

Implementing the Wisdom of Waze *

Shoshana Vasserman
 Harvard University
 svasserman@fas.harvard.edu

Michal Feldman
 Tel Aviv University
 michal.feldman@cs.tau.ac.il

Avinatan Hassidim
 Bar Ilan University
 avinatanh@gmail.com

Abstract

We study a setting of non-atomic routing in a network of m parallel links with asymmetry of information. While a central entity (such as a GPS navigation system) — a *mediator* hereafter — knows the cost functions associated with the links, they are unknown to the individual agents controlling the flow. The mediator gives *incentive compatible* recommendations to agents, trying to minimize the total travel time. Can the mediator do better than when agents minimize their travel time selfishly without coercing agents to follow his recommendations? We study the *mediation ratio*: the ratio between the *mediated equilibrium* obtained from an incentive compatible mediation protocol and the social optimum. We find that mediation protocols can reduce the efficiency loss compared to the full revelation alternative, and compared to the non mediated Nash equilibrium. In particular, in the case of two links with affine cost functions, the mediation ratio is at most $8/7$, and remains strictly smaller than the price of anarchy of $4/3$ for any fixed m . Yet, it approaches the price of anarchy as m grows. For general (monotone) cost functions, the mediation ratio is at most m , a significant improvement over the unbounded price of anarchy.

1 Introduction

Popular navigation services such as Waze are used by drivers both to plan out routes and to optimally navigate real time road congestion. Waze collects aggregate traffic information in areas of interest and so can take real time traffic conditions, which are incalculable to individual drivers, into account when computing optimal route recommendations. As

*The work of S. Vasserman and M. Feldman was partially supported by the European Research Council under the European Union’s Seventh Framework Programme (FP7/2007-2013) / ERC grant agreement number 337122 and the People Programme (Marie Curie Actions) of the European Unions Seventh Framework Programme (FP7/2007-2013) under REA grant agreement number 274919. The work of A. Hassidim is supported by ISF grant number 1241/12 and by BSF 2012344. We thank Itai Ashlagi for his great comments and discussion.

such, it is likely that Waze can recommend a route to shorten the travel time of each individual driver. However, it may not be possible for each individual user to avoid traffic without creating congestion on the clearer roads, and it might even be that such a recommendation leads to longer aggregate routes.

Consider the following simplistic example. Suppose that a thousand drivers want to route from city Source to city Destination, which is reachable from Source through two parallel roads. The travel time in each of these roads depends on whether or not an accident has occurred. Specifically, suppose that in each of these roads, in the absence of accidents, each driver’s trip takes $x/1000$ hours, where x is the number of drivers on the road (e.g., if half of the drivers take a road with no accident, then the travel time of each driver is half an hour). However, if an accident occurs, then the road becomes clogged, and each driver’s trip takes one hour, independently of the number of drivers on the road. Suppose that an accident occurs on each road with some probability p , which is known to all drivers, but whether or not an accident has occurred on a given road is unknown.

If Waze did not exist, then each driver would choose a road at random. To a first order approximation, we assume that exactly half of the drivers would take each road. In this case, one can easily verify that the average time spent per driver is $p^2 + 3p(1-p)/2 + (1-p)^2/2$.

Now suppose that Waze knows exactly which road has had an accident, and can route each driver to his individually optimal road. In this case, in every situation except one where no accident occurs on either road, each driver would spend a full hour on the road. The average time spent per driver is $p^2 + 2p(1-p) + (1-p)^2/2$ which is larger than the average time spent in the absence of Waze.

Taking a closer look at this example, we observe that in this scenario, the naive Bayes Nash equilibrium, in which half the players take each road, is actually the socially optimal routing regardless of the status of the roads (that is, for every realization of the prior). The equilibrium generated by Waze corresponds to the full information Nash equilibrium, which is, in general, not optimal.

Surely, Waze could do better. Indeed, by proposing a random route to each driver, Waze could implement the optimal policy. In more complex examples, Waze could implement more sophisticated policies (e.g., never consider a specific road and randomize between the remaining ones). This sce-

nario motivates the main question of this paper - in general, can Waze always implement the socially optimal solution? If not, how close to the optimum can it get?

To study this question we consider a non-atomic selfish routing game, with a source node, a target node, and m parallel links. There is a common prior on the cost functions of each link, known to all players (players do not know the realization). In addition, there is a mediator (Waze), who first commits to a mediation policy, then learns the realization of the prior, and only then sends recommendations to all of the players according to the policy he had committed to. The players (who know the mediator's strategy but not the realization of the prior) decide which route to take, and incur the travel cost realized on their route. Because the mediator cannot inextricably compel agents to follow his recommendations, we restrict attention to *incentive compatible* policies employed by the mediator - that is, policies that induce an equilibrium in which all players are best off by following the mediator's recommendation. We refer to such equilibria as *mediated equilibria*.

Our model is applicable to a number of additional situations considered by the congestion game literature. For example, the mediator might be interpreted as a manager who recommends task allocations to teams or machines. If the difficulty of tasks is realized according to a random process observed by the manager, this paper asks how much faster the manager can get a set of tasks done by giving incentive compatible assignments to each worker, than by revealing his observations and leaving workers to allocate themselves.

In this model, we establish several results. First, we show a revelation principle, which implies that restricting attention to incentive compatible policies is without loss of generality.

Second, we quantify the efficiency loss in this setting using the *mediation ratio*, which is the worst-case ratio (over all possible instances and priors) between the best mediated equilibrium and the socially optimal one. Clearly, the mediator can always commit to implementing the full-information Nash equilibrium of the game (i.e., by the definition of a Nash equilibrium, following the mediator's recommendation is an equilibrium). Therefore, the mediation ratio is always bounded from above by the price of anarchy (PoA). We show that if all cost functions are affine (i.e., of the form $c(x) = ax + b$, where x is the amount of traffic on the road), then for the case of two parallel links, the mediation ratio is $8/7$, and for the case of m links, it is $\frac{4m}{3m+1}$. Recalling that in this case the PoA is $4/3$, we conclude that the mediation ratio is strictly smaller than the PoA for any fixed number of links, but approaches the PoA as the number of links grows. In particular, our results suggest that unlike the PoA, the mediation ratio does depend on the network size, yet the worst case manifests in simple (i.e., parallel-link) networks. For general (non-decreasing) monotone cost functions, we show that the mediation ratio for m parallel links is exactly m . This is in contrast to the PoA, which can be arbitrarily high for non-affine cost-functions, even in the case of two parallel links.

Related work In a full-information routing game, the Price of Anarchy (PoA) is defined as the ratio of the worst Nash

equilibrium and the optimal flow, and the Price of Stability (PoS) is defined as the ratio of the best NE and the optimal flow. These measures are well understood in the full-information case. In particular, for non-atomic routing games, all Nash equilibria have the same cost (thus, the PoA and PoS are equal). If all cost functions are affine, then the PoA is at most $4/3$ [Roughgarden and Tardos, 2000; 2004], and this PoA manifests in the simplest network — one that is composed of two parallel links. In that sense, the PoA in full-information games is independent of the network topology [Roughgarden, 2003].

The notion of mediators is central to the study of multi-agent systems. The simplest form of a mediator introduced in the game theoretic literature is the notion of a correlated equilibrium [Aumann, 1974]. Indeed, a correlated equilibrium is essentially a mediation device that gives recommendation to agents but cannot enforce behavior. Since then, various notions of mediators have been introduced in the literature. These mediators are all reliable entities who behave in a pre-specified way and cannot enforce behavior:

Monderer and Tennenholtz [Monderer and Tennenholtz, 2009] propose the use of mediators to obtain stability against deviations by coalitions. Their mediator plays on behalf of the agents upon their approval and based on their messages. In other mediation models the mediator is assumed to be much stronger. For example, Rozenfeld and Tennenholtz [Rozenfeld and Tennenholtz, 2007] consider a mediator who can observe the agents' actions and impose punishments on those that do not follow his recommendation in routing games. Davis et. al [Joshua R. Davis and Wexler, 2011] consider a mediator who can persuade agents to let him play on their behalf by threatening coalitional punishments in load balancing games. Monderer and Tennenholtz [Monderer and Tennenholtz, 2003] consider the use of mediators to implement desirable outcomes by means of monetary transfers. Mediators in games with incomplete information have been considered by Ashlagi, Monderer and Tennenholtz [Ashlagi et al., 2009] in the context of position auctions. In contrast to our model, where the mediator has more information than the agents, they consider a Bayesian game where the agents have private types that are unknown to the mediator.

Kremer, Mansour and Perry [Kremer et al., 2013] study a sequential setting, where a-priori both the mediator and the agents have only priors over rewards, but the mediator gains more information as new agents explore different alternatives. The authors characterize the optimal incentive-compatible recommendation protocol in this setting.

2 Model and Preliminaries

A nonatomic unit of flow must be passed from a source node to a sink node through a parallel-links network on a set of links $L = \{1 \dots, m\}$. Each infinitesimal unit of flow is controlled by an independent myopic rational agent who chooses one link to be routed through. The aggregate decisions of all agents yield a feasible flow, $\alpha = (\alpha_1 \dots \alpha_m)$ in which a measure of α_k is routed through link k , with $\alpha_k \geq 0$ for each k , and $\sum_{k=1}^m \alpha_k = 1$. Each link k is associated with a non-decreasing cost function $c_k : [0, 1] \rightarrow R^+$ which maps every

fraction routed on link k to the cost incurred by each agent flowing on that link. The *social cost* of a given tuple of cost functions and flow is given by $cost(c, \alpha) = \sum_{i=1}^m \alpha_k c_k(\alpha_k)$, where $c = (c_1, \dots, c_m)$ is a vector of cost functions, and $\alpha \in \Delta_m$ is a feasible flow. We consider a Bayesian framework in which agents have incomplete information regarding the cost functions on the links in the network. We call all games that follow the description above, *Bayesian Congestion Games* (BCGs), and focus on two models in particular: a *permutation* model and an *IID* model.

In a *permutation* model, a set of cost functions $\{c^1, \dots, c^m\}$ are first chosen without repetitions from a multiset of cost functions $\{c_i : i \in I\}$ for some set of indices I , and then a permutation function $\pi : [m] \rightarrow [m]$ assigns costs to the links, such that for every link $i \in [m]$, $c_i = c^{\pi(i)}$. A special case of the permutation model is one where there is a known set of m cost functions $\{c^1, \dots, c^m\}$ and cost functions on the links are assigned based on the permutation π . In the permutation model, the agents do not know π , but they know the set of cost functions and the prior p over its distribution.

The *IID* model is characterized by a probability distribution over cost functions F such that for every link k , the cost function c_k is drawn identically and independently from F . We will denote the realization of cost functions by $c = (c_1, \dots, c_m)$ and the common prior regarding c by p .

We further endow the game with an informed, benevolent mediator who observes the realization of cost functions before any flow is routed, and can communicate with agents. The mediator's action is a policy function, σ , that maps a realization of all cost functions into a set of arbitrary signals, such that each agent i receives a private signal σ_i . The mediator's policy is known to all agents, but the realization is known only to the mediator and no agent has access to the signals received by other agents.

The strategy of each agent i is a (possibly random) strategy τ_i that maps a tuple of the common prior over cost functions and the signal received by the mediator into one of the m links. As is standard in the analysis of one-shot Bayesian games, we focus on strategies that constitute a Bayes Nash Equilibrium for all agents.

Definition 2.1. [BNE] Given a mediator's policy σ , a strategy profile τ is a *Bayesian Nash Equilibrium* if for each agent i , playing $\tau_i(p, \sigma_i)$ minimizes agent i 's cost when i receives the signal σ_i and every agent $j \neq i$, plays according to τ_j . Note that every BNE τ essentially induces a 'pure' flow $\alpha = (\alpha_1, \dots, \alpha_m)$.¹

The mediator may be able to use his knowledge of the cost realizations to incur a routing flow that is socially preferred to the flows that would be chosen in his absence, but he cannot compel agents to take his advice. A mediator's policy is said to be *implementable* if following the mediator's advice is *Bayesian incentive compatible*, or incentive compatible (IC) for short.

¹This is because our model considers a continuum of identical infinitesimal players so that the flow α does not distinguish between the composition of each component α_k and it is without loss of generality to assume that agents play deterministic strategies only.

Definition 2.2. [IC] A mediator's policy σ is *incentive compatible* if:

- (1) For every cost realization c , and every agent i , $\sigma_i \in [m]$ (i.e., the signal is a recommendation of which route to follow).
- (2) There exists a BNE τ such that $\tau_i(p, \sigma_i) = \sigma_i$ for all i .

When the mediator invokes an incentive compatible policy σ , we assume that the BNE in which his recommendations are followed is played. When the mediator does not invoke any policy, the players have no information. Since the links are a-priori indistinguishable to the players, the only BNE is the one in which $1/m$ flow is routed on each link. We call this strategy of the players the uninformed BNE, and call the equilibrium an *uninformed BNE*.

In what follows we present a revelation principle for our setting, implying that we can, without loss of generality, restrict attention to IC policies. We omit the full proof, which employs a standard simulation argument, due to space constraints.

Theorem 2.3. [Revelation principle] For every mediator's policy σ and corresponding BNE τ that induces a flow $\alpha = (\alpha_1, \dots, \alpha_m)$, there exists an incentive compatible policy σ' that induces the same flow α .

A flow that is induced by an IC policy is referred to as a *mediated equilibrium*.

Definition 2.4. [Mediated Equilibrium] Given an IC policy σ , a *mediated equilibrium* is a flow $\alpha^\sigma = (\alpha_1^\sigma, \dots, \alpha_m^\sigma)$ such that α_k^σ is the measure of players receiving signal k by the policy σ .

The cost of an IC policy σ is thus given by

$$cost(\sigma, p) = E_{c \sim p} [cost(c, \alpha^\sigma)]$$

where α^σ is the flow induced by σ under prior p .

As we shall soon show, in some cases the unconstrained social optimal flow cannot be implemented as a Mediated Equilibrium. To measure the difference from the optimal solution, we introduce the *mediation ratio* (MR), defined as the ratio of the expected costs of best mediated equilibrium flow and the globally optimal flow.

Definition 2.5. [Mediation Ratio] Given a prior p , the mediation ratio (MR) with respect to p is defined as:

$$MR(p) = \frac{\min_{\sigma: \sigma \text{ is IC}} E_{c \sim p} [cost(c, \alpha^\sigma)]}{E_{c \sim p} \left[\min_{\alpha} cost(c, \alpha) \right]}.$$

For a family \mathcal{P} of priors, the mediation ratio with respect to \mathcal{P} is defined as :

$$MR(\mathcal{P}) = \text{Sup}_{p \in \mathcal{P}} MR(p).$$

The mediator can always calculate a Nash equilibrium of the full information game (as he knows the realizations of the cost functions) and direct the agents accordingly. Incentive compatibility follows from the properties of Nash equilibria. Consequently, the mediation ratio is upper bounded by the price of stability (note that in a non-atomic routing game all NE of the full information game have the same social cost, thus the price of anarchy equals the price of stability). We show several more properties of our models:

Lemma 2.6. *The MR in an IID BCG with m links and arbitrary cost distribution F can be no larger than the case in which F has positive density on only m cost functions: $f(c) > 0$ for $c \in C$ for some set of costs C with cardinality $|C| = m$, and $f(c) = 0$ otherwise.*

The proof of this lemma is omitted due to space constraints and can be found in the full version of this paper. It is based on the fact that if the drivers benefit from following the advice of the mediator in some realization, then the mediator can incentivize them to follow the advice in general, even if there are cases in which following the advice hurts them (as long as these cases are rare enough). For any F , the MR maximizes the average ratio between the mediated equilibrium and the Nash equilibrium across all realization. Thus, if F has support on more than m cost functions, the MR is bounded from above by the MR from a truncation of F that has support on the worst m costs.

Lemma 2.7. *For every m , the MR of a BCG game with m links in the permutation model is weakly larger than the MR of a BCG game with m links in the IID model.*

Lemma 2.7 follows from the additional randomness that is endemic to IID games. Given any IID game, we can compute the MR of the permutation game induced from every set of cost functions, realizable under the IID model. The highest MR among these will be higher than the MR of the IID game as the latter randomizes over socially preferred outcomes and (as a consequence) has weaker incentive constraints.

Based on the last lemma, lower bounds on the MR that are established with respect to the IID model apply to both the IID and the permutation models. We now establish a general IC condition that determines whether a policy σ is implementable in a Permutation BCG.

Lemma 2.8. *A policy σ that induces the equilibrium flow $\alpha_1, \dots, \alpha_m$ in a Permutation BCG with m parallel links whose costs are $\{c_i\}_{i=1}^m$ is implementable if and only if*

$$\sum_{i=1}^m \alpha_i c_i(\alpha_i) \leq \frac{1}{m} \sum_{i=1}^m c_i(\alpha_i). \quad (1)$$

Proof. Let σ be a policy that generates the flow $\alpha_1, \dots, \alpha_m$. We assume that when an agent receives a signal σ_i and defies the mediator, he randomizes among the remaining links $j \neq i$. Under a uniform prior, this implies that σ is incentive compatible if and only if

$$\sum_{i=1}^m \alpha_i \left(c_i(\alpha_i) - \frac{1}{m-1} \sum_{j \neq i} c_j(\alpha_j) \right) \leq 0. \quad (2)$$

Rearranging, and noting that

$$\frac{1}{m-1} \sum_{i=1}^m \alpha_i \sum_{j \neq i} c_j(\alpha_j) = \frac{1}{m-1} \sum_{i=1}^m c_i(\alpha_i) (1 - \alpha_i),$$

we have that the left-hand side of equation (2) is equal to:

$$\sum_{i=1}^m \alpha_i c_i(\alpha_i) - \frac{1}{m-1} \sum_{i=1}^m c_i(\alpha_i) + \frac{1}{m-1} \sum_{i=1}^m \alpha_i c_i(\alpha_i)$$

and simplifying, we obtain the desired condition. \square

3 Affine cost functions

In the full-information case, if all costs are affine functions, then the PoA is $4/3$ in the case of two links, and this is the worst-possible ratio across all networks. We show that the mediation ratio is bounded away from the PoA in the case of two links, and for any fixed number of links. However, as m grows, the mediation ratio approaches $4/3$. We begin with the case of two parallel links.

Proposition 3.1. *The MR of BCG permutation games with affine cost functions on two links is at least $8/7$.*

Proof. Consider a prior p that permutes the costs $\{c_1(x) = x, c_2(x) = 1/2\}$. Note that the global optimum solves $\alpha^* \in \arg \max_{\alpha} \{\alpha^2 + \frac{1}{2}(1 - \alpha)\}$, and so $\alpha^* = \frac{1}{4}$, and the optimal social cost is $7/16$. To find the best mediated policy, note that for any policy σ that sends α^σ flow to the link with cost $c_1(x)$, an agent who receives a signal σ_i expects to have been sent to the link with cost $c_1(x)$ with probability α^σ . Thus he faces an expected cost of $(\alpha^\sigma)^2 + \frac{1}{2}(1 - \alpha^\sigma)$ if he follows the mediator's policy, or an expected cost of $\frac{\alpha^\sigma}{2} + (1 - \alpha^\sigma)\alpha^\sigma$ if he defies the mediator and instead uses the link he was not sent to. By definition, σ is incentive compatible if and only if the expected cost of following σ is lower than the expected cost of defying it for each agent. This reduces to the constraint that $(\alpha^\sigma)^2 - \alpha^\sigma + 1/4 < 0$, which is not satisfied for any $\alpha^\sigma \in [0, 1]$. Thus, the only policy that the mediator can implement is an uninformative one, under which, agents randomize uniformly between the two links, yielding a social cost of $1/2$. Putting these together, $MR(p)$ is given by the ratio of the cost of the uninformed BNE and the cost of the unconstrained optimum: $MR(p) = \frac{1/2}{7/16} = \frac{8}{7}$. \square

Proposition 3.2. *The MR of BCG permutation games with affine cost functions on two links is at most $8/7$.*

Proof. Let \mathcal{P} be the family of priors on permutations of costs $\{c_1(x) = a_1(x) + b_1, c_2(x) = a_2(x) + b_2\}$ with $\{a_1, a_2, b_1, b_2\} \in \mathbb{R}_+^4$. As we are interested in $\sup_{p \in \mathcal{P}} MR(p)$, we consider priors for which the globally optimal flow is not implementable by the mediator (so that $MR(p) > 1$), so that the majority of agents pays a higher cost than the minority under the globally optimal flow.² That is, if for the parameters, $\{a_1, a_2, b_1, b_2\}$, α^* is the globally optimal flow that is not implementable in a mediated equilibrium, then Lemma 2.8 implies that

$$a_1 \alpha^* + b_1 \geq a_2 (1 - \alpha^*) + b_2$$

whenever $\alpha^* > 1/2$ on link with cost $c_1(x)$. Notice that when both a_1 and a_2 or both b_1 and b_2 are 0, the globally optimal flow never satisfies this, and so is implementable, a contradiction. Thus, this restriction imposes a constraint on the parameters $\{a_1, a_2, b_1, b_2\}$ in the support of p : namely, that (without loss of generality) $a_1 > 0$ and $b_2 > 0$. The maximal $MR(p)$ is then obtained from the constrained optimization problem that maximizes $MR(p)$, given $a_1 > 0$

²That is, $c(\alpha^*) > c(1 - \alpha^*)$ whenever $\alpha^* > 1/2$ in the globally optimal flow.

and $b_2 > 0$.³ The solution to this problem yields the costs $c_1(x) = x$ and $c_2(x) = 1/2$ with $MR(p) = 8/7$. \square

We complement this positive result by a negative one, showing that when m is large, the MR goes to $4/3$ in the *IID* model, and hence also in the *permutation* model. Since the MR is upper bounded by the price of anarchy (which is bounded by $4/3$), this settles the question for large values of m . Before presenting this result, we introduce several helpful lemmas.

Lemma 3.3. *Consider a set of cost vectors c^1, c^2, \dots, c^ℓ , where for every i , $c^i = (c_1^i, \dots, c_m^i)$ is a vector of cost functions on the m links. If for every i the best mediated equilibrium is the uninformed BNE, then for every distribution of these costs the best mediated equilibrium is the uninformed BNE (where $1/m$ is routed on every link).*

The full proof, which can be found in the full version of this paper, shows that if a distribution of cost functions admits an implementable policy that is superior to the uninformed BNE, a realization of costs in the support can be found that allows for a superior implementable policy independently. Toward establishing our negative result, we first show that in a very simple BCG game, mediation gives no advantage, and the best mediated equilibrium is the trivial one in which agents choose which edge to follow uniformly at random. The proof, which follows from Lemma 2.8, is omitted due to space constraints and can be found in the full version.

Lemma 3.4. *Consider an instance where for some k (between 0 and m), there are k cost functions of the form $c(x) = x$, and $(m - k)$ cost functions of the form $1/m$. Then, the best mediated equilibrium is the uninformed BNE (where $1/m$ is routed on every link), independent of k .*

We are now ready to present our negative result, by constructing a simple BCG for which the MR goes to $4/3$ in the limit as the number of links goes to infinity.

Proposition 3.5. *The MR in an IID BCG with m parallel links and affine cost functions is at least*

$$\frac{4}{3 + 2 \log(m)/m + (1 - 2 \log(m)/m)^m}. \quad (3)$$

The proof of Proposition 3.5 considers an IID BCG on m links, in which each link has cost $c(x) = 1/m$ with probability $p = 2 \log(m)/m$, and has cost $c(x) = x$ with the remaining probability. It then shows that the MR of this BCG is given by equation (3), which approaches $4/3$ as m approaches infinity. As $4/3$ is the maximal PoA, it is also the maximum possible MR , and so this bound is tight. The full proof can be found in the full version of this paper.

4 General cost functions

As established in the previous section, when the cost functions are restricted to the set of affine functions, the MR in

³That is, for any set of parameters such that $a_1 > 0$ and $b > 0$, OPT is given by the solution to the first order condition of social cost, and the best mediated policy is the uninformed BNE. The maximal $MR(p)$ maximizes the ratio of these two costs over the parameter space $\{a_1, b_1, a_2, b_2\}$.

m -link networks converges to the PoA as the number of links m increases.

In this section we show that this is not the case for general nondecreasing cost functions. In particular, while the PoA can be arbitrarily high in this case, the MR never exceeds m .

Theorem 4.1. *The MR in a BCG with m links with general non-decreasing cost functions is at most m .*

Before proving Theorem 4.1, we introduce a reduced condition under which a flow can be implemented. We then show that whenever the optimal flow of a BCG is not implementable, there exists an implementable flow with a social cost that is within a factor of m of the optimal social cost.

Definition 4.2. A flow $\alpha_1, \dots, \alpha_m$ can be divided into buckets, if there exists a partition of the links in $[m]$ to sets S_1, \dots, S_k such that:

1. For any $i \leq k$ and any two links $s, t \in S_i$ we have $c_s(\alpha_s) = c_t(\alpha_t)$.
2. If $i < j$, $s \in S_i$ and $t \in S_j$ then $c_s(\alpha_s) < c_t(\alpha_t)$.
3. For every i , let $\sum_{j=i}^k |S_j| = m_i$, and let $S_{\geq i} = S_i \cup S_{i+1} \cup \dots \cup S_k$. We have

$$\frac{\sum_{s \in S_i} \alpha_s}{\sum_{s \in S_{\geq i}} \alpha_s} > \frac{|S_i|}{m_i}$$

4. Denoting $S_i = \{\alpha_{i1}, \dots, \alpha_{ir}\}$ then $\alpha_{i1}, \dots, \alpha_{ir}$ can be implemented on r links.

Note that Definition 4.2 did not require that $\sum_{i=1}^m \alpha_i = 1$. Before we show that every flow that can be divided into buckets can be implemented, we introduce the following lemma.

Lemma 4.3. *Let $\alpha_1, \dots, \alpha_m$ be a flow that can be divided into buckets, and let $\sum_{i=1}^m \alpha_i = \gamma$. Then $\sum_{i=1}^m \alpha_i c_i(\alpha_i) \leq \frac{\gamma}{m} \sum_{i=1}^m c_i(\alpha_i)$.*

Proof. We proceed inductively over the number of links. When $m = 1$, the statement holds trivially. Let $\tilde{\gamma} = \sum_{i=2}^m \alpha_i = \gamma - \alpha_1$, and suppose, for the inductive hypothesis, that

$$\sum_{i=2}^m \alpha_i c_i(\alpha_i) \leq \frac{\tilde{\gamma}}{m-1} \sum_{i=2}^m c_i(\alpha_i).$$

Since α_i are divided into buckets, by property (2), $c_1(\alpha_1) \geq \frac{1}{m-1} \sum_{i=2}^m c_i(\alpha_i)$ and by property (3), $\alpha_1 \leq \frac{\gamma}{m} \leq \frac{1}{m}$, and $\frac{\tilde{\gamma}}{m-1} \leq \frac{1}{m}$. Thus, using the inductive hypothesis, we have that:

$$\begin{aligned} \sum_{i=1}^m \alpha_i c_i(\alpha_i) &= \alpha_1 c_1(\alpha_1) + \sum_{i=2}^m \alpha_i c_i(\alpha_i) \\ &\leq \frac{\tilde{\gamma}}{m-1} \sum_{i=2}^m c_i(\alpha_i) + \alpha_1 c_1(\alpha_1) \end{aligned} \quad (4)$$

$$\leq \left(\frac{\gamma - \frac{1}{m}}{m-1}\right) \sum_{i=2}^m c_i(\alpha_i) + \frac{1}{m} c_1(\alpha_1) \quad (5)$$

$$\leq \frac{1}{m} \sum_{i=2}^m c_i(\alpha_i) + \frac{1}{m} c_1(\alpha_1) = \frac{1}{m} \sum_{i=1}^m c_i(\alpha_i), \quad (6)$$

where the inequality in (3) comes from plugging in the inductive hypothesis and, since $\alpha_1 \leq \frac{1}{m}$, the inequality in (4) is equivalent to:

$$\begin{aligned} & \left[\gamma - \frac{1}{m} - (\gamma - \alpha_1) \right] \sum_{i=2}^m \frac{c_i(\alpha_i)}{m-1} + \left[\frac{1}{m} - \alpha_1 \right] c_1(\alpha_1) \\ &= \alpha_1 \left(\sum_{i=2}^m \frac{c_i(\alpha_i)}{m-1} - c_1(\alpha_1) \right) + \frac{1}{m} \left(c_1(\alpha_1) - \sum_{i=2}^m \frac{c_i(\alpha_i)}{m-1} \right) \geq 0. \end{aligned}$$

The inequality in (5) comes from substituting $\frac{\tilde{\gamma}}{m-1} \leq \frac{1}{m}$. \square

Corollary 4.4. *Let $\alpha_1, \dots, \alpha_m$, be a flow that can be divided into buckets. Then it is implementable as a mediated equilibrium.*

We now prove a stronger statement which implies Theorem 4.1.

Lemma 4.5. *Let $\alpha_1, \dots, \alpha_m$ be an optimal flow on m links. Then either it is implementable, or there exists another flow β_1, \dots, β_m on the links, such that $\sum \beta_i = \sum \alpha_i$, the cost is $\sum_i \beta_i c_i(\beta_i) \leq m \sum_i \alpha_i c_i(\alpha_i)$, the minimum cost per link doesn't decrease, $\min_i c_i(\beta_i) \geq \min_i c_i(\alpha_i)$, and β_1, \dots, β_m can be divided into buckets.*

Proof. Toward a proof by induction, assume that the lemma holds for every number of links $n < m$. Now for m links, if the optimum is implementable, then we are done. Otherwise, the IC constraint must be violated for α so that

$$\sum_{i=1}^m \alpha_i c_i(\alpha_i) > \frac{1}{m} \sum_{i=1}^m c_i(\alpha_i)$$

and so there must exist some k such that $\alpha_k \geq \frac{1}{m}$ and $c_k(\alpha_k) \geq \frac{1}{m} \sum_{i=1}^m c_i(\alpha_i)$. Choose k to maximize $c_k(\alpha_k)$.

Let $J = \{i : c_i(\alpha_i) > c_k(\alpha_k)\}$, the set of links with costs strictly higher than $c_k(\alpha_k)$ under α . By the definition of k , for every $i \in J$ we have that $\alpha_i < 1/m$. Since $|J| < m$, we apply the induction hypothesis on J , and get flow $\beta_{i1} \dots \beta_{ir}$ which can be divided into buckets. Let S_1, \dots, S_b denote those buckets. Let $C = \{i : i \notin J\}$, where C stands for cheap links. Note that:

$$\frac{\sum_{i \in C} \alpha_i}{\sum_{i \in m} \alpha_i} \geq \frac{|C|}{m}.$$

Let $\{\beta_i\}_{i \in C}$ be a full information Nash equilibrium that routes $\sum_{i \in C} \alpha_i$ flow on the links of C , so that $c_i(\beta_i) = c_j(\beta_j)$ for every pair of links i, j . We have that:

1. The minimal cost didn't decrease $\min_{i \in C} c_i(\beta_i) \geq \min_{i \in C} c_i(\alpha_i)$.
2. The total cost increased by at most a factor of m :

$$\sum_{i \in C} \beta_i c_i(\beta_i) \leq m \sum_{i \in C} \alpha_i c_i(\alpha_i).$$

3. For every $i, j \in C$ we have that $c_i(\beta_i) = c_j(\beta_j)$.

We can now take the flow $\{\beta_i\}_{i \in J} \cup \{\beta_i\}_{i \in C}$, and get a flow β_1, \dots, β_m . This new flow can be partitioned to $b+1$ buckets, where the cheapest bucket is $\{\beta_i\}_{i \in C}$ and the other buckets are S_1, \dots, S_b . We also have:

1. The minimal cost did not decrease $\min_{i \in [m]} c_i(\beta_i) = \min_{i \in C} c_i(\beta_i) \geq \min_{i \in C} c_i(\alpha_i) = \min_{i \in [m]} c_i(\alpha_i)$, since we chose C to contain the cheap links.

2. As for the social cost, we have

$$\begin{aligned} \sum_{i \in [m]} \beta_i c_i(\beta_i) &= \sum_{i \in C} \beta_i c_i(\beta_i) + \sum_{i \in J} \beta_i c_i(\beta_i) \\ &\leq m \sum_{i \in C} \alpha_i c_i(\alpha_i) + m \sum_{i \in J} \alpha_i c_i(\alpha_i) = m \sum_{i \in [m]} \alpha_i c_i(\alpha_i). \end{aligned}$$

Which completes the proof. \square

Moreover, the upper bound on $MR(\mathcal{P})$ in games with general cost functions is tight. As we demonstrate in the following lemma, there exist BCGs on m links and priors p with $MR(p) = m$. Thus, $MR(\mathcal{P})$ over all such priors is exactly m .

Theorem 4.6. *The MR in a BCG with m links congestion game and general non-decreasing cost functions is at least m .*

Proof. Consider a game with m links, each of which has a step cost function that assigns cost 0 if the amount of flow on it is $x \in [0, \frac{1}{m})$ and assigns cost 1 if $x \in [\frac{1}{m}, 1]$. The optimal policy is to route slightly less than $\frac{1}{m}$ of the flow on each one of $m-1$ links, and the remaining flow through the remaining link. This way, only the flow that goes through the last link incurs a positive cost of 1, thus the social cost is slightly more than $\frac{1}{m}$.

We now claim that the only implementable policy is one that sends exactly $1/m$ of the flow on each one of the m links, yielding an expected cost of 1. Suppose toward contradiction that there were a policy σ that induced an equilibrium α^σ which produced a social cost in which only k out of m links had flow greater or equal to $1/m$, for some $k < m$. For at least one of these links i , it holds that $\alpha_i > 1/m$; therefore the sum of flows on these k links is strictly greater than k/m . Consequently, an agent who follows the mediator's recommendation incurs an expected cost that is strictly greater than k/m . On the other hand, choosing a link uniformly at random yields an expected cost of k/m . We conclude that the only implementable flow routes $1/m$ of the flow through each link, yielding a social cost of 1. \square

5 Conclusion

We study a class of Bayesian non-atomic parallel-links routing games in which agents have incomplete information about the costs of the links in the game, but a knowledgeable mediator can correlate outcomes by giving incentive compatible recommendations in order to minimize the total travel time of the system. We define the *mediation ratio*: the ratio between the mediated equilibrium arising from IC recommendations and the social optimum, which is always bounded from above by the price of anarchy. We find that for general monotone cost functions, while the price of anarchy is unbounded, the mediation ratio for m parallel links is exactly m . The main open question left by our work is proving mediation ratio bounds for general road networks. For price of anarchy, the worst case manifests in the parallel link networks that we have studied.

References

- [Ashlagi *et al.*, 2009] Itai Ashlagi, Dov Monderer, and Moshe Tennenholtz. Mediators in position auctions. *Games and Economic Behavior*, 67(1):2–21, September 2009.
- [Aumann, 1974] Robert J. Aumann. Subjectivity and correlation in randomized strategies. *Journal of Mathematical Economics*, 1(1):67–96, March 1974.
- [Joshua R. Davis and Wexler, 2011] Alexa Sharp Joshua R. Davis, David Liben-Nowell and Tom Wexler. Mediated equilibria in Load-Balancing Games. *Chicago Journal of Theoretical Computer Science*, 2011(5), November 2011.
- [Kremer *et al.*, 2013] Ilan Kremer, Yishay Mansour, and Motty Perry. Implementing the "Wisdom of the Crowd", February 2013.
- [Monderer and Tennenholtz, 2003] Dov Monderer and Moshe Tennenholtz. k-Implementation. In *Proceedings of the 4th ACM conference on Electronic commerce - EC '03*, pages 19–28, New York, New York, USA, June 2003. ACM Press.
- [Monderer and Tennenholtz, 2009] Dov Monderer and Moshe Tennenholtz. Strong mediated equilibrium. *Artificial Intelligence*, 173(1):180–195, January 2009.
- [Roughgarden and Tardos, 2000] Tim Roughgarden and Éva Tardos. How bad is selfish routing? In *41st Annual Symposium on Foundations of Computer Science, FOCS 2000, 12-14 November 2000, Redondo Beach, California, USA*, pages 93–102, 2000.
- [Roughgarden and Tardos, 2004] T. Roughgarden and É. Tardos. Bounding the inefficiency of equilibria in nonatomic congestion games. *Games and Economic Behavior*, 49(2):389–403, 2004.
- [Roughgarden, 2003] Tim Roughgarden. The price of anarchy is independent of the network topology. *J. Comput. Syst. Sci.*, 67(2):341–364, 2003.
- [Rozenfeld and Tennenholtz, 2007] Ola Rozenfeld and Moshe Tennenholtz. Routing Mediators. In *IJCAI International Joint Conference on Artificial Intelligence*, pages 1488–1493, 2007.