Signaling in Bayesian Network Congestion Games: the Subtle Power of Symmetry

Matteo Castiglioni, Andrea Celli, Alberto Marchesi, Nicola Gatti
Politecnico di Milano, Piazza Leonardo da Vinci 32, I-20133, Milan, Italy
{name.surname}@polimi.it

Abstract

Network congestion games are a well-understood model of multi-agent strategic interactions. Despite their ubiquitous applications, it is not clear whether it is possible to design information structures to ameliorate the overall experience of the network users. We focus on Bayesian games with atomic players, where network vagaries are modeled via a (random) state of nature which determines the costs incurred by the players. A third-party entity—the *sender*—can observe the realized state of the network and exploit this additional information to send a *signal* to each player. A natural question is the following: *is it possible for an informed sender to reduce the overall social cost via the strategic provision of information to players who update their beliefs rationally?* The paper focuses on the problem of computing optimal *ex ante* persuasive signaling schemes, showing that *symmetry* is a crucial property for its solution. Indeed, we show that an optimal *ex ante* persuasive signaling scheme can be computed in polynomial time when players are symmetric and have affine cost functions. Moreover, the problem becomes NP-hard when players are asymmetric, even in non-Bayesian settings.

Introduction

Network congestion games, where players seek to minimize their own costs selfishly, are a canonical example of a setting where externalities may induce socially inefficient outcomes [Roughgarden 2005]. In real-world problems, the state of the network may be uncertain, and not known to its users (*e.g.*, drivers may not be aware of road works and accidents in a road network). This setting is modeled via *Bayesian network congestion games* (BNCGs). We investigate whether providing players with partial information about the state of the network may mitigate inefficiencies.

We model this information-structure design problem through the Bayesian persuasion framework by Kamenica and Gentzkow [2011]. At its core, this framework involves an informed *sender* trying to influence the behavior of a set of self-interested players—the *receivers*—via the provision of payoff-relevant information. In the specific case of BNCGs, the sender is informed about the realized state of the network. The model assumes that the sender has the ability to commit to a publicly disclosed *signaling scheme*, which is a randomized mapping from network states to action recommendations (*i.e.*, routes suggestions) for the players. The commitment assumption is realistic in many settings [Dughmi 2017]. These include BNCGs, where signaling schemes are usually implemented as software (*e.g.*, real-time traffic apps) whose features are publicly revealed by the sender. One argument to that effect is that reputation and credibility may be a key factor for the long-term utility of the sender [Rayo and Segal 2010]. We focus on the notion of *ex ante persuasiveness*, as introduced by Xu [2020] and Celli, Coniglio, and Gatti [2020], where the receivers are incentivized to follow the sender's recommendations having observed only the signaling scheme. This assumes credible receivers' commitments to following recommendations, which is reasonable in practice. In our setting, it is natural to assume that each player decides to either follow the signaling scheme (*i.e.*, adopting the traffic app) or act based on her prior belief about the network state. In some cases, the receivers could also be forced to stick to their *ex ante* commitment by contractual agreements or penalties.

Related Works Arnott, De Palma, and Lindsey [1991] and Acemoglu et al. [2018] study the impact of information on traffic congestion. Several recent works focus on *non-atomic* games [Das, Kamenica, and Mirka 2017; Massicot and Langbort 2019; Wu, Amin, and Ozdaglar 2018; Vasserman, Feldman, and Hassidim 2015]. The majority of these works focus on *ex interim* persuasiveness, where the receivers are incentivized to follow recommendations after receiving them. Bhaskar et al. [2016] study the inapproximability of finding optimal *ex interim* persuasive signaling schemes in non-atomic games. Liu and Whinston [2019] focus on atomic games with costs uncertainties and study *ex interim* persuasion by placing stringent constraints on the network structure. To the best of our knowledge, the present work is the first studying *ex ante* persuasion in general atomic BNCGs. Other related works study the simpler problem of finding (coarse) correlated equilibria in non-Bayesian congestion games [Christodoulou and Koutsoupias 2005; Papadimitriou and Roughgarden 2008]. The closest to our work is that of Jiang and Leyton-Brown [2011], who provide a polynomial-time algorithm to find an optimal coarse correlated equilibrium (*i.e.*, an *ex ante* persuasive signaling scheme in the non-Bayesian setting) in simple congestion games with symmetric players selecting a single resource (a.k.a. *singleton* congestion games).

Original Contributions We investigate whether it is possible to efficiently compute optimal (*i.e.*, minimizing the social cost) *ex ante* persuasive signaling schemes in BNCGs. First, we show that an optimal *ex ante* persuasive signaling scheme can be computed in polynomial time in symmetric BNCGs (*i.e.*, where all the players share the same source and destination pair) with edge costs defined as affine functions of the edge congestion. To prove this result, we exploit the ellipsoid algorithm by designing a sophisticated polynomial-time separation oracle based on a suitably defined min-cost flow problem. Then, we show that *symmetry* is a crucial property for efficient signaling by proving that it is NP-hard to compute an optimal *ex ante* persuasive signaling scheme in *asymmetric* BNCGs. Our reduction proves an even stronger hardness result, as it works for non-Bayesian singleton congestion games with affine costs, which is arguably the simplest class of asymmetric congestion games. Furthermore, in such setting, a solution to our problem is an optimal coarse correlated equilibrium and, thus, computing optimal coarse correlated equilibria is NP-hard.

Signaling in Network Congestion Games

We study atomic network congestion games where edge costs depend on a stochastic state of nature. In this section, we introduce the main elements of our model.

Network Congestion Game (NCG) A *network congestion game* [Fabrikant, Papadimitriou, and Talwar 2004] is defined as a tuple $(N, G, \{c_e\}_{e\in E}, \{(s_p, t_p)\}_{p\in N})$, where:

- $N := \{1, \ldots, n\}$ denotes the set of players;
- $G := (V, E)$ is the directed graph underlying the game, with V being its set of nodes and each $e = (v, v') \in E$ representing a directed edge from v to v';
- $\{c_e\}_{e\in E}$ are the edge costs, with each $c_e : \mathbb{N} \to \mathbb{R}_+$ defining the cost of edge $e \in E$ as a function of the number of players traveling through e;
- $\{(s_p, t_p)\}_{p\in N}$, with $s_p, t_p \in V$, denote the source-destination pairs for all the players.

In an NCG, the set A_p of actions available to a player $p \in N$ is implicitly defined by the graph G, the source s_p, and the destination t_p. Formally, A_p is the set of all directed paths from s_p to t_p in the graph G. In this work, we use $a_p \in A_p$ to denote a player p's path and we write $e \in a_p$ whenever the path contains the edge $e \in E$. An action profile $a \in A$, where $A := \times_{p\in N} A_p$, is a tuple of s_p-t_p directed paths $a_p \in A_p$, one per player $p \in N$. Sometimes, we denote an action profile $a \in A$ as $a = (a_p, a_{-p})$, where $a_p \in A_p$ is the action played by player $p \in N$ and a_{-p} collectively denotes the actions of the other players. For the ease of notation, given an action profile $a \in A$, we let f_e^a be the congestion of edge $e \in E$ in a, *i.e.*, the number of players selecting a path passing thorough e in a; formally, $f_e^a := |\{p \in N \mid e \in a_p\}|$. Thus, $c_e(f_e^a)$ denotes the cost of edge e in a. Finally, the cost incurred by player $p \in N$ in an action profile $a \in A$ is denoted by $c_p(a) := \sum_{e\in a_p} c_e(f_e^a)$.

Bayesian Network Congestion Game (BNCG) We define a *Bayesian network congestion game* as a tuple $(N, G, \Theta, \mu, \{c_{e,\theta}\}_{e\in E,\theta\in\Theta}, \{(s_p, t_p)\}_{p\in N})$, where, differently from the basic setting, the edge cost functions $c_{e,\theta} : \mathbb{N} \to \mathbb{R}_+$ also depend on a state of nature θ drawn from a finite set of states Θ. Moreover, μ encodes the prior beliefs that the players have over the states of nature, *i.e.*, $\mu \in \text{int}(\Delta_\Theta)$ is a fully-supported probability distribution over the set Θ, with μ_θ denoting the prior probability that the state of nature is $\theta \in \Theta$. All the other components are defined as in non-Bayesian NCGs. Notice that, in BNCGs, the cost experienced by player $p \in N$ in an action profile $a \in A$ also depends on the state of nature $\theta \in \Theta$, and, thus, it is defined as $c_{p,\theta}(a) := \sum_{e\in a_p} c_{e,\theta}(f_e^a)$. A BNCG is *symmetric* if all the players share the same (s_p, t_p) pair, *i.e.*, whenever they all have the same set of actions (paths). For the ease of notation, in such settings we let $s, t \in V$ be the common source and destination. Moreover, we focus on BNCGs with *affine costs*, *i.e.*, for all $e \in E$ and $\theta \in \Theta$, there exist constants $\alpha_{e,\theta}, \beta_{e,\theta} \in \mathbb{R}_+$ such that the edge cost function can be expressed as $c_{e,\theta}(f_e^a) := \alpha_{e,\theta} f_e^a + \beta_{e,\theta}$. [1]

Signaling in BNCGs Suppose that a BNCG is employed to model a road network subject to vagaries. It is reasonable to assume that third-party entities (*e.g.*, the road management company) may have access to the realized state of nature. We call one such entity *the sender*. We focus on the following natural question: *is it possible for an informed sender to mitigate the overall costs through the strategic provision of information to players who update their beliefs rationally?* The sender can publicly commit to a *signaling scheme* which maps the realized state of nature to a signal for each player. The sender can exploit general *private* signaling schemes, sending different signals to each player through private communication channels. In this setting, a simple revelation-principle-style argument shows that it is enough to employ players' actions as signals [Arieli and Babichenko 2016; Kamenica and Gentzkow 2011]. Therefore, a private signaling scheme is a function $\phi : \Theta \to \Delta_A$ which maps any state of nature to a probability distribution over action profiles (signals). For the ease of notation, the probability of recommending an action profile $a \in A$ given the state of nature $\theta \in \Theta$ is denoted by $\phi_{\theta,a}$. Then, it has to hold $\sum_{a\in A} \phi_{\theta,a} = 1$, for each $\theta \in \Theta$. After observing the state of nature $\theta \in \Theta$, the sender draws an action profile $a \in A$ according to $\phi_{\theta,a}$ and recommends action a_p to each player $p \in N$. A signaling scheme is *persuasive* if following recommendations is an equilibrium of the underlying *Bayesian game* [Bergemann and Morris 2016a,b]. We focus

[1] We focus on affine costs since: *(i)* the assumption is reasonable in many applications [Vasserman, Feldman, and Hassidim 2015], and *(ii)* the problem is trivially NP-hard when generic costs are allowed (see Section).

on the notion of *ex ante persuasiveness* as defined by Xu [2020] and Celli, Coniglio, and Gatti [2020].

Definition 1. *A signaling scheme* $\phi : \Theta \to \Delta_A$ *is* ex ante persuasive *if, for each* $p \in N$ *and* $a_p \in A_p$*, it holds:*

$$\sum_{\theta \in \Theta} \mu_\theta \sum_{a'=(a'_p, a_{-p}) \in A} \phi_{\theta,a'} \Big(c_{p,\theta}(a_p, a_{-p}) - c_{p,\theta}(a') \Big) \geq 0.$$

Then, a *coarse correlated equilibrium* (CCE) [Moulin and Vial 1978] may be seen as an *ex ante* persuasive signaling scheme in non-Bayesian NCGs in which there are no states of nature, *i.e.*, when $|\Theta| = 1$. Finally, a sender's *optimal ex ante* persuasive signaling scheme ϕ^* is such that it minimizes the *expected social cost* of the solution, *i.e.*:

$$\phi^* \in \operatorname*{argmin}_{\phi} \sum_{\theta \in \Theta} \mu_\theta \sum_{a \in A} \phi_{\theta,a} \sum_{p \in N} c_{p,\theta}(a).$$

The following example illustrates the interaction flow between the sender and the players (receivers).

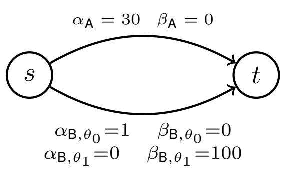

	Signals			
	(B,B,B)	(A,A,B)	(A,B,A)	(B,A,A)
θ_0	1	0	0	0
θ_1	0	1/3	1/3	1/3

Figure 1: *Left*: BNCG for Example 1. *Right*: An *ex ante* persuasive signaling scheme for the case with $n = 3$. The table displays only those $a \in A$ such that $\phi_{\theta,a} > 0$ for some state of nature $\theta \in \Theta = \{\theta_0, \theta_1\}$.

Example 1. *Figure 1 (Left) describes a simple BNCG modeling the road network between the JFK International Airport (node s), and Manhattan (node t). It is late at night and three lone researchers have to reach the AAAI venue. They are following navigation instructions from the same application, whose provider (the sender) has access to the current state of the roads (called* A *and* B*, respectively). Roads costs (i.e., travel times) are depicted in Figure 1 (Left). In normal conditions (state* θ_0*), road* B *is extremely fast (*$\alpha_{\mathsf{B}} = 1$ *and* $\beta_{\mathsf{B}} = 0$*). However, it requires frequent road works for maintenance (state* θ_1*), which increase the travel time. Moreover, it holds* $\mu_{\theta_0} = \mu_{\theta_1} = 1/2$*. The interaction between the sender and the three players goes as follows: (i) the sender commits to a signaling scheme* ϕ*; (ii) the players observe* ϕ *and decide whether to adhere to the navigation system or not; (iii) the sender observes the realized state of nature and exploits this knowledge to compute recommendations. Figure 1 (Right) describes an* ex ante *persuasive signaling scheme. In this case, when the state of nature is* θ_1*, one of the players is randomly selected to take road* B*, even if it is undergoing maintenance. In expectation, following sender's recommendations is strictly better than congesting road* A*.*

A simple variation of Example 1 is enough to show that the introduction of signaling allows the sender to reach solutions with arbitrarily better expected social cost than what can be achieved via the optimal Bayes-Nash equilibrium in absence of signaling. Specifically, consider the BNCG in Figure 1 (Left) with the following modifications: $n = 1$, β coefficients always equal to zero, $\alpha_{\mathsf{A},\theta_0} = \infty$, $\alpha_{\mathsf{A},\theta_1} = 0$, $\alpha_{\mathsf{B},\theta_0} = 0$, and $\alpha_{\mathsf{B},\theta_1} = \infty$. Without signaling, the optimal choice yields an expected social cost of ∞. However, a perfectly informative signal (*i.e.*, one revealing the realized state of nature) allows the player to avoid any cost.

The Power of Symmetry

We design a polynomial-time algorithm to compute an optimal *ex ante* persuasive signaling scheme in symmetric BNCGs with affine cost functions. Our algorithm exploits the ellipsoid method. We first formulate the problem as an LP (Problem (1)) with polynomially many constraints and exponentially many variables. Then, we show how to find an optimal solution to the LP in polynomial time by applying the ellipsoid algorithm to its dual (Problem (2)), which features polynomially many variables and exponentially many constraints. This calls for a polynomial-time separation oracle for Problem (2), which is not readily available since the problem has an exponential number of constraints. We prove that, in our setting, a polynomial-time separation oracle can be implemented by solving a suitably defined min-cost flow problem. The proof of this result crucially relies on the symmetric nature of the problem and the assumption that the costs are affine functions of the edge congestion.

The following lemma shows how to formulate the problem as an LP. [2] For the ease of presentation, we use $I_{\{e \notin a_p\}}$ to denote the indicator function for the event $e \notin a_p$, *i.e.*, it holds $I_{\{e \notin a_p\}} = 1$ if $e \notin a_p$, while $I_{\{e \notin a_p\}} = 0$ otherwise.

Lemma 1. *Given a symmetric BNCG, an optimal* ex ante *persuasive signaling scheme* ϕ *can be found with the LP:*

$$\min_{\phi \geq 0, x} \sum_{\theta \in \Theta} \mu_\theta \sum_{a \in A} \phi_{\theta,a} \sum_{p \in N} c_{p,\theta}(a) \quad \text{s.t.} \tag{1a}$$

$$\sum_{\theta \in \Theta} \mu_\theta \sum_{a \in A} c_{p,\theta}(a) \phi_{\theta,a} \leq x_{p,s} \qquad \forall p \in N \tag{1b}$$

$$x_{p,v} \leq \sum_{\theta \in \Theta} \mu_\theta \sum_{a \in A} c_{e,\theta} \left(f_e^a + I_{\{e \notin a_p\}} \right) \phi_{\theta,a} + x_{p,v'} \qquad \forall p \in N, \forall e = (v, v') \in E \tag{1c}$$

$$x_{p,t} = 0 \qquad \forall p \in N \tag{1d}$$

$$\sum_{a \in A} \phi_{\theta,a} = 1 \qquad \forall \theta \in \Theta \tag{1e}$$

Proof. Clearly, Objective (1a) is equivalent to minimizing the social cost, while Constraints (1e) imply that ϕ is well formed. Constraints (1b) enforce *ex ante* persuasiveness for every player $p \in N$: the expression on the left-hand side represents player p's expected cost, while $x_{p,s}$ is the cost of her best deviation (*i.e.*, a cost-minimizing path given μ and ϕ). This is ensured by Constraints (1c) and (1d). In particular, for every player $p \in N$ and node $v \in V \setminus \{t\}$, the former guarantee that $x_{p,v}$ is the minimum cost of a path from v to

[2] LPs analogous to Problem (1) and Problem (2) can also be derived for the asymmetric setting. However, the separation problem for the dual is solvable in poly-time only in the symmetric case.

t. This is shown by noticing that (given that $x_{p,t} = 0$) such cost can be inductively defined as follows:

$$\min_{\substack{v' \in V: \\ e=(v,v') \in E}} \left\{ \sum_{\theta \in \Theta} \mu_\theta \sum_{a \in A} c_{e,\theta} \left(f_e^a + I_{\{e \notin a_p\}} \right) \phi_{\theta,a} + x_{p,v'} \right\},$$

where $f_e^a + I_{\{e \notin a_p\}}$ accounts for the fact that the congestion of edge e must be incremented by one if player p does not select a path containing e in the action profile a. □

Lemma 2. *The dual of Problem* (1) *reads as follows:*

$$\max_y \sum_{\theta \in \Theta} y_\theta \quad \text{s.t.} \tag{2a}$$

$$\mu_\theta \left(\sum_{p \in N} c_{p,\theta}(a) y_p - \sum_{p \in N} \sum_{e \in E} c_{e,\theta} \left(f_e^a + I_{\{e \notin a_p\}} \right) y_{p,e} \right) + y_\theta \le \mu_\theta \sum_{p \in N} c_{p,\theta}(a) \qquad \forall \theta \in \Theta, \forall a \in A \tag{2b}$$

$$\sum_{v' \in V: e=(v,v') \in E} y_{p,e} - \sum_{v' \in V: e=(v',v) \in E} y_{p,e} = 0 \qquad \forall p \in N, \forall v \in V \setminus \{s,t\} \tag{2c}$$

$$\sum_{v \in V: e=(s,v) \in E} y_{p,e} - y_p = 0 \qquad \forall p \in N \tag{2d}$$

$$y_{p,t} - \sum_{v \in V: e=(v,t) \in E} y_{p,e} = 0 \qquad \forall p \in N \tag{2e}$$

$$y_p \le 0 \qquad \forall p \in N \tag{2f}$$

$$y_{p,e} \le 0 \qquad \forall p \in N, \forall e \in E. \tag{2g}$$

Proof. It directly follows from LP duality, by letting y_p (for $p \in N$), $y_{p,e}$ (for $p \in N$ and $e \in E$), $y_{p,t}$ (for $p \in N$), and y_θ (for $\theta \in \Theta$) be the dual variables associated to, respectively, Constraints (1b), (1c), (1d), and (1e). □

Since $|A|$ is exponential in the size of the game, Problem (1) features exponentially many variables, while its number of constraints is polynomial. Conversely, Problem (2) has polynomially many variables and exponentially many constraints, which enables the use of the ellipsoid algorithm to find an optimal solution to Problem (2) in polynomial time.This requires a polynomial-time separation oracle for Problem (2), *i.e.*, a procedure that, given a vector y of dual variables, it either establishes that y is feasible for Problem (2) or, if not, it outputs a hyperplane separating y from the feasible region. In the following, we focus on a particular type of separation oracles: those generating violated constraints of Problem (2).

Given that Problem (2) has an exponential number of constraints, a polynomial-time separation oracle is not readily available. It turns out that, in our setting, we can design one by leveraging the symmetry of the players and the fact that the cost functions are affine, as described in the following.

First, we prove that Problem (2) always admits an optimal player-symmetric solution, *i.e.*, a vector y such that, for each pair of players $p, q \in N$, it holds that $y_p = y_q$, $y_{p,e} = y_{q,e}$ for all $e \in E$, and $y_{p,t} = y_{q,t}$. This result allows us to restrict the attention to player-symmetric vectors y.

Lemma 3. *Problem* (2) *always admits an optimal player-symmetric solution.*

Proof. Given any optimal solution y to Problem (2), we can always recover, in polynomial time, a player-symmetric optimal solution $\tilde{y}$. Specifically, for every $p \in N$, let $\tilde{y}_p = \frac{\sum_{p \in N} y_p}{n}$, $\tilde{y}_{p,e} = \frac{\sum_{p \in N} y_{p,e}}{n}$ for all $e \in E$, and $\tilde{y}_{p,t} = \frac{\sum_{p \in N} y_{p,t}}{n}$, while $\tilde{y}_\theta = y_\theta$ for every $\theta \in \Theta$. Let us remark that $\tilde{y}$ is player-symmetric since: *(i)* for every $e \in E$, it holds $\tilde{y}_{p,e} = \tilde{y}_{q,e}$ for each pair of players $p, q \in N$; and *(ii)* $\tilde{y}_p = \tilde{y}_q$ and $\tilde{y}_{p,t} = \tilde{y}_{q,t}$ for each $p, q \in N$. First, notice that y and $\tilde{y}$ provide the same objective value, as $\tilde{y}_\theta = y_\theta$ for all $\theta \in \Theta$. Thus, we only need to prove that $\tilde{y}$ satisfies all the constraints of Problem (2). For $a \in A$ and $i \in [n]$, let us denote with $\pi_i(a)$ an action profile $a' \in A$ such that $a'_p = a_{((p+i) \mod n)}$, *i.e.*, a permutation of a in which each player $p \in N$ takes on the role of player $(p+i) \mod n$. Moreover, let $\pi(a) := \bigcup_{i \in [n]} \pi_i(a)$. Constraints (2b) are satisfied by $\tilde{y}$, since, for every $\theta \in \Theta$ and $a \in A$, it holds:

$$\mu_\theta \left(\sum_{p \in N} c_{p,\theta}(a) \tilde{y}_p - \sum_{p \in N} \sum_{e \in E} c_{e,\theta} \left(f_e^a + I_{\{e \notin a_p\}} \right) \tilde{y}_{p,e} \right) + \tilde{y}_\theta$$

$$= \frac{1}{n} \sum_{a' \in \pi(a)} \mu_\theta \left(\sum_{p \in N} c_{p,\theta}(a') y_p - \sum_{p \in N} \sum_{e \in E} c_{e,\theta} \left(f_e^{a'} + I_{\{e \notin a'_p\}} \right) y_{p,e} \right) + y_\theta$$

$$\le \frac{1}{n} \sum_{a' \in \pi(a)} \mu_\theta \sum_{p \in N} c_{p,\theta}(a') = \mu_\theta \sum_{p \in N} c_{p,\theta}(a).$$

Similar arguments show that $\tilde{y}$ satisfies all the other constraints, concluding the proof. □

Notice that any polynomial-time separation oracle for Problem (2) can explicitly check whether each member of the polynomially many Constraints (2c), (2d), and (2e) is satisfied for the given y. Thus, we focus on the separation problem restricted to the exponentially many Constraints (2b), which, using Lemma 3, can be formulated as stated in the following lemma.

Lemma 4. *Given a player-symmetric y, solving the separation problem for Constraints* (2b) *amounts to finding $\theta \in \Theta$ and $a \in A$ that are optimal for the following problem:*

$$\min_{\theta \in \Theta, a \in A} \mu_\theta \left((1 - \bar{y}) \sum_{p \in N} c_{p,\theta}(a) - \sum_{p \in N} \sum_{e \in E} c_{e,\theta} \left(f_e^a + I_{\{e \notin a_p\}} \right) \bar{y}_e \right) - y_\theta, \tag{3}$$

where we let $\bar{y} = y_1$ and $\bar{y}_e = y_{1,e}$ for all $e \in E$.

Next, we show how Problem (3) can be equivalently formulated avoiding the minimization over the exponentially-sized set A. Intuitively, we rely on the fact that, for a fixed

$\theta \in \Theta$, we can exploit the symmetry of the players to equivalently represent action profiles $a \in A$ as integer vectors q of edge congestions $q_e \in [n]$, for all $e \in E$.

Lemma 5. *Problem* (3) *can be formulated as* $\min_{\theta\in\Theta} \chi(\theta)$*, where* $\chi(\theta)$ *is the optimal value of the following problem:*

$$\min_{q\in\mathbb{Z}_+^{|E|}} \ (1-\bar{y}) \sum_{e\in E} \alpha_{e,\theta} q_e^2 + \beta_{e,\theta} q_e -$$

$$- \sum_{e\in E} \bar{y}_e \Big(n\alpha_{e,\theta} q_e + (n-q_e)\alpha_{e,\theta} + n\beta_{e,\theta} \Big) \quad \text{s.t.} \quad (4a)$$

$$\sum_{v\in V: e=(s,v)\in E} q_e = n \quad (4b)$$

$$\sum_{v\in V: e=(v,t)\in E} q_e = n \quad (4c)$$

$$\sum_{\substack{v'\in V:\\ e=(v',v)\in E}} q_e = \sum_{\substack{v'\in V:\\ e=(v,v')\in E}} q_e \quad \forall v \in V \setminus \{s,t\}. \quad (4d)$$

Proof. First, given a state $\theta \in \Theta$, Problem (3) reduces to computing $\chi(\theta) := \min_{a\in A}(1-\bar{y})\sum_{p\in N} c_{p,\theta}(a) - \sum_{p\in N}\sum_{e\in E} c_{e,\theta}\left(f_e^a + I_{\{e\notin a_p\}}\right)\bar{y}_e$, where the function to be minimized only depends on the number of players selecting each edge $e \in E$ in a, rather than the identity of the players who are choosing e (since they are symmetric). Letting $q_e \in [n]$ be the congestion level of edge $e \in E$ and using $c_{e,\theta} = \alpha_{e,\theta} q_e + \beta_{e,\theta}$ (affine costs), it holds $\sum_{p\in N} c_{p,\theta}(a) = \sum_{e\in E} \alpha_{e,\theta} q_e^2 + \beta_{e,\theta} q_e$, and, for every $e \in E$, $\sum_{p\in N} c_{e,\theta}\left(f_e^a + I_{\{e\notin a_p\}}\right) = n\alpha_{e,\theta} q_e + (n - q_e)\alpha_{e,\theta} + n\beta_{e,\theta}$. This gives Objective (4a). Moreover, Constraints (4b), (4c), and (4d) ensure that q is well defined. □

Let us remark that computing an optimal integer solution to Problem (4) is necessary in order to (possibly) find a violated constraint for a given y; otherwise, we would not be able to easily recover an action profile $a \in A$ from q.

Now, we show that an optimal integer solution to Problem (4) can be found in polynomial time by reducing it to an instance of integer min-cost flow problem. Intuitively, it is sufficient to consider a modified version of the original graph G in which each edge $e \in E$ is replaced with n parallel edges with unit capacity and increasing unit costs. This is possible given that the Objective (4a) is a convex function of q, which is guaranteed by the fact that costs are affine.

Lemma 6. *An optimal integer solution to Problem* (4) *can be found in polynomial time by solving a suitably defined instance of integer min-cost flow problem.*

Proof. First, notice that Objective (4a) is a sum edge costs, in which the cost of each edge $e \in E$ is a convex function of the edge congestion q_e, as the only quadratic term is $(1-\bar{y})a_{e,\theta} q_e^2$, where the multiplying coefficient is always positive, given $\bar{y} \leq 0$ and $\alpha_{e,\theta} \geq 0$. This allows us to formulate Problem (4) as an instance of integer min-cost flow problem. We build a new graph where each $e \in E$ is replaced with n parallel edges, say e_i for $i \in [n]$. For $e \in E$ and $i \in [n]$, let us define $g(e,i) := (1-\bar{y})\left(\alpha_{e,\theta} i^2 + \beta_e i\right) - \bar{y}_e\left(n\alpha_{e,\theta} i + (n-i)\alpha_{e,\theta} + n\beta_{e,\theta}\right)$. Each (new) edge e_i has unit capacity and a per-unit cost equal to $\delta(e_i) := g(e,i) - g(e,i-1)$. Clearly, finding an integer min-cost flow is equivalent to minimizing Objective (4a). Notice that, since the original edge costs are convex, it holds $\delta(e_i) \geq \delta(e_j)$ for all $j < i \in [n]$. Thus, an edge e_i is used (*i.e.*, it carries a unit of flow) only if all the edges e_j, for $j < i \in [n]$, are already used. This allows us to recover an integer vector q from a solution to the min-cost flow problem. Finally, let us recall that we can find an optimal solution to the integer min-cost flow problem in polynomial time by solving its LP relaxation. □

The last lemma allows us to prove our main result:

Theorem 1. *Given a symmetric BNCG, an optimal* ex-ante *persuasive signaling scheme can be computed in poly-time.*

Proof. The algorithm applies the ellipsoid algorithm to Problem (2). At each iteration, we require that the vector of dual variables y given to the separation oracle be player-symmetric, which can be easily obtained by applying the symmetrization technique introduced in the proof of Lemma 3. The separation oracle needs to solve an instance of integer min-cost flow problem for every $\theta \in \Theta$ (see Lemmas 5 and 6). Notice that an integer solution is required in order to be able to identify a violated constraint. Finally, the polynomially many violated constraints generated by the ellipsoid algorithm can be used to compute an optimal ϕ. □

The Curse of Asymmetry

In this section, we provide our hardness result on asymmetric BNCGs. Our proof is split into two intermediate steps: *(i)* we prove a hardness result for a simple class of asymmetric non-Bayesian congestion games in which each player selects only one resource (Lemma 7); and *(ii)* we show that such games can be represented as NCGs with only a polynomial blow-up in the representation size (Lemma 8). Our main result reads as follows:

Theorem 2. *The problem of computing an optimal* ex ante *persuasive signaling scheme in BNCGs with asymmetric players is* NP*-hard, even with affine costs.* [3]

The proof of Theorem 2 is based on a reduction that maps an instance of 3SAT (a well-known NP-hard problem, see [Garey and Johnson 1979]) to a game in the class of *singleton congestion games* (SCGs) [Ieong et al. 2005], where each player can select only one resource at a time. A (non-Bayesian) SCG is described by a tuple $(N, R, \{A_p\}_{p\in N}, \{c_r\}_{r\in R})$, where R is a finite set of resources, each player $p \in N$ selects a single resource from the set $A_p \subseteq R$ of available resources, and resource $r \in R$ has a cost $c_r : \mathbb{N} \to \mathbb{R}_+$. Another way of interpreting SCGs is as games played on parallel-link graphs, where each player can select only a subset of the edges.

[3] Without affine costs, computing an optimal *ex ante* persuasive signaling scheme is trivially NP-hard even in symmetric BNCGs. This directly follows from [Meyers and Schulz 2012], which shows that even finding an optimal action profile (that is also an optimal Nash equilibrium) is NP-hard in symmetric (non-Bayesian) NCGs.

First, let us provide the following definition and notation.

Definition 2 (3SAT). *Given a finite set C of three-literal clauses defined over a finite set V of variables, is there a truth assignment to the variables satisfying all the clauses?*

We denote with $l \in \varphi$ a literal (*i.e.*, a variable or its negation) appearing in a clause $\varphi \in C$. Moreover, we let m and s be, respectively, the number of clauses and variables, *i.e.*, $m := |C|$ and $s := |V|$. W.l.o.g., we assume that $m \geq s$.

Lemma 7 introduces our main reduction, proving that finding a social-cost-minimizing CCE is NP-hard in SCGs with asymmetric players, *i.e.*, whenever the resource sets A_p are different among each other. [4] Notice that the games used in the reduction are *not* Bayesian; this shows that the hardness fundamentally resides in the asymmetry of the players.

Lemma 7. *The problem of computing a social-cost-minimizing CCE in SCGs with asymmetric players is NP-hard, even with affine costs.*

Proof. Our 3SAT reduction shows that the existence of a polynomial-time algorithm for computing a social-cost-minimizing CCE in SCGs would allow us to solve any 3SAT instance in polynomial time. Given (C, V), let $z := m^{40}$, $u := m^{12}$, and $\epsilon := \frac{1}{m^4}$. We build an SCG $\Gamma(C, V)$ admitting a CCE with social cost smaller than or equal to $\gamma := z^2+(4us+s+3m)(z-u)+\frac{3z}{m^9}$ iff (C, V) is satisfiable.

Mapping. $\Gamma(C, V)$ is defined as follows (for every $r \in R$, the cost c_r is an affine function with coefficients α_r and β_r).

- $N = \{p_v \mid v \in V\} \cup \{p_{\varphi,q} \mid \varphi \in C, q \in [3]\} \cup \{p_{v,j}, p_{\bar{v},j} \mid v \in V, j \in [2u]\} \cup \{p_i \mid i \in [z]\}$;
- $R = \{r_t\} \cup \{r_v, r_{\bar{v}}, r_{v,1}, r_{v,2}, r_{\bar{v},1}, r_{\bar{v},2} \mid v \in V\}$;
- $A_{p_v} = \{r_v, r_{\bar{v}}, r_t\} \ \forall v \in V$;
- $A_{p_{\varphi,q}} = \{r_l \mid l \in \varphi\} \ \forall \varphi \in C, \forall q \in [3]$;
- $A_{p_{v,j}} = \{r_v, r_{v,1}, r_{v,2}\} \ \forall v \in V, \forall j \in [2u]$;
- $A_{p_{\bar{v},j}} = \{r_{\bar{v}}, r_{\bar{v},1}, r_{\bar{v},2}\} \ \forall v \in V, \forall j \in [2u]$;
- $A_{p_i} = \{r_t\} \ \forall i \in [z]$;
- $\alpha_{r_v} = \alpha_{r_{\bar{v}}} = \epsilon$ and $\beta_{r_v} = \beta_{r_{\bar{v}}} = z + 1 - \epsilon \ \forall v \in V$;
- $\alpha_{r_{v,1}} = \alpha_{r_{v,2}} = \alpha_{r_{\bar{v},1}} = \alpha_{r_{\bar{v},2}} = 1 \ \forall v \in V$;
- $\beta_{r_{v,1}} = \beta_{r_{v,2}} = \beta_{r_{\bar{v},1}} = \beta_{r_{\bar{v},2}} = z + 1 - u \ \forall v \in V$;
- $\alpha_{r_t} = 1$ and $\beta_{r_t} = 0$.

Figure 2 shows a picture representing how the players' action sets are constructed in games $\Gamma(C, V)$, where, for simplicity, only the part referring to a single variable $v \in V$ and a single clause $\varphi \in C$ is reported.

Overview. Intuitively, in games $\Gamma(C, V)$ the social cost is small if players p_i (for $i \in [z]$) are the only ones selecting resource r_t. Then, each player p_v (for $v \in V$) must choose either r_v or $r_{\bar{v}}$ (rather than r_t), representing the fact that variable v is set to either false or true, respectively. At the same time, players p_v do not deviate to resource r_t only if they are the only players selecting their resources. This implies that all the players $p_{\varphi,q}$ (for $\varphi \in C$ and $q \in [3]$) must play a resource not selected by any player p_v. Hence, each player $p_{\varphi,q}$ plays a resource r_l whose corresponding literal l is true, which results in φ being satisfied. The action profile defined thus far does *not* constitute an equilibrium, as players $p_{\varphi,q}$ have an incentive to deviate to resources r_l with l evaluating to false. Players $p_{v,j}$ and $p_{\bar{v},j}$ are used to avoid such deviations. They are told to play resources r_l with very small probability, so that other players do not deviate to them.

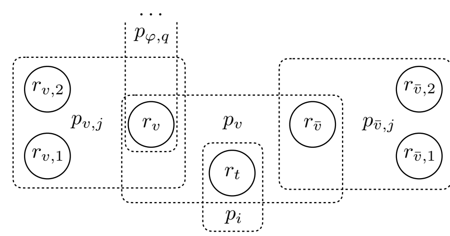

Figure 2: Example of players' action sets in a game instance $\Gamma(C, V)$ used for the reduction in the proof of Lemma 7.

If. Suppose (C, V) is satisfiable, and let $\tau : V \to \{\mathsf{T}, \mathsf{F}\}$ be a truth assignment satisfying all the clauses in C. For the ease of presentation, we let $\tau(l) \in \{\mathsf{T}, \mathsf{F}\}$ be the truth value of literal $l \in \{v, \bar{v} \mid v \in V\}$ under τ. Using τ, we recover a CCE $\phi \in \Delta_A$ with social cost smaller than or equal to γ. This selects the action profiles $\{a^1, a^2, a^3\} \subseteq \times_{p \in N} A_p$ defined in the following with probabilities $\phi_{a^1} = \phi_{a^2} = \frac{1}{2} - \frac{1}{2m^{10}}$ and $\phi_{a^3} = \frac{1}{m^{10}}$. First, we determine actions for players $p_{\varphi,q}$ (the same in a^1, a^2, and a^3). Each player $p_{\varphi,q}$ (for $\varphi \in C$ and $q \in [3]$) plays a resource r_l with $l \in \varphi$ such that $\tau(l) = \mathsf{T}$, so that none of these players has an incentive to deviate to another resource r_l with $\tau(l) = \mathsf{T}$. Moreover, players' actions are such that each r_l with $\tau(l) = \mathsf{T}$ has at least one player using it, which is useful to avoid that other players deviate on the resource. To formally define players $p_{\varphi,q}$' actions, we consider a congestion game Γ_{R} restricted to the players $\{p_{\varphi,q} \mid \varphi \in C, q \in [3]\}$ with action spaces limited to resources $r_l \in A_{p_{\varphi,q}}$ with $\tau(l) = \mathsf{T}$ (since τ satisfies all clauses, each player has at least one action). Clearly, Γ_{R} admits a pure NE [Rosenthal 1973]. We show that, in any pure NE, each resource is selected by at least one player. By contradiction, suppose that there exists a resource r_l such that no player chooses it. Then, there must be at least two players $p_{\varphi,q}$ (with $l \in \varphi$) selecting some resource different from r_l. As a result, there must be one player with an incentive to deviate to the empty resource (as she would pay $z+1$ rather than something $\geq z + 1 + \epsilon$), contradicting the NE assumption. In conclusion, for every $\varphi \in C$ and $q \in [3]$, we let $a^1_{p_{\varphi,q}}$, $a^2_{p_{\varphi,q}}$, and $a^3_{p_{\varphi,q}}$ all be equal to the resource played by the corresponding player in some pure NE of Γ_{R}. Now, we define actions for players $p_{v,j}$ and $p_{\bar{v},j}$. Each player $p_{l,j}$ plays r_l in a^3 (drawn with a small probability of $\frac{1}{m^{10}}$) only if $\tau(l) = \mathsf{F}$, while this never happens in a^1 and a^2. Intuitively, this avoids that other players deviate to a resource r_l with $\tau(l) = \mathsf{F}$. Moreover, players $p_{l,j}$ are split into two

[4] The reduction in Lemma 7 does *not* rely on standard constructions, as most of the reductions for congestion games only work with action profiles, while ours needs randomization. Indeed, in asymmetric SCGs, a social-cost-minimizing action profile can be computed in poly-time by solving an instance of min-cost flow problem. This also prevents the use of other techniques for proving the hardness of CCEs, *e.g.*, those by Barman and Ligett [2015].

groups alternating between resources $r_{l,1}$ and $r_{l,2}$ in action profiles a^1 and a^2. This prevents deviations to either $r_{l,1}$ or $r_{l,2}$ (as there are at least u players using the resource with high probability). Formally, for every $l \in \{v, \bar{v} \mid v \in V\}$:

- for $j \in [u]$, we let $a^1_{p_{l,j}} = r_{l,1}$, $a^2_{p_{l,j}} = r_{l,2}$, and $a^3_{p_{l,j}} = r_l$ if $\tau(l) = \mathsf{F}$, while $a^3_{p_{l,j}} = r_{l,1}$ if $\tau(l) = \mathsf{T}$;
- for $j \in [2u] : j > u$, we let $a^1_{p_{l,j}} = r_{l,2}$, $a^2_{p_{l,j}} = r_{l,1}$, and $a^3_{p_{l,j}} = r_l$ if $\tau(l) = \mathsf{F}$, while $a^3_{p_{l,j}} = r_{l,2}$ if $\tau(l) = \mathsf{T}$.

Finally, we introduce players p_v' actions. In a^1 and a^2 (selected with high probability $1 - \frac{1}{m^{10}}$), each player p_v uses r_v if $\tau(v) = \mathsf{F}$, while $r_{\bar{v}}$ otherwise. Instead, in a^3 (drawn with a small probability of $\frac{1}{m^{10}}$), player p_v selects r_t so as to keep the cost of players $p_{l,j}$ small. Thus, for every $v \in V$, we let $a^3_{p_v} = r_t$ and $a^1_{p_v} = a^2_{p_v} = r_v$ if $\tau(v) = \mathsf{F}$, while $a^1_{p_v} = a^2_{p_v} = r_{\bar{v}}$ if not. Next, we show that players have no incentive to defect from ϕ, *i.e.*, ϕ is a CCE. Given that player $p_{\varphi,q}$'s action (for $\varphi \in C$ and $q \in [3]$) is determined by a pure NE of Γ_{R}, she does not have any incentive to deviate to another resource $r_l \in A_{\varphi,q}$ with $\tau(l) = \mathsf{T}$ (as these resources are not selected by players not participating to Γ_{R} and the players in Γ_{R} are at an NE). Moreover, in ϕ, player $p_{\varphi,q}$'s expected cost is at most $z+1+3\epsilon m$, while she would pay at least $(z+1+\epsilon)(1-\frac{1}{m^{10}})+(z+1+2u\epsilon)\frac{1}{m^{10}} \geq z+1+2\epsilon m^2$ by selecting a resource $r_l \in A_{\varphi,q}$ with $\tau(l) = \mathsf{F}$. Each player p_v (for $v \in V$) does not defect from ϕ, since her expected cost is $(z+1)(1-\frac{1}{m^{10}}) + (z+s)\frac{1}{m^{10}}$, while she would pay:

- the same amount by switching to resource r_t;
- at least $z+1+\epsilon$ by playing r_l with $l \in \{v, \bar{v}\}$ and $\tau(l) = \mathsf{T}$ (as there is at least one player $p_{\varphi,q}$ on r_l);
- at least $(z+1)(1-\frac{1}{m^{10}})+(z+1+2u\epsilon)\frac{1}{m^{10}} = z+1+2\frac{1}{m^2}$ by selecting r_l with $l \in \{v, \bar{v}\}$ and $\tau(l) = \mathsf{F}$.

Each player $p_{l,j}$ (for $l \in \{v, \bar{v} \mid v \in V\}$ and $j \in [2u]$) with $\tau(l) = \mathsf{F}$ does not deviate, since her cost is $(z+1)(1 - \frac{1}{m^{10}}) + (z+1-\epsilon+2u\epsilon)\frac{1}{m^{10}}$, while she would pay:

- at least $(z+1)(\frac{1}{2}-\frac{1}{2m^{10}})+(z+2)(\frac{1}{2}-\frac{1}{2m^{10}})$ by switching to either $r_{l,1}$ or $r_{l,2}$;
- at least $(z+1+\epsilon)(1-\frac{1}{m^{10}}) + (z+1-\epsilon+2u\epsilon)\frac{1}{m^{10}}$ by selecting resource $r_l \in A_{p_{l,j}}$.

Moreover, each player $p_{l,j}$ with $\tau(l) = \mathsf{T}$ does not deviate either, as her cost is $(z+1)$, while she would pay:

- at least $z+1+\epsilon$ by playing resource r_l;
- at least $(z+1)(\frac{1}{2}+\frac{1}{2m^{10}})+(z+2)(\frac{1}{2}-\frac{1}{2m^{10}})$ by switching to either $r_{l,1}$ or $r_{l,2}$.

Finally, players p_i must select resource r_t; thus, they experience a cost of $z(1-\frac{1}{m^{10}}) + (z+s)\frac{1}{m^{10}}$. Moreover, since the maximum cost of a resource different from r_t is $z+1+u$, players p_v incur a cost at most of $(z+1+u)(1-\frac{1}{m^{10}}) + (z+s)\frac{1}{m^{10}}$, while all the other players pay at most $z+1+u$. Then, the CCE ϕ provides a social cost smaller than or equal to $z\left[z(1-\frac{1}{m^{10}}) + (z+s)\frac{1}{m^{10}}\right] + s\left[(z+1+u)(1-\frac{1}{m^{10}}) + (z+s)\frac{1}{m^{10}}\right] + (4us+3m)(z+1+u) \leq z^2+\frac{zs}{m^{10}}+(z+1+u)(s+4us+3m)+s(z+s)\frac{1}{m^{10}} = z^2 + (s+4us+3m)(z-u) + (2u+1)(s+4us+3m) + s(2z+s)\frac{1}{m^{10}} \leq \gamma$, where the last inequality follows from $(2u+1)(s+4us+3m) + s(2z+s)\frac{1}{m^{10}} \leq (2m^{12}+1)(m+4m^{13}+3m) + m(2z+m)\frac{1}{m^{10}} \leq \frac{3z}{m^9}$ for m large enough.

Only if. Suppose there exists a CCE $\phi \in \Delta_A$ with social cost smaller than or equal to γ. First, we prove that, with probability at most $\frac{1}{m^8}$, at least one player p_v plays r_t. By contradiction, assume that this is not the case. Then, the social cost would be at least $(z^2+(4us+s+3m)(z-u))(1-\frac{1}{m^8}) + ((z+1)^2 + (4us+s+3m-1)(z-u))\frac{1}{m^8} \geq z^2 + (4us+s+3m)(z-u) + (2z-z)\frac{1}{m^8} > \gamma$. This implies that each player p_v is playing either r_v or $r_{\bar{v}}$ with probability at least $1-\frac{1}{m^8}$. Then, we prove that p_v is the only player on that resource with probability at least $1-\frac{1}{m^8}-\frac{1}{m^2}$. Otherwise, by contradiction, her cost would be at least $z+1+\frac{\epsilon}{m^2} = z+1+\frac{1}{m^6}$, while by playing r_t she would pay at most $(z+1)(1-\frac{1}{m^8})+(z+s)\frac{1}{m^8} \leq z+1+\frac{1}{m^7}$. By a union bound, there exists an action profile $a \in \times_{p \in N} A_p$ played with probability at least $1-s(\frac{1}{m^8}+\frac{1}{m^2}) > 0$ in which all the players p_v are alone on their resources (either r_v or $r_{\bar{v}}$). Let $\tau : V \to \{\mathsf{T}, \mathsf{F}\}$ be a truth assignment such that $\tau(v) = \mathsf{T}$ if $a_{p_v} = r_{\bar{v}}$ and $\tau(v) = \mathsf{F}$ if $a_{p_v} = r_v$. Then, τ satisfies all the clauses, since all the players $p_{\varphi,q}$ play r_l with $\tau(l) = \mathsf{T}$ and, thus, each clause has at least a true literal. □

The following lemma concludes the proof of Theorem 2.

Lemma 8. *Any SCG can be represented as an NCG of size polynomial in the size of the original SCG.*

Proof. Given an SCG $(N, R, \{A_p\}_{p \in N}, \{c_r\}_{r \in R})$ we build an NCG $(N, G, \{c_e\}_{e \in E}, \{(s_p, t_p)\}_{p \in N})$ as follows. The graph $G = (V, E)$ has two nodes $v_{r,1}, v_{r,2} \in V$ for each resource $r \in R$, and, additionally, for every player $p \in N$, there is a source node $s_p \in V$ and a destination one $t_p \in V$. Moreover, there is an edge $(v_{r,1}, v_{r,2}) \in E$ for every $r \in R$ and, for every $p \in N$ and $r \in A_p$, there two edges $(s_p, v_{r,1}) \in E$ and $(v_{r,2}, t_p) \in E$. Finally, for the edges $e = (v_{r,1}, v_{r,2})$, we let $c_e = c_r$, while $c_e = 0$ for all the other edges. Clearly, the size of the NCG is polynomially bounded by that of the original SCG, proving the result. □

Discussion and Future Works

The paper studies information-structure design problems in atomic BNCGs, where an informed sender can observe the actual state of the network and commit to a signaling scheme. We focus on the problem of computing optimal *ex ante* persuasive signaling schemes in such setting. We show that, with affine costs, *symmetry* is the property marking the transition from polynomial-time tractability to NP-hardness.

In the future, we are interested in studying the problem of approximating optimal *ex ante* persuasive signaling schemes, and in the design of practical algorithms for real-world network signaling problems. Moreover, to make the framework even more applicable, it would be interesting to explore how the sender can handle uncertainty about receivers' payoffs [Castiglioni et al. 2020].

Acknowledgments

This work has been partially supported by the Italian MIUR PRIN 2017 Project ALGADIMAR "Algorithms, Games, and Digital Market".

References

Acemoglu, D.; Makhdoumi, A.; Malekian, A.; and Ozdaglar, A. 2018. Informational Braess' paradox: The effect of information on traffic congestion. *Operations Research* 66(4): 893–917.

Arieli, I.; and Babichenko, Y. 2016. Private Bayesian persuasion. *Available at SSRN 2721307* .

Arnott, R.; De Palma, A.; and Lindsey, R. 1991. Does providing information to drivers reduce traffic congestion? *Transportation Research Part A: General* 25(5): 309–318.

Barman, S.; and Ligett, K. 2015. Finding Any Nontrivial Coarse Correlated Equilibrium is Hard. *SIGecom Exch.* 14(1): 76–79.

Bergemann, D.; and Morris, S. 2016a. Bayes correlated equilibrium and the comparison of information structures in games. *Theoretical Economics* 11(2): 487–522.

Bergemann, D.; and Morris, S. 2016b. Information design, Bayesian persuasion, and Bayes correlated equilibrium. *American Economic Review* 106(5): 586–91.

Bhaskar, U.; Cheng, Y.; Ko, Y. K.; and Swamy, C. 2016. Hardness results for signaling in bayesian zero-sum and network routing games. In *Proceedings of the 2016 ACM Conference on Economics and Computation*, 479–496.

Castiglioni, M.; Celli, A.; Marchesi, A.; and Gatti, N. 2020. Online Bayesian Persuasion. *Advances in Neural Information Processing Systems* 33.

Celli, A.; Coniglio, S.; and Gatti, N. 2020. Private Bayesian Persuasion with Sequential Games. In *The Thirty-Fourth AAAI Conference on Artificial Intelligence*, 1886–1893.

Christodoulou, G.; and Koutsoupias, E. 2005. On the Price of Anarchy and Stability of Correlated Equilibria of Linear Congestion Games. In *Proceedings of the 13th Annual European Conference on Algorithms*, 59–70.

Das, S.; Kamenica, E.; and Mirka, R. 2017. Reducing congestion through information design. In *2017 55th Annual Allerton Conference on Communication, Control, and Computing (Allerton)*, 1279–1284. IEEE.

Dughmi, S. 2017. Algorithmic information structure design: a survey. *SIGecom Exch.* 15(2): 2–24.

Fabrikant, A.; Papadimitriou, C.; and Talwar, K. 2004. The complexity of pure Nash equilibria. In *Proceedings of the thirty-sixth annual ACM symposium on Theory of computing*, 604–612.

Garey, M.; and Johnson, D. 1979. *Computers and Intractability: A Guide to the Theory of NP-completeness*. WH Freeman and Company.

Ieong, S.; McGrew, R.; Nudelman, E.; Shoham, Y.; and Sun, Q. 2005. Fast and compact: A simple class of congestion games. In *AAAI*, volume 5, 489–494.

Jiang, A. X.; and Leyton-Brown, K. 2011. A general framework for computing optimal correlated equilibria in compact games. In *International Workshop on Internet and Network Economics*, 218–229. Springer.

Kamenica, E.; and Gentzkow, M. 2011. Bayesian persuasion. *American Economic Review* 101(6): 2590–2615.

Liu, Y.; and Whinston, A. B. 2019. Efficient Real-Time Routing for Autonomous Vehicles Through Bayes Correlated Equilibrium: An Information Design Framework. *Information Economics and Policy* 47: 14–26.

Massicot, O.; and Langbort, C. 2019. Public Signals and Persuasion for Road Network Congestion Games under Vagaries. *IFAC-PapersOnLine* 51(34): 124–130.

Meyers, C. A.; and Schulz, A. S. 2012. The complexity of welfare maximization in congestion games. *Networks* 59(2): 252–260.

Moulin, H.; and Vial, J.-P. 1978. Strategically zero-sum games: the class of games whose completely mixed equilibria cannot be improved upon. *International Journal of Game Theory* 7(3-4): 201–221.

Papadimitriou, C. H.; and Roughgarden, T. 2008. Computing Correlated Equilibria in Multi-Player Games. *Journal of the ACM* 55(3).

Rayo, L.; and Segal, I. 2010. Optimal information disclosure. *Journal of political Economy* 118(5): 949–987.

Rosenthal, R. W. 1973. A class of games possessing pure-strategy Nash equilibria. *International Journal of Game Theory* 2(1): 65–67.

Roughgarden, T. 2005. *Selfish routing and the price of anarchy*, volume 174. MIT press Cambridge.

Vasserman, S.; Feldman, M.; and Hassidim, A. 2015. Implementing the wisdom of waze. In *Proceedings of the Twenty-Fourth International Joint Conference on Artificial Intelligence*, 660–666.

Wu, M.; Amin, S.; and Ozdaglar, A. E. 2018. Value of Information Systems in Routing Games. *arXiv preprint arXiv:1808.10590* .

Xu, H. 2020. On the Tractability of Public Persuasion with No Externalities. In *Proceedings of the 2020 ACM-SIAM Symposium on Discrete Algorithms*, 2708–2727.