

Information Signalling with Concurrent Monetary Incentives in Bayesian Congestion Games

Bryce L. Ferguson, Philip N. Brown, and Jason R. Marden

Abstract—The uncertainty held by a system’s users can cause ineffective decision making. Nowhere is this more apparent than in transportation networks, where drivers’ uncertainty over current road/traffic conditions can negatively alter their opinion of which route to take. To alleviate this, an informed system operator may elect to *signal* information to uninformed users with the hope of persuading them to take more preferable actions (e.g., Google/Apple maps providing live traffic updates). In this work, we study public signalling mechanisms in the context of Bayesian congestion games. We find that though revealing information can reduce system cost in some settings, it can also be detrimental and cause worse performance than not signalling at all. However, we find an important relationship between information signalling and monetary incentives: by utilizing both signalling and incentive mechanisms concurrently, the system operator can guarantee that revealing information does not worsen performance. We prove these findings in a general class of Bayesian congestion games and further exemplify them with analytically characterized bounds on the change in system cost from signalling in parallel networks with polynomial latency functions. Finally, we consider the problem of solving for optimal signals with and without the concurrent use of monetary incentives of two types: those that update with the revealed information and those that do not. We construct solvable optimization problems whose solutions give optimal signalling policies even when the signalling policy is limited in its support; we then quantify the benefit of these, and other, signalling mechanisms in numerical examples.

Index Terms—traffic congestion control, routing decisions, information design, monetary incentives, game theory

I. INTRODUCTION

THE degree of traffic congestion on highways and roads in busy city areas is inherently caused by the collective route choices of the drivers [2]. Though drivers often choose routes that minimize their own travel time, the system behavior that emerges from this decision making need not be optimal [3]. This inefficiency can be further exacerbated by drivers’ uncertainty over the state of the system [4], [5], e.g., uncertainty on current weather conditions, traffic rates, or on-road collisions. With the deployment of new sensing and communication technologies (e.g., vehicle-to-device and vehicle-to-infrastructure), the traffic engineers overseeing these systems gain the opportunity to learn these unknown system parameters; however, the effect of revealing this information to drivers is not well understood.

This research was supported by ONR Grant #N00014-20-1-2359, AFOSR Grant #FA9550-20-1-0054, and NSF Grant #ECCS-2013779

B. L. Ferguson (corresponding author) and J. R. Marden are with the Department of Electrical and Computer Engineering, University of California, Santa Barbara, CA, {blferguson, jrmarden}@ece.ucsb.edu.

P.N. Brown is with the Department of Computer Science, University of Colorado at Colorado Springs, {philip.brown}@uccs.edu.

The conference version of this work appears in [1].

The emergence of new sensing and communication technologies opens the door to new methods for coordinating driving behavior and improving traffic patterns. One such method is that of *information signalling* by a well informed central authority [6]–[9]. By partially revealing their information about system parameters to uninformed users, the signaller allows the system users the opportunity to form new beliefs about their environment. If the signaller reveals this information strategically, they may alter user behavior in such a way that the overall system performance is improved; for example, Google/Apple maps can share the travel times of certain routes to guide driver decision making in a way that can alter aggregate driving patterns and improve performance for the user population [10]. One may initially think that all information should be shared with the users; however, we find that this need not be optimal and could further degrade system performance. The main focus of this work is determining what capabilities a system operator has in improving congestion via information signalling and identifying when and how this information should be shared.

We study the principles of information signalling in the context of *Bayesian congestion games*, where a group of users (syn. drivers) must route themselves through a congestable network while the exact congestion characteristics of each path is unknown. Deterministic congestion games have been used in the transportation literature to model driver decision making and its effect on road traffic [2], [11]–[19]. Recently, Bayesian congestion games have emerged as a generalized model in which the edge latency functions are random variables. The users possess a common prior belief over the random latencies of each route (for example, the belief there was an accident on a road or the chance weather has affected driving conditions) but do not observe the realization. The informed system operator does observe the realized values of the random parameters and can strategically signal information about them to the system users. This model for uncertain driving conditions has been used to study how information signalling policies should be designed [20]–[22], what behavior is likely to emerge [23], [24], and the associated performance of specific signalling structures [10], [25], [26]. The results are typically limited to computational methods for finding signalling policies or identifying whether or not revealing the state exactly is optimal. Additionally, for ease of analysis, much of the work in this area often assumes the signals are private (sent to individual users) which does not give relevant insights on public signals (sent to all users) which we consider in this work.

Signalling mechanisms are becoming a topic of increasing

research in their ability to influence user behavior, however, this is not the only influencing mechanism at a system operator's disposal. Incentive mechanisms, where users are assessed monetary penalties or rewards based on their actions, have long been studied as an effective means of coordinating system behavior [27]–[31]. In transportation settings, these incentives may manifest as road/bridge tolls or transit prices. The interplay between incentives and signalling is an emerging area of study, and has up until now been limited to studying mechanisms where users must pay to acquire information [32]–[34]. To the best of our knowledge, no current work has studied how monetary incentives and information signalling can be used concurrently to improve system performance.

In this work, we provide general insights on the benefit information signalling can provide with and without the concurrent use of monetary incentives. Through example, we demonstrate two key observations: (1) signalling, on its own, can worsen system cost, and (2) co-designing signal and incentive mechanisms offer opportunities for improvement that were not present when using each separately. In Theorem 1, we highlight the a benefit of co-designing these mechanisms: with appropriate monetary incentives, information signalling will not worsen system performance, essentially making signalling robust. To further understand the benefit of utilizing both mechanisms in tandem, in Section IV we focus on the class of parallel networks with polynomial latency functions and derive bounds on the possible benefit a signalling policy can provide with and without concurrent incentives.

Finally, in Section V, we address the problem of finding optimal signal-incentive pairs, including when the possible number of signals is bounded. We show how one can create solvable optimization problems to find the optimal signalling policies with and without concurrent incentives that may or may not be allowed to update with the sent signal.

We bolster the conclusions of this work with numerical examples in Section IV-C and Section V-D, in which we find that concurrent signal-incentive mechanisms offer notable improvements.

II. SYSTEM MODEL

A. Congestion Games

Consider a directed graph (V, E) with vertex set V , edge set $E \subseteq (V \times V)$, and k origin-terminal pairs (o_i, t_i) . Denote by $\mathcal{P}_i$ the set of all simple paths connecting origin o_i to destination t_i . Further, let $\mathcal{P} = \bigcup_{i=1}^k \mathcal{P}_i$ denote the set of all paths in the graph. A *flow* on the graph is a vector $f \in \mathbb{R}_{\geq 0}^{|\mathcal{P}|}$, where f_P expresses the mass of traffic utilizing path P . The mass of traffic on an edge $e \in E$ is thus $f_e = \sum_{P: e \in P} f_P$, and we say $f = \{f_e\}_{e \in E}$. A flow f is *feasible* if it satisfies $\sum_{P \in \mathcal{P}_i} f_P = r_i$ for each source-destination pair, where r_i is the mass of traffic traveling from origin o_i to terminal t_i .

When a larger number of users traverse the same path, the congestion (and thus transit delay) on that path increases. To characterize this, each edge e is endowed with a *latency function* $\ell_e : \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}_{\geq 0}$ that maps the mass of traffic on an edge to the delay users on that edge observe. We assume each latency function is positive, convex, non-decreasing, and




Figure 1: System model diagram. System users (drivers) must travel from their source to destination through a congestible network with uncertain congestion rates. The users do not know the current network state, however, the system operator does. Leveraging its greater information, the system operator can devise a signalling mechanism to send messages of partial information to the users to alter their beliefs and ultimately their actions.

continuously differentiable. The system cost of a flow f is the *total latency*,

$$\mathcal{L}(f) = \sum_{e \in E} f_e \cdot \ell_e(f_e). \quad (1)$$

A *routing problem* is specified by the tuple $G = (V, E, \{\ell_e\}_{e \in E}, \{r_i, (o_i, t_i)\}_{i=1}^k)$, and we let $\mathcal{F}(G)$ denote the set of all feasible flows. We define the optimal flow f^{opt} as one that minimizes the total latency, i.e.,

$$f^{\text{opt}} \in \arg \min_{f \in \mathcal{F}(G)} \mathcal{L}(f). \quad (2)$$

Though this flow is desirable, it need not emerge from the self interested decision making of users. To model the setting where users are free to choose their own paths (such as drivers selecting their own routes), let $x \in [0, r_i]$ denote the index of an infinitesimal agent who uses a path $P_x \in \mathcal{P}_i$. When agents are left to choose routes that minimize their own observed travel time, i.e., each agent possesses a cost function $J_x(P_x; f) = \sum_{e \in P_x} \ell_e(f_e)$, then a plausible behavior that can emerge in the system is that of a *Nash flow* f^{Nf} [2], [11], [35], which satisfies

$$J_x(P_x; f) \leq J_x(P'; f), \forall P' \in \mathcal{P}_i, x \in [0, r_i], i \in [k]. \quad (3)$$

These system states are those where no user has an incentive to change their action and need not be optimal [36]; additionally, the total latency in any Nash flow in a game of this form is the same [37].

B. Bayesian Congestion Games & Information Signalling

We consider a setting where the exact traffic conditions are unknown to drivers but known precisely by a central system operator (e.g., Google Maps, Waze, Apple Maps, etc.). Let each latency function take the form $\ell_e(f_e) = \sum_{d \in \mathcal{D}} \alpha_{e,d} \cdot \ell_d(f_e)$, where $\mathcal{D} = \{\ell_1, \dots, \ell_D\}$ is a set of basis latency functions¹ and $\alpha_{e,d} \geq 0$ is the weight of the basis function

¹This formulation can capture many models of congestion including polynomial ($\mathcal{D} = \{x^0, x^1, x^2, \dots, x^D\}$), exponential ($\mathcal{D} = \{e^{0x}, e^{0.5x}, e^{2x}, \dots\}$), and the Bureau of Public Roads (BPR) latency functions ($\mathcal{D} = \{x^0, x^4\}$), commonly used to model the congestion characteristics of physical roads [38], [39].




Figure 2: Two-link, parallel, Bayesian congestion game. One edge possesses a linear latency function, the other a constant latency function. The coefficients of each of these latency functions $\alpha_{1,0}, \alpha_{2,1}$ are unknown but distributed with prior μ_0 over $A = [0, 2]^2$. At right, is an illustration of a truthful signalling policy $\pi : A \rightarrow \{s_1, s_2, s_3\}$, which partitions A to map realizations to signals. After receiving a signal s , the agents compute their posterior μ_s , as illustrated by the posterior beliefs with support defined by the subset of A to which it is associated. In general, signals need not be deterministic/truthful, and we may choose a signalling policy $\pi : A \rightarrow \Delta(S)$, such that the signal s is drawn from a distribution conditioned on α ; the posterior beliefs are computed the same, but now need not be of the partitioned form shown in this example.

$\ell_d(\cdot)$ on the edge $e \in E$. To capture the idea of uncertainty in this problem, let $\alpha \in \mathbb{R}_{\geq 0}^{|E| \cdot |D|}$ (whose elements contain the weight of each basis latency function on each edge) be a random variable with distribution $\mu_0(x) = \mathbb{P}[\alpha = x]$ and support A . We assume the system operator observes the realization of this parameter, but the system users do not. If the users reach a flow f , we extend (1) to be the *expected total latency* over a distribution μ ,

$$\mathcal{L}(f; \mu) = \mathbb{E}_{\alpha \sim \mu} \left[\sum_{e \in E} f_e \cdot \ell_e(f_e) \right].$$

Because $\ell_e(\cdot)$ is determined by α , it is a random variable.

As a method to coordinate behavior and induce more desirable system states, the system operator may choose to signal relevant information to the users so they may update their beliefs. To do so, a system operator selects a *signalling policy* $\pi : A \rightarrow \Delta(S)$ that maps realizations of the system state $\alpha \in A$ to a distribution, from which a signal $s \in S$ is sampled² which may reveal information to system users. We assume these signals are *public*, in that every user receives the same message, but need not be truthful or reveal the exact realization. At the reception of signal $s \in S$, users infer the posterior distribution over the system state α as

$$\mu_s(\alpha) = \frac{\pi(s|\alpha) \cdot \mu_0(\alpha)}{\int_{\alpha' \in A} \pi(s|\alpha') \cdot \mu_0(\alpha') d\alpha'},$$

where $\pi(s|\alpha)$ is the probability of sending signal s when the system state realization is α . Fig. 2 illustrates how agents beliefs may be shaped in a simple example.

Under a signalling policy π , agents may change their chosen path based on which signal they receive. Let $\mathbf{f} = \{f(s)\}_{s \in S}$ denote the tuple containing the flow that occurs at the reception of each signal, and let $\sigma_x = \{\sigma_x(s) \in \mathcal{P}_i\}_{s \in S}$ denote the path user $x \in [0, r_i]$ selects after receiving each signal. When each agent adopts a strategy based on the information system's

signals, the system designer now cares about the expected total latency, expressed as

$$\mathcal{L}(\mathbf{f}; \mu_0, \pi) = \sum_{s \in S} \psi(s) \cdot \mathcal{L}(f(s); \mu_s), \quad (4)$$

where $\psi(y) = \mathbb{P}[s = y] = \int_{\alpha' \in A} \pi(s|\alpha') \cdot \mu_0(\alpha') d\alpha'$ denotes the distribution over signals. An agent's cost will now be their expected travel time,

$$J_x(\sigma_x; \mathbf{f}, \mu_0, \pi) = \sum_{s \in S} \psi(s) \cdot \mathbb{E}_{\alpha \sim \mu_s} \left[\sum_{e \in \sigma_x(s)} \ell_e(f_e(s)) \right]$$

We can now define a *Bayes-Nash Equilibrium* as a tuple $(\mathbf{f}^{\text{BNf}}, \sigma^{\text{BNf}})$ as a set of strategies where no agent elects to unilaterally change, i.e.,

$$J_x(\sigma_x^{\text{BNf}}; \mathbf{f}^{\text{BNf}}, \mu_0, \pi) \leq J_x(\sigma'; \mathbf{f}^{\text{BNf}}, \mu_0, \pi), \quad \forall \sigma' \in (\mathcal{P}_i)^m, \quad x \in [0, r_i], i \in \{1, \dots, k\}. \quad (5)$$

Our main focus in this work is understanding what opportunities a system designer has in lowering the expected total latency by way of information provisioning, i.e., comparing $\mathcal{L}(\mathbf{f}^{\text{BNf}}; \mu, \pi)$ with $\mathcal{L}(\mathbf{f}^{\text{BNf}}; \mu, \emptyset)$ (where the use of $\emptyset$ where a π was expected denotes the case where no information is shared with users). To quantify this improvement in system performance, we define the *benefit to system cost* as our performance metric

$$\mathbf{B}(\pi; \mu) = \mathcal{L}(\mathbf{f}^{\text{BNf}}; \mu, \emptyset) - \mathcal{L}(\mathbf{f}^{\text{BNf}}; \mu, \pi), \quad (6)$$

which measures the reduction in system cost from utilizing a signal policy π . The system operator's objective is to institute a signalling structure that reduces the system cost or, equivalently, has a positive benefit. Several works have shown encouraging results on the capabilities of information signalling and identified situations in which system cost can be significantly reduced [21], [23], [25]. However, the consequences of information signalling need not always be positive. In the following example, we identify that this may be the case even in simple settings.

²For ease of notation, we will often treat the set of signals S as finite; however, the set of signals can be generalized to include a unique signal for each realization of the system state, i.e., $S = A$, which can be uncountable.

Example 1 (Consequences From Signalling). In this example, consider a population of drivers tasked with selecting one of two commute options. One of the routes is always delayed (either from natural hazards, uncertain demand, or irregular maintenance), while the other has free flowing traffic but congests as the number of drivers on that route increases. However, the drivers are uncertain which route will be delayed.




Figure 3: Bayesian congestion game used in example 2.

To model this, consider a congestion game with two parallel edges $E = \{e_1, e_2\}$. One edge has a linear latency function, and the other a constant; each edge is the linear congestable edge with probability $1/2$. Let $\mu_0(\alpha^1) = 1/2$ be the prior belief that the first edge is the linear congestable edge (state α^1). When no information is revealed, users split over the two edges equally ($f_1 = f_2 = 1/2$), and the expected system cost is $\mathcal{L}(\mathbf{f}^{\text{BNf}}; \mu_0, \emptyset) = 0.75$.

If an information signal s is sent to the users, let $q := \mathbb{P}[\alpha = \alpha^1 \mid s]$ be the posterior belief that the first edge is the linear congestable edge. With this posterior, $f_1 = q$ users utilize the first edge, and the expected cost is $\mathcal{L}(f; \mu_s) = q^2 - q + 1$. For any value of $q \neq 1/2$, the expected system cost is greater than not revealing information; as such, any signalling policy that causes users beliefs to differ from the prior will increase cost, i.e., $\mathbf{B}(\pi; \mu) = \mathcal{L}(\mathbf{f}^{\text{BNf}}; \mu, \emptyset) - \mathcal{L}(\mathbf{f}^{\text{BNf}}; \mu, \pi) < 0$ for all $\pi : A \rightarrow \Delta(S)$. This demonstrates our first observation:

Observation: Revealing information to users can have negative consequences and increase system cost.

Example 1 highlights that signalling, on its own, may not be capable of reducing system cost. However, this is not the only influencing mechanism at a traffic engineers disposal. Another mechanism to influence user behavior is that of *monetary incentives* which have been well studied in transportation and alleviating congestion [27]–[31], but, to the authors' best knowledge, the use of information signalling and monetary incentives in tandem has yet to be studied in the context of traffic networks.

C. Monetary Incentives

Consider a congestion game G ; an incentive designer can apply an incentive $\tau_e \in \mathbb{R}$ to each edge $e \in E$ to change the cost experienced by users utilizing that edge, i.e.,

$$J_x(e; f) = \ell_e(f_e) + \tau_e.$$

When a signal s is sent to the users, the expected cost to a user x on path P_x in flow f becomes $J_x(P_x; f) = \mathbb{E}[\sum_{e \in P_x} \ell_e(f_e) + \tau_e \mid s]$. This change in cost affects the users' decision-making and ultimately leads to new Nash flows, ideally with lower total latency. Monetary incentives are

a well studied and highly utilized method of controlling congestion in transportation [15], [16]. However, the relationship between incentives and information signalling is not currently well understood; studying these two mechanisms concurrently is the main focus of this work.

To model the interplay of these two influencing mechanisms, note that at each signal the selected tolls will alter the Bayes-Nash Equilibrium³ by propagating this new cost into (5). With incentives τ and signalling policy π , the equilibrium system cost will be written $\mathcal{L}(\mathbf{f}^{\text{BNf}}; \mu_0, \pi, \tau)$. One can identify scenarios where either mechanism is capable of reducing congestion, however, this is not true in general. For a given Bayesian congestion game, it is not immediately apparent if either influencing mechanism can independently reduce system cost at all. In the following example, we will see that even in a simple setting, there exist quantifiable benefits to designing these mechanisms concurrently.

Example 2 (The Need for Co-Design). In this example, again, consider a population of drivers tasked with choosing between two commutes. The traffic rates on one route is always known, but the second sometimes contains a significant delay (perhaps caused by routine closures and detours).




Figure 4: Bayesian congestion game used in example 3.

To model this situation, consider a congestion game with two parallel edges $E = \{e_1, e_2\}$. The first edge has a deterministic latency function $\ell_1(f_1) = f_1 + 1$, while the second edge has a latency function $\ell_2(f_2) = f_2 + \zeta$ where $\zeta = 2$ with probability $1/2$ and $\zeta = 0$ otherwise. First, we show that no toll can reduce system cost alone. When no information is revealed, each edge has the same expected cost and the Bayes-Nash flow is $\mathbf{f}^{\text{BNf}} = \{(1/2, 1/2)\}$; the optimal flow is the same, i.e., $f^{\text{opt}} = (1/2, 1/2)$. As the unincentivized equilibrium is already optimal, clearly no toll can reduce system cost.

Now, consider some information signalling policy π with signal set $S = \{s_1, \dots, s_m\}$. Note that $\mathcal{L}^{\text{Nf}}(\zeta) = 1 + \zeta/2$; from Lemma 1,

$$\mathcal{L}(\mathbf{f}^{\text{BNf}}; \mu_0, \pi) = \sum_{s \in S} \psi(s)(1 + \mathbb{E}[\zeta \mid s]/2) = 1 + \mathbb{E}[\zeta]/2,$$

and $\mathcal{L}(\mathbf{f}^{\text{BNf}}; \mu_0, \emptyset) = 1 + \mathbb{E}[\zeta]/2$, thus, any signalling policy π does not reduce system cost.

Finally, let π be the full-reveal signalling policy and $\tau_1 = 0.5$ and $\tau_2 = 0$. The expected system cost with this signal/incentive pair is $1.4375 < 1.5 = 1 + \mathbb{E}[\zeta]/2$. This demonstrates our second observation:

³Note that the Bayes-Nash equilibrium flow $\mathbf{f}^{\text{BNf}}$ is now inherently dependent on the selected incentives.

Observation: There exist situations where signalling alone cannot reduce system cost, tolling alone cannot reduce system cost, but signalling and tolling together can reduce system cost.

Example 2 points to an important relationship between signals and incentives: there exist opportunities in designing the two together, but the benefits are not readily obvious. A co-design of the two mechanisms can be accomplished in two ways (1) by creating a larger optimization problem in which signals and incentives are both decision variables (see Section V), and (2) designing an incentive policy that can update with the sent signal. We call incentives that can update with the signal *signal-aware* and incentives that cannot *signal-agnostic*. This subsection highlighted the limitations of signal-agnostic incentives, in Section III and Section IV, we will largely focus on the benefit of signal-aware incentives.

D. Summary of Contributions

The main contributions of this work come in characterizing the interplay of two influencing mechanisms: information signalling and monetary incentives. Further, we describe how these mechanisms can be designed concurrently to provide increased benefits in reducing total latency. We propose two methods to this co-design. The first is utilizing signal-aware incentives designed for a given signalling policy. In Proposition 1, we characterize the optimal signal-aware toll for any signalling policy. One insight this work provides is that these incentives make the signalling policy robust; in Theorem 1 we show that while using the optimal signal-aware incentives, no information signalling policy can worsen system cost. To further illustrate the advantage of co-designing signals and incentives, in Section IV we consider the sub-class of problems with parallel networks and polynomial latency functions, in which we find analytical bounds for how much a signalling policy can change system cost. The insights from this section follow the more general results and show that signalling can still provide significant reduction in system cost when incentives are also used.

The second method of co-design involves directly solving for the optimal signal-incentive pairs. To do so, we leverage existing results on Generalized Moment Problems to solve for optimal signalling mechanisms. In Section V, we survey the existing literature and show how the optimal signal-incentive pairs (with either signal-agnostic or signal-aware incentives) can be transcribed and solved as GMPs. Additionally, we amend the problem to handle the case where there is a limited number of signals that can be sent (i.e., $|S|$ is bounded).

Finally, in Section V-D, we offer a numerical simulation to quantify the above results of signalling and incentive mechanisms concurrent use. This experiment demonstrates several of the insights from this work, including that co-designed incentive mechanisms offer notable performance improvements.

III. ADVANTAGE OF INCENTIVES

Example 1 and Example 2 highlighted an opportunity to design signals and incentives in tandem. In this section, we will focus on the qualities of incentives that can update with

the sent signal. Consider a *signal-aware incentive mechanism* $T(s; \pi, \mu_0)$ that assigns tolls $\{\tau_e(s)\}_{s \in S}$ dependent on the signal broadcast by the information provider. A player $x \in [0, r_i]$ with the strategy σ_x now observes an expected cost of

$$J_x(\sigma_x; \mathbf{f}, \mu_0, \pi, T) = \sum_{s \in S} \psi(s) \cdot \mathbb{E}_{\alpha \sim \mu_s} \left[\sum_{e \in \sigma_x(s)} \ell_e(f_e(s)) + \tau_e(s) \right].$$

The Bayes-Nash flow definition remains as shown in (5), but now with users' tolled cost. We now seek to understand the effectiveness of jointly implementing a signalling policy π and an incentive mechanism T . As such, we extend the definition of (6) which quantifies the gain in system performance to include the effect of an incentive mechanism T , i.e.,

$$\mathbf{B}(\pi; \mu, T) = \mathcal{L}(\mathbf{f}^{\text{BNf}}; \mu, \emptyset, T) - \mathcal{L}(\mathbf{f}^{\text{BNf}}; \mu, \pi, T). \quad (7)$$

We measure the benefit of a signalling policy by comparing the system cost with incentives and signalling and incentives alone. We do this because we largely want to focus on the value that information signalling can provide on its own.

First, we must decide what monetary incentives to use. In Proposition 1, we characterize an optimal signal-aware incentive mechanism for a given signalling policy.

Proposition 1. *Let μ_0 be a prior on the latency coefficients α in a Bayesian congestion game G with positive, convex, non-decreasing, and continuously differentiable latency functions that are of the form $\ell_e(f_e) = \sum_{d \in \mathcal{D}} \alpha_{e,d} \cdot \ell_d(f_e)$ where $\ell_d \in \mathcal{D}$, and let $\pi : A \rightarrow \Delta(S)$ be a signalling policy. An optimal signal-aware incentive mechanism T^* (i.e., maximizes $\mathcal{L}(\mathbf{f}^{\text{BNf}}; \mu_0, \pi, T)$) assigns tolls according to*

$$\tau_e^*(s) = \sum_{d \in \mathcal{D}} \mathbb{E}_{\alpha_{e,d} \sim \mu_s} [\alpha_{e,d}] x_e \ell'_d(x_e), \quad (8)$$

where $x \in \arg \min_{f \in \mathcal{F}(G)} \mathcal{L}(f; \mathbb{E}_{\alpha \sim \mu_s} [\alpha])$.

The proof appears in the appendix. Proposition 1 provides a mechanism for computing the optimal incentives for any signalling policy π . The use of these incentives in tandem with a signalling policy will alter the equilibrium flow and thus the system cost. Motivated by the observed negative consequences of information signalling shown in Example 1, Theorem 1 states the possible consequences of information signalling but now while concurrently using the optimal monetary incentives T^* ; we find that the concurrent use of monetary incentives makes it such that signalling can never worsen system performance, i.e., have no negative benefit.

Theorem 1. *Let μ_0 be a prior on the latency coefficients α in a Bayesian congestion game G with positive, convex, non-decreasing, and continuously differentiable latency functions that are of the form $\ell_e(f_e) = \sum_{d \in \mathcal{D}} \alpha_{e,d} \cdot \ell_d(f_e)$ where $\ell_d \in \mathcal{D}$. While using the signal-aware incentive policy T^* (as defined in Proposition 1, any signalling policy $\pi : A \rightarrow \Delta(S)$ has non-negative benefit to system cost, i.e.,*

$$\mathbf{B}(\pi; \mu_0, T^*) \geq 0. \quad (9)$$

Proof of Theorem 1: Consider a realization of a congestion game G with latency coefficients α . Let $\mathcal{L}^*(\alpha)$ denote the total latency in a Nash flow while using the incentive mechanism T^* as defined in Proposition 1. First, we characterize the Bayesian-Nash flow with incentives T^* . If the signal $s \in S$ is sent to users, they update their belief via Bayesian inference to $\mu_s(x) = \frac{\pi(s|\alpha) \cdot \mu_0(\alpha)}{\psi(s)}$. In a flow f , user $x \in [0, r_i]$ taking path $P_x \in \mathcal{P}_i$ experiences an expected cost of

$$\begin{aligned} J_x(P_x; f, \mu_s) &= \mathbb{E}_{\alpha \sim \mu_s} \left[\sum_{e \in P_x} \sum_{d \in \mathcal{D}} \alpha_{e,d} \cdot \ell_d(f_e) + \tau_e^*(s) \right] \\ &= \sum_{e \in P_x} \sum_{d \in \mathcal{D}} \mathbb{E}[\alpha_{e,d} | s] \ell_d(f_e) + \tau_e^*(s). \end{aligned}$$

Note that if f were not a Nash flow in the congestion game with coefficients $\mathbb{E}[\alpha | s]$, then by (3) at least one user x would be able to deviate their strategy $\sigma_x(s)$ and experience lower cost. Therefore, the only Bayes-Nash flows occur when $f(s)$ is a Nash flow with respect to $\mathbb{E}[\alpha | s]$ and tolls $\tau_e^*(s)$ for all $s \in S$. From Proposition 1, this is the optimal flow in the network with coefficients $\mathbb{E}[\alpha | s]$.

Next, consider the prior distribution μ_0 on α , and let f be a flow in the network. The expected total latency

$$\begin{aligned} \mathcal{L}(f; \mu) &= \mathbb{E}_{\alpha \sim \mu} \left[\sum_{e \in E} f_e \cdot \sum_{d \in \mathcal{D}} \alpha_{e,d} \ell_d(f_e) \right] \\ &= \sum_{e \in E} f_e \cdot \sum_{d \in \mathcal{D}} \mathbb{E}_{\alpha \sim \mu} [\alpha_{e,d}] \ell_d(f_e) \\ &= \mathcal{L}(f; \mathbb{E}_{\alpha \sim \mu} [\alpha]), \end{aligned}$$

which follows from linearity of expected value.

Combining the previous two observations, we obtain that the total latency in a Nash flow in the congestion game G with latency coefficients α when using T^* , can be expressed as

$$\mathcal{L}(\mathbf{f}^{\text{BNf}}; \mu_0, \pi, T^*) = \sum_{s \in S} \psi(s) \mathcal{L}^*(\mathbb{E}[\alpha | s]). \quad (10)$$

Next, we observe that $\mathcal{L}^*(\alpha)$ is concave. $\mathcal{L}^*(\alpha)$ can be expressed as the pointwise infimum over $f \in \mathcal{F}(G)$ for a given α ,

$$\mathcal{L}^*(\alpha) = \inf_{f \in \mathcal{F}(G)} \mathcal{L}(f; \alpha) = \inf_{f \in \mathcal{F}(G)} \sum_{e \in E} \sum_{d \in \mathcal{D}} \alpha_{e,d} f_e \cdot \ell_d(f_e).$$

Observing that $\sum_{e \in E} \sum_{d \in \mathcal{D}} \alpha_{e,d} f_e \ell_d(f_e)$ is affine in α (and thus concave in α) for each f , we can invoke that the pointwise infimum over a class of functions that are each concave is itself, concave [40]. Thus $\mathcal{L}^*(\alpha)$ is concave though need not be affine.

Now, consider the total latency in a Bayes-Nash flow with signal policy π and incentive T^* ,

$$\mathcal{L}(\mathbf{f}^{\text{BNf}}; \mu_0, \pi, T^*) = \sum_{s \in S} \psi(s) \cdot \mathcal{L}^*(\mathbb{E}[\alpha | s]) \quad (11a)$$

$$\leq \mathcal{L}^* \left(\sum_{s \in S} \psi(s) \cdot \mathbb{E}[\alpha | s] \right) \quad (11b)$$

$$= \mathcal{L}^*(\mathbb{E}[\alpha]) = \mathcal{L}(\mathbf{f}^{\text{BNf}}; \mu_0, \emptyset, T^*), \quad (11c)$$

where the (11a) holds from (10) and (11b) holds from the concavity of $\mathcal{L}^*$. From this, we can see that $\mathbf{B}(\pi; \mu_0, T^*) \geq 0$, i.e., while utilizing incentive scheme T^* , signalling cannot increase system cost. $\square$

IV. POLYNOMIAL ROUTING GAMES ON PARALLEL NETWORKS

Section III provided insights into the signalling problem in general Bayesian congestion games, showing that signalling can be detrimental to system performance unless concurrently utilized with monetary incentives. In this section, we seek to further understand the connection between signalling and incentivizing by characterizing closed form bounds on the benefit a signalling policy can provide with and without incentives. We do this in the context of Bayesian congestion games on parallel networks (i.e., one source-terminal pair with n directed edges directly connecting them) with polynomial latency functions (i.e., $\ell_d(x) = (x)^d$). These latency functions are positive, increasing, convex polynomials where $\mathcal{D} = (d_1, \dots, d_k)$ and $d_i \in \mathbb{Z}_{\geq 0}$ expresses the possible degrees⁴, e.g., $\mathcal{D} = \{0, 1\}$ represents affine congestion rates [28], $\mathcal{D} = \{0, 4\}$ can represent the well-known Bureau of Public Roads (BPR) latency functions, commonly used to model the congestion characteristics of physical roads [38], [39], and $\mathcal{D} = \{0, \dots, D\}$ can represent any positive, convex, increasing polynomial up to degree D [31]. Additionally, we note from Remark 2 in the appendix that, without loss of generality, we can normalize to a unit demand, i.e., $r = 1$.

We highlight two important possible realizations of the random variable α that will be used throughout: $\tilde{\alpha} \in \mathbb{R}_{\geq 0}^{|\mathcal{E}| \cdot |\mathcal{D}|}$ such that $\tilde{\alpha}_{e,d} = \inf\{\text{supp}(\alpha_{e,d})\}$ (where $\text{supp}(\cdot)$ denotes the support of α) in which each parameter takes its lowest value, and $\hat{\alpha} \in \mathbb{R}_{\geq 0}^{|\mathcal{E}| \cdot |\mathcal{D}|}$ such that $\hat{\alpha}_{e,d} = \sup\{\text{supp}(\alpha_{e,d})\}$ in which each parameter takes its largest value. Note that $\tilde{\alpha}$ and $\hat{\alpha}$ need not be in the support of α , but rather represent the corners of the smallest box that contains the support of α that are closest and furthest from the origin respectively. Further, to avoid degenerate cases, we institute the following assumption on Bayesian congestion games.

Assumption 1. In a Bayesian Congestion Game G with prior μ_0 , $0, 1 \in \mathcal{D}$ is always satisfied, and, for each edge $e \in E$, $\tilde{\alpha}_{e,0}, \tilde{\alpha}_{e,1} > 0$.

This assumption prevents cases where traffic can be routed with zero delay and has zero affect on congestion.

In the context of these games, we address the possible benefit of a signalling policy.

A. Signalling Alone

When a system designer seeks to improve system performance by solely using a public information-signalling system, Theorem 2 provides bounds on the benefit a signalling policy can provide.

Theorem 2. Consider the class of parallel Bayesian congestion games with polynomial latency functions whose degrees

⁴We assume that $0 \in \mathcal{D}$ is always satisfied.

come from the set $\mathcal{D}$. For any distribution over the latency coefficients μ_0 and any signalling policy π , the benefit in the expected total latency of a Bayes-Nash flow from signalling satisfies

$$-\Theta \|\mathbb{E}[\alpha] - \check{\alpha}\|_2 \leq \mathbf{B}(\pi; \mu_0) \leq \Theta \|\mathbb{E}[\alpha] - \check{\alpha}\|_2, \quad (12)$$

where $\Theta := |\mathcal{D}| + \frac{\rho^+ - \rho_0^-}{2\rho_1^-} (|E| + |\mathcal{D}| - 1)$, $\rho_0^- = \min_{e \in E} \check{\alpha}_{e,0}$, $\rho_1^- = \min_{e \in E} \check{\alpha}_{e,1}$, $\rho^+ = \max_{e \in E} \sum_{d \in \mathcal{D}} (d+1) \hat{\alpha}_{e,d}$, $\mathbb{E}[\alpha] = \int_{x \in A} x \mu_0(x) dx$, and $\check{\alpha} \in \mathbb{R}_{\geq 0}^{|E| \cdot |\mathcal{D}|}$ such that $\check{\alpha}_{e,d} = \inf \{\text{supp}(\alpha_{e,d})\}$ for each $e \in E$, $d \in \mathcal{D}$. Additionally, there exists a μ_0 such that for any truthful $\pi \neq \emptyset$ (i.e., $\pi : A \rightarrow S$ is deterministic),

$$\mathbf{B}(\pi; \mu_0) = \sqrt{|\mathcal{D}|} \cdot \|\mathbb{E}[\alpha] - \check{\alpha}\|_2. \quad (13)$$

Further, if $d \in \mathcal{D}$ where $d > 0$, then there exists a μ_0 such that for any truthful $\pi \neq \emptyset$,

$$\mathbf{B}(\pi; \mu_0) = -\|\mathbb{E}[\alpha] - \check{\alpha}\|_2. \quad (14)$$

The proof of Theorem 2 appears at the end of this section.

Theorem 2 reveals the capabilities a signalling policy has in improving system performance. It also highlights the reality that revealing information can make system performance worse. The bounds for the benefit of a signalling policy depends on the number of terms considered in each latency function $|\mathcal{D}|$, the size of the network $|E|$, as well as the distance between the average system state and the edge of its support $\|\mathbb{E}[\alpha] - \check{\alpha}\|_2$ and other terms that change with the support. One can think that the number of latency terms $|\mathcal{D}|$ characterizes the complexity of the model of network congestion while $\|\mathbb{E}[\alpha] - \check{\alpha}\|_2$ measures the amount of uncertainty about the system parameters. Additionally, (13) and (14) shows that there exist situations where regardless of what signalling policy is chosen, revealing information can greatly benefit or hinder system performance.

For many of the proofs in this sections, we will utilize the following Lemma, demonstrated in the proof of Theorem 1 and proven in the appendix.

Lemma 1. *With a prior μ_0 and a signalling policy π , the Bayes-Nash flow $\mathbf{f}^{\text{BNf}}$ can be characterized by $\{\bar{f}^{\text{Nf}}(s)\}_{s \in S}$, where $\bar{f}^{\text{Nf}}(s)$ is the Nash flow in the network G with coefficients $\bar{\alpha}_s = \mathbb{E}[\alpha|s]$. Additionally, for a given flow f , the expected total latency with distribution μ over the coefficients α is equal to the total latency in the network with the expected coefficients, i.e., $\mathcal{L}(f; \mu) = \mathcal{L}(f; \mathbb{E}_{\alpha \sim \mu}[\alpha])$. Together, these facts show that the expected total latency in a Bayes-Nash flow is equal to the weighted average of the total latency in the expected network after receiving each signal, i.e.,*

$$\mathcal{L}(\mathbf{f}^{\text{BNf}}; \mu_0, \pi) = \sum_{s \in S} \psi(s) \mathcal{L}^{\text{Nf}}(\mathbb{E}_{\alpha \sim \mu_s}[\alpha]), \quad (15)$$

where $\mathcal{L}^{\text{Nf}}(\alpha)$ denotes the total latency in a Nash flow in the deterministic congestion game G with latency coefficients α .

To prove Theorem 2, in Lemma 2 we provide a fact about the function $\mathcal{L}^{\text{Nf}}(\alpha)$ which bounds the difference in total latency between any two realizations of edge latency coefficients.

Lemma 2. *Consider the class of parallel congestion games with polynomial latency functions with degrees drawn from the set $\mathcal{D}$ with coefficients $\alpha \in A$. Let $a, b \in \mathbb{R}_{\geq 0}^{|E| \cdot |\mathcal{D}|}$ be two possible sets of coefficients for a congestion game with edge set E , then*

$$|\mathcal{D}| + \frac{\rho^+ - \rho_0^-}{2\rho_1^-} (|E| + |\mathcal{D}| - 1) \geq \frac{\mathcal{L}^{\text{Nf}}(a) - \mathcal{L}^{\text{Nf}}(b)}{\|a - b\|_2}, \quad (16)$$

where $\rho_0^- = \min_{e \in E} \check{\alpha}_{e,0}$, $\rho_1^- = \min_{e \in E} \check{\alpha}_{e,1}$, and $\rho^+ = \max_{e \in E} \sum_{d \in \mathcal{D}} (d+1) \hat{\alpha}_{e,d}$.

Proof of Theorem 2: We start by proving the lower bound on $\mathbf{B}(\pi; \mu_0)$, which quantifies how much the use of a signalling policy can worsen the system performance. Consider the prior μ_0 and signalling policy $\pi : A \rightarrow \Delta(S)$. If μ_s is the posterior formed from receiving signal s , let $\bar{\alpha}_s = \mathbb{E}_{\alpha \sim \mu_s}[\alpha]$.

To prove the lower bound, first, define the set $\hat{A} = \prod_{e \in E, d \in \mathcal{D}} [\inf \{\text{supp}(\alpha_{e,d}), \sup \{\text{supp}(\alpha_{e,d})\}]$ that is the smallest box in $\mathbb{R}_{\geq 0}^{|E| \cdot |\mathcal{D}|}$ that contains $A = \text{supp}(\alpha)$. Note that $\check{\alpha}$ is the corner of this box that is closest to the origin. Let, $\hat{\mathcal{L}}^{\text{Nf}}$ be the concave closure of the function $\mathcal{L}^{\text{Nf}}$ over $\hat{A}$, i.e.,

$$\hat{\mathcal{L}}^{\text{Nf}} = \sup \{z | (\alpha, z) \in \text{Conv}_{\hat{A}}(\mathcal{L}^{\text{Nf}})\},$$

where $\text{Conv}_{\hat{A}}(\mathcal{L}^{\text{Nf}})$ denotes the convex hull of the graph of $\mathcal{L}^{\text{Nf}}$ over the domain $\hat{A}$. With this, we show

$$\mathbf{B}(\pi; \mu_0) = \mathcal{L}(\mathbf{f}^{\text{BNf}}; \mu_0, \emptyset) - \mathcal{L}(\mathbf{f}^{\text{BNf}}; \mu_0, \pi) \quad (17a)$$

$$= \mathcal{L}^{\text{Nf}}(\bar{\alpha}_0) - \sum_{s \in S} \psi(s) \cdot \mathcal{L}^{\text{Nf}}(\bar{\alpha}_s) \quad (17b)$$

$$\geq \mathcal{L}^{\text{Nf}}(\check{\alpha}) - \hat{\mathcal{L}}^{\text{Nf}}(\bar{\alpha}_0) \quad (17c)$$

$$= \hat{\mathcal{L}}^{\text{Nf}}(\check{\alpha}) - \hat{\mathcal{L}}^{\text{Nf}}(\bar{\alpha}_0) \quad (17d)$$

$$\geq -\Theta \|\mathbb{E}[\alpha] - \check{\alpha}\|_2, \quad (17e)$$

where (17b) holds from Lemma 1, (17c) holds from $\mathcal{L}^{\text{Nf}}$ monotonically increasing in α and $\check{\alpha}_{e,d} \leq \alpha_{e,d}$ for all $e \in E$, $d \in \mathcal{D}$, and $\alpha \in A$, as well as the concave closure $\hat{\mathcal{L}}^{\text{Nf}}$ being greater than any concave combination of points in A and $\sum_{s \in S} \psi(s) \bar{\alpha}_s = \bar{\alpha}_0$, (17d) holds from $\mathcal{L}^{\text{Nf}}(\check{\alpha}) = \hat{\mathcal{L}}^{\text{Nf}}(\check{\alpha})$ due to $\check{\alpha}$ being a corner of $\hat{A}$, and (17e) holds from Lemma 2 and the definition of Θ from the theorem statement along with the observation that the maximum gradient in the concave closure $\hat{\mathcal{L}}^{\text{Nf}}$ must also occur in the original function $\mathcal{L}^{\text{Nf}}$ by the intermediate value theorem.

Now, we prove the upper bound using similar methods:

$$\mathbf{B}(\pi; \mu_0) = \mathcal{L}(\mathbf{f}^{\text{BNf}}; \mu_0, \emptyset) - \mathcal{L}(\mathbf{f}^{\text{BNf}}; \mu_0, \pi) \quad (18a)$$

$$= \mathcal{L}^{\text{Nf}}(\bar{\alpha}_0) - \sum_{s \in S} \psi(s) \cdot \mathcal{L}^{\text{Nf}}(\bar{\alpha}_s) \quad (18b)$$

$$\leq \mathcal{L}^{\text{Nf}}(\bar{\alpha}_0) - \mathcal{L}^{\text{Nf}}(\check{\alpha}) \quad (18c)$$

$$\leq \Theta \|\mathbb{E}[\alpha] - \check{\alpha}\|_2. \quad (18d)$$

Together, these two bounds show the range of attainable performance improvements by utilizing a signalling policy π . The fact that $\mathbf{B}(\pi; \mu_0)$ can be negative shows that providing information to users need not always help. In fact, for any signalling policy π , there exist scenarios where adding information can be detrimental to system performance.

To see (13), consider a two link parallel network where $\ell_1(f_1) = 1$, and $\ell_2(f_2) = \sum_{d \in \mathcal{D}} \beta(f_2)^d$, i.e., $\alpha_{2,d} = \beta$ for all $d \in \mathcal{D}$. When $\beta \leq 1/|\mathcal{D}|$, $f_2 = 1$ and $\mathcal{L}^{\text{Nf}}(\beta) = \beta \cdot |\mathcal{D}|$. When $\beta > 1$, $f_1 > 0$ and $\mathcal{L}^{\text{Nf}}(\beta) = 1$. Consider a distribution μ_0 where $\beta = 0$ with probability $1 - \epsilon$ and $\frac{1}{\epsilon \cdot |\mathcal{D}|}$ with probability ϵ . The expected value of β is thus $1/|\mathcal{D}|$. The expected total latency without signalling is $\mathcal{L}^{\text{Nf}}(\beta = 1/|\mathcal{D}|) = 1$. Because any truthful signal π will reveal the two possible realizations of β , the expected total latency with π is

$$\mathcal{L}^{\text{Nf}}(0)(1 - \epsilon) + \mathcal{L}^{\text{Nf}}\left(\frac{1}{\epsilon \cdot |\mathcal{D}|}\right)\epsilon \rightarrow \mathcal{L}^{\text{Nf}}(0) = 0$$

as $\epsilon \rightarrow 0$. The benefit of π is thus

$$\mathcal{L}(\mathbf{f}^{\text{BNf}}; \mu_0, \emptyset) - \mathcal{L}(\mathbf{f}^{\text{BNf}}; \mu_0, \pi) = 1 = -\sqrt{|\mathcal{D}|} \cdot \|\mathbb{E}[\alpha] - \check{\alpha}\|_2,$$

where α is a vector with $|\mathcal{D}|$ entries of value $\frac{1}{|\mathcal{D}|}$.

To see (14), consider the two link parallel network where $\ell_1(f_1) = (f_1)^d + \beta$ and $\ell_2(f_2) = \beta(f_2)^d + 1$ and $f_1 + f_2 = 1$; in this congestion game, $\beta = \alpha_{1,0} = \alpha_{2,d}$ is a single parameter that represents two, correlated coefficients. It is difficult to characterize the Nash flow in closed form, however, we can utilize the following two facts (1) $\frac{\partial}{\partial \beta} \mathcal{L}^{\text{Nf}}(\beta)|_{\beta=0} = 0$ and (2) $\frac{\partial}{\partial \beta} \mathcal{L}^{\text{Nf}}(\beta)|_{\beta \rightarrow \infty} = 1$. Let μ_0 be a distribution on β . From the first fact, $\mathcal{L}(\mathbf{f}^{\text{BNf}}; \mu_0, \emptyset) = \mathcal{L}^{\text{Nf}}(\bar{\beta}) \rightarrow \mathcal{L}^{\text{Nf}}(0)$ as $\bar{\beta} \rightarrow 0$. Now, consider the prior distribution $\mu_0(x) = \{0, \text{w.p. } 1 - \epsilon, \bar{\beta}/\epsilon, \text{w.p. } \epsilon\}$. Any signalling policy π will reveal which β as 0 or $\bar{\beta}/\epsilon$, as such,

$$\mathcal{L}(\mathbf{f}^{\text{BNf}}; \mu_0, \pi) = \mathcal{L}^{\text{Nf}}(0)(1 - \epsilon) + \mathcal{L}^{\text{Nf}}(\bar{\beta}/\epsilon)\epsilon.$$

From fact (2) above, as $\epsilon \rightarrow 0$, $\mathcal{L}(\mathbf{f}^{\text{BNf}}; \mu_0, \pi) \rightarrow \bar{\beta}$. From these two facts, with sufficiently small $\bar{\beta}$,

$$\mathcal{L}(\mathbf{f}^{\text{BNf}}; \mu_0, \emptyset) - \mathcal{L}(\mathbf{f}^{\text{BNf}}; \mu_0, \pi) \rightarrow -\bar{\beta} = -\|\mathbb{E}[\beta] - \underline{\beta}\|_2,$$

where $\underline{\beta} = 0$ to match (14). $\square$

B. Signals & Incentives

Theorem 2 showed that revealing information has the possibility of increasing or decreasing system cost. It is already known from Theorem 1 that concurrently utilizing appropriate monetary incentives removes the possibility of worsening performance; however, it is not yet clear how these incentives affect a signalling policy's ability to improve performance.

Theorem 3 provides bounds on the benefit a signalling policy can provide while also utilizing the signal-aware incentive mechanism T^* . We see that by concurrently utilizing incentives and signalling, the system designer can guarantee the benefit of signalling is non-negative and still have room for significant improvement.

Theorem 3. *Consider the class of parallel Bayesian congestion games with polynomial latency functions whose degrees come from the set $\mathcal{D}$. For any distribution over the latency coefficients μ_0 and any signalling policy π , the decrease in the expected total latency of a Bayes-Nash flow from signalling satisfies*

$$0 \leq \mathbf{B}(\pi; \mu_0, T^*) \leq \Xi \|\mathbb{E}[\alpha] - \check{\alpha}\|_2, \quad (19)$$

where $\Xi := |\mathcal{D}| + \frac{\rho^+ - \rho_0^-}{4\rho_1^-} \left(|E| + \sum_{d \in \mathcal{D} \setminus \{0\}} (d+1)^d \right)$, $\rho_0^- = \min_{e \in E} \check{\alpha}_{e,0}$, $\rho_1^- = \min_{e \in E} \check{\alpha}_{e,1}$, $\rho^+ = \max_{e \in E} \sum_{d \in \mathcal{D}} (d+1)\hat{\alpha}_{e,d}$, $\mathbb{E}[\alpha] = \int_{x \in A} x \mu_0(x) dx$, and $\check{\alpha} \in \mathbb{R}_{\geq 0}^{|E| \cdot |\mathcal{D}|}$ such that $\check{\alpha}_{e,d} = \inf \{\text{supp}(\alpha_{e,d})\}$ for each $e \in E$, $d \in \mathcal{D}$.

Comparing the bounds on the benefit of a signalling policy with and without the use of incentives (i.e., (12) and (19)), we see that incentives can make the use of signals more robust (non-negative benefit) while allowing for similar opportunities to improve performance. We further support this conclusion in Section IV-C by providing a numerical example and comparing the benefit of revealing information with and without incentives.

Before we prove Theorem 3, we state the following lemma that is similar to Lemma 2 but applies to $\mathcal{L}^*$.

Lemma 3. *Consider the class of parallel congestion games with polynomial latency functions with degrees coming from the set $\mathcal{D}$ with coefficients $\alpha \in A$. Let $a, b \in \mathbb{R}_{\geq 0}^{|E| \cdot |\mathcal{D}|}$ be two possible sets of coefficients for a congestion game with edge set E , then*

$$|\mathcal{D}| + \frac{\rho^+ - \rho_0^-}{4\rho_1^-} \left(|E| + \sum_{d \in \mathcal{D} \setminus \{0\}} (d+1)^d \right) \geq \frac{\mathcal{L}^*(a) - \mathcal{L}^*(b)}{\|a - b\|_2}. \quad (20)$$

where $\rho_0^- = \min_{e \in E} \check{\alpha}_{e,0}$, $\rho_1^- = \min_{e \in E} \check{\alpha}_{e,1}$, $\rho^+ = \max_{e \in E} \sum_{d \in \mathcal{D}} (d+1)\hat{\alpha}_{e,d}$.

The proof of Lemma 3 is in the appendix.

Proof of Theorem 3: Consider the prior μ_0 and signalling policy $\pi : A \rightarrow \Delta(S)$. If μ_s is the posterior formed from receiving signal s , then let $\bar{\alpha}_s = \mathbb{E}_{\alpha \sim \mu_s}[\alpha]$. When utilizing the incentive mechanism T^* which assigns incentives as stated in (8), then Proposition 1 states that the equilibrium flow that emerges when signal s is received will be $f^*(s) \in \arg \min \mathcal{L}(f; \bar{\alpha}_s)$; as such $\mathcal{L}^*(\bar{\alpha}_i)$ is the total latency that occurs.

The lower bound of (19) is immediate from Theorem 1. For the upper bound, we show that

$$\begin{aligned} \mathbf{B}(\pi; \mu_0, T^*) &= \mathcal{L}(\mathbf{f}^{\text{BNf}}; \mu_0, \emptyset, T^*) - \mathcal{L}(\mathbf{f}^{\text{BNf}}; \mu_0, \pi, T^*) \\ &= \mathcal{L}^*(\bar{\alpha}_0) - \sum_{s \in S} \psi(s) \cdot \mathcal{L}^*(\bar{\alpha}_s) \end{aligned} \quad (21a)$$

$$\leq \mathcal{L}^*(\bar{\alpha}_0) - \mathcal{L}^*(\check{\alpha}) \quad (21b)$$

$$\leq \Xi \|\mathbb{E}[\alpha] - \check{\alpha}\|_2, \quad (21c)$$

where the (21a) holds from Lemma 1, (21b) from $\mathcal{L}^*$ non-decreasing with α , and (21c) from Lemma 3 and the definition of Ξ in the theorem statement.

The bound can be proven tight by considering an example in which $d_i = 0$ for each $d \in \mathcal{D}$ (i.e., all latency terms are constant). Consider an Bayesian congestion game in a two link parallel network in which the first edge has latency $\ell_1(f_1) = 1$ and the second has $\ell_2(f_2) = \sum_{d \in \mathcal{D}} \zeta$, where $\zeta = \alpha_{2,d} \geq 0$ for each $d \in \mathcal{D}$ is a single unknown latency parameter that represents $|\mathcal{D}|$, perfectly correlated coefficients. Let μ_0 be a prior distribution on ζ such that $\mu_0(\zeta = 0) = 1 - \epsilon$ and $\mu_0(\zeta = 1/(|\mathcal{D}| \cdot \epsilon)) = \epsilon$; as such $\mathbb{E}_{\zeta \sim \mu_0}[\zeta] = 1/|\mathcal{D}|$ and using Lemma 1 tells us the expected total latency without




Figure 5: The benefit of revealing information with and without the concurrent use of incentives. π is the uniform-grid signal structure where the support A is partitioned into a grid with granularity b ; as b increases, more information is revealed to the users. At left, the benefit of using the uniform-grid signalling policy π^b is shown for the setting described in Section IV-C with and without the concurrent use of the incentive mechanism T^* . When incentives are used, revealing information provides a positive benefit and improves performance, which is shown to be generally true in Theorem 1. With no incentives, the benefit becomes negative and revealing information worsens system cost, which was shown to be possible in Example 1 and Theorem 2.

signalling is $\mathcal{L}(\mathbf{f}^{\text{BNF}}; \mu_0, \emptyset, T^*) = 1$. Now, consider the use of a signalling policy π that reveals the state to the users; again using Lemma 1 we see the total latency with the signalling policy π is $\mathcal{L}(\mathbf{f}^{\text{BNF}}; \mu_0, \pi, T^*) = 0$. To see this matches our bound in (19), we note that the mean coefficient vector is $\mathbb{E}[\alpha] = [1, 0, \dots, 0, 1/|\mathcal{D}|, \dots, 1/|\mathcal{D}|]^T$ and the bottom of our support is $\tilde{\alpha} = [1, 0, \dots, 0]^T$. Substituting this in, we get that the bound in (19) equates to 1. $\square$

C. Benefit of Truthful Signalling

To understand how the benefit of signalling changes as more, truthful information is revealed, we offer the following numerical example. Consider a Bayesian congestion game with two edges, $\ell_1(f_1) = f_1^2 + \alpha_{1,0}$ and $\ell_2(f_2) = \alpha_{2,2}f_2^2 + 1$, where $\alpha_{1,0}$ and $\alpha_{2,2}$ are parameters unknown to the user. We consider that these two parameters are drawn from a truncated normal distribution, i.e., let

$$z \sim \mathcal{N}\left(\begin{bmatrix} 30 \\ 30 \end{bmatrix}, 180 \cdot \begin{bmatrix} 2 & 1 \\ 1 & 2 \end{bmatrix}\right),$$

and define the prior over $\alpha_{1,0}, \alpha_{2,2}$ as $\mu_0(x) = \mathbb{P}[x = z | z \in A]$ where $A = [0, 60]^2$ is their support.

Now, we analyze the benefit of the *uniform-grid signalling policy* with and without the concurrent use of the signal-aware incentive mechanism T^* as defined in Proposition 1. This signalling policy is truthful in that each user is informed accurately of what partition the realization of the latency coefficient parameter α is in, and reveals more information as the number of partitions increases. Let b be an integer representing the granularity of the signalling mechanism, i.e., the number of times A is partitioned along each dimension as shown in Fig. 5, i.e.,

$$\pi_{i,j} = \left[\frac{60}{b}(i-1), \frac{60}{b}i\right] \times \left[\frac{60}{b}(j-1), \frac{60}{b}j\right],$$

essentially forming a uniform grid over A .

In Fig. 5, we plot the benefit for using the uniform-grid signalling policy with and without the concurrent use of the incentive mechanism T^* , i.e., $\mathbf{B}(\pi; \mu_0)$ and $\mathbf{B}(\pi; \mu_0, T^*)$.

Observe that when no incentives are used, increasing the amount of information revealed to users (i.e., larger b) causes the benefit to become increasingly negative; meaning as more information is revealed, the signalling policy makes the system performance worse. Conversely, while using incentive mechanism T^* , as more information is added, the benefit becomes increasingly positive and revealing information now improves performance.


V. OPTIMAL SIGNAL DESIGN

In the preceding sections of this paper, we showed the range of possible benefit a signalling policy can provide with and without the concurrent use of monetary incentives. In this section, we address how one can compute an optimal signalling policy π^* . In general, the optimal signal can be NP-hard to find [41]; however, in many problems, this is not the case [42]. We remain in the context of parallel, polynomial-latency congestion games. However, if the signalling platform is limited (e.g., physical signs or discrete UI options), the designer may be limited to a small/finite number of possible signals, i.e., $|S| = \Lambda \leq |A|$. In Section V-A, we survey the result in [21] which shows a method of transcribing the optimal signalling policy problem as a generalized moment problem (GMP) which can be approximately solved with existing solvers [43]. Because we have observed in this work that the concurrent use of monetary incentives can aid in information signalling, we propose two extensions to solve for co-designed signal/incentive pairs with both signal-aware and signal-agnostic incentives. In Section V-B, we show that the signal-agnostic co-design problem can be done by an expansion of the decision variables and the problem remains a GMP. In Section V-C, we show that the signal-aware co-design allows for a simplification where the polynomial constraints can be removed, making the program geometric and solvable via convex programming techniques.

A. Computing Optimal Signals

We assume that α is realized from a prior distribution μ_0 with finite support $A = \{\alpha^1, \dots, \alpha^m\}$. A signalling policy

Latency Functions		
State 1 w.p. 0.3	State 2 w.p. 0.4	State 3 w.p. 0.3
$\ell_1(f_1) = 25f_1 + 5$	$\ell_1(f_1) = 30f_1 + 25$	$\ell_1(f_1) = 30f_1 + 25$
$\ell_2(f_2) = 17f_2 + 10$	$\ell_2(f_2) = 35f_2 + 13$	$\ell_2(f_2) = 17f_2 + 10$
$\ell_3(f_3) = 30f_3 + 25$	$\ell_3(f_3) = 13f_3 + 15$	$\ell_3(f_3) = 30f_3 + 25$
$\ell_4(f_4) = 10f_4 + 25$	$\ell_4(f_4) = 10f_4 + 25$	$\ell_4(f_4) = 11f_4 + 35$



Info./Incentive Setting	Total Lat.
No signal / No tolls	23.99
True signal / No tolls	22.25
Opt. signal / No tolls	21.91*
No signal / w/ ag. tolls	23.41
Opt. signal / w/ ag. tolls	21.35*
No signal / w/ aw. tolls	23.41
Opt. signal / w/ aw. tolls	21.29*

Figure 6: Simulation results for the system cost (expected total latency) in a four-link parallel congestion game with affine latency functions. The comparison is made between seven information/incentive settings: no, true/full, and optimal signalling as well as no signalling and optimal signalling with the use of concurrent signal-agnostic (ag. tolls) and signal-aware (aw. tolls) incentives, respectively. The optimal signals and associated total latency are found using (PT) and (P) with and without constraints (23); the asterisk denotes that the solution is approximate, found using the GloptiPoly solver. We find that the optimal signals provide notable improvements over signalling naively (true signal) and that both types of tolls further aid the benefit of signalling.

$\pi : A \rightarrow \Delta(S)$ can now be represented by an $m \times \Lambda$ column stochastic matrix, where $\Lambda \leq m$ defines the designer's constraint on their available signals⁵, and $\pi(s, k) = \mathbb{P}[s|\alpha^k]$. A signal-dependent flow tuple $\mathbf{f}$ can be represented by a $\Lambda \times n$ matrix, where $\mathbf{f}(s, e)$ is the flow on edge e in the flow that emerges after receiving signal s . The expected system cost of a signalling policy π with flows $\mathbf{f}$ can thus be written as

$$\begin{aligned} \mathcal{L}(\mathbf{f}; \mu_0, \pi) &= \sum_{k=1}^m \sum_{s=1}^{\Lambda} \mathcal{L}(\mathbf{f}(s, -); \alpha^k) \cdot \mathbb{P}[s|\alpha^k] \\ &= \sum_{k=1}^m \sum_{s=1}^{\Lambda} \sum_{e \in E} \sum_{d \in \mathcal{D}} \alpha_{e,d}^k (\mathbf{f}(s, e))^{d+1} \cdot \pi(s, k) \mu_0(k). \end{aligned} \quad (22)$$

Note that (22) is polynomial in $\mathbf{f}$ and π .

In order find the optimal signalling mechanism, we must introduce a constraint that $\mathbf{f}$ is a Bayes-Nash equilibrium; from Lemma 1, we can do this by requiring $\mathbf{f}(s, -)$ to be a Nash flow with the expected latency coefficients given the signal, i.e.,

$$\sum_{k=1}^m \ell_{e,k}(\mathbf{f}(s, e)) \mu_s(k) \leq \sum_{k=1}^m \ell_{e',k}(\mathbf{f}(s, e')) \mu_s(k), \quad \forall e \in E \text{ s.t. } \mathbf{f}(s, e) > 0, \quad e' \in E, \quad s \in S,$$

where $\mu_s(k) = \frac{\pi(s, k) \mu_0(k)}{\psi(s)}$; this can be rewritten as

$$\mathbf{f}(s, e) \cdot \sum_{k=1}^m (\ell_{e,k}(\mathbf{f}(s, e)) - \ell_{e',k}(\mathbf{f}(s, e'))) \pi(s, k) \mu_0(k) \leq 0, \quad \forall e, e' \in E, \quad s \in S. \quad (23)$$

Using (22) and (23), along with other constraints, we can write the following optimization problem, whose solution is the signalling mechanism which minimizes expected total latency in a Bayes-Nash flow:

$$\begin{aligned} &\text{minimize} && \mathcal{L}(\mathbf{f}; \mu_0, \pi) \\ &\mathbf{f} \in \mathbb{R}_{\geq 0}^{\Lambda \times n}, \quad \pi \in \mathbb{R}_{\geq 0}^{\Lambda \times m} \\ &\text{subject to} && (23), \\ &&& \mathbb{1}_{\Lambda}^T \pi = \mathbb{1}_m^T, \\ &&& \mathbf{f} \mathbb{1}_n = \mathbb{1}_{\Lambda} \end{aligned} \quad (P)$$

Note that (P) has a polynomial objective, polynomial inequality constraints, and linear equality constraints. Problems

⁵From [6], we need not consider signal sets with more than m signals.

of this form can be cast as instances of the generalized problem of moments and solved approximately using a semidefinite programming approach [43], as discussed in [21]. In Section V-D, we will use the solution to (P) to numerically investigate the benefit of optimal signals.

B. Optimal Signal-Agnostic Co-design

Example 2 showed that signals and tolls may be less effective when designed separately. We will show how the Co-design with signal-agnostic incentives can be done with little more complication than the signal design case. Let $\tau \in \mathbb{R}_{\geq 0}^n$ be a vector for the signal-agnostic incentive levied on each edge. The objective of the optimization problem will remain the same as in (22); however the equilibrium constraint will be affected by τ . The new equilibrium constraints become

$$\begin{aligned} \mathbf{f}(s, e) \sum_{k=1}^m (\ell_{e,k}(\mathbf{f}(s, e)) + \tau_e - \ell_{e',k}(\mathbf{f}(s, e')) - \tau_{e'}) \pi(s, k) \mu_0(k) \\ \leq 0, \quad \forall e, e' \in E, \quad s \in S. \end{aligned} \quad (24)$$

Using (22) and (24), we get the new program

$$\begin{aligned} &\text{minimize} && \mathcal{L}(\mathbf{f}; \mu_0, \pi) \\ &\mathbf{f} \in \mathbb{R}_{\geq 0}^{\Lambda \times n}, \quad \pi \in \mathbb{R}_{\geq 0}^{\Lambda \times m}, \quad \tau \in \mathbb{R}_{\geq 0}^n \\ &\text{subject to} && (24), \\ &&& \mathbb{1}_{\Lambda}^T \pi = \mathbb{1}_m^T, \\ &&& \mathbf{f} \mathbb{1}_n = \mathbb{1}_{\Lambda} \end{aligned} \quad (PT)$$

The program Eq. (PT) belongs to the same class of GMPs as Eq. (P), but with more constraints. In Section V-D and Fig. 6, we discuss how the co-designed mechanisms may improve performance.

C. Optimal Signal-Aware Co-design

In this subsection, we seek to solve for optimal signals/incentive pairs with signal-aware incentives. The incentive design portion of this task is handled by Proposition 1, which states that the incentive mechanism T^* is optimal for any π . We thus look for how to design a signalling policy while concurrently using these incentives.

Remark 1. The optimal signal-incentive pair (π^*, T^*) uses monetary incentives from Proposition 1 and signalling policy from the solution to (P) without constraints (23).

Remark 1 follows from the fact that T^* causes the Bayes-Nash flow $f(s, \cdot)$ after receiving a signal s to be one that minimizes the expected total latency given s . As such, removing (23) from (P) allows $f(s, \cdot)$ to be any feasible flow; in the minimization problem $f(s, \cdot)$ thus becomes one that minimizes the expected total latency, or be what emerges from using T^* .

Additionally, after removing (23) from (P), the problem has only linear equality constraints and a posynomial objective. This problem is thus a geometric program and can be transformed to a convex optimization problem [40]. We will use the solution to this program to compare the effectiveness of the optimal signal-incentive pair with signalling-aware incentive to other settings in Section V-D.

D. Value of Optimal Signalling

To quantify the performance of optimal signalling mechanisms, we discuss a generated numerical example and draw several conclusions. The example is described in Fig. 6 and depicts a setting in which users must traverse a parallel network with four edges whose travel delays grow in an affine manner. Users are uncertain of these latency functions but known they come from three possible states (potentially caused by road accidents or weather related hazards). In this problem, we compute the expected total latency in seven different settings: no signalling, truthful signalling which reveals the exact state, and the optimal signal, as well as no signalling and optimal signalling alongside the optimal signal-agnostic and signal-aware incentives respectively. The optimal signals are found using the polynomial optimization solver GloptiPoly [44], which casts the problem as a generalized moment problem and finds an approximate solution via semi-definite programming (the asterisk in Fig. 6 is to denote the signalling mechanisms are found by this approximate solution method). We identify the following observation from the simulation:

- 1) Signalling can offer notable performance improvements. Simply revealing the truth offered a 7.25% reduction in system cost and signalling optimally offered an 8.67% reduction.
- 2) Incentives can further aid in the capabilities of signalling. The optimal signal-incentive pairs – for both signal-aware and signal-agnostic incentives – offered the most significant performance improvements over signalling alone or incentivizing alone.
- 3) Using signal-aware incentives gives the best performance and makes optimal mechanisms easiest to compute. This is apparent from the last row of the right table in Fig. 6 and 1.

VI. CONCLUSION

In this paper, we study the effectiveness of information signalling in the context Bayesian congestion games. Our main observations are that designing signalling mechanisms and monetary incentives concurrently can offer improvements that cannot be offered by either alone; one such improvement is that concurrently using appropriate monetary incentives and information signals can help avoid cases where revealing information worsens expected total travel latency. To further this understanding, we derive bounds on the possible benefit

of signalling with and without the concurrent use of monetary incentives and provide methods to compute the optimal signalling policies.

Future work may investigate the capabilities of a system operator with less reliable mechanisms (e.g., uncertainty of their own about the system state and heterogeneity in users' beliefs and responses to incentives). Additionally, further studies may uncover if these conclusions exist in settings outside of Bayesian congestion games.

REFERENCES

- [1] B. L. Ferguson, P. N. Brown, and J. R. Marden, "Avoiding Unintended Consequences: How Incentives Aid Information Provisioning in Bayesian Congestion Games," *ArXiv preprint arXiv:2204.06046*, apr 2022. [Online]. Available: <https://arxiv.org/abs/2204.06046v1>
- [2] A. C. Pigou, *The Economics of Welfare*. Macmillan, New York, 1920.
- [3] G. Hardin, "The Tragedy of the Commons," *Science*, vol. 162, no. 3859, pp. 1243–1248, 1968. [Online]. Available: <https://www.science.org/doi/abs/10.1126/science.162.3859.1243>
- [4] J. Liu, S. Amin, and G. Schwartz, "Effects of information heterogeneity in Bayesian routing games," *arXiv preprint arXiv:1603.08853*, 2016.
- [5] D. Acemoglu, A. Makhdoumi, A. Malekian, and A. Ozdaglar, "Informational Braess' Paradox: The Effect of Information on Traffic Congestion," *Operations Research*, vol. 66, no. 4, 2018.
- [6] E. Kamenica and M. Gentzkow, "Bayesian Persuasion," *American Economic Review*, vol. 101, no. 6, pp. 2590–2615, oct 2011. [Online]. Available: <https://www.aeaweb.org/articles?id=10.1257/aer.101.6.2590>
- [7] I. Goldstein and C. Huang, "Bayesian Persuasion in Coordination Games," *American Economic Review*, vol. 106, no. 5, pp. 592–596, may 2016. [Online]. Available: <https://www.aeaweb.org/articles?id=10.1257/aer.p20161047>
- [8] D. Bergemann and S. Morris, "Information design: A unified perspective," *Journal of Economic Literature*, vol. 57, no. 1, pp. 44–95, 2019.
- [9] F. Sezer, H. Khazaei, and C. Eksin, "Social Welfare Maximization and Conformism via Information Design in Linear-Quadratic-Gaussian Games," *arXiv preprint arXiv:2102.13047*, 2021.
- [10] R. Lindsey, T. Daniel, E. Gisches, and A. Rapoport, "Pre-trip information and route-choice decisions with stochastic travel conditions: Theory," *Transportation Research Part B: Methodological*, vol. 67, no. C, pp. 187–207, 2014. [Online]. Available: <https://ideas.repec.org/a/eee/transb/v67y2014icp187-207.html>
- [11] J. Wardrop, "Some Theoretical Aspects of Road Traffic Research," in *Proc. of the Institution of Civil Engineers, Part II, Vol. 1, No. 36*, 1952, pp. 352–362.
- [12] R. W. Rosenthal, "A class of games possessing pure-strategy Nash equilibria," *International Journal of Game Theory*, vol. 2, no. 1, pp. 65–67, dec 1973.
- [13] C. S. Fisk, "Game theory and transportation systems modelling," *Transportation Research Part B: Methodological*, vol. 18, no. 4-5, pp. 301–313, aug 1984.
- [14] D. Braess, A. Nagurney, and T. Wakolbinger, "On a Paradox of Traffic Planning," *Transportation Science*, vol. 39, no. 4, pp. 446–450, nov 2005. [Online]. Available: <http://transci.journal.informs.org/cgi/doi/10.1287/trsc.1050.0127>
- [15] D. Levinson, "Micro-foundations of congestion and pricing: A game theory perspective," *Transportation Research Part A: Policy and Practice*, vol. 39, no. 7-9, pp. 691–704, aug 2005.
- [16] F. Farokhi and K. H. Johansson, "A Study of Truck Platooning Incentives Using a Congestion Game," *IEEE Transactions on Intelligent Transportation Systems*, vol. 16, no. 2, pp. 581–595, 2015.
- [17] A. Yekkehkhany and R. Nagi, "Risk-Averse Equilibria for Vehicle Navigation in Stochastic Congestion Games," *IEEE Transactions on Intelligent Transportation Systems*, pp. 1–17, 2022.
- [18] S. Spana, L. Du, and Y. Yin, "Strategic Information Perturbation for an Online In-Vehicle Coordinated Routing Mechanism for Connected Vehicles Under Mixed-Strategy Congestion Game," *IEEE Transactions on Intelligent Transportation Systems*, vol. 23, no. 5, pp. 4541–4555, 2022.
- [19] Y. Chen and V. Gupta, "Toll Design for Routing Games With Stochastic Demands," *IEEE Control Systems Letters*, vol. 6, pp. 3445–3450, 2022.
- [20] M. Castiglioni, A. Celli, A. Marchesi, and N. Gatti, "Signaling in Bayesian Network Congestion Games: The Subtle Power of Symmetry," feb 2020. [Online]. Available: <http://arxiv.org/abs/2002.05190>

- [21] Y. Zhu and K. Savla, "A Semidefinite Approach to Information Design in Non-atomic Routing Games," *CoRR*, vol. abs/2005.0, 2020. [Online]. Available: <https://arxiv.org/abs/2005.03000>
- [22] M. Wu and S. Amin, "Information design for regulating traffic flows under uncertain network state," in *2019 57th Annual Allerton Conference on Communication, Control, and Computing (Allerton)*. IEEE, 2019, pp. 671–678.
- [23] M. Wu, S. Amin, and A. E. Ozdaglar, "Value of Information in Bayesian Routing Games," *Operations Research*, vol. 69, no. 1, pp. 148–163, jan 2021. [Online]. Available: <http://pubsonline.informs.org/doi/10.1287/opre.2020.1999>
- [24] J. Verbree and A. Cherukuri, "Inferring the prior in routing games using public signalling," 2021. [Online]. Available: <https://arxiv.org/abs/2109.05895>
- [25] H. Tavaafoghi and D. Teneketzis, "Informational incentives for congestion games," in *55th Annual Allerton Conference on Communication, Control, and Computing, Allerton 2017*, vol. 2018-Janua. Institute of Electrical and Electronics Engineers Inc., jan 2018, pp. 1285–1292.
- [26] M. H. Lotfi, R. J. La, and N. C. Martins, "Bayesian Congestion Game with Traffic Manager: Binary Signal Case," in *2018 IEEE Conference on Decision and Control (CDC)*, 2018, pp. 327–333.
- [27] L. Fleischer, K. Jain, and M. Mahdian, "Tolls for Heterogeneous Selfish Users in Multicommodity Networks and Generalized Congestion Games," in *Proc. 45th IEEE Symp. on Foundations of Computer Science*, Rome, Italy, 2004, pp. 277–285. [Online]. Available: http://ieeexplore.ieee.org/lpdocs/epic03/wrapper.htm?arnumber=1366247http://ieeexplore.ieee.org/xpls/abs/_all.jsp?arnumber=1366247
- [28] B. L. Ferguson, P. N. Brown, and J. R. Marden, "Utilizing Information Optimally to Influence Distributed Network Routing," in *Proceedings of the IEEE Conference on Decision and Control*, vol. 2019-Decem. Institute of Electrical and Electronics Engineers Inc., dec 2019, pp. 5008–5013.
- [29] —, "The Effectiveness of Subsidies and Tolls in Congestion Games," *IEEE Transactions on Automatic Control*, vol. 67, no. 6, pp. 2729–2742, 2022. [Online]. Available: <https://ieeexplore.ieee.org/document/9451652/>
- [30] R. Cole, Y. Dodis, and T. Roughgarden, "Pricing network edges for heterogeneous selfish users," in *Proc. of the 35th ACM symposium on Theory of computing*, New York, New York, USA, 2003, pp. 521–530. [Online]. Available: <http://portal.acm.org/citation.cfm?doid=780542.780618>
- [31] D. Paccagnan, R. Chandan, B. L. Ferguson, and J. R. Marden, "Optimal Taxes in Atomic Congestion Games," *ACM Trans. Econ. Comput.*, vol. 9, no. 3, aug 2021. [Online]. Available: <https://doi.org/10.1145/3457168>
- [32] A. de Palma, R. Lindsey, and N. Picard, "Risk Aversion, the Value of Information, and Traffic Equilibrium," *Transportation Science*, vol. 46, 2012.
- [33] G. Egorov and K. Sonin, "Persuasion on Networks," National Bureau of Economic Research, Working Paper 27631, jul 2020. [Online]. Available: <http://www.nber.org/papers/w27631>
- [34] N. Heydari Beni and A. Anastasopoulos, "Joint Information and Mechanism Design for Queues with Heterogeneous Users," in *2021 60th IEEE Conference on Decision and Control (CDC)*, 2021, pp. 3806–3811.
- [35] W. H. Sandholm, "Potential Games with Continuous Player Sets," *Journal of Economic Theory*, vol. 97, pp. 81–108, 2001. [Online]. Available: <http://linkinghub.elsevier.com/retrieve/pii/S0022053100926966%5Cnhttp://www.sciencedirect.com/science/article/pii/S0022053100926966>
- [36] T. Roughgarden, "How Bad Is Selfish Routing?" *J. ACM*, vol. 49, no. 2, pp. 236–259, 2002.
- [37] A. Mas-Colell, "On a Theorem of Schmeidler," *Mathematical Economics*, vol. 13, pp. 201–206, 1984.
- [38] F. Farokhi and K. H. Johansson, "A piecewise-constant congestion taxing policy for repeated routing games," *Transportation Research Part B: Methodological*, vol. 78, pp. 123–143, 2015. [Online]. Available: <http://dx.doi.org/10.1016/j.trb.2015.04.003>
- [39] R. Singh and R. Dowling, "Improved Speed-Flow Relationships: Application To Transportation Planning Models," *Seventh TRB Conference on the Application of Transportation Planning Methods*, no. 4, pp. 340–349, 2002. [Online]. Available: <http://trid.trb.org/view.aspx?id=734743%5Cnhttp://docs.trb.org/00939750.pdf>
- [40] S. Boyd, S. P. Boyd, and L. Vandenberghe, *Convex optimization*. Cambridge university press, 2004.
- [41] U. Bhaskar, Y. Cheng, Y. K. Ko, and C. Swamy, "Hardness Results for Signaling in Bayesian Zero-Sum and Network Routing Games," 2015. [Online]. Available: <https://arxiv.org/abs/1512.03543>
- [42] S. M. Griesbach, M. Hoefer, M. Klimm, and T. Koglin, "Public Signals in Network Congestion Games," 2022. [Online]. Available: <https://arxiv.org/abs/2205.09823>
- [43] J. B. Lasserre, "A semidefinite programming approach to the generalized problem of moments," *Mathematical Programming* 2007 112:1, vol. 112, no. 1, pp. 65–92, feb 2007. [Online]. Available: <https://link.springer.com/article/10.1007/s10107-006-0085-1>
- [44] D. Henrion, J. B. Lasserre, and J. Lofberg, "GloptiPoly 3: Moments, optimization and semidefinite programming," *Optimization Methods and Software*, vol. 24, no. 4-5, pp. 761–779, aug 2009.
- [45] I. Milchtaich, "Network topology and the efficiency of equilibrium," *Games and Economic Behavior*, vol. 57, no. 2, pp. 321–346, 2006. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0899825605001284>

APPENDIX

Remark 2. *In parallel, polynomial Bayesian congestion games, without loss of generality, we can assume a unit traffic rate, $r = 1$, even when $r \sim \nu$ is a random variable.*

proof: Consider a congestion game G with demand r . Define a mapping $Q(G, \gamma)$ that outputs a new congestion game $\hat{G}$ with latency functions $\hat{\ell}_e(x) = \sum_{d \in \mathcal{D}} \frac{\alpha_{e,d}}{\gamma^{d+1}} (f_e)^d$. Let f be a flow in G with total traffic r . Now, consider the flow $\gamma f = \{\gamma \cdot f_e\}_{e \in E}$ in $\hat{G}$. Each edge $e \in E$ will have latency

$$\hat{\ell}_e(\gamma f_e) = \sum_{d \in \mathcal{D}} \frac{\alpha_{e,d}}{\gamma^{d+1}} (\gamma f_e)^d = \frac{1}{\gamma} \ell_e(f_e).$$

Notice that latency on each edge is scaled by $1/\gamma$ and the preference structure is preserved; therefore, if f is a Nash flow in G then γf is a Nash flow in $\hat{G}$. Further,

$$\mathcal{L}(\gamma f; \hat{G}) = \sum_{e \in E} \gamma f_e \hat{\ell}_e(\gamma f_e) = \sum_{e \in E} f_e \ell_e(f_e) = \mathcal{L}(f; G),$$

and the two networks will have the same total latency.

If $(\alpha, r) \sim \mu_0$, e.g., $\mu_0(x, y) = \mathbb{P}[\alpha = x, r = y]$, then consider that $\hat{\alpha} \sim \hat{\mu}_0$, where $\hat{\mu}_0(z) = \sum_{x,y|Q(x,1/y)=z} \mu_0(x, y)$. Now $\hat{\alpha}$ has the same distribution over total latency. $\square$

Proof of Lemma 1: To prove the first claim, consider the prior μ_0 on α and the signalling policy $\pi : A \rightarrow \Delta(S)$. If the signal $s \in S$ is sent to users, they update their belief via Bayesian inference to $\mu_s(x) = \frac{\pi(s|\alpha) \cdot \mu_0(\alpha)}{\psi(s)}$. In a flow f , user $x \in [0, r_i]$ taking path $P_x \in \mathcal{P}_i$ experiences an expected cost of

$$\begin{aligned} J_x(P_x; f, \mu_s) &= \mathbb{E}_{\alpha \sim \mu_s} \left[\sum_{e \in P_x} \sum_{d \in \mathcal{D}} \alpha_{e,d} \cdot \ell_d(f_e) \right] \\ &= \sum_{e \in P_x} \sum_{d \in \mathcal{D}} \mathbb{E}[\alpha_{e,d} | s] \ell_d(f_e). \end{aligned}$$

Note that if f were not a Nash flow in the congestion game with coefficients $\mathbb{E}[\alpha | s]$, then by (3) at least one user x would be able to deviate their strategy $\sigma_x(s)$ and experience lower cost. Therefore, the only Bayes-Nash flows occur when $f(s)$ is a Nash flow with respect to $\mathbb{E}[\alpha | s]$ for all $s \in S$. Further, because the total latency in a Nash flow is unique, so too is the expected total latency in a Bayes-Nash flow.

To prove the second claim, consider the distribution μ_0 on α , and let f be a flow in the network. The expected total latency

$$\begin{aligned}\mathcal{L}(f; \mu) &= \mathbb{E}_{\alpha \sim \mu} \left[\sum_{e \in E} f_e \cdot \sum_{d \in \mathcal{D}} \alpha_{e,d} \ell_d(f_e) \right] \\ &= \sum_{e \in E} f_e \cdot \sum_{d \in \mathcal{D}} \mathbb{E}_{\alpha \sim \mu} [\alpha_{e,d}] \ell_d(f_e) \\ &= \mathcal{L}(f; \mathbb{E}_{\alpha \sim \mu} [\alpha]),\end{aligned}$$

which follows from linearity of expected value. $\square$

Proof of Proposition 1: Consider that the users receive signal s from signalling policy π and prior μ_0 (forming posterior μ_s), and reach a flow of f . From Lemma 1 the expected total latency in this flow is equal to the total latency of this flow in the expected network, i.e., $\mathcal{L}(f; \mu_s) = \mathcal{L}(f; \bar{\alpha}_s)$, where $\bar{\alpha}_s = \mathbb{E}_{\alpha \sim \mu_s} [\alpha]$. Thus, an optimal flow at the reception of signal s is one that satisfies $f^{\text{opt}}(s) \in \arg \min_f \mathcal{L}(f; \bar{\alpha}_s)$.

We now look for incentives that will influence users such that $f^{\text{opt}}(s)$ becomes a Nash flow in a congestion game with latency coefficients $\bar{\alpha}_s$. To do so, we note that G with flow-varying incentive functions $\tau_e(f_e)$ is a potential game [12] with potential function

$$\Phi(f; \alpha) = \sum_{e \in E} \int_0^{f_e} \ell_e(x) + \tau_e(x) dx.$$

As such, the flow in $\arg \min_f \Phi(f; \alpha)$ is a Nash equilibrium. For the polynomial latency functions considered in this work, let $\tau_e(x) = \sum_{d \in \mathcal{D}} x \alpha_{e,d} \ell'_d(x)$. Now, the potential function becomes

$$\begin{aligned}\Phi(f; \alpha) &= \sum_{e \in E} \int_0^{f_e} \sum_{d \in \mathcal{D}} \alpha_{e,d} \ell_d(x) + x \alpha_{e,d} \ell'_d(x) dx \\ &= \sum_{e \in E} \sum_{d \in \mathcal{D}} f_e \alpha_{e,d} \ell_d(f_e) = \mathcal{L}(f; \alpha),\end{aligned}$$

and the Nash flow that minimizes Φ also minimizes $\mathcal{L}$; as such, $f^{\text{opt}}(s)$ becomes a Nash equilibrium in the game with the coefficients $\bar{\alpha}_s$.

Finally, notice that by selecting the fixed incentive $\tau_e^*(s) = \tau_e(f_e^{\text{opt}}(s))$, the equilibrium conditions do not change and $f^{\text{opt}}(s)$ remains a Nash flow. Nash flows retain the same uniqueness properties under fixed incentives, and thus assigning $\tau^*(s)$ minimizes the expected total latency when s is sent. If this is done for each signal, the total latency with each signal will be minimal and so too will the overall expected total latency, making T^* an optimal incentive mechanism. $\square$

Proof of Lemma 2 We assume $r = 1$, which is without loss of generality from Remark 2. We note that $\mathcal{L}^{\text{Nf}}(\alpha)$ is continuous, but need not be continuously differentiable; as such we look for the largest gradient in the differentiable regions of the support. Let $f^{\text{Nf}}(\alpha)$ be the Nash flow in the parallel congestion game with polynomial coefficients α , i.e., $\mathcal{L}^{\text{Nf}}(\alpha) = \mathcal{L}(\alpha, f^{\text{Nf}}(\alpha)) = \sum_{e \in E} \sum_{d \in \mathcal{D}} \alpha_{e,d} (f_e^{\text{Nf}})^{d+1}$. First, we seek to bound the partial derivative of $\mathcal{L}^{\text{Nf}}(\alpha)$ with respect to some parameter $\alpha_{e,d}$. Clearly, $\frac{\partial}{\partial \alpha_{e,d}} \mathcal{L}^{\text{Nf}}(\alpha) \geq 0$ in parallel networks as no Braess's paradox type example can exist [45]. To upper-bound this partial derivative, we will consider a case

where by increasing $\alpha_{e,d}$ any mass of traffic that chooses to leave edge e will all choose the edge e' ; in general, this may not occur with every change in $\alpha_{e,d}$, as users may disperse over multiple edges, however, if we consider that users do all move to the same edge, and we pick edge e' as the one that increases the total latency most rapidly, then the following upper-bound will hold. With this in mind, it implies that $\frac{\partial}{\partial \alpha_{e,d}} f_e^{\text{Nf}} = -\frac{\partial}{\partial \alpha_{e,d}} f_{e'}^{\text{Nf}}$, and that we can evaluate the partial derivative as

$$\begin{aligned}\frac{\partial}{\partial \alpha_{e,d}} \mathcal{L}^{\text{Nf}}(\alpha) &= (f_e^{\text{Nf}})^{d+1} + \left(\sum_{d'' \in \mathcal{D}} \alpha_{e',d''} (d'' + 1) (f_{e'}^{\text{Nf}})^{d''+1} \right. \\ &\quad \left. - \sum_{d' \in \mathcal{D}} \alpha_{e,d'} (d' + 1) (f_e^{\text{Nf}})^{d'} \right) \frac{\partial}{\partial \alpha_{e,d}} f_{e'}^{\text{Nf}}\end{aligned}\quad (25)$$

Now, we note that latency on edges e and e' must be the same in a Nash flow, thus

$$\begin{aligned}\ell_e(f_e^{\text{Nf}}) &= \ell_{e'}(f_{e'}^{\text{Nf}}) \\ \frac{\partial}{\partial \alpha_{e,d}} \ell_e(f_e^{\text{Nf}}) &= \frac{\partial}{\partial \alpha_{e,d}} \ell_{e'}(f_{e'}^{\text{Nf}}).\end{aligned}$$

Using this equality, and the fact that $\frac{\partial}{\partial \alpha_{e,d}} f_e^{\text{Nf}} = -\frac{\partial}{\partial \alpha_{e,d}} f_{e'}^{\text{Nf}}$, we can evaluate the derivative and rearrange to get

$$\frac{\partial}{\partial \alpha_{e,d}} f_{e'}^{\text{Nf}} = \frac{(f_e^{\text{Nf}})^d}{\ell'_e(f_e^{\text{Nf}}) + \ell'_{e'}(f_{e'}^{\text{Nf}})} \leq \frac{(f_e^{\text{Nf}})^d}{2\rho_1^-}, \quad (26)$$

where $\rho_1^- = \min_{e \in E} \bar{\alpha}_{e,1}$. Substituting (26) into (25) gives us

$$\frac{\partial}{\partial \alpha_{e,d}} \mathcal{L}^{\text{Nf}}(\alpha) \leq (f_e^{\text{Nf}})^{d+1} + \frac{\rho^+ - \rho_0^-}{2\rho_1^-} (f_e^{\text{Nf}})^d, \quad (27)$$

where $\rho_0^- = \min_{e \in E} \bar{\alpha}_{e,0}$ and $\rho^+ = \max_{e \in E} \sum_{d \in \mathcal{D}} (d + 1) \bar{\alpha}_{e,d}$.

Now, the gradient of $\mathcal{L}^{\text{Nf}}(\alpha)$ must satisfy

$$\begin{aligned}\|\nabla \mathcal{L}^{\text{Nf}}(\alpha)\|_2 &\leq \sqrt{\sum_{e \in E} \sum_{d \in \mathcal{D}} \left((f_e^{\text{Nf}})^{d+1} + \frac{\rho^+ - \rho_0^-}{2\rho_1^-} (f_e^{\text{Nf}})^d \right)^2} \\ &\leq \sqrt{\left(\sum_{e \in E} \sum_{d \in \mathcal{D}} (f_e^{\text{Nf}})^{d+1} + \frac{\rho^+ - \rho_0^-}{2\rho_1^-} (f_e^{\text{Nf}})^d \right)^2} \\ &\leq \sum_{d \in \mathcal{D}} \left(\sum_{e \in E} f_e^{\text{Nf}} \right)^{d+1} + \frac{\rho^+ - \rho_0^-}{2\rho_1^-} \\ &\quad \left(\sum_{e \in E} (f_e^{\text{Nf}})^0 \sum_{d \in \mathcal{D} \setminus \{0\}} \left(\sum_{e \in E} f_e^{\text{Nf}} \right)^d \right) \\ &= |\mathcal{D}| + \frac{\rho^+ - \rho_0^-}{2\rho_1^-} (|E| + |\mathcal{D}| - 1),\end{aligned}$$

where the first inequality holds from (27), the second and third hold from the super-additivity of convex monomials of positive terms, and the final equality holds from the assumption that $r = 1$.

Finally, consider two sets of coefficients $a, b \in A$. Observe

$$\begin{aligned}
& \frac{\mathcal{L}^{\text{Nf}}(a) - \mathcal{L}^{\text{Nf}}(b)}{\|a - b\|_2} \\
&= \frac{1}{\|a - b\|_2} \int_{\lambda=0}^1 (a - b)^T \nabla \mathcal{L}^{\text{Nf}}(\lambda a + (1 - \lambda)b) d\lambda \\
&\leq \frac{1}{\|a - b\|_2} \int_{\lambda=0}^1 \|a - b\|_2 \cdot \|\nabla \mathcal{L}^{\text{Nf}}(\lambda a + (1 - \lambda)b)\|_2 d\lambda \\
&\leq \int_{\lambda=0}^1 |\mathcal{D}| + \frac{\rho^+ - \rho_0^-}{2\rho_1^-} (|E| + |\mathcal{D}| - 1) d\lambda \\
&= |\mathcal{D}| + \frac{\rho^+ - \rho_0^-}{2\rho_1^-} (|E| + |\mathcal{D}| - 1).
\end{aligned}$$

Where the first inequality holds from Cauchy-Schwarz, and the second holds from our observation above on the norm of the gradient of $\mathcal{L}^{\text{Nf}}(\alpha)$. $\square$

Proof of Lemma 3: This proof follows very similarly to the proof of Lemma 2, but now, in an optimal flow $f^*(\alpha)$, the latency on each edge is not equal, however, the marginal-cost on each edge is [2]. Let $\nu_e(f_e) = \sum_{d \in \mathcal{D}} (d + 1) \alpha_{e,d} (f_e)^d$ be the marginal cost on edge e with flow f_e . Now, in the optimal flow f^* (which emerges from using the tolls T^*),

$$\begin{aligned}
\nu_e(f_e^*) &= \nu_{e'}(f_{e'}^*) \\
\frac{\partial}{\partial \alpha_{e,d}} \nu_e(f_e^*) &= \frac{\partial}{\partial \alpha_{e',d}} \nu_{e'}(f_{e'}^*).
\end{aligned}$$

Evaluating and rearranging these derivatives gives

$$\frac{\partial}{\partial \alpha_{e,d}} f_{e'}^* = \frac{(d + 1)(f_e^*)^d}{\nu_e'(f_e^*) + \nu_{e'}'(f_{e'}^*)} \leq \frac{(d + 1)(f_e^*)^d}{4\rho_1^-}. \quad (28)$$

With this, we can upper bound the partial derivative of $\mathcal{L}^*$ as

$$\frac{\partial}{\partial \alpha_{e,d}} \mathcal{L}^*(\alpha) \leq (f_e^*)^{d+1} + \frac{\rho^+ - \rho_0^-}{4\rho_1^-} (d + 1)(f_e^*)^d. \quad (29)$$

Following the same steps as in the proof of Lemma 2, the gradient of $\mathcal{L}^*$ must satisfy

$$\|\nabla \mathcal{L}^*(\alpha)\|_2 \leq |\mathcal{D}| + \frac{\rho^+ - \rho_0^-}{4\rho_1^-} \left(|E| + \sum_{d \in \mathcal{D} \setminus \{0\}} (d + 1)^d \right).$$

Finally, we can use this bound as in the proof of Lemma 2 to complete the proof. $\square$



Philip N. Brown is an Assistant Professor in the Department of Computer Science at the University of Colorado, Colorado Springs. Philip received the Bachelor of Science in Electrical Engineering in 2007 from Georgia Tech, after which he spent several years designing control systems and process technology for the biodiesel industry. He received the Master of Science in Electrical Engineering in 2015 from the University of Colorado at Boulder under the supervision of Jason R. Marden, where he was a recipient of the University of Colorado Chancellor's Fellowship. He received the PhD in Electrical and Computer Engineering from the University of California, Santa Barbara under the supervision of Jason R. Marden. He was finalist for the Best Student Paper Award at the 2016 and 2017 IEEE Conferences on Decision and Control, received the 2018 CCDC Best PhD Thesis Award from UCSB, and received the AFOSR Young Investigator Award. Philip is interested in the interactions between engineered and social systems.



Jason R. Marden is a Professor in the Department of Electrical and Computer Engineering at the University of California, Santa Barbara. Jason received the Bachelor of Science in Mechanical Engineering in 2001 from UCLA, and the PhD in Mechanical Engineering in 2007, also from UCLA, under the supervision of Jeff S. Shamma, where he was awarded the Outstanding Graduating PhD Student in Mechanical Engineering. After graduating from UCLA, he served as a junior fellow in the Social and Information Sciences Laboratory at the California Institute of Technology until 2010, and then as an Assistant Professor at the University of Colorado until 2015. Jason is a recipient of an ONR Young Investigator Award (2015), NSF Career Award (2014), AFOSR Young Investigator Award (2012), SIAM CST Best Sicon Paper Award (2015), and the American Automatic Control Council Donald P. Eckman Award (2012). Jason's research interests focus on game theoretic methods for the control of distributed multiagent systems.



Bryce L. Ferguson is a PhD student in the Electrical and Computer Engineering Department at the University of California, Santa Barbara. Bryce received his BS and MS in Electrical Engineering from the University of California, Santa Barbara in June 2018 and March 2020 respectively, and his A.A. in Mathematics from Santa Rosa Junior College in 2016. He was named a 2022 CPS Rising star and was a finalist for the Best Student Paper Award at the 2020 American Controls Conference. Bryce's research interests focus on using game theoretic methods for describing and controlling both societal and engineered multiagent systems.