

When Efficiency meets Equity in Congestion Pricing and Revenue Refunding Schemes

Devansh Jalota, Kiril Solovey, *Fellow, IEEE*, Karthik Gopalakrishnan, Stephen Zoepf, Hamsa Balakrishnan, *Member, IEEE*, and Marco Pavone, *Member, IEEE*

Abstract—Congestion pricing has long been hailed as a means to mitigate traffic congestion; however, its practical adoption has been limited due to the resulting social inequity issue, e.g., low-income users are priced out of certain roads. This issue has spurred interest in the design of equitable mechanisms that aim to refund the collected toll revenues as lump-sum transfers to users. Although revenue refunding has been extensively studied for over three decades, there has been no thorough characterization of how such schemes can be designed to simultaneously achieve system efficiency and equity objectives. In this work, we bridge this gap through the study of *congestion pricing and revenue refunding* (CPRR) schemes in non-atomic congestion games. We first develop CPRR schemes, which in comparison to the untolled case, simultaneously increase system efficiency without worsening wealth inequality, while being *user-favorable*: irrespective of their initial wealth or values-of-time (which may differ across users), users would experience a lower travel cost after the implementation of the proposed scheme. We then characterize the set of optimal user-favorable CPRR schemes that simultaneously maximize system efficiency and minimize wealth inequality. Finally, we provide a concrete methodology for computing optimal CPRR schemes and also highlight additional equilibrium properties of these schemes under different models of user behavior. Overall, our work demonstrates that through appropriate refunding policies we can design user-favorable CPRR schemes that maximize system efficiency while reducing wealth inequality.

Index Terms—Congestion Games, Traffic Routing, Wealth Inequality

I. INTRODUCTION

Road congestion pricing, which typically involves users paying for the externalities they impose on other road users, has been widely accepted as a mechanism to alleviate traffic congestion. However, the practical

adoption of congestion pricing has been limited [1] primarily due to the resultant inequity concerns, e.g., high income users are likely to get the most benefit with shorter travel times while low income users suffer large travel times since they avoid the high toll roads. Several empirical works have noted the regressive nature of congestion pricing [2] and a recent theoretical work [3] has also characterized the influence of tolls on wealth inequality. In particular, [3] developed an *Inequity Theorem* for users travelling between the same origin-destination (O-D) pair, and proved that any form of tolls would increase the wealth inequality. These rigorous critiques are complemented by opinions in the popular press that congestion fees amount to “a tax on the working class [4].”

The lack of support for congestion pricing due to its social inequity issues [5] has led to a growing interest in designing equitable congestion-pricing schemes [6]. One approach that has been proposed to alleviate the inequity issues of congestion pricing is direct revenue redistribution, i.e., refunding the toll revenues to users in the form of lump-sum transfers. The idea of revenue refunding is analogous to that of *feebates*, where refunds are used as a means to induce desirable behavior in society. Our work is centered on the design of congestion pricing and revenue refunding (CPRR) schemes that improve system performance without reducing wealth inequality, and benefit every user irrespective of their wealth or value-of-time. We view our work as paving the way for the design of practical, sustainable, and publicly acceptable congestion pricing schemes.

a) Contributions: In this work, we present the first study of the wealth-inequality effects of CPRR schemes in non-atomic congestion games, with a focus on devising CPRR schemes that simultaneously reduce the total system cost, i.e., the sum of the travel times on all edges of the network weighted by the corresponding values-of-time of users, without increasing the level of wealth (or income) inequality. We consider the setting of heterogeneous users, with differing values-of-time and income, who seek to minimize their individual travel cost, which is a linear function of their travel times, tolls, and refunds, in the system. As in previous work [3], we incorporate the income elasticity of travel time, i.e., increased travel time corresponds to lost income, to reason about the income distribution of users before and after the imposition of a CPRR scheme.

Manuscript received September 12, 2022. This work was supported in part by the National Science Foundation, CAREER Award CMMI1454737, and by the Stanford Institute for Human-Centered AI.

Devansh Jalota, Karthik Gopalakrishnan, and Marco Pavone are with the Autonomous Systems Laboratory, Stanford University, Stanford, CA 94305 USA (email: djalota@stanford.edu; kgopalakrishnan@stanford.edu; pavone@stanford.edu).

Kiril Solovey is with the Faculty of Electrical and Computer Engineering at Technion - Israel Institute of Technology, Haifa, 3200003, Israel (e-mail: kirilsol@technion.ac.il).

Stephen Zoepf is with the U.S. Department of Transportation, Washington, D.C 20590, USA (e-mail: stephen.zoepf@dot.gov).

Hamsa Balakrishnan is with the Department of Aeronautics and Astronautics, Massachusetts Institute of Technology, Cambridge, MA 02139 USA (e-mail: hamsa@mit.edu).

The views expressed in this paper are those of the authors and do not necessarily represent the views of the U.S. Department of Transportation or the Federal government.

To capture the behavior of selfish users, we study the effect of the Nash equilibria induced by CPRR schemes on wealth inequality. We begin with the study of exogenous equilibria, which is the standard Nash equilibrium model with heterogeneous users [7], wherein users minimize a linear function of their travel time and tolls, without considering refunds. In this setting, in Section IV, we establish the existence of a *Pareto-improving* CPRR scheme that, compared with the untolled outcome, (i) is user-favorable, i.e., every user group, irrespective of their initial wealth, has a lower travel cost after the implementation of the scheme, (ii) lowers total system cost, and (iii) does not increase wealth inequality (see Fig. 1). When all travel demand is between a single O-D pair and each user's value-of-time is proportional to their income, we further show that the same CPRR scheme does not increase wealth inequality relative to the ex-ante income distribution, i.e., the users' income profiles prior to making their trips. Thus, our results show that it is possible to reverse the wealth-inequality effects of congestion pricing established in the *Inequity Theorem* in [3] through appropriate revenue refunding schemes.

Next, we characterize the set of optimal CPRR schemes that are favorable to all users in the exogenous equilibrium setting. In particular, in Section V, we establish the existence of CPRR schemes that simultaneously minimize total system cost and wealth inequality among all CPRR schemes that are favorable to any user (see Fig. 1). Further, we develop a method to compute the optimal CPRR scheme in Sections VI-A and VI-B and show for a commonly used wealth inequality measure, the discrete Gini coefficient, that a simple max-min allocation of the refunds among user groups with different incomes is optimal. We further present numerical experiments in Appendix D to demonstrate the efficacy of optimal CPRR schemes and show that the benefits of CPRR can even be realized when users' values of time are not exactly known to the central planner.

Finally, in Section VI-C, we consider the endogenous equilibrium, a new notion we introduce, wherein users additionally consider refunds in their travel cost minimization. In this setting, we show that the optimal CPRR scheme is robust to coalitions, i.e., any exogenous equilibrium induced by an optimal CPRR scheme is also an endogenous equilibrium with coalitions.

We remark that in line with prior literature on traffic routing with heterogeneous users [3], [8], [9], we assume a complete information setting wherein the different attributes (i.e., the income, value-of-time, and O-D pair) of the user groups are known and can be used to design CPRR schemes. To this end, our results can be interpreted as the theoretical limits of what is achievable in terms of the efficiency and equity outcomes given perfect state information. However, we remark that even though we consider the complete information setting wherein the tolls and refunds are computed in a centralized manner, the developed optimal CPRR schemes induce selfish users to distributedly optimize their individual objectives

and collectively enforce a traffic pattern that minimizes both total system cost and the level of wealth inequality. Further, since our results contribute to the literature on designing intervention and control schemes under perfect state information [10]–[12], our proposed approach to designing CPRR schemes serves as more of a design module rather than an end-to-end solution to the equity problem associated with congestion pricing. We do note, however, that there are several methods to estimate user attributes, e.g., their values of time or preferences, that have been explored in the empirical literature [13]–[15], which can be used to inform the inputs that are necessary for the design of optimal CPRR schemes.

Overall, our work demonstrates that if we appropriately refund the collected toll revenues, we can achieve system efficiency without increasing inequality. Further, in doing so, we ensure that our designed schemes are publicly acceptable as we guarantee that each user is at least as well off as before the introduction of the CPRR scheme. Thus, we view our work as a significant step in shifting the discussion around congestion pricing from one focused on the societal inequity impacts of tolls to one that centers around *how* to best preserve equity through the distribution of toll revenues.

II. RELATED WORK

The design of mechanisms that satisfy both system efficiency and user fairness desiderata has been a centerpiece of algorithm design for a range of applications. For instance, in resource allocation settings, [16] quantified the loss in efficiency when the allocation outcomes are required to satisfy certain fairness criteria. In machine learning classification tasks, [17] studied group-based fairness notions to prevent discrimination against individuals belonging to disadvantaged groups. In the context of traffic routing, [18] introduced a fairness-constrained traffic-assignment problem to achieve a balance between the total travel time of a traffic assignment and its level of fairness. Here, fairness is measured through the maximum ratio between the travel times of users travelling between the same O-D pair.

Resolving the efficiency and equity trade-off is particularly important for allocation mechanisms involving monetary transfers given their impact on low-income groups. Although achieving system efficiency involves allocating goods to users with the highest willingness to pay, in many settings, e.g., cancer treatment, the needs of users are not well expressed by their willingness to pay [19]. Since Weitzman's seminal work on accounting for agent's needs in allocation decisions [19], there has been a rich line of work on taking into account redistributive considerations [20] in resource allocation problems. For instance, [21] analyzed the free provision of a low-quality public good to low-income users by taxing individuals that consume the same good of a higher quality in the private market. More recently, [22] studied the allocation of objects to agents with the objective of maximizing agent's values that may be different from their willingness to pay.

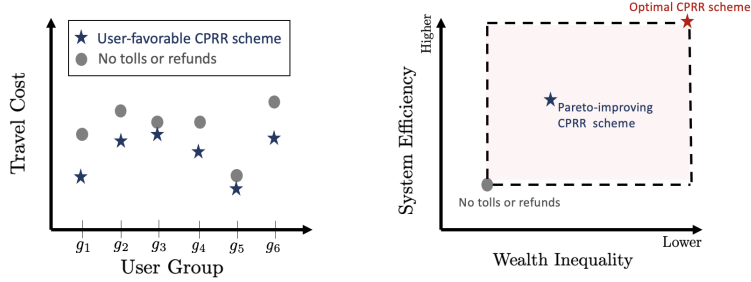


Fig. 1. Depiction of *user-favorable*, *Pareto-improving*, and *optimal* congestion pricing and revenue refunding (CPRR) schemes.

In the context of congestion pricing, revenue redistribution has long been considered as a means to alleviate the inequity issues of congestion pricing [23]. Several revenue redistribution strategies have been proposed in the literature, such as the lump-sum transfer of toll revenues to users [24]. In the setting of Vickrey’s bottleneck congestion model [25]—a benchmark representation of peak-period traffic congestion on a single lane— [26] investigated how a uniform lump-sum payment of toll revenues can be used to make heterogeneous users better off than prior to the implementation of the tolls and refunds. In more general networks with a single O-D pair, [27] established the existence of a tolling mechanism with uniform revenue refunds that reduced the travel cost for each user while decreasing the total system travel time as compared to before the tolling reform. The extension of this result to general road networks with a multiple O-D pair travel demand and heterogeneous users was investigated by [7]. While [7] characterize conditions for the CPRR scheme to be user-favorable, our work studies the influence of such schemes by characterizing their influence on wealth inequality.

III. PRELIMINARIES

In this section, we introduce basic definitions and concepts regarding traffic flow, congestion pricing and revenue refunding (CPRR) schemes, and metrics for system efficiency and wealth-inequality.

A. Elements of Traffic Flow

We model the road network as a directed graph $G = (V, E)$, with the vertex and edge sets denoted by V and E , respectively. Each edge $e \in E$ has a flow-dependent travel-time function $t_e : \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}_{\geq 0}$, which maps x_e , the traffic flow rate on edge e , to the travel time $t_e(x_e)$. The flow rate x_e on edge e represents the average number of vehicles traversing through that edge during a fixed time interval (e.g., over an hour). As is standard in the literature, we assume that the function t_e , for each $e \in E$, is differentiable, convex and monotonically increasing. While we assume that the edge travel times have infinite capacities, as is common in the non-atomic congestion game and transportation literature [3], [28], we note that our model can be extended to the setting with soft capacity constraints for appropriate choices of the travel time functions that grow very steeply once the road capacities have been exceeded.

Users make trips in the road network and belong to a discrete set of user groups based on their (i) value-of-time, (ii) income, and (iii) O-D pair. Let \mathcal{G} denote the set of all user groups, and let $v_g > 0$, $q_g > 0$, and $w_g = (s_g, u_g)$ denote the value-of-time, income, and O-D pair represented by an origin s_g and destination u_g , respectively, for each user in group $g \in \mathcal{G}$. Each user belonging to a group g makes a trip on a path, which is a sequence of directed edges beginning at s_g and ending at u_g (without visiting any node more than once). The set of all possible paths between OD-pair w_g is denoted as \mathcal{P}_g and the travel demand d_g of user group g represents the total flow to be routed through paths in \mathcal{P}_g .

A path flow pattern $\mathbf{f} = \{f_{P,g} : g \in \mathcal{G}, P \in \mathcal{P}_g\}$ specifies for each user group g , the amount of flow $f_{P,g} \geq 0$ routed on a path $P \in \mathcal{P}_g$. In particular, a flow \mathbf{f} must satisfy the user demand, i.e., $\sum_{P \in \mathcal{P}_g} f_{P,g} = d_g$, for all $g \in \mathcal{G}$. We denote the set of all non-negative flows that satisfy this constraint as Ω .

Each path flow $\mathbf{f} = \{f_{P,g} : g \in \mathcal{G}, P \in \mathcal{P}_g\}$ is associated with a corresponding edge flow $\mathbf{x} = \{x_e\}_{e \in E}$ and group specific edge flows $\mathbf{x}^g = \{x_e^g\}_{e \in E}$ for all $g \in \mathcal{G}$, where x_e^g represents the flow of users in group g on edge e . The relationship between the path and edge flows is given by $\sum_{P \in \mathcal{P}_g : e \in P} f_{P,g} = x_e^g$, for all $e \in E, g \in \mathcal{G}$ and $\sum_{g \in \mathcal{G}} x_e^g = x_e$, for all $e \in E$. Here $P \in \mathcal{P}_g : e \in P$ denotes the set of paths $P \in \mathcal{P}_g$ that include edge e .

B. CPRR Schemes

A congestion pricing and revenue refunding (CPRR) scheme is defined by a tuple $(\boldsymbol{\tau}, \mathbf{r})$, where (i) $\boldsymbol{\tau} = \{\tau_e : e \in E\}$ is a vector of edge prices (or tolls), and (ii) $\mathbf{r} = \{r_g : g \in \mathcal{G}\}$ is a vector of group-specific revenue refunds, where each user in group g receives a lump-sum transfer of r_g . In other words, everybody pays the same toll for using an edge independent of their group, and all users with the same income, value-of-time and O-D pair get the same refund, irrespective of the actual path they take between the O-D pair w_g . We note that the vector of refunds \mathbf{r} , in general, need not be non-negative and can take on any real values. Under the CPRR scheme $(\boldsymbol{\tau}, \mathbf{r})$ and a vector of edge flows \mathbf{x} , the total value of tolls collected is given by $\Pi := \sum_{e \in E} \tau_e x_e$. In this work we consider CPRR schemes such that the sum of the revenue refunds equals the sum of the revenue collected from the edge tolls, i.e., $\sum_{g \in \mathcal{G}} r_g d_g = \Pi$. In addition, we consider revenue refunding schemes that depend only

on the groups \mathcal{G} and the total revenue Π induced by a flow \mathbf{f} , but not on the specific paths taken by users under \mathbf{f} . We leave the study of more complex path-dependent refunding schemes for future work (see Section VII).

The total travel cost incurred by the user consists of two components: (i) a linear function of their travel time and tolls, which is a commonly-used modelling approach [8], [9], and (ii) the refund received. The overall model we use, which is formally defined below, has been previously considered in the literature [7].

Definition 1 (User Travel Cost). *Consider a CPRR scheme $(\boldsymbol{\tau}, \mathbf{r})$ and a flow pattern \mathbf{f} with edge flow \mathbf{x} . Then, the total cost incurred by a user belonging to a group $g \in \mathcal{G}$ when traversing a path $P \in \mathcal{P}_g$ with $f_{P,g} > 0$ is given by $\mu_P^g(\mathbf{f}, \boldsymbol{\tau}, \mathbf{r}) := \sum_{e \in P} (v_g t_e(x_e) + \tau_e) - r_g$.*

With slight abuse of notation, we will denote $\mu_P^g(\mathbf{f}, \boldsymbol{\tau}, \mathbf{0})$ as a travel cost that does not include refunds, and $\mu_P^g(\mathbf{f}, \mathbf{0}, \mathbf{0})$ as a travel cost that does not account for tolls or refunds, where $\mathbf{0}$ is a vector of zeros.

C. System Efficiency and Wealth Inequality Metrics

We evaluate the quality of a CPRR scheme using two metrics: (i) system efficiency, which is measured through the total system cost, and (ii) wealth inequality.

a) *Total System Cost*: For any feasible path flow \mathbf{f} with edge flows \mathbf{x} and group specific edge flows \mathbf{x}^g , the total system cost $C(\mathbf{f})$, is the sum of travel times weighted by the users' values-of-time across all edges [7] i.e., $C(\mathbf{f}) := \sum_{e \in E} \sum_{g \in \mathcal{G}} v_g x_e^g t_e(x_e)$. We denote by $C^* := \min_{\mathbf{f} \in \Omega} C(\mathbf{f})$ the widely studied cost-based system optimum.

b) *Wealth Inequality*: We measure the impact of a CPRR scheme on wealth inequality in the following manner. For a profile of incomes $\mathbf{q} = \{q_g : g \in \mathcal{G}\}$, we let a function $W : \mathbb{R}_{>0}^{|\mathcal{G}|} \rightarrow \mathbb{R}_{\geq 0}$ measure the level of wealth inequality of society. We say that an income distribution $\tilde{\mathbf{q}}$ has a lower level of wealth inequality than \mathbf{q} if and only if $W(\tilde{\mathbf{q}}) \leq W(\mathbf{q})$.

In this work, we assume that the wealth-inequality measure $W(\cdot)$ satisfies the following properties:

- 1) *Scale Independence*: The wealth-inequality is unchanged after re-scaling incomes by a positive constant, i.e., $W(\lambda \mathbf{q}) = W(\mathbf{q})$ for any $\lambda > 0$.
- 2) *Constant Income Transfer Property*: If the initial income distribution is \mathbf{q} and each user is transferred a non-negative (non-positive) amount of money λ ($-\lambda$) where $0 \leq \lambda < \min_{g \in \mathcal{G}} q_g$, then the wealth inequality cannot increase (decrease). That is, $W(\mathbf{q} + \lambda \mathbf{1}) \leq W(\mathbf{q})$ and $W(\mathbf{q} - \lambda \mathbf{1}) \geq W(\mathbf{q})$, where $\mathbf{1}$ is a vector of ones.

The above properties are well defined for any wealth inequality distribution when the incomes of all users are strictly positive, which we assume in this work. These properties, including scale independence [29], [30], are fairly natural [3] and hold for commonly used wealth-inequality measures, such as the discrete Gini coefficient, which we elucidate in detail in Section VI-B. Further,

the constant income transfer property is a direct consequence of the fact that regressive (progressive) taxes increase (decrease) wealth inequality, as elucidated in the extended version of this work [31].

When using the wealth inequality measure W , we are interested in understanding the influence of a flow \mathbf{f} for a given CPRR scheme $(\boldsymbol{\tau}, \mathbf{r})$ on the income distribution of users. To this end, we define the income profile of users before making their trip as the *ex-ante income distribution* $\mathbf{q}^0 > \mathbf{0}$ and that after making their trip as the *ex-post income distribution*, which is defined as follows.

Definition 2 (Ex-Post Income Distribution). *For a given CPRR scheme $(\boldsymbol{\tau}, \mathbf{r})$ and an equilibrium flow \mathbf{f} , the induced ex-post income distribution of users is denoted by $\mathbf{q}(\mathbf{f}, \boldsymbol{\tau}, \mathbf{r})$ and is defined as follows. For a group g , $q_g(\mathbf{f}, \boldsymbol{\tau}, \mathbf{r}) := q_g^0 - \beta \mu^g(\mathbf{f}, \boldsymbol{\tau}, \mathbf{r})$, where \mathbf{q}^0 is the ex-ante income distribution and β is a small constant such that the ex-post income of users is strictly positive and represents the relative importance of the congestion game under consideration to an individual's well-being [3].*

We reiterate that the small constant β does not depend on the type of trip being made or the importance of that trip to the user but solely reflects the importance of the congestion game under consideration to an individual's well-being, as in [3]. The positive income assumption ensures that the above defined wealth inequality properties (including scale independence) hold, which would not be the case if users have negative incomes.

We note that in this paper we consider time-invariant travel demand that is fixed for all user groups and assume fractional flows, both of which are standard assumptions in the literature [3], [7]. In line with prior work by [7], we assume that users are refunded based on their income, value-of-time, and O-D pair. Furthermore, similar to much of the prior literature in traffic routing with heterogeneous users [3], [8], [9], we assume that the different attributes (i.e., the income, value-of-time, and O-D pair) of the user groups are known, and can be used in the design of CPRR schemes. In practice, such centralized information on user attributes may not be known and we defer the problem of dealing with incomplete information settings to future work.

IV. PARETO IMPROVING CPRR SCHEMES

In this section, we show that if the tolls collected from congestion pricing are refunded to users in an appropriate way then the wealth inequality effects of congestion pricing can be reversed. Throughout this section and the next we assume that user behavior is characterized through the *exogenous equilibrium* model wherein users minimize a linear function of their travel time and tolls, without considering refunds.

After formally defining exogenous equilibrium below, we develop a CPRR scheme that simultaneously decreases the total system cost of all users while not increasing the level of wealth inequality relative to the untolled outcome, a property which we refer to as *Pareto improving*. Moreover, when designing the scheme, we

ensure that it is politically acceptable by guaranteeing that each user is at least as well off in terms of the travel cost μ^g , which includes travel time, tolls, and refunds, under the CPRR scheme than that without the implementation of congestion pricing or refunds.

Next, we consider the important special case when users travel between the same O-D pair, and have values-of-time proportional to their income. In this setting, we establish the existence of a Pareto improving CPRR scheme that results in an ex-post income distribution that has a lower wealth inequality as compared to that of the ex-ante income distribution. Note that this result is stronger than the more general case with multiple O-D pairs, as the wealth-inequality measure of the ex-ante income distribution is lower than that of the ex-post income distribution for the untolled case.

A. Exogenous Equilibrium

To capture the strategic behavior of users, we present below the standard model of Nash equilibrium with heterogeneous users, which we call exogenous equilibrium. The exogenous setting is commonly studied in the context of non-atomic congestion games without [8], [9] or with refunds [7]. As the name suggests, in an exogenous equilibrium the revenue refunds are assumed to be *exogenous* and do not influence the behavior and route choice of users in the transportation network. That is, users minimize a linear function of their travel time and tolls, without considering refunds.

We note that such a model of user behavior can be quite realistic in certain settings, since accounting for refunds when making route choices may often involve quite sophisticated decision making for users. Furthermore, for users to reason about how their path choice will influence their refund, they must know the refunding policy, which may typically not be known in practice, thereby making the notion of an exogenous equilibrium more appropriate in such settings. We do consider the more sophisticated *endogenous* setting in Section VI-C and demonstrate that our results obtained in the exogenous setting also extend to endogenous setting as well.

The following definition formalizes the notion of an exogenous equilibrium, which only depends on the toll component τ of a CPRR scheme (τ, r) .

Definition 3 (Exogenous Equilibrium). *For a given congestion-pricing scheme τ , a path flow pattern \mathbf{f} is an exogenous equilibrium if for each group $g \in \mathcal{G}$ it holds that $f_{P,g} > 0$ for some path $P \in \mathcal{P}_g$ if and only if $\mu_P^g(\mathbf{f}, \tau, \mathbf{0}) \leq \mu_Q^g(\mathbf{f}, \tau, \mathbf{0})$, for all $Q \in \mathcal{P}_g$. We say that such an \mathbf{f} is an exogenous τ -equilibrium.*

We reiterate that the above notion of an exogenous equilibrium is the standard Nash equilibrium concept used in non-atomic congestion games and follows since users are infinitesimal, unlike equilibrium concepts in atomic congestion games or in the presence of coalitions (see Definition 6). In this work, we refer to this equilibrium concept as *exogenous* to explicitly distinguish it

from the *endogenous* setting when users also account for refunds when making travel decisions. A key property of any exogenous τ -equilibrium \mathbf{f} is that all users within a given group $g \in \mathcal{G}$ incur the same travel cost without refunds, irrespective of the path on which they travel. Hence, we drop the path dependence in the notation and denote the user travel cost without refunds for any user in group g at flow \mathbf{f} as $\mu^g(\mathbf{f}, \tau, \mathbf{0})$. Additionally, since the refund r_g is the same for all users in group g , the travel cost with refunds is denoted as $\mu^g(\mathbf{f}, \tau, r)$.

Another useful property of an exogenous equilibrium is that for a given congestion-pricing scheme τ , the resulting total system cost, user travel cost, and ex-post income distribution are invariant under the different τ -equilibria (see the extended version of our paper [31] for a discussion). That is for any two τ -equilibria \mathbf{f} and \mathbf{f}' it holds that $C(\mathbf{f}) = C(\mathbf{f}')$, $\mu^g(\mathbf{f}, \tau, \mathbf{0}) = \mu^g(\mathbf{f}', \tau, \mathbf{0})$, and $q(\mathbf{f}, \tau, r) = q(\mathbf{f}', \tau, r)$. Thus, we will use the simplified notation $C_\tau := C(\mathbf{f})$, $\mu^g(\tau, r) := \mu^g(\mathbf{f}, \tau, r)$, and $q(\tau, r) := q(\mathbf{f}, \tau, r)$ for any exogenous τ -equilibrium \mathbf{f} , when considering the exogenous equilibrium model. In this context, note that C_0 corresponds to the untolled total system cost, and this quantity is identical for both the exogenous and endogenous equilibrium (we consider the latter in Section VI-C).

B. User-Favorable Pareto Improving CPRR Schemes

To ensure that the CPRR schemes we develop are politically acceptable, we consider schemes, as in [7], that result in equilibrium outcomes wherein each individual user is at least as well off as compared to that under the untolled user equilibrium outcome, a property we refer to as user-favorable (see Fig. 1). We note that the definition below readily extends to the setting of endogenous equilibria as well.

Definition 4 (User-Favorable CPRR Schemes). *A CPRR scheme (τ, r) is user-favorable if for any (exogenous) τ -equilibrium the travel cost of any user group g does not increase with respect to any untolled 0-equilibrium \mathbf{f}^0 , i.e., $\mu^g(\tau, r) \leq \mu^g(\mathbf{0}, \mathbf{0})$.*

We mention that the the above definition can readily be extended to incorporate the notion of a user-favorable CPRR scheme relative to any status-quo traffic equilibrium pattern, which is not necessarily equal to the untolled case, e.g., the traffic pattern in a city that has already implemented some form of congestion pricing. Thus, considering the untolled user equilibrium \mathbf{f}^0 in the above definition is without loss of generality.

We now present the main result of this section, which establishes that any pricing scheme τ that improves the system efficiency compared to the untolled case can be paired with a refunding scheme r such that the wealth inequality relative to the ex-post income distribution under the untolled setting is not increased, i.e., the CPRR scheme (τ, r) is Pareto improving (see Fig. 1) and user-favorable. Note that designing CPRR schemes that achieve a lower wealth inequality and total system cost as compared to that of the untolled user equilibrium

outcome is desirable since this implies that the CPRR scheme improves upon both the efficiency and equity metrics relative to the status-quo equilibrium pattern.

Proposition 1 (Existence of Pareto Improving CPRR Scheme). *Let τ be a congestion-pricing scheme such that $C_\tau \leq C_0$, where C_0 is the untolled total system cost. Then, there exists a refund scheme r such that (τ, r) is user-favorable and does not increase wealth inequality, i.e., $W(q(\tau, r)) \leq W(q(\mathbf{0}, \mathbf{0}))$. That is, the scheme (τ, r) is Pareto improving.*

For a proof of Proposition 1, see Appendix A. Both Definition 4 and Proposition 1 can readily be extended to incorporate the notions of user-favorable and Pareto improving CPRR schemes relative to any status-quo traffic equilibrium pattern beyond the untolled user equilibrium. For simplicity, we prove those properties relative to the untolled user equilibrium setting.

We now present a consequence of this result for single O-D pair travel demand when all users have values-of-time proportional to their incomes. In this setting, we show the existence of a revenue refunding scheme that decreases the wealth inequality relative to the ex-ante income distribution. Note that this is a stronger result than Proposition 1 since the wealth inequality of the ex-ante income distribution is lower than that of the ex-post income distribution for the untolled case.

Corollary 1 (CPRR Decreases Wealth Inequality for Single O-D Pair). *Consider the setting where all users travel between the same O-D pair and have values-of-time proportional to their incomes, i.e., $v_g = \omega q_g^0$ for some $\omega > 0$ for each group g . Let τ be road tolls such that $C_\tau \leq C_0$. Then, there exists a revenue refunding scheme r such that the CPRR scheme (τ, r) is user-favorable and $W(q(\tau, r)) \leq W(q^0)$.*

For a proof of Corollary 1, see Appendix B. Corollary 1 indicates that appropriate refunding can reverse the negative consequences of tolls on wealth inequality, as established in the ‘‘Inequity Theorem’’ [3]. In particular, the ‘‘Inequity Theorem’’ asserts that for the setting considered in Corollary 1, any form of tolls increases the level of wealth inequality compared with the ex-ante income distribution q^0 in the absence of refunds.

A main ingredient in Corollary 1 is that the wealth inequality of the ex-ante income distribution q^0 is equal to that of the ex-post income distribution under the untolled user equilibrium. This result holds when users travel between the same O-D pair and have values-of-time scaling proportionally with their incomes. However, it does not hold in general for users travelling between different O-D pairs, since in such a case, users may incur different travel times at the untolled user equilibrium. For multiple O-D pairs, we show through an example in the extended version of our work [31] that there are travel demand instances when no CPRR scheme can reduce income inequality relative to that of the ex-ante income distribution. Thus, for the rest of this paper we devise

CPRR schemes that do not increase the wealth inequality relative to the ex-post income distribution under the untolled user equilibrium outcome rather than relative to the ex-ante income distribution. Note that doing so is reasonable, as we look to design CPRR schemes that improve on the status quo traffic pattern, which is typically described by the untolled user equilibrium.

V. OPTIMAL CPRR SCHEMES

In the previous section, we established the existence of a user-favorable CPRR scheme that simultaneously reduces total system cost without increasing wealth inequality relative to an untolled outcome. In this section, we prove the existence of optimal CPRR schemes that achieve a total system cost and wealth inequality that cannot be improved by any other user-favorable CPRR scheme. In particular, we establish that the optimal CPRR schemes are those that induce exogenous equilibrium flows with the minimum total system cost while also resulting in ex-post income distributions with the lowest level of wealth inequality among the class of all user-favorable CPRR schemes (see Fig. 1). We further show in Section VI that these optimal CPRR schemes induce equilibria even when coalitions of users endogenize the effect of refunds on their travel decisions.

We first present the main result of this section, which characterizes the set of optimal CPRR schemes.

Theorem 1 (Optimal CPRR Scheme). *There exists a user-favorable CPRR scheme (τ^*, r^*) such that for any user-favorable CPRR scheme (τ, r) it holds that $C_{\tau^*} \leq C_\tau$ and $W(q(\tau^*, r^*)) \leq W(q(\tau, r))$.*

The proof of this theorem is constructive as it provides a recipe for computing the optimal CPRR scheme (τ^*, r^*) . The proof relies on two intermediate results of independent interest. The first lemma shows that under any user-favorable CPRR scheme, each user’s income is at least their ex-post income under the untolled case.

Lemma 1 (Ex-post Income Distribution). *Let τ be road tolls such that $C_\tau \leq C_0$. Then, under any refunds r such that the CPRR scheme (τ, r) is user-favorable, the ex-post income of any user in group g is $q_g(\tau, r) = q_g(\mathbf{0}, \mathbf{0}) + \beta c_g$, where the transfer value $c_g \geq 0$ and satisfies the relation $\sum_{g \in \mathcal{G}} c_g d_g = C_0 - C_\tau$.*

Proof. Denote the ex-post income of group g as $\hat{q}_g = q_g(\tau, r)$. We now prove the ex-post income relation using the definition of a user-favorable CPRR scheme. In particular, for any user-favorable CPRR scheme (τ, r) the user travel cost does not increase from the untolled case, i.e., $\mu^g(\tau, r) \leq \mu^g(\mathbf{0}, \mathbf{0})$. As it holds that $\mu^g(\tau, r) = \mu^g(\tau, \mathbf{0}) - r_g$, we observe that for some $c_g \geq 0$ the following relation must hold for each user in group g : $\mu^g(\tau, \mathbf{0}) - r_g + c_g = \mu^g(\mathbf{0}, \mathbf{0})$. Then, for an ex-ante income distribution q^0 , the ex-post income of each user belonging to group g is given by

$$\hat{q}_g \stackrel{(a)}{=} q_g^0 - \beta \mu^g(\mathbf{0}, \mathbf{0}) + \beta c_g \stackrel{(b)}{=} q_g(\mathbf{0}, \mathbf{0}) + \beta c_g,$$

where (a) follows as $\mu^g(\boldsymbol{\tau}, \mathbf{0}) - r_g = \mu^g(\mathbf{0}, \mathbf{0}) - c_g$, and (b) follows as the ex-post income of users in group g for the untolled setting is $q_g(\mathbf{0}, \mathbf{0}) = q_g^0 - \beta\mu^g(\mathbf{0}, \mathbf{0})$.

Next, to show that $\sum_{g \in \mathcal{G}} c_g d_g = C_0 - C_\tau$ we characterize the quantities C_0 and C_τ . In particular, observe that by definition $C_0 = C(\mathbf{f}^0)$ and $C_\tau = C(\mathbf{f})$, where \mathbf{f}^0 is the untolled $\mathbf{0}$ -equilibrium and \mathbf{f} is an exogenous τ -equilibrium. Now, note that both flows \mathbf{f}^0 and \mathbf{f} can be expressed in closed form. In particular, for a given congestion-pricing scheme τ' the exogenous τ' -equilibrium $\mathbf{h}(\tau')$ can be written as

$$\mathbf{h}(\tau') = \arg \min_{\mathbf{h}' \in \Omega} \sum_{e \in E} \int_0^{x(\mathbf{h}')_e} t_e(\omega) d\omega + \sum_{e \in E} \sum_{g \in \mathcal{G}} \frac{1}{v_g} x(\mathbf{h}')_e \tau_e, \quad (1)$$

where $x(\mathbf{f}')$ denotes the edge representation of a path flow \mathbf{f}' . We note that this program corresponds to the *multi-class user-equilibrium optimization problem* [28].

Given this representation of $\mathbf{h}(\tau')$, we derive the following relation between the total system cost $C_{\tau'}$ and collected revenues, by analyzing the KKT conditions of this minimization problem. In particular, it holds that

$$C_{\tau'} = \sum_{g \in \mathcal{G}} \mu^g(\tau', \mathbf{0}) d_g - \sum_{e \in E} \tau'_e x(\mathbf{h}(\tau'))_e. \quad (2)$$

Note that the edge flow $x(\mathbf{h}(\tau'))$ is unique by the strict convexity of the travel-time function. We defer the proof of (2) to the extended version of this paper [31].

We now leverage (2) to obtain that $C_\tau = \sum_{g \in \mathcal{G}} \mu^g(\tau, \mathbf{0}) d_g - \sum_{e \in E} \tau_e x(\mathbf{f})_e$, where $x(\mathbf{f}) = x(\mathbf{h}(\tau))$. Furthermore, from (2) for the untolled setting, we obtain that $C_0 = \sum_{g \in \mathcal{G}} \mu^g(\mathbf{0}, \mathbf{0}) d_g$. Finally, using these two relations and leveraging the fact that $c_g = \mu^g(\mathbf{0}, \mathbf{0}) - \mu^g(\tau, \mathbf{0}) + r_g$ we get

$$\begin{aligned} \sum_{g \in \mathcal{G}} c_g d_g &\stackrel{(a)}{=} C_0 - \sum_{g \in \mathcal{G}} \mu^g(\tau, \mathbf{0}) d_g + \sum_{e \in E} \tau_e x(\mathbf{f})_e, \\ &\stackrel{(b)}{=} C_0 - C_\tau, \end{aligned}$$

where (a) follows as $\sum_{g \in \mathcal{G}} r_g d_g = \sum_{e \in E} \tau_e x(\mathbf{f})_e$ and $C_0 = \sum_{g \in \mathcal{G}} \mu^g(\mathbf{0}, \mathbf{0}) d_g$, and (b) follows as $C_\tau = \sum_{g \in \mathcal{G}} \mu^g(\tau, \mathbf{0}) d_g - \sum_{e \in E} \tau_e x(\mathbf{f})_e$. This proves our claim. \square

The second result required to prove Theorem 1 relies on the observation that there is a monotonic relationship between the minimum achievable wealth-inequality measure and the total system cost.

Lemma 2 (Monotonicity of Refunds). *Suppose that there are two congestion-pricing schemes τ_A and τ_B with total system costs satisfying $C_{\tau_A} \leq C_{\tau_B} \leq C_0$. Then there exists a revenue refunding scheme \mathbf{r}_A such that (τ_A, \mathbf{r}_A) is user-favorable and achieves a lower wealth inequality measure than any user-favorable CPRR scheme (τ_B, \mathbf{r}_B) for any revenue refunds \mathbf{r}_B , i.e., $W(\mathbf{q}(\tau_A, \mathbf{r}_A)) \leq W(\mathbf{q}(\tau_B, \mathbf{r}_B))$.*

Proof. We prove this claim by constructing for each revenue refunding scheme \mathbf{r}_B under the tolls τ_B , a

revenue refunding scheme \mathbf{r}_A under the tolls τ_A that achieves a lower wealth inequality. To this end, we first introduce some notation. Let c_g^A and c_g^B be non-negative transfers for each group g as in Lemma 1, where $\sum_{g \in \mathcal{G}} c_g^A d_g = C_0 - C_{\tau_A}$ and $\sum_{g \in \mathcal{G}} c_g^B d_g = C_0 - C_{\tau_B}$ must hold for the feasibility of the scheme.

Then, by Lemma 1 we have that the ex-post income of users in group g can be expressed as: $q_g(\tau_A, \mathbf{r}_A) = q_g(\mathbf{0}, \mathbf{0}) + \beta c_g^A$ and $q_g(\tau_B, \mathbf{r}_B) = q_g(\mathbf{0}, \mathbf{0}) + \beta c_g^B$. Let $c_g^A = c_g^B + \frac{1}{\sum_{g \in \mathcal{G}} d_g} (C_{\tau_B} - C_{\tau_A})$. We now show that the refunding \mathbf{r}_A is feasible by observing that $\sum_{g \in \mathcal{G}} c_g^A d_g = \sum_{g \in \mathcal{G}} c_g^B d_g + C_{\tau_B} - C_{\tau_A} = C_0 - C_{\tau_A}$. Here we leveraged the fact that $\sum_{g \in \mathcal{G}} c_g^B d_g = C_0 - C_{\tau_B}$.

Under the above defined non-negative transfer c_g^A , we observe that the ex-post income distribution under the CPRR scheme (τ_A, \mathbf{r}_A) is the same as the ex-post income distribution under the CPRR scheme (τ_B, \mathbf{r}_B) plus a constant positive transfer, which is equal for all users. That is, we have $\mathbf{q}(\tau_A, \mathbf{r}_A) = \mathbf{q}(\tau_B, \mathbf{r}_B) + \lambda \mathbf{1}$ for $\lambda = \frac{\beta}{\sum_{g \in \mathcal{G}} d_g} (C_{\tau_B} - C_{\tau_A}) \geq 0$. Finally, by the constant income transfer property (Section III) it follows that $W(\mathbf{q}(\tau_A, \mathbf{r}_A)) \leq W(\mathbf{q}(\tau_B, \mathbf{r}_B))$. \square

The above result establishes a very natural property of any user-favorable revenue-refunding policy for which the total refund remaining after satisfying the user-favorable condition is $C_0 - C_\tau$. In particular, a smaller total system cost yields a larger amount of remaining refund $C_0 - C_\tau$, which, in turn, results in a greater degree of freedom in distributing these refunds to achieve an overall lower level of wealth inequality.

Finally, Theorem 1 follows directly by the monotonicity relation established in Lemma 2, and prescribes a two-step procedure to find an optimal CPRR scheme that is also user-favorable. In particular, choose a congestion pricing scheme τ^* such that the total travel cost is minimized, i.e., $C_{\tau^*} = C^*$. Next, select the revenue refunding scheme \mathbf{r}^* to be such that the expression $W(\mathbf{q}(\tau^*, \mathbf{r}^*))$ is minimized and (τ^*, \mathbf{r}^*) is user-favorable through an appropriate selection of transfers c_g . For more details on computing an optimal CPRR scheme, we refer to Sections VI-A and VI-B.

a) *Significance of Theorem 1:* Theorem 1 establishes that the optimal CPRR scheme is one that simultaneously achieves the highest efficiency whilst also reducing wealth inequality to the maximum degree possible among all user-favourable CPRR schemes. This finding is counter-intuitive since equity and efficiency are typically at odds but Theorem 1 establishes that no such tradeoff between system efficiency and wealth inequality exists. The reason for this is that the remaining refund after satisfying the user-favourable condition increases as the total system cost decreases (Lemma 2), thereby giving greater leverage in the design of the refunding mechanism to achieve a lower wealth inequality. We further present numerical experiments in Appendix D to demonstrate the efficacy of optimal CPRR schemes and also show that the benefits of CPRR can even be realized

in the setting when users' values of time are not exactly known to the central planner.

VI. COMPUTATIONAL AND EQUILIBRIUM PROPERTIES OF OPTIMAL CPRR SCHEMES

Having established the existence of optimal CPRR schemes, we now show how such schemes can be computed and highlight additional equilibrium properties of these schemes. To this end, in Sections VI-A and VI-B, we provide a concrete recipe for computing the optimal CPRR scheme (τ^*, r^*) for a commonly used wealth inequality measure, the *discrete Gini coefficient*. Then, in Section VI-C, we consider the endogenous equilibrium setting, wherein users minimize a linear function of not only their travel times and tolls but also refunds. In this setting, we show that the optimal CPRR scheme is robust to user coalitions, i.e., optimal CPRR schemes induce equilibria even when *coalitions* of users endogenize the effect of refunds on their travel decisions.

A. Computing Optimal Tolls

The problem of computing optimal tolls τ^* has been widely studied [28]. In particular, [28] showed by analysing the KKT conditions of the minimum total system cost problem, presented in Section III-C, that the optimal toll on each edge is given by $\tau_e^* = \left(\sum_{g \in \mathcal{G}} \frac{x_e^g}{x_e} v_g \right) x_e t'_e(x_e)$, where edge flows x and the group specific edge flows x^g correspond to the edge decomposition of the optimal path flow f of the minimum total system cost problem: $f = \arg \min_{f \in \Omega} C(f)$. Observe that the optimal tolls τ^* to minimize the total system cost is akin to marginal cost prices, given by $x_e t'_e(x_e)$ for each edge e , when all users have the same values of time. In particular, the optimal toll on each edge is given by the travel time externality, i.e., the marginal cost prices, of users multiplied by the average value of time of users on that edge.

B. Computing Optimal Revenue Refunds

Given the method to compute optimal tolls τ^* , as elucidated in the previous section, we now focus our attention on deriving the optimal revenue refunding policy r^* for a commonly used wealth inequality measure, the discrete Gini coefficient. In particular, we show in this section that the optimal revenue refunding scheme for the discrete Gini coefficient measure corresponds to a natural max-min refunding scheme wherein the refunds are given to users belonging to the lowest income groups.

We first present the discrete Gini coefficient measure and discuss some of its properties.

Definition 5 (Discrete Gini Coefficient [6]). *Let the mean income corresponding to the income distribution \mathbf{q} with a demand vector $\mathbf{d} = \{d_g : g \in G\}$ be $\Delta(\mathbf{q}) = \frac{\sum_{g \in \mathcal{G}} q_g d_g}{\sum_{g \in \mathcal{G}} d_g}$. Then, the discrete Gini coefficient W is given by $W(\mathbf{q}) = \frac{1}{2(\sum_{g \in \mathcal{G}} d_g)^2 \Delta(\mathbf{q})} \sum_{g_1, g_2 \in \mathcal{G}} d_{g_1} d_{g_2} |q_{g_1} - q_{g_2}|$.*

A few comments about the discrete Gini coefficient as a wealth inequality measure are in order. First, the discrete Gini coefficient satisfies the scale independence and constant income transfer properties (presented in Section III-C) required for it to be a valid wealth inequality measure and we present a proof of this claim in the extended version of our work [31]. Next, the discrete Gini coefficient is zero if all users have the same income, i.e., there is perfect equality in society. Furthermore, due to the absolute value of the difference between user incomes in the numerator, the discrete Gini coefficient is larger if the dispersion of incomes between different user groups is greater. Finally, note that we do not write the discrete Gini coefficient measure as a function of the vector of demands $\mathbf{d} = \{d_g : g \in G\}$ as we assume that user demands are fixed in this work.

For the discrete Gini coefficient, we now present a mathematical program for computing the revenue refunding policy r^* . To this end, we first observe by Lemma 1 that for any user-favorable CPRR scheme (τ^*, r^*) each user's ex-post income is given by $q_g(\tau^*, r^*) = q_g(\mathbf{0}, \mathbf{0}) + c_g$ (where, for ease of exposition, we let $\beta = 1$) for some $c_g \geq 0$, where $\sum_{g \in \mathcal{G}} c_g d_g = C_0 - C_{\tau^*}$. Thus, the choice of the optimal revenue refunds r^* can be reduced to computing the optimal transfers c_g . In particular, we formulate the computation of the optimal transfers c_g to minimize the discrete Gini coefficient through the following optimization problem:

$$\min_{c_g \geq 0, \forall g \in \mathcal{G}} W(\mathbf{q}(\mathbf{0}, \mathbf{0}) + \mathbf{c}) \text{ s.t. } \sum_{g \in \mathcal{G}} c_g d_g = C_0 - C_{\tau^*},$$

where $\mathbf{c} = \{c_g : g \in \mathcal{G}\}$ and $\mathbf{q}(\mathbf{0}, \mathbf{0}) + \mathbf{c}$ represents the income distribution of users after receiving the revenue refunds. Furthermore, noting that $\Delta(\mathbf{q}(\mathbf{0}, \mathbf{0}) + \mathbf{c}) = \frac{C_0 - C_{\tau^*} + \sum_{g \in \mathcal{G}} q_g(\mathbf{0}, \mathbf{0}) d_g}{\sum_{g \in \mathcal{G}} d_g}$ is a fixed quantity, the above problem corresponds to a linear program (see Chapter 6 in [32]). The optimal revenue refunding policy r^* corresponding to the above optimization problem results in a natural max-min outcome and we present further formalism and a proof of this claim in the extended version of our paper [31]. In particular, users in the lowest income groups are provided refunds until their incomes equal that of the second lowest income groups, and this process is repeated until all the refunds are exhausted. We note here that this greedy process of refunding revenues to the lowest income groups is reminiscent of Rawl's difference principle of giving the greatest benefit to the most disadvantaged groups of society [33].

C. CPRR Schemes and Endogenous Equilibria

In this section, we consider the setting of the *endogenous* equilibrium, wherein users minimize a linear function of not only their travel times and tolls but also refunds. In particular, we consider two equilibrium notions (without and with user coalitions) in this endogenous setting and show that the optimal CPRR scheme induces equilibria in both settings. To this end, we first consider endogenous equilibria without coalitions and show that

any endogenous equilibrium is an exogenous equilibrium. Next, in the setting of endogenous equilibria *with* coalitions, we show that while, in general, endogenous equilibria do not coincide with exogenous equilibria, the optimal CPRR scheme is robust to coalitions, i.e., any exogenous equilibrium induced by an optimal CPRR scheme is also an endogenous equilibrium.

1) *Endogenous Equilibria without Coalitions*: We begin by considering the setting of an endogenous equilibrium without user coalitions and show that endogenous and exogenous equilibria are equivalent. In this setting without user coalitions, the definition of an exogenous equilibrium (Definition 3) can be readily extended to the setting when users additionally account for refunds in their travel cost minimization, as is elucidated by the following definition. In particular, for a given CPRR scheme (τ, \mathbf{r}) , a path flow pattern \mathbf{f} is an endogenous (τ, \mathbf{r}) -equilibrium without coalitions if for each group $g \in \mathcal{G}$ it holds that $f_{P,g} > 0$ for some path $P \in \mathcal{P}_g$ if and only if $\mu_P^g(\mathbf{f}, \tau, \mathbf{r}) \leq \mu_Q^g(\mathbf{f}, \tau, \mathbf{r})$, for all $Q \in \mathcal{P}_g$.

Given this notion of an endogenous equilibrium without coalitions, we show in the extended version of our paper [31] that any exogenous equilibrium is also an endogenous equilibrium without coalitions and vice versa, i.e., the two equilibrium concepts are equivalent. This result follows naturally since we are in the setting of a non-atomic congestion game, wherein users are infinitesimal, and thus a unilateral deviation by any user will not influence their overall refunds since the flow of users remains unchanged and the tolls are fixed.

2) *Endogenous Equilibria with Coalitions*: We now consider the stronger endogenous equilibrium notion wherein *coalitions* of users minimize a linear function of not only their travel times and tolls but also refunds. In particular, we consider the setting wherein each user group is a coalition. Note that unlike the setting without coalitions, in this setting, a change in the strategy of the entire group, i.e., the flow sent on each feasible path, will likely result in a change in the overall network flow and correspondingly the revenues obtained by users in the group. In the presence of coalitions, we show that while exogenous equilibria and endogenous equilibria with coalitions do not agree in general, any exogenous equilibrium induced by an optimal CPRR scheme is also an endogenous equilibrium with coalitions.

To this end, we begin by introducing the notion of an endogenous equilibrium with coalitions.

Definition 6 (Endogenous Equilibrium with Coalitions). *Let (τ, \mathbf{r}) be a CPRR scheme, and \mathbf{f} be a flow pattern. Then \mathbf{f} is an endogenous (τ, \mathbf{r}) -equilibrium with coalitions if for each group $g \in \mathcal{G}$, every path $P \in \mathcal{P}_g$ such that $f_{P,g} > 0$, and any flow pattern \mathbf{f}' such that $f'_{P',g'} = f_{P',g'}$, for all $g' \in \mathcal{G} \setminus \{g\}, P' \in \mathcal{P}_{g'}$, it holds that $\mu_P^g(\mathbf{f}, \tau, \mathbf{r}(\mathbf{f}, \tau)) \leq \mu_Q^g(\mathbf{f}', \tau, \mathbf{r}(\mathbf{f}', \tau))$, for all $Q \in \mathcal{P}_g$. Here \mathbf{f}' denotes a flow that results from \mathbf{f} where exactly one group changes its path assignment.*

A few comments about the above definition are in

order. First, it is clear that the above definition of endogenous equilibrium is a stronger notion than the standard Nash equilibrium considered in non-atomic congestion games. This is because every endogenous equilibrium is a Nash equilibrium when users minimize their travel costs including refunds but not every Nash equilibrium is necessarily an endogenous equilibrium.

Next, we restrict the set of possible coalitions to those corresponding to strategies for a given user group. This is often reasonable, since users belonging to similar income levels that make similar trips, i.e., travel between the same O-D pair, are more likely to be socially connected with each other and share travel information as compared to users across groups. As a result, we do not consider the setting of equilibrium formation that is robust to any arbitrary set of coalitions [34], and defer this as an interesting direction for future research.

Furthermore, we can view the endogenous equilibrium as a non-atomic analogue of the atomic equilibrium setting, wherein each group g controls a flow of d_g . In atomic settings, each group only sends its flow on one path, whereas in the non-atomic setting, the flows can be dispersed across multiple paths with equal travel costs.

a) *Endogenous Equilibria with Coalitions Differ from Exogenous Equilibria*: We first show that, in general, the endogenous equilibria with coalitions and exogenous equilibria are not the same. To this end, we first recall that an exogenous equilibrium only depends on the tolling scheme τ and is independent of the refunds \mathbf{r} . On the other hand, since users take into account revenue refunds in the case of the endogenous equilibrium, each user must know the refunding policy \mathbf{r} to reason about their strategies when making travel decisions. In particular, each user (and coalition of users within a group) must be able to reason about how a change in their strategy, i.e., the path(s) on which they travel, will change the total amount of refund they receive, and in effect their travel cost. Thus, for this section, we restrict our attention to revenue-refunding schemes resulting from the max-min refunding policy described in Section VI-B. That is, users are given refunds through a process analogous to a max-min allocation. We now construct an example to show that an exogenous τ -equilibrium flow may no longer be an equilibrium when users take into account refunds in their travel cost minimization.

Proposition 2 (Non-Equivalence of Equilibria). *There exists a setting with (i) a two-edge parallel network, (ii) three income classes, and (iii) tolls τ , such that the induced exogenous τ -equilibrium is not an endogenous (τ, \mathbf{r}) -equilibrium with coalitions, where \mathbf{r} results from the max-min revenue refunding policy in Section VI-B.*

For a proof of Proposition 2, see the extended version of our paper [31]. The above proposition is quite natural, since low-income users may take routes that were previously unaffordable when taking into account revenue refunds in their route selection process.

b) *Endogenous Equilibria with Coalitions Coincide with Exogenous Equilibria at the Optimal Solution*:

While Proposition 2 indicates that, in general, the exogenous equilibria and endogenous equilibria with coalitions do not coincide, we now establish that any exogenous equilibrium induced by an optimal user-favorable CPRR scheme (τ^*, r^*) , where the refund satisfies a mild condition, is also an endogenous equilibrium. In particular, we have the following lemma:

Lemma 3 (Optimal CPRR Scheme under Endogenous Equilibria). *Let (τ^*, r^*) be an optimal user-favorable CPRR scheme under the exogenous equilibrium model and let \mathbf{f}^* be its exogenous equilibrium. In addition, let \mathbf{f}^0 be a 0-equilibrium and C^* be the minimum total system cost. Further, suppose that the refunding scheme r^* is defined as $r_g^* := \mu^g(\tau^*, \mathbf{0}) - \mu^g(\mathbf{0}, \mathbf{0}) + c_g(C^*)$, where the non-negative transfer $c_g(C(\mathbf{f}))$ for each group g is monotonically non-increasing in the total system cost $C(\mathbf{f})$ for a given flow \mathbf{f} . Then \mathbf{f}^* is also an endogenous (τ^*, r^*) -equilibrium with coalitions.*

For a proof of Lemma 3, see Appendix C. We note that the condition in Lemma 3 that the non-negative transfer c_g for any group g is monotonically non-increasing in the total system cost is not demanding. For instance, the optimal refunding scheme, i.e., the one minimizing wealth inequality, for the discrete Gini coefficient respects this monotonicity relation, as described in Section VI-B.

VII. DISCUSSION AND FUTURE WORK

In this paper, we studied user-favorable CPRR schemes that mitigate the regressive wealth inequality effects of congestion pricing. Our work demonstrated that if we look at congestion pricing from the lens of refunding the collected tolls, then we can simultaneously achieve the *economic* and *equity* goals of sustainable transportation. Thus, we view our work as a significant step in shifting the discussion around congestion pricing from one focused on the inequity impacts of tolls to one that centers around how to best distribute the revenues collected to different sections of society. For a more in-depth discussion on how our work paves the way for the design of sustainable, publicly-acceptable congestion-pricing schemes and its associated practical challenges, we refer to the extended version of our paper [31].

There are several interesting directions for further research. The first would be to relax some of the commonly-used assumptions in transportation research and game theory, e.g., considering time-varying travel demand or travel cost functions that are non-linear in the travel times, tolls, and refunds. Next, since we only consider direct refunds to road users, it would be worthwhile to extend our framework to analyze system designs with cross subsidies across multiple forms of transport, e.g., subsidies to improve the transit infrastructure. Finally, it would be interesting to go beyond the direct lump-sum transfers of the collected revenues studied in this work and investigate more general group-specific differential congestion pricing mechanisms wherein the price on a given path may differ by user group.

APPENDIX

A. Proof of Proposition 1

Consider the refunds $r_g = \mu^g(\tau, \mathbf{0}) - \mu^g(\mathbf{0}, \mathbf{0}) + \frac{1}{\sum_{g \in \mathcal{G}} d_g} (C_0 - C_\tau)$ for each user in group g . Through an argument similar to that in [7, Theorem 1], it can be shown that the corresponding CPRR scheme is user-favorable, which we present in the extended version of this paper [31]. We now show that under this revenue refunding scheme, the ex-post income distribution $\hat{\mathbf{q}} = \mathbf{q}(\tau, \mathbf{r})$ has a lower wealth inequality measure relative to the untolled user equilibrium ex-post income distribution $\tilde{\mathbf{q}} = \mathbf{q}(\mathbf{0}, \mathbf{0})$. That is, we show that $W(\hat{\mathbf{q}}) \leq W(\tilde{\mathbf{q}})$.

To see this, we begin by considering the ex-ante income distribution \mathbf{q}^0 . Under the untolled user equilibrium, users in group g incur a travel cost $\mu^g(\mathbf{0}, \mathbf{0})$, and thus the ex-post income distribution of users in group g is given by $\tilde{q}_g = q_g^0 - \beta \mu^g(\mathbf{0}, \mathbf{0})$, where β is the scaling factor as in Definition 2. On the other hand, under the CPRR scheme (τ, \mathbf{r}) , the ex-post income distribution of users in group g is given by

$$\hat{q}_g = q_g^0 - \beta (\mu^g(\tau, \mathbf{0}) - r_g) = \tilde{q}_g + \beta \frac{1}{\sum_{g \in \mathcal{G}} d_g} (C_0 - C_\tau),$$

where we used that $\tilde{q}_g = q_g^0 - \beta \mu^g(\mathbf{0}, \mathbf{0})$. Since the above relation holds for all groups g , $\hat{\mathbf{q}} = \tilde{\mathbf{q}} + \lambda \mathbf{1}$, where $\lambda = \frac{\beta}{\sum_{g \in \mathcal{G}} d_g} (C_0 - C_\tau) \geq 0$. Finally, the result that $W(\hat{\mathbf{q}}) \leq W(\tilde{\mathbf{q}})$ follows by the constant income transfer property (Section III), establishing our claim. \square

B. Proof of Corollary 1

Consider the same user-favorable CPRR scheme (τ, \mathbf{r}) as in the proof of Proposition 1. We now show that $W(\hat{\mathbf{q}}) \leq W(\mathbf{q})$, where $\hat{\mathbf{q}} = \mathbf{q}(\tau, \mathbf{r})$. To see this, we first show that $W(\mathbf{q}(\mathbf{0}, \mathbf{0})) = W(\mathbf{q}^0)$, which follows from the observation that for any 0-equilibrium flow \mathbf{f}^0 all users incur the same travel time, denoted as γ , since they travel between the same O-D pair. This observation leads to a travel cost of $\mu^g(\mathbf{0}, \mathbf{0}) = \omega q_g^0 \gamma$ for each group g . Then, for the untolled setting, the ex-post income distribution of users in group g is given by

$$\tilde{q}_g = q_g^0 - \beta \mu^g(\mathbf{0}, \mathbf{0}) = q_g^0 - \beta \omega q_g^0 \gamma = q_g^0 (1 - \beta \omega \gamma).$$

From the above, it follows that $\tilde{\mathbf{q}} = \lambda_1 \mathbf{q}^0$ for $\lambda_1 = 1 - \beta \omega \gamma$. Thus, for β small enough it holds that $\lambda_1 > 0$. Under this condition, due to the scale-independence property (Section III) of the wealth-inequality measure it follows that $W(\tilde{\mathbf{q}}) = W(\mathbf{q}^0)$. Finally, since $W(\hat{\mathbf{q}}) \leq W(\tilde{\mathbf{q}})$ by the proof of Proposition 1 it follows that $W(\hat{\mathbf{q}}) \leq W(\tilde{\mathbf{q}}) = W(\mathbf{q}^0)$, which proves our claim. \square

C. Proof of Lemma 3

For any user-favorable CPRR scheme (τ^*, r^*) it holds for some c_g for each group g that the travel cost to users in group g under the exogenous τ^* -equilibrium \mathbf{f}^* is given by $r_g^* = \mu^g(\tau^*, \mathbf{0}) - \mu^g(\mathbf{0}, \mathbf{0}) + c_g$, where $c_g \geq 0$ and $\sum_{g \in \mathcal{G}} c_g d_g = C_0 - C^*$.

We now consider the emerging behavior of users for the endogenous setting. Since $\mu^g(\mathbf{0}, \mathbf{0})$ is a fixed quantity representing the travel cost at the untolled $\mathbf{0}$ -equilibrium \mathbf{f}^0 , the best response of any coalition within a group g under the endogenous equilibrium, when minimizing each user's individual travel cost $\mu^g(\mathbf{0}, \mathbf{0}) - c_g$ (see the analysis in Lemma 1), is to maximize c_g .

Next, since for each user group g , c_g is monotonically non-decreasing in $C_0 - C(\mathbf{f})$, we have that c_g is maximized for each user group g when $C_0 - C(\mathbf{f})$ is maximized. Since C_0 is fixed, $C_0 - C(\mathbf{f})$ is maximized for any flow \mathbf{f}^* with the minimum total system cost C^* . This implies that each user's non-negative transfer c_g is maximized for any flow \mathbf{f}^* with the minimum total system cost. Thus, any exogenous τ^* -equilibrium flow \mathbf{f}^* that achieves the minimum total system cost is also an endogenous equilibrium with coalitions, since a deviation by any coalition of users in group g can never result in a higher non-negative transfer c_g than that at the minimum total system cost solution. \square

D. Numerical Experiments

In this section, we present numerical experiments to demonstrate the efficacy of optimal CPRR schemes in reducing the total system cost without increasing wealth inequality. We also show that the benefits of CPRR can even be realized in the setting when users' values of time are not known to the central planner. To this end, we conducted experiments on four traffic networks and present the corresponding results in Table I, which presents the relative percentage differences of the total system cost and wealth inequality of the ex-post income distribution for the optimal CPRR scheme and the one under incomplete information to the user equilibrium outcome without tolls and refunds. For a detailed discussion on the implementation details of our experiments as well as the chosen network structures, O-D demands, travel-time functions, user values of time, and incomes, we refer to the extended version of our paper [31].

We first note from columns 1 and 3 of Table I that the optimal CPRR scheme, as expected, reduces the total system cost and discrete Gini coefficient compared to the user equilibrium setting with no tolls or refunds, thereby corroborating Proposition 1. In addition, since users' values of time are assumed to be scaled proportions of their incomes for the experiments [35], our results for the optimal CPRR scheme for single O-D pair demand also corroborate Corollary 1 (see our extended paper [31]).

In addition to evaluating the performance of optimal CPRR schemes, we also perform experiments in the incomplete information setting when user specific values of time or incomes may not be known, as is often the case in practice. In this incomplete information setting, we only assume access to the mean values of time and incomes of users and provide all users travelling between a given O-D pair the same refund, i.e., we consider anonymous refunding schemes as in [7]. Our results in Table I indicate that deploying CPRR schemes in this setting generally results in total system costs and

TABLE I

RELATIVE PERCENTAGE DIFFERENCES OF THE TOTAL SYSTEM COST (COLUMNS 1 AND 2) AND WEALTH INEQUALITY (COLUMNS 3 AND 4), EVALUATED BY THE DISCRETE GINI COEFFICIENT, OF THE OPTIMAL CPRR SCHEME (COLUMNS 1 AND 3) AND THAT WITH INCOMPLETE INFORMATION (COLUMNS 2 AND 4) COMPARED TO THE USER EQUILIBRIUM OUTCOME WITHOUT TOLLS ON FOUR TRAFFIC NETWORKS: (I) PIGOU NETWORK, (II) FOUR EDGE PARALLEL NETWORK, (III) SERIES PARALLEL NETWORK, AND (IV) GRID NETWORK. FOR THE GRID NETWORK, TWO O-D PAIRS WERE CONSIDERED FOR THREE SETTINGS DEPENDING ON THE DEGREE OF VARIANCE OF USERS' VALUES OF TIME, I.E., LOW, MEDIUM, OR HIGH. HERE, C_I AND q^I DENOTE THE TOTAL SYSTEM COST AND EX-POST INCOME DISTRIBUTION FOR THE SCHEME WITH INCOMPLETE INFORMATION, $W^* = W(\mathbf{q}(\tau^*, \mathbf{r}^*))$, AND $W^0 = W(\mathbf{q}(\mathbf{0}, \mathbf{0}))$.

Experiment	$\frac{C_0 - C^*}{C_0}$	$\frac{C_0 - C_I}{C_0}$	$\frac{W^0 - W^*}{W^0}$	$\frac{W^0 - W(q^I)}{W^0}$
Pigou (2 edge)	5.1147	5.1029	0.0357	0.0297
Parallel (4 edge)	4.1343	4.1223	0.0167	0.0134
Series-Parallel	4.8809	4.8331	0.0609	0.0554
Grid (Low Var)	0.9910	0.9834	0.0107	0.0071
Grid (Med Var)	1.4824	1.3062	0.0161	0.0070
Grid (High Var)	2.3365	1.6787	0.0253	0.0070

level of wealth inequality that are higher than that of the optimal CPRR schemes in the complete information setting but lower than that corresponding to the user equilibrium setting with no tolls and refunds. Table I also indicates that the performance of the CPRR scheme with incomplete information depends on the variance in the user values of time and income around the mean. In particular, Table I indicates that as the variance in user values of time is decreased, the CPRR scheme with incomplete information achieves a performance closer to that of the optimal CPRR scheme on both total system cost and wealth inequality metrics.

ACKNOWLEDGMENT

The authors thank Alon Eden and Kira Goldner for insightful discussions on non-atomic congestion games.

REFERENCES

- [1] K. A. Small and J. Yan, "The value of "value pricing" of roads: Second-best pricing and product differentiation," *Journal of Urban Economics*, vol. 49, no. 2, pp. 310 – 336, 2001.
- [2] M. Manville and E. Goldman, "Would congestion pricing harm the poor? do free roads help the poor?" *Journal of Planning Education and Research*, vol. 38, no. 3, pp. 329–344, 2018.
- [3] K. Gemici, E. Koutsopoulos, B. Monnot, C. H. Papadimitriou, and G. Piliouras, "Wealth Inequality and the Price of Anarchy," in *International Symposium on Theoretical Aspects of Computer Science*, ser. Leibniz International Proceedings in Informatics, vol. 126. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, 2019, pp. 31:1–31:16.
- [4] A. Paybarah, "Congestion pricing: Mass transit savior or tax on the working class?" *New York Times*, March 2019.
- [5] S. Jaensirisak, M. Wardman, and A. D. May, "Explaining variations in public acceptability of road pricing schemes," *Journal of Transport Economics and Policy*, vol. 39, no. 2, pp. 127–153, 2005.
- [6] D. Wu, Y. Yin, S. Lawphongpanich, and H. Yang, "Design of more equitable congestion pricing and tradable credit schemes for multimodal transportation networks," *Transportation Research Part B: Methodological*, vol. 46, no. 9, pp. 1273–1287, 2012.
- [7] X. Guo and H. Yang, "Pareto-improving congestion pricing and revenue refunding with multiple user classes," *Transportation Research Part B: Methodological*, vol. 44, no. 8, pp. 972–982, 2010.
- [8] R. Cole, Y. Dodis, and T. Roughgarden, "Pricing network edges for heterogeneous selfish users," in *Symposium on Theory of Computing*. Association for Computing Machinery, 2003, p. 521–530.

- [9] L. Fleischer, K. Jain, and M. Mahdian, "Tolls for heterogeneous selfish users in multicommodity networks and generalized congestion games," in *Symposium on Foundations of Computer Science*. IEEE, 2004, pp. 277–285.
- [10] Z. Pan and T. Bacsar, "H_∞-optimal control for singularly perturbed systems. part i: Perfect state measurements," *Automatica*, vol. 29, no. 2, pp. 401–423, 1993.
- [11] E. A. Sofronova and A. I. Diveev, "Traffic flows optimal control problem with full information," in *2020 IEEE 23rd International Conference on Intelligent Transportation Systems (ITSC)*. IEEE, 2020, pp. 1–6.
- [12] O. Sabag, G. Goel, S. Lale, and B. Hassibi, "Regret-optimal controller for the full-information problem," in *2021 American Control Conference (ACC)*. IEEE, 2021, pp. 4777–4782.
- [13] K. Arora, F. Zheng, and K. Girotra, "Private vs. pooled transportation: Customer preference and congestion management," in *INFORMS International Conference on Service Science*. Springer, 2020, pp. 59–74.
- [14] N. Buchholz, "Spatial Equilibrium, Search Frictions, and Dynamic Efficiency in the Taxi Industry," *The Review of Economic Studies*, vol. 89, no. 2, pp. 556–591, 09 2021. [Online]. Available: <https://doi.org/10.1093/restud/rdab050>
- [15] P. Cohen, R. Hahn, J. Hall, S. Levitt, and R. Metcalfe, "Using big data to estimate consumer surplus: The case of uber," National Bureau of Economic Research, Tech. Rep., 2016.
- [16] D. Bertsimas, V. Farias, and N. Trichakis, "The price of fairness," *Operations Research*, vol. 59, no. 1, pp. 17–31, 2011.
- [17] C. Dwork, M. Hardt, T. Pitassi, O. Reingold, and R. Zemel, "Fairness through awareness," in *Innovations in Theoretical Computer Science Conference*. New York, NY, USA: Association for Computing Machinery, 2012, p. 214–226.
- [18] O. Jahn, R. Möhring, A. Schulz, and N. Stier-Moses, "System-optimal routing of traffic flows with user constraints in networks with congestion," *Operations Research*, vol. 53, no. 4, pp. 600–616, 2005.
- [19] M. L. Weitzman, "Is the price system or rationing more efficient in getting a commodity to those who need it most?" *Bell Journal of Economics*, vol. 8, no. 2, pp. 517–524, 1977.
- [20] M. Akbarpour, P. Dworczak, and S. D. Kominers, "Redistributive allocation mechanisms," GRAPE Group for Research in Applied Economics, GRAPE Working Papers 40, 2020. [Online]. Available: <https://ideas.repec.org/p/fme/wpaper/40.html>
- [21] T. Besley and S. Coate, "Public provision of private goods and the redistribution of income," *The American Economic Review*, vol. 81, no. 4, pp. 979–984, 1991.
- [22] D. Condorelli, "Market and non-market mechanisms for the optimal allocation of scarce resources," *Games and Economic Behavior*, vol. 82, pp. 582–591, 2013.
- [23] K. A. Small, "Using the revenues from congestion pricing," *Transportation*, vol. 19, no. 4, pp. 359–381, 1992.
- [24] P. B. Goodwin, "The rule of three: a possible solution to the political problem of competing objectives for road pricing," *Traffic engineering & control*, vol. 30, no. 10, pp. 495–497, 1989.
- [25] W. S. Vickrey, "Congestion theory and transport investment," *The American Economic Review*, vol. 59, no. 2, pp. 251–260, 1969.
- [26] R. Arnott, A. de Palma, and R. Lindsey, "The welfare effects of congestion tolls with heterogeneous commuters," *Journal of Transport Economics and Policy*, vol. 28, no. 2, pp. 139–161, 1994.
- [27] J. Eliasson, "Road pricing with limited information and heterogeneous users: A successful case," *The annals of regional science*, vol. 35, no. 4, pp. 595–604, 2001.
- [28] H. Yang and H.-J. Huang, "The multi-class, multi-criteria traffic network equilibrium and systems optimum problem," *Transportation Research Part B: Methodological*, vol. 38, no. 1, pp. 1–15, 2004.
- [29] T. Sittiyot, P. Budsaratragoon, and K. Holasut, "A scaling perspective on the distribution of executive compensation," *Physica A: Statistical Mechanics and its Applications*, vol. 543, p. 123556, 2020. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0378437119319818>
- [30] F. Bourguignon, "Decomposable income inequality measures," *Econometrica*, vol. 47, no. 4, pp. 901–920, 1979. [Online]. Available: <http://www.jstor.org/stable/1914138>
- [31] D. Jalota, K. Solovey, K. Gopalakrishnan, S. Zoepf, H. Balakrishnan, and M. Pavone, "When efficiency meets equity in congestion pricing and revenue refunding schemes," *CoRR*, vol. abs/2106.10407, 2021.

- [32] J. Bisschop, *AIMMS optimization modeling*. lulu.com, 2016.
- [33] J. Rawls, *A theory of justice*. Harvard university press, 2020.
- [34] N. Andelman, M. Feldman, and Y. Mansour, "Strong price of anarchy," *Games and Economic Behavior*, vol. 65, no. 2, pp. 289–317, 2009.
- [35] D. Jalota, D. Paccagnan, M. Schiffer, and M. Pavone, "Online routing over parallel networks: Deterministic limits and data-driven enhancements," *INFORMS Journal on Computing*, vol. 0, no. 0, p. null, 0. [Online]. Available: <https://doi.org/10.1287/ijoc.2023.1275>



Devansh Jalota received the B.S. degree in civil and environmental engineering and B.A. degree in applied mathematics at University of California, Berkeley, CA, USA. He is currently working toward the Ph.D. degree with the Institute for Computational and Mathematical Engineering at Stanford University, Stanford, CA, USA.



Kiril Solovey received the Ph.D. degree in computer science from Tel Aviv University, Tel Aviv, Israel. He is an Assistant Professor with the Faculty of Electrical and Computer Engineering at Technion - Israel Institute of Technology, Haifa, Israel.



Karthik Gopalakrishnan is a postdoctoral scholar in the Autonomous Systems Lab at Stanford University. His research focuses on developing optimization and control algorithms for networked systems, with a particular focus on transportation systems.



Stephen Zoepf is the Acting Director of the Highly Automated Systems Safety Center of Excellence (HASS COE) and Chief Analyst for the Office of the Assistant Secretary for Research and Technology at the U.S. Department of Transportation. He holds a Ph.D., M.Sc. and B.Sc. from MIT and has two decades of experience in transportation and mobility.



Hamsa Balakrishnan is the William E. Leonhard (1940) Professor of Aeronautics and Astronautics at the Massachusetts Institute of Technology (MIT). Her research is in the design, analysis, and implementation of control and optimization algorithms for large-scale cyber-physical infrastructures, with an emphasis on air transportation systems.



Marco Pavone (Member, IEEE) received the Ph.D. degree in aeronautics and astronautics from the Massachusetts Institute of Technology (MIT), Cambridge, MA, USA, in 2010. He is an Associate Professor of aeronautics and astronautics with Stanford University, Stanford, CA, USA, where he is the Director of the Autonomous Systems Laboratory. Dr. Pavone is currently an Associate Editor for the IEEE Control Systems Magazine.