

Information Design in Bayesian Routing Games

Original

Information Design in Bayesian Routing Games / Cianfanelli, L., Ambrogio, A., Como, G.. - (2023), pp. 3945-3949. (62nd IEEE Conference on Decision and Control, CDC 2023 Singapore (Singapore) 13-15 December 202) [10.1109/cdc49753.2023.10384300].

Availability:

This version is available at: 11583/3009257 since: 2026-03-26T13:57:11Z

Publisher:

Institute of Electrical and Electronics Engineers - IEEE

Published

DOI:10.1109/cdc49753.2023.10384300

Terms of use:

This article is made available under terms and conditions as specified in the corresponding bibliographic description in the repository

Publisher copyright

IEEE postprint/Author's Accepted Manuscript

©2023 IEEE. Personal use of this material is permitted. Permission from IEEE must be obtained for all other uses, in any current or future media, including reprinting/republishing this material for advertising or promotional purposes, creating new collecting works, for resale or lists, or reuse of any copyrighted component of this work in other works.

(Article begins on next page)

Information design in Bayesian routing games

Leonardo Cianfanelli, Alexia Ambrogio and Giacomo Como

Abstract—We study optimal information provision in transportation networks when users are strategic and the network state is uncertain. An omniscient planner observes the network state and discloses information to the users with the goal of minimizing the expected travel time at the user equilibrium. Public signal policies, including full-information disclosure, are known to be inefficient in achieving optimality. For this reason, we focus on private signals and restrict without loss of generality the analysis to signals that coincide with path recommendations that satisfy *obedience constraints*, namely users have no incentive in deviating from the received recommendation according to their posterior belief. We first formulate the general problem and analyze its properties for arbitrary network topologies and delay functions. Then, we consider the case of two parallel links with affine delay functions, and provide sufficient conditions under which optimality can be achieved by information design. Interestingly, we observe that the system benefits from uncertainty, namely it is easier for the planner to achieve optimality when the variance of the uncertain parameters is large. We then provide an example where optimality can be achieved even if the sufficient conditions for optimality are not met.

Index Terms—Information design, Bayesian routing games.

I. INTRODUCTION

Routing games describe the behaviour of a non-atomic set of strategic users travelling in transportation networks [1]. Each link of the network is endowed with a delay function that returns the travel time along the link as a function of the flow, and users aim at selecting paths with minimum travel time. A user equilibrium of the game is a flow such that no user has an incentive in deviating to other paths from the one she has selected. Typically, user equilibria are suboptimal in terms of total travel time on the network. The goal of a network planner is then to influence the user behaviour by incentive mechanisms to steer the equilibrium of the game towards the system optimum. Standard approaches are to charge tolls [2], [3], or to intervene on the network structure [4], [5]. When the state of the network is uncertain (we refer to this class of games by *Bayesian routing games*), this uncertainty can be leveraged by an omniscient planner that observes the state of the network and discloses information to the users. How to disclose this information to minimize the system cost at equilibrium is called in the literature *information design problem*.

The role of information in routing games has been widely studied in the literature. Informational Braess' paradox shows that the travel time at the equilibrium may decrease when a

fraction of users are not aware of some roads of the network [6]. In [7], [8] the authors analyze how the efficiency of the equilibrium varies in terms of the fraction of adopters of information providers, showing that too much information might hurt the system. The information design problem in routing games has been first considered in [9], where the authors show by examples that in routing games revealing full-information and all public signal policies are in general suboptimal and providing private information is preferable. In [10] it is studied an information design problem with two parallel links and affine delay functions where the free-flow delay of one link is a discrete random variable, and sufficient conditions for optimality of private information are provided. In [11], the authors focus on computational perspectives and propose an algorithm to solve efficiently the information design problem with polynomial delay functions and arbitrary network topologies. Dynamical information provision is studied in [12]–[14].

Motivated by the suboptimality of public signal policies, in this paper we consider private information provision in Bayesian routing games. Compared to [10], we formulate Bayesian routing games and the corresponding information design problem for arbitrary network topologies and arbitrary delay functions described by continuous random variables. We show how to equivalently reformulate the problem (due to the revelation principle [15], [16]) by restricting the signal space to path recommendations such that users have no incentive in deviating from the received recommendation, which is known as *obedience constraint*. The policy then specifies the fraction of users that receive each recommendation as a function of the network state. The users are not aware of the recommendations received by other ones, but the prior distribution of the network state and the information policy are publicly known.

After analyzing the properties of the general problem, we focus on the special case of two links with affine delay functions, which is a convex problem. We show that under a restrictive assumption on the support of the random variables that describe the delay functions, there exist sufficient and necessary conditions on the moments of the random variables under which the optimality can be achieved by information design. Interestingly, we show that the uncertainty on the network state is beneficial for the system under the optimal information policy, namely the planner is able to persuade the users to distribute according the system optimum flow, which typically does not happen when users have full-information on the network state. We also remark that the approach can be generalized to arbitrary topologies although the corresponding information design problem is in general

The authors are with Dipartimento di Scienze Matematiche, Politecnico di Torino. Giacomo Como is also with the Department of Automatic Control, Lund University (e-mail: leonardo.cianfanelli@polito.it, alexia.ambrogio@studenti.polito.it, giacomo.como@polito.it, webpage: <https://sites.google.com/view/leonardocianfanelli/>, <https://staff.polito.it/giacomo.como/>)

non-convex, which is left for future research. Finally, we analyze a particular case that allows explicit computation, showing that optimality can be achieved even if the restrictive assumption on the support of the random variables is not met.

The paper is organized as follows. In Section II we present Bayesian routing games, formulate the information design problem, and analyze its properties. We then present the results in Section III. In Section IV we summarize the contribution and discuss future research lines. The proofs of the theoretical results are omitted and can be found in the extended version of this paper [17].

A. Notation

Given a vector x , we let x' denote its transpose. $\mathbf{1}$, \mathbf{I} and $\delta^{(i)}$ denote the vector of all ones, the identity matrix the vector with 1 in i -th entry $\delta_i^{(i)}$ and 0 in all the other positions, whose size may be deduced from the context. For a finite set \mathcal{P} , let $\Delta_{\mathcal{P}}$ denote the simplex on \mathcal{P} , i.e.,

$$\Delta_{\mathcal{P}} = \{x \in \mathbb{R}_+^{\mathcal{P}} : \mathbf{1}'x = 1\}.$$

Given x, a, b in \mathbb{R} , with $a \leq b$, we let

$$[x]_a^b = \max\{a, \min\{x, b\}\} = \begin{cases} a & \text{if } x < a \\ x & \text{if } x \in [a, b] \\ b & \text{if } x > b. \end{cases}$$

II. MODEL AND PROBLEM FORMULATION

We present the model in Section II-A, and formulate the problem in Section II-B.

A. Model

We model the transportation network topology as a directed multigraph $\mathcal{G} = (\mathcal{N}, \mathcal{E})$, where \mathcal{N} is the set of nodes and \mathcal{E} is the set of links. Every link e in \mathcal{E} is directed from its tail node $\sigma(e)$ in \mathcal{N} to its head node $\xi(e)$ in \mathcal{N} . The transportation network topology is fully characterized by its node-link incidence matrix B in $\mathbb{R}^{\mathcal{N} \times \mathcal{E}}$, with entries $B_{ne} = 1$ if $\sigma(e) = n$, -1 if $\xi(e) = n$, or 0 otherwise.

We denote by o and d in \mathcal{N} , with $o \neq d$, the origin and, respectively, the destination nodes. We assume that d is reachable from o . Let a unitary mass of non-atomic users travel from o to d , and define $\nu = \delta^{(o)} - \delta^{(d)}$. We then let

$$\mathcal{F} = \{f \in \mathbb{R}_+^{\mathcal{E}} : Bf = \nu\}$$

be the set of unitary o - d flows, where a flow describes how users are distributed across links.

The state of the network is described by a random variable θ taking values on a probability space $(\Theta, \mathcal{A}, \mathbf{P})$. Each link of the network is endowed with a delay function $\tau_e : \mathbb{R}_+ \times \Theta \rightarrow \mathbb{R}_+$ that depends on the flow and on the realization of the network state. We shall assume that $\tau_e(f_e, \theta)$ is increasing in f_e for every θ in Θ to consider congestion effects.

A network flow is a function $f : \Theta \rightarrow \mathcal{F}$ that assigns to each network state θ in Θ a vector $f(\theta)$ in \mathcal{F} . We define the system cost as the expected travel time on the network, i.e.,

$$C(f) = \int_{\Theta} \sum_{e \in \mathcal{E}} f_e(\theta) \tau_e(f_e(\theta), \theta) d\mathbf{P}(\theta).$$

Network flows depend on the available information on the network state. A setting of interest is the *full-information* setting, where all users know exactly the realization of the network state θ . Given the information setting, we can define two solution concepts, the *system optimum* flow and the *user equilibrium* flow. System optimum flows describe a scenario where each user is forced to select a path with the goal of minimizing the system cost. However, users are typically strategic and aim at minimizing their own travel time. This behavior is captured by the notion of user equilibrium flow.

In full-information setting, flows are allowed to depend on the network state θ , which is known to the users. We can imagine this as the result of an omniscient planner observing the network state and revealing this information to the users. In this setting, the system optimum is a network flow that depends on θ and minimizes the system cost:

$$f^* \in \arg \min_{f: \Theta \rightarrow \mathcal{F}} \int_{\Theta} \sum_{e \in \mathcal{E}} f_e(\theta) \tau_e(f_e(\theta), \theta) d\mathbf{P}(\theta). \quad (1)$$

Observe that in full-information setting, f^* is independent of the prior distribution of θ , since $f^*(\theta)$ is the unitary o - d flow that minimizes $\sum_{e \in \mathcal{E}} f_e(\theta) \tau_e(f_e(\theta), \theta)$ for a given network state θ . To formalize the notion of user equilibrium, let \mathcal{P} denote the set of paths from o to d and let $z : \Theta \rightarrow \Delta_{\mathcal{P}}$ denote a path flow, describing how users distribute on the paths given the network state. Let A in $\mathbb{R}^{\mathcal{E} \times \mathcal{P}}$ be the link-path incidence matrix, with entries $A_{ei} = 1$ if link e belongs to path i , or 0 otherwise. A path flow z induces a unique network flow $f(\theta) = Az(\theta)$, and the cost of a path i for a network state θ is the sum of the delay functions of the links along the path, i.e.,

$$c_i(f(\theta), \theta) = \sum_{e \in \mathcal{E}} A_{ei} \tau_e(f_e(\theta), \theta).$$

The full-information user equilibrium is a network flow $f^W : \Theta \rightarrow \mathcal{F}$ that admits $z^W : \Theta \rightarrow \Delta_{\mathcal{P}}$ such that $f^W(\theta) = Az^W(\theta)$ and such that, for every i in \mathcal{P} and θ in Θ , $z_i^W(\theta) > 0$ implies

$$c_i(f^W(\theta), \theta) \leq c_j(f^W(\theta), \theta) \quad \forall j \in \mathcal{P},$$

In other words, the full-information user equilibrium is a network flow such that, for each network state θ , each used path has optimal cost, which means that no user has an incentive in deviating to other paths.

In general, it is known that revealing full-information is suboptimal for the system, namely the full-information user equilibrium is suboptimal compared to the full-information system-optimum, which is by definition the best network flow in terms of system cost. Moreover, it is known that any public signal is inefficient in achieving optimality [9]–[11]. For this reason, in the next part of the paper we shall consider how to design optimal private signal policies.

B. Problem formulation

We assume that an omniscient planner observes the realization of the network state, and given this observation sends private signals to the users, with the goal of minimizing the

expected total travel time of the user equilibrium flow. A signal policy is a map $\pi : \Theta \rightarrow \Delta_{\mathcal{P}}$ that assigns to each network state the fraction of users that is recommended to take each path. One might consider more general types of signals, however assuming that signals coincide with path recommendations is without loss of generality, as stated by the revelation principle [15], [16].

A fundamental assumption is that the signal policy and the prior distribution of the network state are known to the users. Once the signals have been sent, users update their belief on the network state. Let $d\mathbf{P}_i$ denote the posterior belief of users that receive recommendation i , which is computed according to Bayes' formula, i.e.,

$$d\mathbf{P}_i(\theta) = \frac{\pi_i(\theta)d\mathbf{P}(\theta)}{\int_{\Theta} \pi_i(\omega)d\mathbf{P}(\omega)}, \quad (2)$$

and let $E_i[\cdot] = \int_{\Theta} \cdot d\mathbf{P}_i(\theta)$ denote the expected value after receiving signal i , i.e., according to posterior belief (2).

As the users update their beliefs on the network state, the user equilibrium varies accordingly. Indeed, users follow the recommendation only if it is a best response to do so. Let y in $\mathbb{R}_+^{\mathcal{P} \times \mathcal{P}}$ be a matrix such that $y\mathbf{1} = \mathbf{1}$, whose element y_{ij} denotes the fraction of users that choose path j among those who receive recommendation of taking path i . We can interpret y as the action distribution of a population game with multiple populations that differ in the received signal. Given π and y , the network flow depends on the network state θ via

$$f_e^{\pi, y}(\theta) = \sum_{j \in \mathcal{P}} A_{ej} \sum_{i \in \mathcal{P}} \pi_i(\theta) y_{ij}, \quad \forall e \in \mathcal{E},$$

or, more compactly,

$$f^{\pi, y}(\theta) = Ay'\pi(\theta).$$

Definition 1 (Bayesian user equilibrium): Given a policy π , a flow $f^{\pi, y}(\theta) = Ay'\pi(\theta)$ is a Bayesian user equilibrium if $\pi_i(\theta)y_{ij} > 0$ for at least a θ in Θ and i, j in \mathcal{P} implies that

$$E_i[c_j(f^{\pi, y}(\theta), \theta)] \leq E_i[c_k(f^{\pi, y}(\theta), \theta)] \quad \forall k \in \mathcal{P}.$$

In other words, we require that if some users receive recommendation i and travels on path j , then path j has to be optimal according to the posterior belief of users that have received recommendation i . The next result characterizes Bayesian user equilibria.

Proposition 1: Given a policy π , a flow $f^{\pi, y^*(\pi)}(\theta) = A(y^*(\pi))'\pi(\theta)$ is a Bayesian user equilibrium if and only if $y^*(\pi)$ is a solution of the convex program

$$y^*(\pi) \in \arg \min_{y \in \mathbb{R}_+^{\mathcal{P} \times \mathcal{P}} : y\mathbf{1} = \mathbf{1}} \Phi_{\pi}(y), \quad (3)$$

where

$$\Phi_{\pi}(y) = \int_{\Theta} \sum_{e \in \mathcal{E}} \int_0^{(Ay'\pi(\theta))_e} \tau_e(\theta, s) ds d\mathbf{P}(\theta).$$

Furthermore, the equilibrium $f^{\pi, y^*(\pi)}$ is unique for every policy π .

The idea of the proof is to show that $y^*(\pi)$ is the Nash equilibrium of a population game that admits a weighted potential, where a population corresponds to the set of users that receive different the same recommendation.

For simplicity of notation, from now on we shall denote by f^{π} the Bayesian user equilibrium corresponding to information policy π . We express the cost of a policy π as the expected total travel time at the corresponding Bayesian user equilibrium, i.e.,

$$C(\pi) = \int_{\Theta} \sum_e f_e^{\pi}(\theta) \tau_e(f_e^{\pi}(\theta), \theta) d\mathbf{P}(\theta).$$

The system planner then wants to solve the problem

$$\pi^* \in \arg \min_{\pi : \Theta \rightarrow \Delta_{\mathcal{P}}} C(\pi). \quad (4)$$

Problem (4) can be reformulated by using the *revelation principle* [15], [16], which states that given a policy π , there always exists a policy $\tilde{\pi}$ such that $C(\pi) = C(\tilde{\pi})$ and $y^*(\tilde{\pi}) = \mathbf{I}$, hence $f^{\tilde{\pi}}(\theta) = A\tilde{\pi}(\theta)$. In other words, this means that restricting the attention to policies under which no user wants to deviate from the received recommendation is without loss of generality. This requirement is known as *obedience constraint*. Revelation principle allows to formulate the information design problem as follows.

Problem 1: The optimal information policy is

$$\pi^* = \arg \min_{\pi : \Theta \rightarrow \Delta_{\mathcal{P}}} \int_{\Theta} \sum_{e \in \mathcal{E}} f_e^{\pi}(\theta) \tau_e(f_e^{\pi}(\theta), \theta) d\mathbf{P}(\theta)$$

subject to

$$\pi_i(\omega) E_i[c_i(f^{\pi}(\theta), \theta) - c_j(f^{\pi}(\theta), \theta)] \leq 0, \quad \forall i, j \in \mathcal{P}, \forall \omega \in \Theta$$

$$f^{\pi}(\theta) = A\pi(\theta).$$

Remark 1: Multiplying the obedience constraints by $\pi_i(\omega)$ allows a distinction in two cases: if some user receives signal i , i.e., $\pi_i(\omega) > 0$ for at least a ω in Θ , then we impose that following the recommendation must be convenient; if no user receives signal i , i.e., $\pi_i(\omega) = 0$ for every ω , then the constraint is trivially satisfied.

Remark 2: Observe that (4) is a bi-level program, namely the planner optimizes the policy π , but the performance of the policy depends on $y^*(\pi)$, which in turn is the result of the convex program (3) that depends on π . Instead, Problem 1 is a single-level optimization problem with obedience constraints, which is more tractable. Notice also that, if the delay functions are convex in the first argument, the objective function of Problem 1 is convex in π independently of the network topology and the prior distribution of the network state. In contrast, the obedience constraints are in general non-convex, and they convex if the network has two parallel links and affine delay functions.

We now define the price of anarchy (PoA).

Definition 2 (Price of anarchy): Let π^* be the solution of Problem 1. Then,

$$PoA = \frac{\int_{\Theta} \sum_{e \in \mathcal{E}} f_e^{\pi^*}(\theta) \tau_e(f_e^{\pi^*}(\theta), \theta) d\mathbf{P}(\theta)}{\int_{\Theta} \sum_{e \in \mathcal{E}} f_e^*(\theta) \tau_e(f_e^*(\theta), \theta) d\mathbf{P}(\theta)} \geq 1.$$

In other words, in the context of information design the price of anarchy measures how suboptimal the optimal policy π^* is compared to the full-information system optimum.

III. INFORMATION DESIGN ON TWO PARALLEL LINKS

In this section we consider the information design problem in a network with two parallel links and affine delay functions. For this case, we first establish in Section III-A sufficient conditions on the support and on the moments of the random variables that describe the delay functions under which $PoA = 1$, and apply these conditions to a case of interest in Section III-B. Then, we consider in Section III-C a case-study showing that optimality can still be achieved even if these sufficient conditions are not met.

A. Arbitrary prior

We consider a graph $\mathcal{G} = (\mathcal{N}, \mathcal{E})$ with two nodes $\mathcal{N} = (o, d)$ and two parallel links from o to d . We assume that the delay functions are affine in the flow, i.e.,

$$\tau_e(a_e, b_e, f_e) = a_e f_e + b_e, \quad e \in \{1, 2\},$$

where a_e, b_e are non-negative random variables. In this setting the policy is a map $\pi(a, b)$ such that $\pi_1(a, b) + \pi_2(a, b) = 1$, hence we can formulate the problem in terms of π_1 . With this in mind, by noticing that the policy depends on b via $x = b_1 - b_2$, and rearranging the terms, the problem reads as follows.

Problem 2: Find $\pi_1^* : \Theta \rightarrow [0, 1]$ to minimize

$$\int_{\Theta} [(x - 2a_2)\pi_1(a, x) + (a_1 + a_2)\pi_1^2(a, x)] d\mathbf{P}(a, x)$$

under obedience constraints

$$\int_{\Theta} [(x - a_2)\pi_1(a, x) + (a_1 + a_2)\pi_1^2(a, x)] d\mathbf{P}(a, x) \leq 0, \quad (5)$$

$$\int_{\Theta} [a_2 - x + (x - a_1 - 2a_2)\pi_1(a, x) + (a_1 + a_2)\pi_1^2(a, x)] d\mathbf{P}(a, x) \leq 0. \quad (6)$$

Lemma 1: Consider Problem 2. Then,

$$f_1^*(a, x) = \left[\frac{2a_2 - x}{2(a_1 + a_2)} \right]_0^1.$$

Observe that in networks with parallel links, under obedience constraints it holds $f^\pi = \pi$, because $A = \mathbf{I}$ and $z = \mathbf{I}$. Hence, if the full-information system optimum f^* satisfies the obedience constraints, then $\pi^* = f^*$ and $PoA = 1$, which means that private information design can achieve optimality. This idea is the core of the next result. To better formalize the next result, let us define the support of a random variable. Given a random variable X in \mathbb{R}^n , its support is the set of values that the random variable can take, i.e.,

$$\text{supp}(X) = \{x \in \mathbb{R}^n : \mathbf{P}(B(x, \epsilon)) > 0 \forall \epsilon > 0\},$$

where $B(x, \epsilon)$ is the open ball centered in x with radius ϵ .

Theorem 1: Consider Problem 2. Assume that:

$$\begin{cases} \max(\text{supp}(x)) \leq 2 \min(\text{supp}(a_2)) \\ \min(\text{supp}(x)) \geq -2 \max(\text{supp}(a_1)). \end{cases} \quad (7)$$

Then, $\pi^* = f^*$ and $PoA = 1$ if and only if

$$\begin{cases} E \left[\frac{2a_2x - x^2}{(a_1 + a_2)} \right] \leq 0 \\ E \left[\frac{-2a_1x - x^2}{(a_1 + a_2)} \right] \leq 0 \end{cases} \quad (8)$$

The proof follows from noticing that, under (7), the policy $\pi = f^*$ satisfies the obedience constraints if and only if (8) is satisfied, guaranteeing that $PoA = 1$.

Remark 3: As already observed, Problem 1 is in general non-convex unless the network is composed of two parallel links with affine delay functions. Hence, finding the optimal information policy π^* is not an easy task in general. However, the approach adopted in Theorem 1 to verify whether the system-optimum flow can be induced by information design can be generalized for arbitrary network topologies. Indeed, if there exists a policy π^* that satisfies the obedience constraints and is compatible with the system-optimum flow (i.e., $f^*(\theta) = A\pi^*(\theta)$), then any other policy cannot achieve a better cost by definition of f^* , proving that π^* is the optimal policy and achieves $PoA = 1$.

Observe by Theorem 1 that, under (7), if $E[x] = 0$ (i.e., the two links have same expected free-flow delay), and a and x are independent, i.e., $E[a_e x] = E[a_e]E[x]$, then the two obedience constraints in (8) are always satisfied and optimality is thus achieved by information policy $\pi_1^* = f_1^* = (2a_2 - x)/2(a_1 + a_2)$. Instead, this does not hold true in general if $E[x] = 0$, but a and x are not independent. We remark that condition (8) is necessary and sufficient for optimality assuming that (7) holds. Condition (7) guarantees that the full-information system optimum does not saturate and allows to express conditions for optimality in terms of moments of the random variables. However, (7) is not necessary for optimality. We shall analyze in Section III-C a case where (7) is violated but optimality is achieved.

B. Deterministic linear coefficients and arbitrary prior

In this section we analyze Problem 2 when a is known and b is the only random variable (hence, also $x = b_1 - b_2$ is random). The next result follows from Theorem 1.

Theorem 2: Consider Problem 2 with a known. Assume that:

$$\text{supp}(x) \subseteq [-2a_1, 2a_2]. \quad (9)$$

Then, $\pi^* = f^*$ and $PoA = 1$ if and only if

$$-\frac{E[x^2]}{2a_1} \leq E[x] \leq \frac{E[x^2]}{2a_2}. \quad (10)$$

Remark 4: It is interesting to establish a parallelism between the price of anarchy in the setting of Theorem 2 and in the standard setting of routing games where x and a are

known to the users. To this end, it proves useful to write (10) as

$$\begin{cases} \sigma_x^2 \geq E[x](2a_2 - E[x]) \\ \sigma_x^2 \geq -E[x](2a_1 + E[x]). \end{cases} \quad (11)$$

where σ_x^2 is the variance of x . If $\sigma_x^2 = 0$ (i.e., x is known) and x belongs to $[-2a_1, 2a_2]$, then Theorem 2 states that optimality can be achieved if and only if

$$x(2a_1 + x) \geq 0 \geq x(2a_2 - x)$$

i.e., when either $x = 2a_2$ or $x = -2a_1$ (all users select the same link at user-equilibrium flow and the equilibrium coincides with the system-optimum), or $x = 0$. For other intermediate values of x , if $\sigma_x^2 = 0$ optimality cannot be achieved. Instead, if x is uncertain, optimality is achievable for intermediate values of $E[x]$ also. In particular, (11) states that the larger the variance is, the easier is to achieve optimality through information design. This shows that the uncertainty on the network state is beneficial for the system, in the sense that an omniscient planner that discloses privately the information on the network state can leverage the uncertainty to persuade the users to distribute according to the system-optimum flow.

C. Deterministic linear coefficients and uniform prior

The theoretical results in Theorems 1 and 2 assume that the support of the random variables is such that under the system-optimum flow f^* the users always distribute on both the links. This assumption indeed allows us to state sufficient and necessary conditions for optimality in terms of moments of the random variables that characterize whether the obedience constraints are active or inactive. In this section we consider a case-study that allows explicit computation even relaxing the assumption on the support. In particular, we assume that

$$d\mathbf{P}(b_1, b_2) = \begin{cases} 1 & \text{if } b_1, b_2 \in [0, 1]^2 \\ 0 & \text{otherwise,} \end{cases} \quad (12)$$

and assume that a is given and non-negative. The next result characterizes the solution of the information design problem.

Theorem 3: Consider Problem 2. Let the prior distribution of b be as in (12), and let a be known. Then, $PoA = 1$ for every a in \mathbb{R}_+^2 .

Notice that specifying the prior distribution of b determines the distribution of x . Moreover, notice that (12) implies that $\text{supp}(x) = [-1, 1]$. Hence, (9) is satisfied if $a_1, a_2 \geq 1/2$. Nonetheless, Theorem 2 states that optimality is achieved for every value of a , extending the results of Theorem 2.

IV. CONCLUSION

In this paper we study optimal information design in Bayesian routing games. The problem investigates how a planner that observes the stochastic state of a transportation network should privately disclose information to strategic users in order to influence their routing behavior, with the goal of steering the system towards the system optimum.

Due to the revelation principle, the analysis focuses without loss of generality on private path recommendations under the assumption that users have no incentive in deviating from the recommendation they receive (obedience constraints). We formalize and analyze the general problem, and then focus on the case of two links and affine delay functions, establishing sufficient conditions on the moments and on the support of the random variables under which the system optimum can be achieved by information design. We then analyze a special case where optimality is achieved even if these sufficient conditions are not satisfied.

This work focuses on when optimality is achievable. Future research lines focus on quantifying how suboptimal is the best information policy when optimality cannot be achieved. Other research lines aim at expanding the analysis to more general network topologies and non-linear delay functions. Finally, it would be interesting to analyze the outcome of multiple information providers competing for customers.

REFERENCES

- [1] J. G. Wardrop, "Road paper. some theoretical aspects of road traffic research." *Proceedings of the institution of civil engineers*, vol. 1, no. 3, pp. 325–362, 1952.
- [2] R. Cole, Y. Dodis, and T. Roughgarden, "How much can taxes help selfish routing?" *Journal of Computer and System Sciences*, vol. 72, no. 3, pp. 444–467, 2006.
- [3] R. Maffiastro and G. Como, "Stability and optimality of multi-scale transportation networks with distributed dynamic tolls," in *2018 IEEE Conference on Decision and Control (CDC)*. IEEE, 2018, pp. 211–216.
- [4] L. Cianfanelli, G. Como, A. Ozdaglar, and F. Parise, "Optimal intervention in transportation networks," *IEEE Transactions on Automatic Control*, 2023.
- [5] R. Z. Farahani, E. Miandoabchi, W. Y. Szeto, and H. Rashidi, "A review of urban transportation network design problems," *European journal of operational research*, vol. 229, no. 2, pp. 281–302, 2013.
- [6] D. Acemoglu, A. Makhdoomi, A. Malekian, and A. Ozdaglar, "Informational braess' paradox: The effect of information on traffic congestion," *Operations Research*, vol. 66, no. 4, pp. 893–917, 2018.
- [7] M. Wu, J. Liu, and S. Amin, "Informational aspects in a class of bayesian congestion games," in *2017 American Control Conference (ACC)*. IEEE, 2017, pp. 3650–3657.
- [8] M. Wu, S. Amin, and A. E. Ozdaglar, "Value of information in bayesian routing games," *Operations Research*, vol. 69, no. 1, pp. 148–163, 2021.
- [9] S. Das, E. Kamenica, and R. Mirka, "Reducing congestion through information design," *55th Annual Allerton Conference*, 2017.
- [10] H. Tavaafoghi and D. Teneketzis, "Informational incentives for congestion games," in *2017 55th Annual Allerton Conference on Communication, Control, and Computing (Allerton)*. IEEE, 2017, pp. 1285–1292.
- [11] Y. Zhu and K. Savla, *Information design in non-atomic routing games with partial participation: computation and properties*. University of Southern California, 2021.
- [12] S. Dughmi and H. Xu, "Algorithmic bayesian persuasion," in *Proceedings of the forty-eighth annual ACM symposium on Theory of Computing*, 2016, pp. 412–425.
- [13] E. Meigs, F. Parise, A. Ozdaglar, and D. Acemoglu, "Optimal dynamic information provision in traffic routing," *arXiv preprint arXiv:2001.03232*, 2020.
- [14] H. Tavaafoghi and D. Teneketzis, "Strategic information provision in routing games," 2019.
- [15] D. Bergemann and S. Morris, "Bayes correlated equilibrium and the comparison of information structures in games," 2016.
- [16] —, "Information design: A unified perspective," *Journal of Economic Literature*, vol. 57, no. 1, pp. 44–95, 2019.
- [17] L. Cianfanelli, A. Ambrogio, and G. Como, "Information design in bayesian routing games," *arXiv preprint arXiv:2309.06859*, 2023.