

When Congestion Games Meet Mobile Crowdsourcing: Selective Information Disclosure

Hongbo Li, Lingjie Duan

Singapore University of Technology and Design
hongbo_li@mymail.sutd.edu.sg, lingjie_duan@sutd.edu.sg

Abstract

In congestion games, users make myopic routing decisions to jam each other, and the social planner with the full information designs mechanisms on information or payment side to regulate. However, it is difficult to obtain time-varying traffic conditions, and emerging crowdsourcing platforms (e.g., Waze and Google Maps) provide a convenient way for mobile users travelling on the paths to learn and share the traffic conditions over time. When congestion games meet mobile crowdsourcing, it is critical to incentive selfish users to change their myopic routing policy and reach the best exploitation-exploration trade-off. By considering a simple but fundamental parallel routing network with one deterministic path and multiple stochastic paths for atomic users, we prove that the myopic routing policy's price of anarchy (PoA) can be arbitrarily large as the discount factor approaches 1. To remedy such huge efficiency loss, we propose a selective information disclosure (SID) mechanism: we only reveal the latest traffic information to users when they intend to over-explore the stochastic paths, while hiding such information when they want to under-explore. We prove that our mechanism reduces PoA to less than 2. Besides the worst-case performance, we further examine our mechanism's average-case performance by using extensive simulations.

Introduction

In transportation networks of limited bandwidth, mobile users are selfish to choose routing decisions myopically and aim to minimize their own travel costs on the way. Traditional congestion games study such selfish routing to understand the efficiency loss using the concept of the price of anarchy (PoA) (Roughgarden and Tardos 2002; Cominetti et al. 2019; Bilò and Vinci 2020; Hao and Michini 2022). To regulate atomic or non-atomic users' selfish routing and reduce social cost, various incentive mechanisms are designed by using monetary payments to penalize users travelling on undesired paths (Brown and Marden 2017; Ferguson, Brown, and Marden 2022; Li and Duan 2023). As it may be difficult to implement such payments on users, non-monetary mechanisms are also designed to provide information restriction on selfish users to change their routing decisions to approach the social optimum (Tavafoghi and

Teneketzis 2017; Sekar et al. 2019; Castiglioni et al. 2021). However, these works largely assume that the social planner has full information of all traffic conditions, and limit attentions to an one-shot static scenario to regulate.

In common practice, the traffic information dynamically changes over time and is difficult to predict in advance (Nikolova and Stier-Moses 2011). To obtain such time-varying information, emerging traffic navigation platforms (e.g., Waze and Google Maps) crowdsource mobile users to learn and share their observed traffic conditions on the way (Vasserman, Feldman, and Hassidim 2015; Zhang et al. 2018). However, such platforms make all information public, and current users still make selfish routing decisions to the path with shortest travel latency, instead of choosing diverse paths to learn more information for future users. As a stochastic path's traffic condition alternates between congestion states over time, the platforms may miss enough exploration to reduce the social cost.


There are some recent works studying information sharing among users in a dynamic scenario. For example, Meigs, Parise, and Ozdaglar (2017) and Wu and Amin (2019) make use of former users' observation to help learn the future travel latency and converge to the Wardrop Equilibrium under full information. Similarly, Vu, Antonakopoulos, and Mertikopoulos (2021) design an adaptive information learning framework to accelerate convergence rates to Wardrop equilibrium for stochastic congestion games. However, these works cater to users' selfish interests and do not consider mechanism design to motivate users to reach social optimum. To study the social cost minimization, multi-armed bandit (MAB) problems are also formulated to derive the optimal exploitation-exploration policy among multiple stochastic arms (paths) (Gittins, Glazebrook, and Weber 2011; Krishnasamy et al. 2021). Recently, Bozorgchenani et al. (2022) apply MAB models to predict the network congestion in a fast changing vehicular environment. However, all of these MAB works strongly assume that users upon arrival always follow the social planner's recommendations and overlook users' deviation to selfish routing.

When congestion games meet mobile crowdsourcing, how to incentive selfish users to listen to the social planner's optimal recommendations is our key question in this paper. As traffic navigation platforms seldom charge users, we target at non-monetary mechanism design which satisfies


budget balance in nature. Yet we cannot borrow those information mechanisms from the literature in mobile crowdsourcing, as their considered traffic information is exogenous and does not depend on users' routing decisions (Kremer, Mansour, and Perry 2014; Papanastasiou, Bimpikis, and Savva 2018; Li, Courcoubetis, and Duan 2017, 2019). For example, Li, Courcoubetis, and Duan (2019) consider a simple two-path transportation network, one with deterministic travel cost and the other alternates over time between a high and a low stochastic cost states due to external weather conditions. In their finding, a selfish user is always found to under-explore the stochastic path to learn latest information there for future users. In our congestion problem, however, a user will add himself to the traffic flow and change the congestion information in the loop. Thus, we imagine users may not only under-explore but also over-explore stochastic paths over time. Furthermore, since the congestion information (though random) depends on users' routing decisions, it is easier for a user to reverse-engineer the system states based on the platform's optimal recommendation. In consequence, the prior information hiding mechanisms (Tavafoghi and Teneketzis 2017; Li, Courcoubetis, and Duan 2019; Zhu and Savla 2022) become no longer efficient.

We summarize our key novelty and main contributions in this paper as follows.

- *Mechanism design when congestion games meet mobile crowdsourcing:* To our best knowledge, this paper is the first to regulate atomic users' routing over time to reach the best exploitation-exploration trade-off by providing incentives. In Section 2, we model a dynamic congestion game in a transportation network of one deterministic path and multiple stochastic paths to learn by users themselves. When congestion games meet mobile crowdsourcing, our study extends the traditional congestion games fundamentally to create positive information learning generated by users themselves.
- *POMDP formulation and PoA analysis:* In Section 3, we formulate users' dynamic routing problems using the partially observable Markov decision process (POMDP) according to hazard beliefs of risky paths. Then in Section 4, we analyze both myopic and socially optimal policies to learn stochastic paths' states, and prove that the myopic policy misses both exploration (when strong hazard belief) and exploitation (when weak hazard belief) as compared to the social optimum. Accordingly, we prove that the resultant price of anarchy (PoA) is larger than $\frac{1}{1-\rho}$, which can be arbitrarily large as discount factor $\rho \rightarrow 1$.
- *Selective information disclosure (SID) mechanism to remedy efficiency loss:* In Section 5, we first prove that the prior information hiding mechanism in congestion games makes PoA infinite in our problem. Alternatively, we propose a selective information disclosure mechanism: we only reveal the latest traffic information to users when they over-explore the stochastic paths, while hiding such information when they under-explore. We prove that our mechanism reduces PoA to be less than $\frac{1}{1-\frac{1}{2}}$, which is no larger than 2. Besides the worst-case performance,



(a) A typical parallel transportation network.



(b) The partially observable Markov chain for modelling $\alpha_i(t)$ dynamics of stochastic path i .

Figure 1: At the beginning of each time slot $t \in \{1, 2, \dots\}$, a user arrives to choose a path among $N + 1$ paths according to travel latency $\ell_i(t)$ of each path $i \in \{0, 1, \dots, N\}$ in Figure 1a. Path 0 is a safe route and its latency has a fixed correlation coefficient $\alpha \in (0, 1)$ to change from the last round. Yet any risky path $i \in \{1, \dots, N\}$ has a stochastic correlation coefficient $\alpha_i(t)$, which alternates according to the partially observable Markov chain in Figure 1b.

we further examine our mechanism's average-case performance by using extensive simulations.

Due to the page limit, we move the lengthy proofs of the paper to our supplementary material and online technical report (Li and Duan 2022), and also provide code here.¹

System Model

As illustrated in Figure 1a, we consider a dynamic congestion game lasting for infinite discrete time horizon. At the beginning of each time epoch $t \in \{1, 2, \dots\}$, an atomic user arrives to travel on one out of $N + 1$ paths from origin O to destination D . Similar to the existing literature of congestion games (e.g., Kremer, Mansour, and Perry 2014; Tavafoghi and Teneketzis 2017; Li, Courcoubetis, and Duan 2019), in Figure 1a the top path 0 as a safe route has a fixed traffic condition α that is known to the public, while the other N bottom paths are risky/stochastic to alternate between traffic conditions α_L and α_H over time. Thus, the crowdsourcing platform expects users to travel to risky paths from time to time to learn the actual traffic information and plan better routing advisory for future users.

In the following, we first introduce the dynamic congestion model for the transportation network, and then introduce the users' information learning and sharing in the crowdsourcing platform.

Dynamic Congestion Model

Let $\ell_i(t)$ denote the travel latency of path $i \in \{0, 1, \dots, N\}$ estimated by a new user arrival on path i at the beginning of each time slot $t \in \{1, 2, \dots\}$. Then the current user decides the best path $i \in \{0, 1, \dots, N\}$ to

¹<https://github.com/redglassli/Congestion-games-SID>

choose by comparing the travel latencies among all paths. We denote a user's routing choice at time t as $\pi(t) \in \{0, 1, \dots, N\}$. For this user, he predicts $\ell_i(t)$ based on the latest latency $\ell_i(t-1)$ and the last user's decision $\pi(t-1)$.

Some existing literature of delay pattern estimation (e.g., Ban et al. 2009; Alam, Farid, and Rossetti 2019) assumes that $\ell_i(t+1)$ is linearly dependent on $\ell_i(t)$. Thus, for safe path 0 with the fixed traffic condition, its next travel latency $\ell_0(t+1)$ changes from $\ell_0(t)$ with constant correlation coefficient α . Here $\alpha \in (0, 1)$ measures the leftover flow to be serviced over time. Yet, if the current atomic user chooses this path (i.e., $\pi(t) = 0$), he will introduce an addition $\Delta\ell$ to the next travel latency $\ell_0(t+1)$, i.e.,

$$\ell_0(t+1) = \begin{cases} \alpha\ell_0(t) + \Delta\ell, & \text{if } \pi(t) = 0, \\ \alpha\ell_0(t), & \text{if } \pi(t) \neq 0. \end{cases} \quad (1)$$

Differently, on any risky path $i \in \{1, \dots, N\}$, its correlation coefficient $\alpha_i(t)$ in this round is stochastic due to the random traffic condition (e.g., accident and weather change) at each time slot t . Similar to the congestion game literature (Meigs, Parise, and Ozdaglar 2017), we suppose $\alpha_i(t)$ alternates between low coefficient state $\alpha_L \in [0, 1)$ and high state $\alpha_H \in [1, +\infty)$ below:

$$\alpha_i(t) = \begin{cases} \alpha_L, & \text{if path } i \text{ has a good traffic condition at } t, \\ \alpha_H, & \text{if path } i \text{ has a bad traffic condition at } t. \end{cases}$$

Note that we consider $\alpha_L < \alpha < \alpha_H$ such that each path can be chosen by users and we also allow jamming on risky paths with $\alpha_H \geq 1$. The transition of $\alpha_i(t)$ over time is modeled as the partially observable Markov chain in Figure 1b, where the self-transition probabilities are q_{LL} and q_{HH} with $q_{LL} + q_{LH} = 1$ and $q_{HH} + q_{HL} = 1$. Then the travel latency $\ell_i(t+1)$ of any risky path $i \in \{1, \dots, N\}$ is estimated as

$$\ell_i(t+1) = \begin{cases} \alpha_i(t)\ell_i(t) + \Delta\ell, & \text{if } \pi(t) = i, \\ \alpha_i(t)\ell_i(t), & \text{if } \pi(t) \neq i. \end{cases} \quad (2)$$

To obtain this $\alpha_i(t)$ realization for better estimating future $\ell_i(t+1)$ in (2), the platform may expect current user to travel on this risky path i to learn and share his observation.

Crowdsourcing Model for Learning

After choosing a risky path $i \in \{1, \dots, N\}$ to travel, in practice a user may not obtain the whole path information when making local observation and reporting to the crowdsourcing platform. Two different users travelling on the same path may have different experiences. Similar to Li, Courcoubetis, and Duan (2019), we model $\alpha_i(t)$ dynamics as the partially observable two-state Markov chain in Figure 1b from the user point of view. We define a random observation set $\mathbf{y}(t) = \{y_1(t), \dots, y_N(t)\}$ for N risky paths, where $y_i(t) \in \{0, 1, \emptyset\}$ denotes the traffic condition of path i as observed by the current user there during time slot t . More specifically, $y_i(t) = 1$ tells that the current user at time t observes a hazard (e.g., 'black ice' segments, poor visibility, jamming) after choosing path $\pi(t) = i$. $y_i(t) = 0$ tells that the user does not observe any hazard on path i . Finally, $y_i(t) = \emptyset$ tells that this user travels on another path with $\pi(t) \neq i$, without making any observation of path i .

Given $\pi(t) = i$, the chance for the user to observe $y_i(t) = 1$ or 0 depends on the random correlation coefficient $\alpha_i(t)$. Under the correlation state $\alpha_i(t) = \alpha_H$ or α_L at time t , we respectively denote the probabilities for the user to observe a hazard as:

$$\begin{aligned} p_H &= \Pr(y_i(t) = 1 | \alpha_i(t) = \alpha_H), \\ p_L &= \Pr(y_i(t) = 1 | \alpha_i(t) = \alpha_L). \end{aligned} \quad (3)$$

Note that $p_L < p_H$ because a risky path in bad traffic condition ($\alpha_i(t) = \alpha_H$) has a larger probability for the user to observe a hazard (i.e., $y_i(t) = 1$). Even if path i has good traffic condition ($\alpha_i(t) = \alpha_L$), it is not entirely hazard free and there is still some probability p_L to face a hazard.

As users keep learning and sharing traffic conditions with the crowdsourcing platform, the historical data of their observations ($\mathbf{y}(1), \dots, \mathbf{y}(t-1)$) and routing decisions ($\pi(1), \dots, \pi(t-1)$) before time t keep growing in the time horizon. To simplify the ever-growing history set, we equivalently translate these historical observations into a hazard belief $x_i(t)$ for seeing bad traffic condition $\alpha_i(t) = \alpha_H$ at time t , by using the Bayesian inference:

$$x_i(t) = \Pr(\alpha_i(t) = \alpha_H | x_i(t-1), \pi(t-1), \mathbf{y}(t-1)). \quad (4)$$

Given the prior probability $x_i(t)$, the platform will further update it to a posterior probability $x'_i(t)$ after a new user with routing decision $\pi(t)$ shares his observation $y_i(t)$ during the time slot:

$$x'_i(t) = \Pr(\alpha_i(t) = \alpha_H | x_i(t), \pi(t), \mathbf{y}(t)). \quad (5)$$

Below, we explain the dynamics of our information learning model.

- At the beginning of time slot t , the platform publishes any risky path i 's hazard belief $x_i(t)$ in (4) about coefficient $\alpha_i(t)$ and the latest expected latency $\mathbb{E}[\ell_i(t) | x_i(t-1), y_i(t-1)]$ to summarize observation history ($\mathbf{y}(1), \dots, \mathbf{y}(t-1)$) till $t-1$.
- During time slot t , a user arrives to choose a path (e.g., $\pi(t) = i$) to travel and reports his following observation $y_i(t)$. Then the platform updates the posterior probability $x'_i(t)$, conditioned on the new observation $y_i(t)$ and the prior probability $x_i(t)$ in (5). For example, if $y_i(t) = 0$, by Bayes' Theorem, $x'_i(t)$ for the correlation coefficient $\alpha_i(t) = \alpha_H$ is

$$\begin{aligned} x'_i(t) &= \Pr(\alpha_i(t) = \alpha_H | x_i(t), \pi(t) = i, y_i(t) = 0) \quad (6) \\ &= \frac{x_i(t)(1 - p_H)}{x_i(t)(1 - p_H) + (1 - x_i(t))(1 - p_L)}. \end{aligned}$$

Similarly, if $y_i(t) = 1$, we have

$$x'_i(t) = \frac{x_i(t)p_H}{x_i(t)p_H + (1 - x_i(t))p_L}. \quad (7)$$

Besides this traveled path i , for any other path $j \in \{1, \dots, N\}$ with $y_j(t) = \emptyset$, we keep $x'_j(t) = x_j(t)$ as there is no added observation to this path at t .

- At the end of this time slot, the platform estimates the posterior correlation coefficient:

$$\begin{aligned}\mathbb{E}[\alpha_i(t)|x'_i(t)] &= \mathbb{E}[\alpha_i(t)|x_i(t), y_i(t)] \\ &= x'_i(t)\alpha_H + (1 - x'_i(t))\alpha_L.\end{aligned}\quad (8)$$

By combining (8) with (2), we can obtain the expected travel latency on stochastic path i for time $t + 1$ as

$$\begin{aligned}\mathbb{E}[\ell_i(t+1)|x_i(t), y_i(t)] &= \\ \begin{cases} \mathbb{E}[\alpha_i(t)|x'_i(t)]\mathbb{E}[\ell_i(t)|x_i(t-1), y_i(t-1)] + \Delta\ell, & \text{if } \pi(t) = i, \\ \mathbb{E}[\alpha_i(t)|x'_i(t)]\mathbb{E}[\ell_i(t)|x_i(t-1), y_i(t-1)], & \text{if } \pi(t) \neq i. \end{cases}\end{aligned}\quad (9)$$

Based on the partially observable Markov chain in Figure 1b, the platform updates each path i 's hazard belief from $x'_i(t)$ to $x_i(t+1)$ below:

$$x_i(t+1) = x'_i(t)q_{HH} + (1 - x'_i(t))q_{LH}. \quad (10)$$

Finally, the new time slot $t + 1$ begins and repeats the process since above.

POMDP Problem Formulations for Myopic and Socially Optimal Policies

Based on the dynamic congestion and crowdsourcing models in the last section, we formulate the problems of myopic policy (for guiding myopic users' selfish routing) and the socially optimal policy (for the social planner/platform's best path advisory), respectively.

Problem Formulation for Myopic Policy

In this subsection, we consider the myopic policy (e.g. used by Waze and Google Maps) that the selfish users will naturally follow. First, we summarize the dynamics of expected travel latencies among all $N + 1$ paths and the hazard beliefs of N stochastic paths into vectors:

$$\begin{aligned}\mathbf{L}(t) &= \{\ell_0(t), \mathbb{E}[\ell_1(t)|x_i(t-1), y_i(t-1)], \dots, \\ &\quad \mathbb{E}[\ell_N(t)|x_N(t-1), y_N(t-1)]\}, \\ \mathbf{x}(t) &= \{x_1(t), \dots, x_N(t)\},\end{aligned}\quad (11)$$

which are obtained based on (9) and (10). For a user arrival at time t , the platform provides him with $\mathbf{L}(t)$ and $\mathbf{x}(t)$ to help make his routing decision. We define the best stochastic path $\hat{i}(t)$ to be the one out of N risky paths to provide the shortest expected travel latency at time t below:

$$\hat{i}(t) = \arg \min_{i \in \{1, \dots, N\}} \mathbb{E}[\ell_i(t)|x_i(t-1), y_i(t-1)]. \quad (12)$$

The selfish user will only choose between safe path 0 and this path $\hat{i}(t)$ to minimize his own travel latency.

We formulate this problem as a POMDP, where the time correlation state $\alpha_i(t)$ of each stochastic path i is partially observable to users in Figure 1b. Thus, the states here are $\mathbf{L}(t)$ and $\mathbf{x}(t)$ in (11). Under the myopic policy, define $C^{(m)}(\mathbf{L}(t), \mathbf{x}(t))$ to be the long-term discounted cost function with discount factor $\rho < 1$ to include social cost of all users since t . Then its dynamics per user arrival has the following two cases. If $\mathbb{E}[\ell_{\hat{i}(t)}(t)|x_{\hat{i}(t)}(t-1), y_{\hat{i}(t)}(t-1)] \geq$

$\ell_0(t)$, a selfish user will choose path 0 and add $\Delta\ell$ to path 0 to have latency $\ell_0(t+1) = \alpha\ell_0(t) + \Delta\ell$ in (1). Since no user enters stochastic path i , there is no information reporting (i.e., $y_i(t) = \emptyset$) and $x'_i(t)$ in (5) equals $x_i(t)$ in (4) for updating $x_i(t+1)$ in (10). The expected travel latency of stochastic path i in the next time slot is updated to $\mathbb{E}[\ell_i(t+1)|x_i(t), y_i(t) = \emptyset]$ according to (9). In consequence, the travel latency and hazard belief sets at the next time slot $t + 1$ are updated to

$$\begin{aligned}\mathbf{L}(t+1) &= \{\alpha\ell_0(t) + \Delta\ell, \mathbb{E}[\ell_1(t+1)|x_1(t), y_1(t) = \emptyset], \\ &\quad \dots, \mathbb{E}[\ell_N(t+1)|x_N(t), y_N(t) = \emptyset]\}, \\ \mathbf{x}(t+1) &= \{x_1(t+1), \dots, x_N(t+1)\}.\end{aligned}\quad (13)$$

Then the cost-to-go $Q_0^{(m)}(t+1)$ since the next user is

$$Q_0^{(m)}(t+1) = C^{(m)}(\mathbf{L}(t+1), \mathbf{x}(t+1)|y_{i(t)}(t) = \emptyset). \quad (14)$$

If $\mathbb{E}[\ell_{\hat{i}(t)}(t)|x_{\hat{i}(t)}(t-1), y_{\hat{i}(t)}(t-1)] < \ell_0(t)$, the user will choose the best stochastic path $\hat{i}(t)$ in (12). Then the platform updates the expected travel latency on path $\hat{i}(t)$ to $\mathbb{E}[\ell_{\hat{i}(t)}(t)|x_{\hat{i}(t)}(t), y_{\hat{i}(t)}(t)]$ in (9), depending on whether $y_{\hat{i}(t)}(t) = 1$ or 0. Note that according to (3),

$$\Pr(y_{\hat{i}(t)}(t) = 1) = (1 - x_{\hat{i}(t)}(t))p_L + x_{\hat{i}(t)}(t)p_H. \quad (15)$$

While path 0's latency in next time changes to $\alpha\ell_0(t)$, and path $i \neq \hat{i}(t)$ has no exploration and its expected latency at time $t + 1$ becomes $\mathbb{E}[\ell_i(t+1)|x_i(t), y_i(t) = \emptyset]$. Then the expected cost-to-go since the next user in this case is

$$\begin{aligned}Q_{\hat{i}(t)}^{(m)}(t+1) &= \\ \Pr(y_{\hat{i}(t)}(t) = 1)C^{(m)}(\mathbf{L}(t+1), \mathbf{x}(t+1)|y_{\hat{i}(t)}(t) = 1) \\ + \Pr(y_{\hat{i}(t)}(t) = 0)C^{(m)}(\mathbf{L}(t+1), \mathbf{x}(t+1)|y_{\hat{i}(t)}(t) = 0).\end{aligned}\quad (16)$$

To combine (14) and (16), we formulate the ρ -discounted long-term cost function since time t under myopic policy as

$$\begin{aligned}C^{(m)}(\mathbf{L}(t), \mathbf{x}(t)) &= \\ \begin{cases} \ell_0(t) + \rho Q_0^{(m)}(t+1), & \text{if } \mathbb{E}[\ell_{\hat{i}(t)}(t)|x_{\hat{i}(t)}(t-1), y_{\hat{i}(t)}(t-1)] \geq \ell_0(t), \\ \mathbb{E}[\ell_{\hat{i}(t)}(t)|x_{\hat{i}(t)}(t-1), y_{\hat{i}(t)}(t-1)] + \rho Q_{\hat{i}(t)}^{(m)}(t+1), & \text{otherwise.} \end{cases}\end{aligned}\quad (17)$$

A selfish user is not willing to explore any stochastic path i with longer expected travel latency, and the next arrival may not know the fresh congestion information. On the other hand, selfish users may keep choosing the path with the shortest latency and jamming this path for future users.

Socially Optimal Policy Problem Formulation

Different from the myopic policy that focuses on the one-shot to minimize the current user's immediate travel cost, the goal of the social optimum is to find optimal policy $\pi^*(t)$ at any time t to minimize the expected social cost over an infinite time horizon.

Denote the long-term ρ -discounted cost function by $C^*(\mathbf{L}(t), \mathbf{x}(t))$ under the socially optimal policy. The optimal policy depends on which path choice yields the minimal long-term social cost. If the platform asks the current user to choose path 0, this user will bear cost $\ell_0(t)$ to travel this path. Due to no information observation (i.e., $\mathbf{y}(t) = \emptyset$), the cost-to-go $Q_0^*(t+1)$ from the next user can be similarly determined as (14) with $\mathbf{L}(t+1)$ and $\mathbf{x}(t+1)$ in (13).

If the platform asks the user to explore a stochastic path i , this choice is not necessarily path $\hat{i}(t)$ in (12). Then the platform updates $\mathbf{x}(t+1)$, depending on whether the user's observation on this path is $y_i(t) = 1$ or $y_i(t) = 0$. Similar to (16), the optimal expected cost function from next user is denoted as $Q_i^*(t+1)$. Then we are ready to formulate the social cost function under socially optimal policy below:

$$C^*(\mathbf{L}(t), \mathbf{x}(t)) = \min_{i \in \{1, \dots, N\}} \{ \ell_0(t) + \rho Q_0^*(t+1), \ell_i(t) + \rho Q_i^*(t+1) \}. \quad (18)$$

Problem (18) is non-convex and its analysis will cause the curse of dimensionality in the infinite time horizon (Bellman 1966). Though it is difficult to solve, we still analytically compare the two policies by their structural results below.

Comparing Myopic Policy to Social Optimum for PoA Analysis

In this section, we first prove that both myopic and socially optimal policies to explore stochastic paths are of threshold-type with respect to expected travel latency. Then we show that the myopic policy may both under-explore and over-explore risky paths.² Finally, we prove that the myopic policy can perform arbitrarily bad.

Lemma 1. *The cost functions $C^{(m)}(\mathbf{L}(t), \mathbf{x}(t))$ in (17) and $C^*(\mathbf{L}(t), \mathbf{x}(t))$ in (18) under both policies increase with any path's expected latency $\mathbb{E}[\ell_i(t)|x_i(t-1), y_i(t-1)]$ in $\mathbf{L}(t)$ and $\mathbf{x}(t)$ in (11).*

With this monotonicity result, we next prove that both policies are of threshold-type.

Proposition 1. *Provided with $\mathbf{L}(t)$ and $\mathbf{x}(t)$ in (11), the user arrival at time t under the myopic policy keeps staying with path 0, until the expected latency of the best stochastic path $\hat{i}(t)$ in (12) reduces to be smaller than the following threshold:*

$$\ell^{(m)}(t) = \ell_0(t). \quad (19)$$

Similarly, the socially optimal policy will choose stochastic path i instead of path 0 if $\mathbb{E}[\ell_i(t)|x_i(t-1), y_i(t-1)]$ is less than the following threshold:

$$\ell_i^*(t) = \arg \max_z \{ z | z \leq \rho Q_i^*(t+1) - \rho Q_0^*(t+1) - \ell_0(t) \}, \quad (20)$$

which increases with hazard belief $x_i(t)$ of risky path i .

Let $\pi^{(m)}(t)$ and $\pi^*(t)$ denote the routing decisions at time t under myopic and socially optimal policies, respectively. We next compare the exploration thresholds $\ell^{(m)}(t)$ and $\ell_i^*(t)$ as well as their associated social costs.

²Over/under exploration means that myopic policy will choose risky path i more/less often than what the social optimum suggests.

Lemma 2. *If $\pi^{(m)}(t) \neq \pi^*(t)$, then the expected travel latencies on these two chosen paths by the two policies satisfy*

$$\begin{aligned} & \mathbb{E}[\ell_{\pi^*(t)}(t)|\mathbf{x}(t-1), \mathbf{y}(t-1)] \\ & \leq \frac{1}{1-\rho} \mathbb{E}[\ell_{\pi^{(m)}(t)}(t)|\mathbf{x}(t-1), \mathbf{y}(t-1)]. \end{aligned} \quad (21)$$

Intuitively, if the current travel latencies on different paths obviously differ, the two policies tend to make the same routing decision. (21) is more likely to hold for large ρ .

Next, we define the stationary belief $x_i(t)$ of high hazard state α_H as \bar{x} , and we provide it below by using steady-state analysis of Figure 1b:

$$\bar{x} = \frac{1 - q_{LL}}{2 - q_{LL} - q_{HH}}. \quad (22)$$

Based on Proposition 1 and Lemma 2, we analytically compare the two policies below.

Proposition 2. *There exists a belief threshold x^{th} satisfying*

$$\min \left\{ \frac{\alpha - \alpha_L}{\alpha_H - \alpha_L}, \bar{x} \right\} \leq x^{th} \leq \max \left\{ \frac{\alpha - \alpha_L}{\alpha_H - \alpha_L}, \bar{x} \right\}. \quad (23)$$

As compared to socially optimal policy, if risky path $i \in \{1, \dots, N\}$ has weak hazard belief $x_i(t) < x^{th}$, myopic users will only over-explore this path with $\ell^{(m)}(t) \geq \ell_i^(t)$. If strong hazard belief with $x_i(t) > x^{th}$, myopic users will only under-explore this path with $\ell^{(m)}(t) \leq \ell_i^*(t)$.*

Here $\frac{\alpha - \alpha_L}{\alpha_H - \alpha_L}$ in (23) is derived by equating path i 's expected coefficient $\mathbb{E}[\alpha_i(t)|x_i'(t)]$ in (8) to path 0's α . Proposition 2 tells that the myopic policy misses both exploitation and exploration over time. If the hazard belief on path $i \in \{1, \dots, N\}$ is weak (i.e., $x_i(t) < x^{th}$), myopic users choose stochastic path i without considering the congestion to future others on the same path. While the socially optimal policy may still recommend users to safe path 0 to further reduce the congestion cost on path i for the following user. On the other hand, if $x_i(t) > x^{th}$, the socially optimal policy may still want to explore path i to exploit hazard-free state α_L on this path for future use. This result is also consistent with $\ell_i^*(t)$'s monotonicity in $x_i(t)$ in Proposition 1.

In Figure 2, we simulate Figure 1a using a simple two-path transportation network with $N = 1$. We plot exploration thresholds $\ell^{(m)}(t)$ in (19) under myopic policy and optimal $\ell_1^*(t)$ in (20) versus hazard belief $x_1(t)$ of path 1. These two thresholds are very different in Figure 2. Given the belief threshold $x^{th} = 0.45$ here, if the hazard belief $x_1(t) < x^{th}$, we have the myopic exploration threshold $\ell^{(m)}(t) > \ell_1^*(t)$ to over-explore stochastic path. If $x_1(t) > x^{th}$, the myopic exploration threshold satisfies $\ell^{(m)}(t) < \ell_1^*(t)$ to over-explore. This result is consistent with Proposition 2.

After comparing the two policies' thresholds, we are ready to further examine their performance gap. Following Koutsoupias and Papadimitriou (1999), we define the price of anarchy (PoA) to be the maximum ratio between the social cost under myopic policy in (17) and the minimal social cost in (18), by searching all possible system parameters:

$$\text{PoA}^{(m)} = \frac{C^{(m)}(\mathbf{L}(t), \mathbf{x}(t))}{\max_{\alpha, \alpha_H, \alpha_L, q_{LL}, q_{HH}, \mathbf{x}(t), \mathbf{L}(t), \Delta \ell, p_H, p_L} C^*(\mathbf{L}(t), \mathbf{x}(t))}, \quad (24)$$

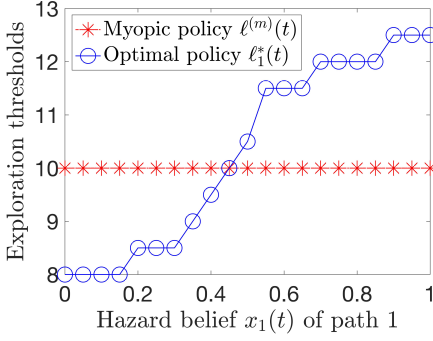


Figure 2: The socially optimal policy’s exploration threshold $l_1^*(t)$ and myopic policy’s threshold $l^{(m)}(t)$ versus hazard belief $x_1(t)$ in a two-path transportation network with $N = 1$. We set $\alpha = 0.6, \alpha_H = 1.2, \alpha_L = 0.2, q_{LL} = 0.5, q_{HH} = 0.5, \Delta\ell = 2, p_H = 0.8, p_L = 0.3, \ell_0(t) = 10$ and $x_1(t) = 0.1$ at current time t .

which is obviously larger than 1. Then we present the lower bound of PoA in the following proposition.

Proposition 3. *As compared to the social optimum in (18), the myopic policy in (17) achieves $\text{PoA}^{(m)} \geq \frac{1}{1-\rho}$, which can be arbitrarily large for discount factor $\rho \rightarrow 1$.*

In this worst-case PoA analysis, we consider a two-path network example, where the myopic policy always chooses safe path 0 but the socially optimal policy frequently explores stochastic path 1 to learn α_L . Here we initially set $\ell_0(0) = \frac{\Delta\ell}{1-\alpha}$ such that the travel latency $\ell_0(1) = \alpha\ell_0(0) + \Delta\ell$ in (1) equals $\frac{\Delta\ell}{1-\alpha}$ all the time for myopic users. Without myopic users’ routing on stochastic path 1, we also keep the expected travel latency on stochastic path 1 unchanged, by setting $x_1(0) = \bar{x}$ in (22) and $\mathbb{E}[\alpha_1(0)|x_1(0) = \bar{x}] = 1$ in (8). Then a myopic user at any time t will never explore the stochastic path 1 given $\ell_1(t) = \ell_0(t)$, resulting in the social cost to be $\frac{\ell_0(0)}{1-\rho}$ in the infinite time horizon. However, the socially optimal policy frequently asks a user arrival to explore path 1 to learn a good condition ($\alpha_L = 0$) for following users. We make $q_{LL} \rightarrow 1$ to maximally reduce the travel latency of path 1, and the optimal social cost is thus no more than $\ell_1(0) + \frac{\rho}{1-\rho}\Delta\ell$. Letting $\frac{\Delta\ell}{\ell_0(0)} \rightarrow 0$, we obtain $\text{PoA}^{(m)} \geq \frac{1}{1-\rho}$.

By Proposition 3, the myopic policy performs worse, as discount factor ρ increases and future costs become more important. As $\rho \rightarrow 1$, PoA approaches infinity and the learning efficiency in the crowdsourcing platform becomes arbitrarily bad to opportunistically reduce the congestion. Thus, it is critical to design efficient incentive mechanism to greatly reduce the social cost.

Selective Information Disclosure

To motivate a selfish user to follow the optimal path advisory at any time, we need to design a non-monetary information mechanism, which naturally satisfies budget balance and is easy to implement without enforcing monetary

payments. Our key idea is to selectively disclose the latest expected travel latency set $\mathbf{L}(t)$ of all paths, depending on a myopic user’s intention to over- or under-explore stochastic paths at time t . To avoid users from perfectly inferring $\mathbf{L}(t)$, we purposely hide the latest hazard belief set $\mathbf{x}(t)$, routing history $(\pi(1), \dots, \pi(t-1))$, and past traffic observation set $(\mathbf{y}(1), \dots, \mathbf{y}(t-1))$, but always provide socially optimal path recommendation $\pi^*(t)$ to any user. Provided with selective information disclosure, we allow sophisticated users to reverse-engineer the path latency distribution and make selfish routing under our mechanism.

Before formally introducing our selective information disclosure in Definition 1, we first consider an information hiding policy $\pi^\theta(t)$ as a benchmark. Similar information hiding mechanisms were proposed and studied in the literature (e.g., Tavaafoghi and Teneketzis 2017 and Li, Courcoubetis, and Duan 2019). In this benchmark mechanism, the user without any information believes that the expected hazard belief $x_i(t)$ of any stochastic path $i \in \{1, \dots, N\}$ has converged to its stationary hazard belief \bar{x} in (22). Then he can only decide his routing policy $\pi^\theta(t)$ by comparing α of safe path 0 to $\mathbb{E}[\alpha_i(t)|\bar{x}]$ in (8) of any path i .

Proposition 4. *Given no information from the platform, a user arrival at time t uses the following routing policy:*

$$\pi^\theta(t) = \begin{cases} 0, & \text{if } \bar{x} \geq \frac{\alpha - \alpha_L}{\alpha_H - \alpha_L}, \\ i \text{ w/ probability } \frac{1}{N}, & \text{if } \bar{x} < \frac{\alpha - \alpha_L}{\alpha_H - \alpha_L}, \end{cases} \quad (25)$$

where $i \in \{1, \dots, N\}$. This hiding policy leads to $\text{PoA}^\theta = \infty$, regardless of discount factor ρ .

Even if we still recommend optimal routing $\pi^*(t)$ in (18), a selfish user sticks to some risky path i given low hazard belief $\bar{x} < \frac{\alpha - \alpha_L}{\alpha_H - \alpha_L}$. This hiding policy can differ a lot from the socially optimal policy in (18) since users cannot observe the latest travel latencies. To tell the $\text{PoA}^\theta = \infty$, we consider the simplest two-path network example: initially safe path 0 has $\ell_0(t=0) = 0$ with $\alpha \rightarrow 1$, and risky path 1 has an arbitrarily large travel latency $\ell_1(0)$ with $\bar{x} = 0$ and $\mathbb{E}[\alpha_1(t)|\bar{x}] = 0$, by letting $q_{LL} = 1$ and $\alpha_L = 0$. Given $\mathbb{E}[\alpha_1(t)|\bar{x}] < \alpha$ or simply $\bar{x} < \frac{\alpha - \alpha_L}{\alpha_H - \alpha_L}$, a selfish user always chooses path $\pi^\theta(t) = 1$, leading to social cost $\ell_1(0) + \frac{\rho\Delta\ell}{1-\rho}$. While letting the first user exploit $\ell_0(0) = 0$ of path 0 to reduce $\mathbb{E}[\ell_1(1)|\bar{x}, \emptyset]$ to 0 for path 1 at time 1, the socially optimal cost is thus $\frac{\rho^2\Delta\ell}{1-\rho}$. Letting $\frac{(1-\rho)\ell_1(0)}{\rho^2\Delta\ell} \rightarrow \infty$, we obtain $\text{PoA}^\theta = \infty$.

This is a PoA^θ example with the maximum-exploration of stochastic paths, which is opposite to the zero-exploration $\text{PoA}^{(m)}$ example after Proposition 3. Given neither information hiding policy $\pi^\theta(t)$ nor myopic policy $\pi^{(m)}(t)$ under full information sharing works well, we need to design an efficient mechanism to selectively disclose information to users to reduce the social cost.

Definition 1. (Selective Information Disclosure (SID) Mechanism:) *If a user arrival at time t is expected to choose a different route $\pi^\theta(t) \neq 0$ in (25) from optimal $\pi^*(t) = 0$ in (18), then our SID mechanism will disclose*

the latest expected travel latency set $\mathbf{L}(t)$ to him. Otherwise, our mechanism hides $\mathbf{L}(t)$ from this user. Besides, our mechanism always provides optimal path recommendation $\pi^*(t)$, without sharing hazard belief set $\mathbf{x}(t)$, routing history $(\pi(1), \dots, \pi(t-1))$, or past observation set $(\mathbf{y}(1), \dots, \mathbf{y}(t-1))$.

According to Definition 1, if $\pi^*(t) = 0$ but a user at time t makes routing decision $\pi^\theta(t) \neq 0$ under $\bar{x} < \frac{\alpha - \alpha_L}{\alpha_H - \alpha_L}$ in (25), our mechanism discloses $\mathbf{L}(t)$ to avoid him from choosing any stochastic path with large expected travel latency. In the other cases, we simply hide $\mathbf{L}(t)$ from any user arrival, as the user already follows optimal routing $\pi^*(t)$.

In consequence, the worst-case for our SID mechanism only happens when $\pi^\theta(t) \neq 0$ and $\pi^*(t) = 0$ under $\bar{x} < \frac{\alpha - \alpha_L}{\alpha_H - \alpha_L}$ in (25). We still consider the same two-path network example with the maximum-exploration after Proposition 4 to show why this SID mechanism works. In this example, our mechanism will provide $\mathbf{L}(t)$, including $\ell_0(0)$ and $\ell_1(0)$, to each user arrival. Observing huge $\ell_1(0)$, the first user turns to choose path 0 with $\ell_0(0) = 0$, which successfully avoids the infinite social cost under $\pi^\theta(t)$. Furthermore, our SID mechanism successfully avoids the worst-cases of $\text{PoA}^{(m)}$ in Proposition 3. Next we prove that our mechanism well bounds the PoA in the following.

Theorem 1. *Our SID mechanism results in $\text{PoA}^{(\text{SID})} \leq \frac{1}{1-\frac{\rho}{2}}$, which is always no more than 2.*

In the worst-case of $\pi^\theta(t) \neq 0$ and $\pi^*(t) = 0$ for our SID mechanism's $\text{PoA}^{(\text{SID})}$, a user knowing $\mathbf{L}(t)$ may deviate to follow the myopic policy $\pi^{(m)}(t) \neq 0$ in (17). To explain the bounded $\text{PoA}^{(\text{SID})}$, we consider a two-path network example with the maximum-exploration under the myopic policy. Here we start with $\ell_0(0) = \ell_1(0) - \varepsilon$ for safe path 0 with $\alpha \rightarrow 1$ to keep the travel latency on path 0 unchanged if no user chooses that path, where ε is positive infinitesimal. We set $\ell_1(0) = \frac{\Delta \ell}{1 - \mathbb{E}[\alpha_1(0)|\bar{x}]}$ for stochastic path 1 with $x_1(0) = \bar{x}$, such that the travel latency $\mathbb{E}[\ell_1(t)|\bar{x}, y_1(t-1)]$ equals $\ell_1(0)$ all the time if all users choose that path. Then in this system, users keep choosing path 1 under myopic policy $\pi^{(m)}(t)$ in (17) to receive social cost $\frac{\ell_1(0)}{1-\rho}$. However, the socially optimal policy may want the first user to exploit path 0 to permanently reduce path 1's expected travel latency for following users there. Thanks to the first user's routing of path 0, the expected travel latency for each following user choosing path 1 at time t is greatly reduced to be less than $\ell_1(0)$ yet is still no less than $\frac{\ell_1(0)}{2}$ for non-zero $\mathbb{E}[\alpha_1(t)|\bar{x}]$. Then the minimum social cost is reduced to be no less than $\ell_0(0) + \frac{\rho \ell_1(0)}{2-2\rho}$, leading to $\text{PoA}^{(\text{SID})} \leq \frac{1}{1-\frac{\rho}{2}}$.

Besides the worst-case performance analysis, we further verify our mechanism's average performance using extensive simulations. Define the following average inefficiency ratio between expected social costs achieved by our SID mechanism and social optimum in (18):

$$\gamma^{(\text{SID})} = \frac{\mathbb{E}[C^{(\text{SID})}(\mathbf{L}(t), \mathbf{x}(t))]}{\mathbb{E}[C^*(\mathbf{L}(t), \mathbf{x}(t))]} \quad (26)$$

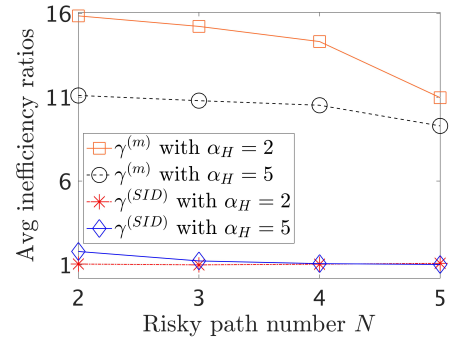


Figure 3: Average inefficiency ratios $\gamma^{(m)}$ under myopic policy in (17) and $\gamma^{(\text{SID})}$ under our selective information disclosure. We vary risky path number N in set $\{2, 3, 4, 5\}$. We set $\alpha = 0.99, \alpha_L = 0, \Delta \ell = 1, p_H = 0.8, p_L = 0.2, q_{HH} = 0.99, q_{LL} = 0.99$ here, and we change $\alpha_H = 2$ and $\alpha_H = 5$ to make comparison. At initial time $t = 0$, we let $\ell_0(0) = 100, \ell_i(0) = 105$ and $x_i(0) = 0.5$ for any path.

To compare, we define $\gamma^{(m)}$ to be the average inefficiency ratio between social costs achieved by the myopic policy in (17) and socially optimal policy in (18). After running 50 long-term experiments for averaging each ratio, we plot Figure 3 to compare $\gamma^{(m)}$ to $\gamma^{(\text{SID})}$ versus risky path number N . Figure 3 shows that our SID mechanism obviously reduces $\gamma^{(m)} > 10$ to $\gamma^{(\text{SID})} < 2$ at $N = 2$, which is consistent with Theorem 1. Figure 3 also shows that the efficiency loss due to users' selfish routing decreases with N , as more choices of risky paths help negate the hazard risk at each path. Here we also vary high hazard state α_H to make a comparison, and we see that a larger α_H causes less efficiency loss due to users' reduced explorations to risky paths.

We can also show using simulations that the average inefficiency ratio under information hiding mechanism in Proposition 4 has a big gap compared to our SID mechanism, especially when users over-explore with $\bar{x} < \frac{\alpha - \alpha_L}{\alpha_H - \alpha_L}$.

Conclusion

In this paper, we studied how to incentive selfish users to reach the best exploitation-exploration trade-off. We use the POMDP techniques to summarize the congestion probability into a dynamic hazard belief. By considering a simple but fundamental parallel routing network with one deterministic path and multiple stochastic paths for atomic users, we proved that the myopic policy's price of anarchy (PoA) is larger than $\frac{1}{1-\rho}$, which can be arbitrarily large as $\rho \rightarrow 1$. To remedy such huge efficiency loss, we proposed a selective information disclosure (SID) mechanism: we only reveal the latest traffic information to users when they intend to over-explore stochastic paths, while hiding such information when they under-explore. We proved that our mechanism reduces PoA to be less than $\frac{1}{1-\frac{\rho}{2}}$. We further examined our mechanism's average-case performance by extensive simulations. We can also extend our system model and key results to a chain road network.

Acknowledgments

The work was supported by the Ministry of Education, Singapore, under its Academic Research Fund Tier 2 Grant under Award MOE- T2EP20121-0001.

References

- Alam, I.; Farid, D. M.; and Rossetti, R. J. 2019. The prediction of traffic flow with regression analysis. In *Emerging Technologies in Data Mining and Information Security*, 661–671. Springer.
- Ban, X.; Herring, R.; Hao, P.; and Bayen, A. M. 2009. Delay pattern estimation for signalized intersections using sampled travel times. *Transportation Research Record*, 2130(1): 109–119.
- Bellman, R. 1966. Dynamic programming. *Science*, 153(3731): 34–37.
- Bilò, V.; and Vinci, C. 2020. The price of anarchy of affine congestion games with similar strategies. *Theoretical Computer Science*, 806: 641–654.
- Bozorgchenani, A.; Maghsudi, S.; Tarchi, D.; and Hossain, E. 2022. Computation Offloading in Heterogeneous Vehicular Edge Networks: On-Line and Off-Policy Bandit Solutions. *IEEE Transactions on Mobile Computing*, 21(12): 4233–4248.
- Brown, P. N.; and Marden, J. R. 2017. Optimal mechanisms for robust coordination in congestion games. *IEEE Transactions on Automatic Control*, 63(8): 2437–2448.
- Castiglioni, M.; Celli, A.; Marchesi, A.; and Gatti, N. 2021. Signaling in bayesian network congestion games: the subtle power of symmetry. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, 5252–5259.
- Cominetti, R.; Scarsini, M.; Schröder, M.; and Stier-Moses, N. E. 2019. Price of anarchy in stochastic atomic congestion games with affine costs. In *Proceedings of the 2019 ACM Conference on Economics and Computation*, 579–580.
- Ferguson, B. L.; Brown, P. N.; and Marden, J. R. 2022. The Effectiveness of Subsidies and Tolls in Congestion Games. *IEEE Transactions on Automatic Control*, 67(6): 2729–2742.
- Gittins, J.; Glazebrook, K.; and Weber, R. 2011. *Multi-armed bandit allocation indices*. John Wiley & Sons.
- Hao, B.; and Michini, C. 2022. The price of Anarchy in series-parallel network congestion games. *Mathematical Programming*, 1–31.
- Koutsoupias, E.; and Papadimitriou, C. 1999. Worst-case equilibria. In *Annual symposium on theoretical aspects of computer science*, 404–413. Springer.
- Kremer, I.; Mansour, Y.; and Perry, M. 2014. Implementing the “wisdom of the crowd”. *Journal of Political Economy*, 122(5): 988–1012.
- Krishnasamy, S.; Sen, R.; Johari, R.; and Shakkottai, S. 2021. Learning unknown service rates in queues: A multiarmed bandit approach. *Operations Research*, 69(1): 315–330.
- Li, H.; and Duan, L. 2022. When Congestion Games Meet Mobile Crowdsourcing: Selective Information Disclosure [Online Technical Report]. *arXiv preprint arXiv:2211.14029*.
- Li, H.; and Duan, L. 2023. Online Pricing Incentive to Sample Fresh Information. *IEEE Transactions on Network Science and Engineering*, 10(1): 514–526.
- Li, Y.; Courcoubetis, C.; and Duan, L. 2019. Recommending paths: Follow or not follow? In *IEEE INFOCOM 2019-IEEE Conference on Computer Communications*, 928–936. IEEE.
- Li, Y.; Courcoubetis, C. A.; and Duan, L. 2017. Dynamic routing for social information sharing. *IEEE Journal on Selected Areas in Communications*, 35(3): 571–585.
- Meigs, E.; Parise, F.; and Ozdaglar, A. 2017. Learning dynamics in stochastic routing games. In *2017 55th Annual Allerton Conference on Communication, Control, and Computing (Allerton)*, 259–266. IEEE.
- Nikolova, E.; and Stier-Moses, N. E. 2011. Stochastic selfish routing. In *International Symposium on Algorithmic Game Theory*, 314–325. Springer.
- Papanastasiou, Y.; Bimpikis, K.; and Savva, N. 2018. Crowdsourcing exploration. *Management Science*, 64(4): 1727–1746.
- Roughgarden, T.; and Tardos, É. 2002. How bad is selfish routing? *Journal of the ACM (JACM)*, 49(2): 236–259.
- Sekar, S.; Zheng, L.; Ratliff, L. J.; and Zhang, B. 2019. Uncertainty in multicommodity routing networks: When does it help? *IEEE Transactions on Automatic Control*, 65(11): 4600–4615.
- Tavaafoghi, H.; and Teneketzis, D. 2017. Informational incentives for congestion games. In *2017 55th Annual Allerton Conference on Communication, Control, and Computing (Allerton)*, 1285–1292. IEEE.
- Vasserman, S.; Feldman, M.; and Hassidim, A. 2015. Implementing the wisdom of waze. In *Twenty-Fourth International Joint Conference on Artificial Intelligence*.
- Vu, D. Q.; Antonakopoulos, K.; and Mertikopoulos, P. 2021. Fast Routing under Uncertainty: Adaptive Learning in Congestion Games via Exponential Weights. *Advances in Neural Information Processing Systems*, 34: 14708–14720.
- Wu, M.; and Amin, S. 2019. Learning an unknown network state in routing games. *IFAC-PapersOnLine*, 52(20): 345–350.
- Zhang, J.; Lu, P.; Li, Z.; and Gan, J. 2018. Distributed trip selection game for public bike system with crowdsourcing. In *IEEE INFOCOM 2018-IEEE Conference on Computer Communications*, 2717–2725. IEEE.
- Zhu, Y.; and Savla, K. 2022. Information Design in Nonatomic Routing Games With Partial Participation: Computation and Properties. *IEEE Transactions on Control of Network Systems*, 9(2): 613–624.