbf{x}), \quad \forall p' \in \mathcal{P}_\theta. \quad (7)$$

For the rest of this work, we will assume that all the flows considered are feasible.

*Remark 1 (User Beliefs):* In order to employ the classical notion of Nash equilibrium, we assume that all uncertainty levels are known by all of the users. That is, a user of type  $\theta$  knows the values  $(r_\theta b_e)_{e \in \mathcal{E}}$  for all  $\theta' \in \mathcal{T}$ . While a user knowing  $(r_\theta b_e)_{e \in \mathcal{E}}$  within its own type  $\theta$  may not be unreasonable, full knowledge of  $(r_\theta b_e)_{e \in \mathcal{E}}$  for all  $\theta \in \mathcal{T}$  is a strong assumption.

This being said, for the types of games we consider, a number of *myopic learning rules*<sup>1</sup> converge to Nash equilibria independent of the beliefs held by users regarding other user types. (see, e.g., [20] and references therein).

#### B. Social Cost and Price of Anarchy

We measure the quality of a solution using its social cost, which is defined to be the aggregate (true) cost incurred by all of the users in the system. Formally, the social cost of a flow  $\mathbf{x}$  is given by

$$C(\mathbf{x}) = \sum_{e \in \mathcal{E}} C_e(x_e) x_e. \quad (8)$$

The social cost is only measured with respect to the true congestion costs and does not reflect users' beliefs.

To capture inefficiencies, we leverage the well-studied notion of the *price of anarchy* which is the ratio of the

<sup>1</sup>By *myopic learning rules*, we mean rules for iterated play that require each player to have minimal-to-no knowledge of other players' cost functions and/or strategies.

social cost of the *worst-case Nash equilibrium* to that of the socially optimal solution [12]. Formally, given an instance  $\mathcal{G}$  of a multi-commodity routing game belonging to some class  $\mathcal{C}$  (a class refers to a set of instances that usually share some property) suppose that  $\mathbf{x}_\mathcal{G}^*$  is the flow that minimizes the social cost  $C(\mathbf{x})$  and that  $\tilde{\mathbf{x}}_\mathcal{G}$  is the Nash equilibrium for the given instance, then the price of anarchy is:

*Definition 2.2 (Price of Anarchy):* Given a class of instances  $\mathcal{C}$ , the *price of anarchy* for this class is

$$\max_{\mathcal{G} \in \mathcal{C}} \frac{C(\tilde{\mathbf{x}}_\mathcal{G})}{C(\mathbf{x}_\mathcal{G}^*)}. \quad (9)$$

Of course, the price of anarchy is always greater than or equal to one.

### III. MAIN RESULTS

The first step to analyzing the multi-commodity game is to characterize the Nash equilibria. If the games fall into the general class of potential games, then the equilibria have "nice properties" in terms of existence, uniqueness, and computability [20]. 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 [10].

Due to the fact that we have linear latencies for each type and the type-dependent uncertainty appears on the  $b_e$  terms for each edge, all game instances of the form we consider admit a potential function and hence, there always exists a Nash equilibrium [10].

*Proposition 3.1:* A feasible flow  $\mathbf{x}$  is a Nash equilibrium for a given instance  $\mathcal{G}$  of a multi-commodity routing game if and only if it minimizes the following potential function:

$$\Phi_r(\mathbf{x}) = \sum_{e \in \mathcal{E}} \left( \frac{a_e x_e^2}{2} + b_e \sum_{\theta \in \mathcal{T}} r_\theta x_e^\theta \right) \quad (10)$$

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

The proof of the proposition follows from standard arguments pertaining to the minimizer of a convex function and from the definition of Nash equilibrium as in (7) and we refer the reader to [20] for more details. The second part of the proposition indicates that the equilibria are essentially unique as the cost on every edge is the same across solutions.

With the above proposition in hand, we now derive three main results on (i) the impact of uncertainty on social cost, (ii) the impact of uncertainty on players, and (iii) bounds on the price of anarchy.

#### 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 under-estimate costs for *every instance* and vice-versa when they over-estimate costs. For a given instance of the multi-commodity routing game, we say that the users are optimistic if  $r_\theta \leq 1$  for all  $\theta \in \mathcal{T}$ . Similarly, the users are pessimistic if  $r_\theta \geq 1$  for all user types.

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 3.2:** 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, whose Nash equilibrium is  $\mathbf{x}^1$ . Suppose all users experience the same uncertainty, where  $r_\theta = r$  for all  $\theta \in \mathcal{T}$ . Then, the following hold:

- 1)  $C(\tilde{\mathbf{x}}) \leq C(\mathbf{x}^1)$  if  $0.5 \leq r \leq 1$ .
- 2)  $C(\tilde{\mathbf{x}}) \geq C(\mathbf{x}^1)$  if  $r \geq 1$ .

We now sketch the key technique involved—a complete proof can be found in the full version of this paper [19]. The main idea behind the proof is that the solutions  $\tilde{\mathbf{x}}$  and  $\mathbf{x}^1$  minimize closely related potential functions, i.e.,  $\Phi_r(\mathbf{x})$  and  $\Phi_1(\mathbf{x})$  respectively. This leads us to derive a sufficient condition highlighting the ‘direction of flow change’ from  $\mathbf{x}^1$  to  $\tilde{\mathbf{x}}$ , namely that  $(r-1) \sum_{e \in \mathcal{E}} b_e (\tilde{x}_e - x_e^1) \leq 0$ . The rest of the proof follows from a case-by-case analysis based on the value of the uncertainty parameter. The following corollary identifies a specific level of (optimistic) uncertainty at which the equilibrium solution is actually optimal.

**Corollary 3.3:** 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  $r_\theta = 0.5$  for all  $\theta \in \mathcal{T}$ , then  $C(\tilde{\mathbf{x}}) = C(\mathbf{x}^*)$ —i.e. the equilibrium is socially optimal.

### B. Impact of Uncertain Users on Those without Uncertainty

Now that we have a better understanding of how uncertainty affects the performance of the entire system as measured by the social cost, we tackle a more nuanced question: *in systems where only some users are uncertain, how does their behavior impact the social cost of the users who do not have uncertainty?* This question is of considerable interest in a number of settings, e.g., in urban transportation networks, where it is believed that [21]–[24] inefficient behavior by the parking users (such as cruising or searching for parking spots) can often cascade into increased congestion for other drivers leading to a detrimental effect on the overall congestion cost.

To answer this question, we restrict our attention to a two-commodity routing game  $\mathcal{G}$ , where both types of users seek to route their flow from a common source node  $s$  to sink node  $t$ . Moreover, we assume that for the first type  $\theta_1$ ,  $r_{\theta_1} = 1$ . For the second type  $\theta_2$ ,  $r_{\theta_2}$  is not necessarily 1. We refer to this as the *two-commodity game having users with and without uncertainty*.

We now define some additional notation. Suppose that  $\mathbf{x}$  denotes a feasible flow for a given instance of the two-commodity game with uncertain and full information users, we use  $C^{\theta_1}(\mathbf{x}) = \sum_{e \in \mathcal{E}} C_e(x_e) x_e^{\theta_1}$  to denote the aggregate cost of users of type  $\theta_1$ . Finally, we also restrict our focus to series-parallel networks, which denote a popular class of network topologies in the literature pertaining to network routing [25], [26]. Informally, a graph is said to be series-parallel if it does not contain an embedded *Wheatstone network* or equivalently if, in the undirected version of this

graph two routes never pass through any edge in opposite directions. The reader is referred to [26] for a more rigorous definition.

We present a surprising result in Theorem 3.4: 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. This result holds for both optimistic and pessimistic behavior.

**Theorem 3.4:** Given an instance  $\mathcal{G}$  of the two-commodity game having users 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,

$$C^{\theta_1}(\tilde{\mathbf{x}}) \leq C^{\theta_1}(\mathbf{x}^1). \quad (11)$$

The proof of Theorem 3.4 is based on Lemma 3 in [26] and can be found in the full version [19].

### C. Price of Anarchy Under Uncertainty

In Theorem 3.2, we showed that the equilibrium cost under uncertainty decreases (resp. increases) when users are optimistic (resp. pessimistic) and all user types have the same level of uncertainty. This naturally raises the question of quantifying the improvement (or degradation in equilibrium) 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 minimum uncertainty in the system and  $\gamma$ , which is the ratio between the minimum and maximum uncertainty among user types.

Let  $r_{\min} = \min_{\theta \in \mathcal{T}} r_\theta$  and  $\gamma = \frac{\min_{\theta} r_\theta}{\max_{\theta} r_\theta}$ .

**Theorem 3.5:** If  $1 < 4r_{\min}\gamma$ , the price of anarchy for multi-commodity routing games is given by

$$\frac{4r_{\min}^2}{4r_{\min}\gamma - 1}. \quad (12)$$

We provide the proof of the above theorem in [19].

The above result broadly validates our message that uncertainty helps equilibrium when users under-estimate their costs and hurts equilibrium when users over-estimate their costs. To understand why, let us consider the case of  $\gamma = 1$ —i.e. the uncertainty is the same across user types. We already know that in the absence of uncertainty, the price of anarchy of multi-commodity routing games with linear costs equals  $\frac{4}{3}$  [13]; this can also be seen by substituting  $r_{\min} = 1$ ,  $\gamma = 1$  in (9). We observe that the price of anarchy is strictly smaller than  $\frac{4}{3}$  for  $r_{\min} < 1$  and reaches the optimum value of one at  $r_{\min} = 0.5$  thereby confirming Corollary 3.3.

Similarly, as  $r_{\min}$  increases from one, the price of anarchy also increases nearly linearly. In fact, our price of anarchy result goes one step beyond Theorem 3.2 as it provides guarantees even when different user types have different uncertainty levels. For example, when  $r_{\min} = 0.6$ , and  $\gamma = 0.9$ —i.e.  $\max_{\theta} r_\theta \approx 0.67$ —the price of anarchy is 1.24, which is still better than the price of anarchy without uncertainty. Furthermore, the price of anarchy result reveals a surprising dichotomy: as long as  $r_{\min} < 1$  and

$\gamma$  is not too large, 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.

#### IV. NUMERICAL EXAMPLES

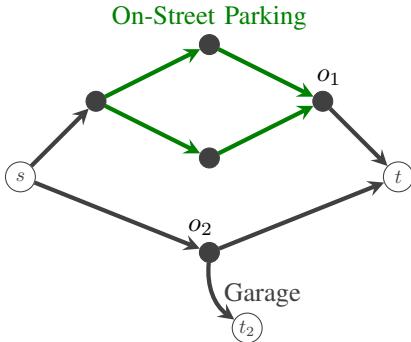


Fig. 1: A special case of our general multi-commodity network with two types of users—parking users and through traffic. All of the network traffic originates at the source node  $s$ . Users belonging to the through traffic simply select a (minimum-cost) path from  $s$  to  $t$  and incur only the latencies 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 (e.g., parking garage).

In this section, we present our main simulation results on a stylized network topology comprising of two types of users (two commodities)—i.e. *through traffic, parking users* (types  $\theta_1, \theta_2$ , respectively). We consider a more general model of uncertainty for our simulations, where the parking users have different uncertainty levels on different parts of the network and the through traffic does not suffer from uncertainty at all. We vary the level of uncertainty faced by the parking users, and observe its effect on the social cost at equilibrium.

In the full version of this paper [19], we present additional simulations on a more realistic urban network topology that captures the downtown area in Seattle. Both our stylized and realistic network simulations validate the theoretical results presented in the previous section.

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

Inspired by the work in [21] 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. 1. The users looking for a parking spot are faced with two options: (i) *on-street parking* which, as in reality, is cheaper but leads to larger wait times due to cruising in search of parking; (ii) an *off-street or a private garage option* that is much easier to access (in terms of wait times) at the expense of a higher price.

To understand the costs faced by the parking users (type  $\theta_2$ ), let  $\mathcal{E}_{os}$  be the set of edges in the on-street parking structure (the green edges in Fig. 1). For parking users that select the on-street parking 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 and  $C_{e,os}^{\theta_2}(x_e) = a_{e,os} x_e + b_{e,os}$  is the parking cost.

Fig. 1 easily transforms into a two-commodity network by creating a *fake* edge  $\tilde{e}$  from node  $o_1$  to  $t_2$ , having the accumulated parking costs from edges  $\mathcal{E}_{os}$ —i.e.  $C_{\tilde{e}}(x_{\tilde{e}}) = \sum_{e \in \mathcal{E}_{os}} C_{e,os}^{\theta_2}(x_e)$ . Then the costs on edges in  $\mathcal{E}_{os}$  are redefined to only contain the travel latency 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  $C_{e'}(x_{e'}) = a_{pg} x_{e'} + b_{pg}$ .

The uncertainty is only faced by the parking users (and only on  $\tilde{e}, e'$ ) who perceive the cost of the two parking structures to be  $\hat{C}_{\tilde{e}}^{\theta_2}(x_{\tilde{e}}) = a_{os} x_{\tilde{e}} + r b_{os}$  and  $\hat{C}_{e'}^{\theta_2}(x_{e'}) = a_{pg} x_{e'} + r b_{pg}$ , respectively, for some parameter  $r > 0$ . The through traffic (type  $\theta_1$ ) does not suffer from uncertainty on any of its edges and therefore,  $r_{\theta_1} = 1$  on all edges. Moreover,  $r_{\theta_2} = 1$  on all edges in the network except  $\tilde{e}, e'$ .

For the simulations, we assume that there is an equal mass of parking and through traffic users originating at the source node  $s$ . The parameters of the edge cost functions were selected uniformly at random from a suitable range. For both on-street and off-street options, the parking costs were set as follows: (i) The parameter  $a_{os}$  was chosen to be inversely proportional to the number of on-street parking slots typically available on roads, whereas we set  $a_{pg} = 0$ , since garages have a large number of parking slots; (ii) The parameters  $b_{os}$  and  $b_{pg}$  equal the price of on-street and garage parking (respectively) commonly used in cities multiplied by a constant that captures the trade-off between time and money.

Fig. 2 shows how the parking users divide themselves among the on-street and garage option (left plot) and how this affects uncertainty as  $r$  varies (right plot). From the left plot, we observe that at the social optimum, approximately 51% of the parking population prefers on-street parking. With no uncertainty (i.e. when  $r = 1$ ), at the Nash equilibrium more parking users (around 80%) gravitate towards the on-street option leading to higher congestion—that is, even without uncertainty, the system is inefficient as is expected.

As  $r$  decreases—that is, as users become more optimistic in their beliefs about prices—more users start flocking to the off-street option as they perceive a multiplicative decrease in price. The result is an improvement in efficiency. On the other hand, for users who tend to over-estimate prices ( $r > 1$ ), the appeal of off-street parking decreases and users move to on-street options leading to increased congestion and poor equilibria. The effect of the above behavior on equilibrium inefficiency is quantified in the right-hand plot in Fig. 2.

#### V. CONCLUSIONS AND FUTURE WORK

In this work, we considered a multi-commodity selfish routing game where different types of users face different levels of uncertainty quantified by a multiplicative parameter  $r_{\theta}$ . Broadly classifying the user attitudes as optimistic and pessimistic, we showed several theoretical and experimental results highlighting the effect that when users under-estimate

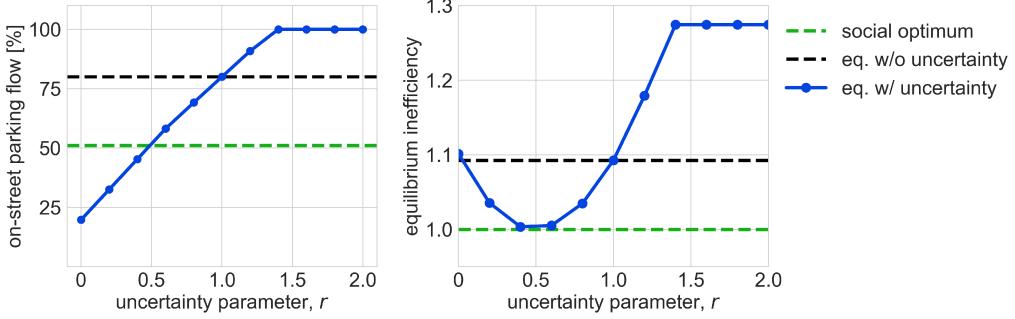


Fig. 2: The left plot shows the on-street parking population mass under the social optimum and the Nash equilibrium as  $r$ , the uncertainty parameter, varies. The right plot shows the equilibrium quality as measured by the ratio of the social cost of the equilibrium solution to that of the optimum for each value of  $r$ . As  $r$ , the level of uncertainty faced by the parking users increases, we observe from the left plot more users move towards on-street parking, which in turn affects the social cost as seen from the right plot. Particularly, when  $0.5 \leq r \leq 1$ , the users under-estimate prices leading to a decrease in social cost as more users choose the garage option. On the other hand, for  $r \geq 1$ , users view the garage option adversely, which leads to more congestion. For a sufficiently large  $r$ , all of the parking users end up using the on-street option and the social cost saturates around  $r = 1.4$  with further increase in uncertainty having no effect on equilibrium cost.

their network costs, equilibrium quality tends to improve and vice-versa when users over-estimate the costs.

Although, we focused on linear congestion cost functions in this work, we remark that *all* of our results extend naturally to polynomial cost functions of the form  $C_e(x_e) = a_e x_e^d + b$ , albeit with different price of anarchy bounds [19]. Finally, on a somewhat abstract note, the central message of this paper that ‘optimism in the face of uncertainty leads to near-optimal solutions’ bears thematic similarities to the popular Upper-Confidence Bound (UCB) algorithm [27] used for Multi-Armed Bandit problems. It would be interesting to see if this connection could be exploited to show similar bounds on equilibrium in routing games with learning.

## 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] 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.
- [5] G. Pierce and D. Shoup, “Getting the prices right: an evaluation of pricing parking by demand in san francisco,” *J. American Planning Association*, vol. 79, no. 1, pp. 67–81, 2013.
- [6] 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.
- [7] P. N. Brown and J. R. Marden, “Fundamental limits of locally-computed incentives in network routing,” in *Proc. American Control Conf.*, 2017, pp. 5263–5268.
- [8] D. Schrank, B. Eisele, T. Lomax, and J. Bak, “2015 urban mobility scorecard,” Texas A&M Transportation Institute, 2015.
- [9] M. Wu, J. Liu, and S. Amin, “Informational aspects in a class of bayesian congestion games,” in *Proc. American Control Conf.*, 2017, pp. 3650–3657.
- [10] F. Farokhi, W. Krichene, A. M. Bayen, and K. H. Johansson, “A heterogeneous routing game,” in *Proc. 51st Annual Allerton Conf. Communication, Control, and Computing*, 2013, pp. 448–455.
- [11] L. Fleischer, K. Jain, and M. Mahdian, “Tolls for heterogeneous selfish users in multicommodity networks and generalized congestion games,” in *Proc. 45th Symp. Found. Comput. Sci.*, 2004, pp. 277–285.
- [12] T. Roughgarden, “Routing games,” in *Algorithmic game theory*. Cambridge University Press New York, 2007, ch. 18, pp. 459–484.
- [13] T. Roughgarden and É. Tardos, “How bad is selfish routing?” *J. ACM*, vol. 49, no. 2, pp. 236–259, 2002.
- [14] 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.
- [15] E. Nikolova and N. E. S. Moses, “The burden of risk aversion in mean-risk selfish routing,” in *Proc. Sixteenth ACM Conf. Economics and Computation*, 2015, pp. 489–506.
- [16] A. Tversky and D. Kahneman, “Judgment under uncertainty: Heuristics and biases,” in *Utility, probability, and human decision making*. Springer, 1975, pp. 141–162.
- [17] 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.
- [18] R. Meir and D. C. Parkes, “Congestion games with distance-based strict uncertainty,” in *Proc. 29th AAAI Conference on Artificial Intelligence*, 2015, pp. 986–992.
- [19] S. Sekar, L. Zheng, L. J. Ratliff, and B. Zhang, “Uncertainty in multi-commodity routing networks: When does it help?” *arXiv:1709.08441 [cs.GT]*, 2017.
- [20] D. Monderer and L. S. Shapley, “Potential games,” *Games and economic behavior*, vol. 14, no. 1, pp. 124–143, 1996.
- [21] 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 Conference on Decision and Control*, 2016, pp. 7605–7610.
- [22] C. Dowling, T. Fiez, L. Ratliff, and B. Zhang, “How much urban traffic is searching for parking?” *arXiv:1702.06156*, 2017.
- [23] D. Shoup and H. Campbell, “Gone parkin’,” *The New York Times*, vol. 29, 2007.
- [24] 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.
- [25] D. Acemoglu and A. E. Ozdaglar, “Competition in parallel-serial networks,” *IEEE J. Selected Areas in Communications*, vol. 25, no. 6, pp. 1180–1192, 2007.
- [26] I. Milchtaich, “Network topology and the efficiency of equilibrium,” *Games and Economic Behavior*, vol. 57, no. 2, pp. 321–346, 2006.
- [27] P. Auer, N. Cesa-Bianchi, and P. Fischer, “Finite-time analysis of the multiarmed bandit problem,” *Machine Learning*, vol. 47, no. 2-3, pp. 235–256, 2002.