

When are Marginal Congestion Tolls Optimal?

Reshef Meir, Technion
reshefm@ie.technion.ac.il

David C. Parkes, Harvard University
parkes@eecs.harvard.edu

Abstract

Marginal tolls are known to provide the existence of an optimal equilibrium in atomic congestion games, but unlike nonatomic games, there might be additional equilibria even with linear cost functions on resources. In this paper, we show that in games with a large number of players, all equilibria are near-optimal.

1 Introduction

It is well known that selfish routing results in suboptimal social behavior and in increased latency [Pigou, 1920]. The modern literature formalizes selfish routing scenarios as *congestion games*, where the inefficiency due to strategic behavior is quantified as the *Price of Anarchy* (PoA)—the ratio between the optimal total latency and the maximal total latency in equilibrium [Roughgarden and Tardos, 2007].

The game theoretic literature on selfish routing can be classified into models of atomic (unsplittable) flow and nonatomic flow, where in the latter, each agent accounts for an infinitesimally small fraction of the total congestion. While in both models a pure equilibrium is guaranteed to exist, and can be found via a simple local best-response dynamics, atomic congestion games are considered more challenging to analyze. Atomic games may have multiple equilibria of different costs, and the price of anarchy can be much higher than in nonatomic games.

The PoA is well understood in congestion games, both atomic and nonatomic, and almost independent from the topology of the network [Roughgarden, 2009]. That is, the inefficiency depends mostly on the edge latency functions, and a simple network of two parallel edges (or roads) is sufficient to create instances with the highest possible PoA.

Still, it is interesting to look to change the behavior of agents by charging them for using a resource. It has been known since [Beckmann *et al.*, 1956] that to enforce optimal behavior in nonatomic games (i.e. such that all equilibria have minimum total latency), it is sufficient to impose *marginal congestion tolls*, i.e., charge each agent based on the latency he *currently* adds to the other agents.¹ Note that we assume tolls are *dynamic* that depend on monitoring the actual congestion on one hand, but can be easily computed. This is in contrast to *static tolls* that typically depend on the

¹It is typically assumed that the tolls themselves are not calculated as part of the total cost, e.g. because they return to the society indirectly, or because the central authority only cares about the latency. Non-refundable tolls are also studied [Cole *et al.*, 2006] but not in this paper.

optimal congestion, and often require extensive computation (see e.g. [Bonifaci *et al.*, 2011]).

For atomic games, it is known that marginal tolls guarantee the existence of at least one optimal equilibrium [Sandholm, 2007], however there may be other inefficient equilibria, even in games with linear latencies [Caragiannis *et al.*, 2010a]. The problem becomes even more involved if we take into account more general notions of equilibrium such as mixed and correlated equilibrium. For a specific classes of atomic routing games, marginal tolls guarantee optimal behavior in any pure equilibrium. This is the case for example for symmetric networks with parallel links (also known as *resource selection games*) since in such networks the equilibrium is unique. The class of networks for which marginal tolls are optimal was extended first in an unpublished (and unfinished) work by Singh [Singh, 2008]. However Singh’s result was very recently refuted by Igal Milchtaich (personal communications) who provided the correct characterization.

Several other papers studied more complicated taxation schemes and how low they can affect the PoA [Fotakis and Spirakis, 2008; Caragiannis *et al.*, 2010a].

Our contribution We show that for any fixed network, if the number of players is sufficiently large, then any equilibrium under marginal tolls is near-optimal. Further, this result extend to mixed, correlated, and coarse correlated equilibria.

We use the smoothness framework [Roughgarden, 2009], which enables the PoA bounds to be established with relatively short and simple proofs.

We also consider agents with variable sensitivity to monetary tolls [Cole *et al.*, 2006; Karakostas and Kolliopoulos, 2004; Fotakis *et al.*, 2010], reflecting how agents trade-off money for time. As discussed in [Yang and Zhang, 2008; Meir and Parkes, 2015b], the parameter may be unobservable, and thus unknown to the central authority setting the tolls. Thus, following [Meir and Parkes, 2015b] and in contrast to most of the mechanism design literature, we assume that a marginal toll is applied, and analyze the equilibrium for a population as the sensitivity parameter varies.

Along the way, we state formally some known results on marginal tolls that seem to have been overlooked in the recent study of atomic congestion and routing games.

2 Preliminaries

For an integer m , $[m] = \{1, 2, \dots, m\}$. We use bold letters to denote vectors, e.g., $\mathbf{a} = (a_1, \dots, a_m)$.

Following the definitions in [Roughgarden, 2007], a routing game is a tuple $G = \langle V, E, N, \mathbf{c}, \mathbf{u}, \mathbf{v} \rangle$, where

- (V, E) are vertices and edges of a directed graph;

- N is a finite set of agents of size n ;
- $\mathbf{c} = (c_e)_{e \in E}$, where $c_e(x) \geq 0$ is a non-decreasing function indicating the cost incurred when x agents use edge e (c_e are called *latency functions*);²
- \mathbf{u}, \mathbf{v} are vectors of n vertices each, where (u_i, v_i) are the source and target nodes of agent i ;

We denote by $A_i \subseteq 2^E$ the set of all directed paths between the pair of nodes (u_i, v_i) in the graph. Thus A_i is the set of *actions* available to agent i . We denote by $A = \cup_i A_i$ the set of all directed source-target paths. A routing game is *symmetric* (also called *single-source-single-target*) if all agents have the same set of actions, i.e., $A_i = A$ for all i .

An *action profile* $\mathbf{a} = (a_i)_{i \in N}$ specifies the path $a_i \in A_i$ of each agent i , and $\mathcal{A} = \times_{i \in N} A_i$ is the set of all action profiles. We denote by $\mathbf{s}_e(\mathbf{a}) \in \mathbb{N}$ the congestion on edge $e \in E$ in profile \mathbf{a} , i.e., $\mathbf{s}_e = \mathbf{s}_e(\mathbf{a}) = |\{i \in N : e \in a_i\}|$ (\mathbf{a} is omitted when clear from context).

The cost for agent i in profile \mathbf{a} is summed over all edges, $C_i(\mathbf{a}) = \sum_{e \in a_i} c_e(s_e)$. The *social cost* in a profile \mathbf{a} in game G is attained by summing over all agents:

$$SC(G, \mathbf{a}) = \sum_{i=1}^n C_i(\mathbf{a}) = \sum_{i=1}^n \sum_{e \in a_i} c_e(s_e) = \sum_{e \in E} s_e c_e(s_e). \quad (1)$$

We denote by $\mathbf{a}^* = \mathbf{a}^*(G) = \operatorname{argmin}_{\mathbf{a} \in \mathcal{A}} SC(G, \mathbf{a})$ the profile that minimizes the social cost (optimal profile).

A profile \mathbf{a} is a *pure Nash equilibrium* if no agent can gain by changing her strategy, i.e. if for all $i \in N, a'_i \in A_i, C_i(\mathbf{a}) \leq C_i(\mathbf{a}_{-i}, a'_i)$, where $\mathbf{a}_{-i} = (a_j)_{j \neq i}$. The definition of equilibrium extends to mixed and correlated strategies. We omit the formal details. Denote by $PNE(G) \subseteq \mathcal{A}$ the sets of pure Nash equilibria of G .

The *price of anarchy* (PoA) of G is the ratio between the social cost of worst equilibrium and the optimal profile, i.e. $\text{PoA}(G) = \frac{\max\{SC(G, \mathbf{a}) : \mathbf{a} \in PNE(G)\}}{SC(G, \mathbf{a}^*)}$ (the definition of mixed- and correlated-POA is similar). It is well known that the PoA can be upper bounded using only the class of latency functions in G , regardless of the structure of (V, E) . For example, if all of c_e are affine functions ($c_e(x) = a_e x + b_e$ for $a_e, b_e \geq 0$) then $\text{PoA}(G) \leq \frac{5}{2}$, and this is true for mixed and correlated-PoA as well [Roughgarden, 2009].

The *price of stability* (PoS) of G is similarly defined as the ratio between the *best* equilibrium and the optimal profile [Christodoulou and Koutsoupias, 2005], i.e. $\text{PoS}(G) = \frac{\min\{SC(G, \mathbf{a}) : \mathbf{a} \in PNE(G)\}}{SC(G, \mathbf{a}^*)}$.

Biased games We are interested in a *biased game*, in our case because of the use of tolls.³ A biased game is a pair (G, \hat{G}) such that G, \hat{G} are identical except in their latency functions. Informally, we assume that players behave according to the “biased costs” $(\hat{c}_e)_{e \in E}$ (e.g. play an equilibrium of \hat{G}), but social cost is measured w.r.t. the “real costs” $(c_e)_{e \in E}$.

²Some authors prefer the term “arc” for directed edges. We stick with the common term in computer science.

³Biased games are also used to model cognitive and behavioral traits such as risk aversion [Ordóñez and Stier-Moses, 2010] or altruism [Caragiannis et al., 2010b].

The *biased price of anarchy/stability* (BPoA/BPoS) compares the equilibria of \hat{G} to the optimum of G , using the real social cost of both. Formally, $\text{BPoA}(G, \hat{G}) = \frac{\max\{SC(G, \mathbf{a}) : \mathbf{a} \in PNE(\hat{G})\}}{SC(G, \mathbf{a}^*)}$, and similarly for BPoS.

The primary bias we will consider in this paper is tolls, and in particular *marginal tolls*. That is, we define $\tau_e(x) = (x-1)[c_e(x) - c_e(x-1)]$, and set $\hat{c}_e^M(x) = c_e(x) + \tau_e(x)$. Toll $\tau_e(x)$ is exactly the marginal cost inflicted upon the remaining $x-1$ agents who use e due to an additional agent. Other toll schemes T can be similarly defined, replacing $\tau_e(x)$ with any other non-negative function $T_e(x)$.

A toll scheme T *strongly enforces* optimal flow in a game G if all equilibria of \hat{G}^T (i.e., the game with biased costs \hat{c}^T) are optimal in G (equivalently, if $\text{BPoA}(G, \hat{G}^T) = 1$) [Fotakis and Spirakis, 2008]. Similarly, a toll scheme *weakly enforces* optimal flows if $\text{BPoS}(G, \hat{G}^T) = 1$.

Marginal tolls in the nonatomic Pigouvian model were suggested by Beckmann [Beckmann et al., 1956], who showed they strongly enforce optimal flows in that model. Our goal is to understand the power of marginal tolls in atomic routing games.

3 Marginal tolls are weakly optimal

The marginal toll scheme for atomic games coincides with the taxes proposed by Sandholm [Sandholm, 2007], albeit Sandholm defined taxes at the strategy level, rather than tolls on particular edges. The observation that marginal tolls weakly enforce optimal flows was also made in an unpublished report by Singh [Singh, 2008].⁴ We state the result for the standard routing games framework.

Theorem 1 ([Sandholm, 2007; Singh, 2008]). *For any atomic congestion game G , there is a pure Nash equilibrium in \hat{G}^M that is optimal in G . Equivalently, $\text{BPoS}(G, \hat{G}^M) = 1$.*

The theorem follows from a simple observation: \hat{G}^M is a potential game [Rosenthal, 1973], whose potential function $\phi(\hat{G}^M, \mathbf{a})$ coincides with the social welfare of G . Thus the optimum of $SC(G, \mathbf{a})$ must be a local minimum of $\phi(\hat{G}^M, \mathbf{a})$, i.e. a pure Nash equilibrium. Quite strikingly, the theorem was extended to a much more general framework where agents have idiosyncratic preferences over strategies, and congestion may depend on agents weight or other features [Sandholm, 2007; Singh, 2008].

Unfortunately, in atomic games there may be additional suboptimal equilibria.

Example 1. *Consider a game with 3 parallel links, $E = \{a, b, c\}$ and 3 agents $N = \{1, 2, 3\}$. $A_1 = \{a, b\}$, $A_2 = \{b, c\}$, and $A_3 = \{c\}$. Latency functions are $c_b(x) = c_c(x) = x$, $c_a \equiv 2$ (see Fig. 1). The modified cost functions under any edge-independent nonnegative tolls can be written as $\hat{c}_b(x) = \hat{c}_c(x) = (1, 2 + T(x), 3 + T'(x))$. The unique optimum is $\mathbf{a}^* = (a, b, c)$ with cost $SC(\mathbf{a}^*) = 2 + 1 + 1 = 4$, which is also a PNE. However, there is another PNE $\mathbf{a}' =$*

⁴Recent works on tolls in routing games seem to be unaware of this observation [Fotakis and Spirakis, 2008; Fotakis et al., 2010; Swamy, 2012].

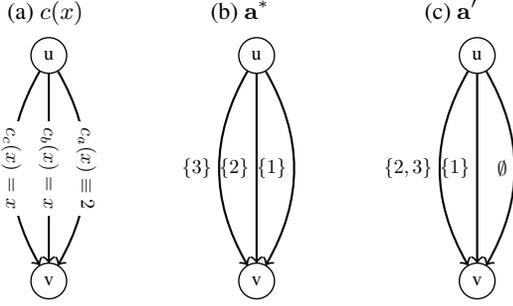


Figure 1: Figure (1a) shows the base game G . The other figures show the optimal state \mathbf{a}^* and the state \mathbf{a}' which is an additional equilibrium of both G and \hat{G}^T .

(b, c, c) with cost $SC(\mathbf{a}') = 1 + 2 + 2 = 5$. This remains a PNE of \hat{G}^T as long as $\hat{c}_b^T(x) = \hat{c}_c^T(x)$: agent 2 is not allowed to use edge a , and agent 1 does not want to use it since $\hat{c}_a(1) = 2 > 1 = \hat{c}_b(1)$.

This means that marginal tolls in atomic games do not, in the general case, strongly enforce optimal flows.

4 Strongly Enforcing Optimal Flows

The prominent technique for proving PoA bounds is *smoothness analysis*. In short, a game G is (λ, μ) -smooth if for all $\mathbf{a} \in \mathcal{A}$ there is $\mathbf{a}' \in \mathcal{A}$ such that $\sum_{i \in N} C_i(\mathbf{a}_{-i}, a'_i) \leq \lambda SC(G, \text{OPT}(G)) + \mu SC(G, \mathbf{a})$. If a game G (not just a routing game) is (λ, μ) -smooth, then $\text{PoA}(G) \leq \frac{\lambda}{1-\mu}$ [Roughgarden, 2009]. Further, this holds for the mixed, correlated, and coarse-correlated PoA as well. For routing games, it is also shown that restricting the class of latency functions results in smooth games. For example, if all cost functions are affine, then G is $(\frac{5}{3}, \frac{1}{3})$ -smooth (thereby showing $\text{PoA}(G) \leq \frac{5}{2}$).

Given a biased game (G, \hat{G}) , we can similarly define the property of *biased smoothness*.

Definition 1. (G, \hat{G}) is $(\hat{\lambda}, \hat{\mu})$ -biased smooth (BS), if there is \mathbf{a}' s.t. for any profile \mathbf{a} ,

$$\sum_{j \in N} (C_j(\mathbf{a}) + \hat{C}_j(\mathbf{a}_{-j}, a'_j) - \hat{C}_j(\mathbf{a})) \leq \hat{\lambda} SC(G, \text{OPT}(G)) + \hat{\mu} SC(G, \mathbf{a}). \quad (2)$$

It is easy to see that if G is (λ, μ) -smooth, then (G, G) is (λ, μ) -BS: we set $\mathbf{a}' = \text{OPT}(G)$, and note that $\sum_{j \in N} (C_j(\mathbf{a}) + C_j(\mathbf{a}_{-j}, a'_j) - C_j(\mathbf{a})) = \sum_{j \in N} C_j(\mathbf{a}_{-j}, a'_j)$.

Theorem 2. Suppose that (G, \hat{G}) is $(\hat{\lambda}, \hat{\mu})$ -BS. Let σ be any equilibrium (pure, mixed, correlated, or coarse-correlated) of the game \hat{G} . Then $SC(G, \sigma) \leq \frac{\hat{\lambda}}{1-\hat{\mu}} SC(G, \text{OPT}(G))$.

The original proof of Roughgarden [Roughgarden, 2009] for the PoA (and coarse-correlated PoA) naturally extends to biased smoothness.⁵ For completeness, we provide the proof

⁵A similar definition of smoothness was applied, for example, for finite congestion games with altruism: when $\hat{C}(\mathbf{a})$ is a combination of $C(\mathbf{a})$ and $SC(\mathbf{a})$, then the BPoA coincides with the “robust PoA” of Chen et al. [Chen et al., 2011].

(almost identical to the ones in [Roughgarden, 2009; Chen et al., 2011]) in the appendix.

In particular, $(1, 0)$ -BS means that $\text{BPoA}(G, \hat{G}) = 1$, i.e. that any PNE of \hat{G} is optimal in G .

We are interested in showing that (G, \hat{G}^M) is BS for some reasonable parameters $\hat{\lambda}, \hat{\mu}$.

4.1 Smoothness in the large

When an atomic game becomes *large*, i.e. when we fix the network and increase the number of players, there is evidence that the game behaves more similarly to a nonatomic game [Feldman et al., 2015]. We show how to extend biased-smoothness analysis (and in particular marginal tolls) to large atomic games. While we can not apply the results of Feldman et al. directly, our techniques are inspired by theirs.

Lemma 3. Let \mathbf{a}, \mathbf{a}' be any two profiles in G with n agents, and let $\epsilon = \epsilon(G) = \max_{e \in E, x \in \mathbb{N}} (c_e(x+1) - c_e(x))$. Then $\sum_{j \in N} C_j(\mathbf{a}_{-j}, a'_j) - C_j(\mathbf{a}) \leq \sum_{e \in E} (s'_e - s_e) c_e(s_e) + O(n\epsilon)$.

Proof.

$$\begin{aligned} & \sum_{j \in N} (C_j(\mathbf{a}_{-j}, a'_j) - C_j(\mathbf{a})) \\ &= \sum_{j \in N} \left(\left(\sum_{e \in a'_j \setminus a_j} c_e(s_e + 1) + \sum_{e \in a_j \cap a'_j} c_e(s_e) \right) - \sum_{e \in a_j} c_e(s_e) \right). \end{aligned}$$

By definition of ϵ , we continue:

$$\begin{aligned} & \leq \sum_{j \in N} \left(\left(\sum_{e \in a'_j \setminus a_j} (c_e(s_e) + \epsilon) + \sum_{e \in a_j \cap a'_j} c_e(s_e) \right) - \sum_{e \in a_j} c_e(s_e) \right) \\ & \leq \sum_{j \in N} \left(\sum_{e \in a'_j} c_e(s_e) - \sum_{e \in a_j} c_e(s_e) \right) + \sum_{j \in N} \sum_{e \in E} \epsilon \\ & = \sum_{j \in N} (s'_e c_e(s_e) - s_e c_e(s_e)) + n|E|\epsilon. \end{aligned}$$

□

That is, we can write the sum of deviations as a function of the aggregate congestion (approximately).

Next, we think of a sequence of atomic games with increasing n : We fix a network (V, E) and continuous quasi-convex cost functions $\mathbf{c} = (c_e)_{e \in E}$, where $c_e : [0, 1] \rightarrow \mathbb{R}^+$. For ease of presentation, we consider symmetric games (i.e. where there is just one source-target pair $u, v \in V$), although similar arguments extend to asymmetric games. This already induces a symmetric nonatomic game $\tilde{G} = (V, E, u, v, \mathbf{c})$. For $n \in \mathbb{N}$, we define G_n by setting $G_n = (V, E, N, u, v, \mathbf{c}^n)$, where $c^n(x) = c(x/n)$. Thus \tilde{G} is the limit of $(G_n)_{n=1,2,\dots}$ (we call it the *limit game*).

Our continuous cost functions can also be subject to biases. Let \hat{c}_e be the biased continuous cost of \tilde{c}_e , and $\hat{c}_e^n(x) = \hat{c}_e(x/n)$. Biased-smoothness for continuous cost functions was defined and explored in [Meir and Parkes, 2015b]: we say that c is $(\hat{\lambda}, \hat{\mu})$ -biased smooth w.r.t. \hat{c} if for all $t, t' \in \mathbb{R}_+$,

$$c(t)t + \hat{c}(t)(t' - t) \leq \hat{\lambda} c(t')t' + \hat{\mu} c(t)t.$$

Clearly, if \tilde{c} is $(\hat{\lambda}, \hat{\mu})$ -biased smooth w.r.t. $\hat{\tilde{c}}$, then c^n is $(\hat{\lambda}, \hat{\mu})$ -biased smooth w.r.t. \hat{c}^n for any n .

Theorem 4. Consider a limit game \tilde{G} , where \tilde{c}_e are quasi-convex and $(\hat{\lambda}, \hat{\mu})$ -biased smooth w.r.t. the bias \hat{c} . Then for any $\delta > 0$ there are $\epsilon > 0, n(\epsilon)$ s.t. for all $n > n(\epsilon)$, the atomic game (G^n, \hat{G}^n) is $((1 + \delta)\hat{\lambda}, \hat{\mu})$ -BS. In particular,

$$BPoA(G^n, \hat{G}^n) \leq (1 + \delta) \frac{\hat{\lambda}}{1 - \hat{\mu}},$$

and this extends to any coarse-correlated equilibrium.

Proof. Let $\mathbf{a}' = \text{OPT}(G^n)$, $Z^n = SC(G^n, \mathbf{a}')$. Since $SC(G^n, \mathbf{a}') = \Omega(n)$ (the cost for each agent is at least some constant), we write $Z^n > \rho n$ for some $\rho > 0$ and $n > n(\rho)$.

Since \tilde{c}_e is bounded and continuous for all $e \in E$,

$$\begin{aligned} \max_{x \in [n]} \{c_e^n(x+1) - c_e^n(x)\} &= \max_{x \in [n]} \{\tilde{c}_e(\frac{x}{n} + \frac{1}{n}) - \tilde{c}_e(\frac{x}{n})\} \\ &\leq \sup_{t \in [0,1]} \{\tilde{c}_e(t + \frac{1}{n}) - \tilde{c}_e(t)\} \xrightarrow{n \rightarrow \infty} 0, \end{aligned}$$

and thus for all $\epsilon > 0$ there is some $n(\epsilon)$ s.t. for all $n > n(\epsilon)$, we have $c_e^n(x+1) - c_e^n(x) < \epsilon$. By Lemma 3

$$\begin{aligned} SC(G^n, \mathbf{a}) + \sum_{j \in N} \hat{C}_j^n(\mathbf{a}_{-j}, a_j) - \hat{C}_j^n(\mathbf{a}) \\ &\leq SC(G^n, \mathbf{a}) + \sum_{e \in E} (s'_e - s_e) \hat{c}_e^n(s_e) + O(n\epsilon) \\ &= \sum_{e \in E} (s_e c_e^n(s_e) + (s'_e - s_e) \hat{c}_e^n(s_e)) + n\epsilon' \\ &\leq \sum_{e \in E} (\hat{\lambda} c_e^n(s'_e) s'_e + \hat{\mu} c_e^n(s_e) s_e) + n\epsilon' \quad (\text{smoothness}) \\ &= \hat{\lambda} Z^n + \hat{\mu} SC(G^n, \mathbf{a}) + n\epsilon' \\ &< \hat{\lambda} SC(G^n, \mathbf{a}') + \hat{\mu} SC(G^n, \mathbf{a}) + \frac{1}{\rho} Z^n \epsilon' \quad (Z^n > \rho n) \\ &= (\hat{\lambda} + \frac{\epsilon'}{\rho}) Z^n + \hat{\mu} SC(G^n, \mathbf{a}) \\ &\leq (1 + \frac{\epsilon'}{\rho}) \hat{\lambda} Z^n + \hat{\mu} SC(G^n, \mathbf{a}). \quad (\hat{\lambda} \geq 1) \end{aligned}$$

Selecting $\epsilon' < \delta \rho$ (and thus sufficiently small $\epsilon > 0$, and $n > \max\{n(\rho), n(\epsilon)\}$), completes the proof. The BPoA bound then follows directly from Theorem 2. \square

Since biased smoothness hold for various pairs of cost functions, Theorem 4 is quite useful. Mainly, we get that marginal tolls strongly enforce near-optimal flow if there are enough players.

Corollary 5. Consider any limit game \tilde{G} , where \tilde{c}_e are quasi-convex. Then for any $\delta > 0$ there is some $n(\delta)$ s.t. for all $n > n(\delta)$, $BPoA(G^n, \hat{G}^n) \leq 1 + \delta$.

Proof. Consider the continuous version of marginal tolls $\hat{c}(t) = \tilde{c}(t) + t \cdot \frac{\partial c(t)}{\partial t}$ [Beckmann *et al.*, 1956].⁶ The proof follows directly from Theorem 4 and from the fact that any quasi-convex function \tilde{c} is $(1, 0)$ -biased smooth w.r.t. \hat{c} [Meir and Parkes, 2015b]. \square

⁶Due to rounding, $\hat{c}^n(x)$ is very close, but not identical to the discrete $\hat{c}^M(x)$ we previously defined.

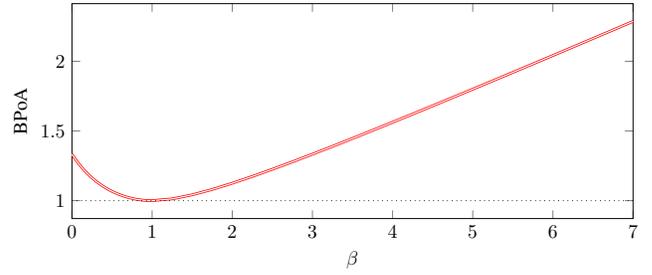


Figure 2: The X-axis shows the tax-sensitivity of agents, where $\beta = 0$ means they ignore the tax.

The double red lines are the tight bounds on the BPoA for large games with affine costs stated in Corollary 6.

5 Tax-sensitivity

We next consider agents with variable sensitivity to monetary tolls, as in [Cole *et al.*, 2006]. Formally, the marginal toll $\tau_e(x)$ is imposed on edge e , but the cost experienced by the agents is $\hat{c}_e^\beta(x) = c(x) + \beta \cdot \tau_e(x)$, where β is a parameter reflecting how agents trade-off money for time. Denote by \hat{G}_β the biased game obtained from G by replacing every cost function $c_e(x)$ with $\hat{c}_e^\beta(x)$. We analyze the equilibrium for a population with parameter β (where $\beta = 1$ means that $\hat{c}_e^\beta(x) = \hat{c}_e^M(x)$).

In [Meir and Parkes, 2015b], various BPoA bounds are derived for nonatomic games with various classes of cost functions (general/convex/polynomial/linear). We show how these bounds extend to large games.

For large atomic games, all the biased smoothness bounds from [Meir and Parkes, 2015b] for tax-sensitivity and other biases immediately apply. These bounds are also known to be tight.

For example, it was shown that affine cost functions (of the form $\tilde{c}(t) = at + b$ for $a, b \geq 0$) are $(1, \frac{(1+\beta)^2}{4} - \beta)$ -biased smooth w.r.t. $\hat{c}(t)$ as defined above for all $\beta \leq 1$ and $(\frac{(1+\beta)^2}{4\beta}, 0)$ -biased smooth for $\beta \geq 1$. We get the following corollary due to Theorem 4:

Corollary 6. Consider any limit game \tilde{G} , where \tilde{c}_e are affine. Then for any $\delta > 0$ there is some $n(\delta)$ s.t. for all $n > n(\delta)$, $BPoA(G^n, \hat{G}^n) \leq \frac{1}{(\beta+1) - \frac{(1+\beta)^2}{4}}$ if $\beta \leq 1$, and $BPoA(G^n, \hat{G}^n) \leq \frac{(1+\beta)^2}{4\beta}$ if $\beta \geq 1$.

Another benefit of smoothness-in-the-large is that the parameters $\hat{\lambda}, \hat{\mu}$ are typically much smaller for classes of continuous functions than for the corresponding class of discrete costs. Indeed, [Feldman *et al.*, 2015] show that the PoA of large games is significantly smaller due to this: for linear costs the PoA drops from $\frac{5}{2}$ to $\frac{4}{3}$, and for polynomials of degree d , the PoA drops from $\Omega(2^d)$ to $O(\frac{d}{\ln d})$. Our result shows that this still holds for large games with biases. For brevity we do not re-state all the results from [Meir and Parkes, 2015b] for large atomic games, however Fig. 2 shows the bounds for affine costs.

6 Discussion

We have studied the problem of strongly enforcing optimal flows in atomic congestion games through marginal congestion tolls. Such tolls always weakly enforce optimal flows, and strongly enforce optimal tolls in large games. Further, our analysis extends to games where agents' tax-sensitivity is not aligned with that of the designer. This is particularly important in the context of mechanism design where we seek to shape drivers' incentives and lead the system to a good equilibrium [Tumer and Agogino, 2006], and when drivers are subject to cognitive and behavioral biases such as risk-aversion [Ordóñez and Stier-Moses, 2010; Nikolova and Stier-Moses, 2015]. One important challenge is to extend the BPoA bounds to games where agents differ in their levels of risk aversion or tax sensitivity. This has been done to some extent in nonatomic games [Meir and Parkes, 2015a; 2015b].

More broadly, this work provides more evidence for the usefulness of "biased-smoothness" analysis, in the line of [Chen *et al.*, 2011; Meir and Parkes, 2015b], and we hope it can lead to a better understanding of routing games where agents are subject to either external influences (like tolls) or behavioral biases.

Acknowledgments

We thank Itai Arieli for pointing us to Sandholm's work.

References

- [Beckmann *et al.*, 1956] M. Beckmann, B. McGuire, and C. Winsten. *Studies in the Economics of Transportation*. Yale University Press, New Haven, 1956.
- [Bonifaci *et al.*, 2011] Vincenzo Bonifaci, Mahyar Salek, and Guido Schäfer. Efficiency of restricted tolls in non-atomic network routing games. In *SAGT'11*, pages 302–313, 2011.
- [Caragiannis *et al.*, 2010a] Ioannis Caragiannis, Christos Kaklamani, and Panagiotis Kanellopoulos. Taxes for linear atomic congestion games. *ACM Transactions on Algorithms*, 7(1):13, 2010.
- [Caragiannis *et al.*, 2010b] Ioannis Caragiannis, Christos Kaklamani, Panagiotis Kanellopoulos, Maria Kyropoulou, and Evi Papaioannou. The impact of altruism on the efficiency of atomic congestion games. In *TGC*, pages 172–188. Springer, 2010.
- [Chen *et al.*, 2011] Po-An Chen, Bart De Keijzer, David Kempe, and Guido Schäfer. The robust price of anarchy of altruistic games. In *WINE'11*, pages 383–390. Springer, 2011.
- [Christodoulou and Koutsoupias, 2005] George Christodoulou and Elias Koutsoupias. On the price of anarchy and stability of correlated equilibria of linear congestion games. In *Algorithms-ESA 2005*, pages 59–70. Springer, 2005.
- [Cole *et al.*, 2006] Richard Cole, Yevgeniy Dodis, and Tim Roughgarden. How much can taxes help selfish routing? *Journal of Computer and System Sciences*, 72(3):444–467, 2006.
- [Feldman *et al.*, 2015] Michal Feldman, Nicole Immorlica, Brendan Lucier, Tim Roughgarden, and Vasilis Syrgkanis. The price of anarchy in large games. *arXiv preprint arXiv:1503.04755*, 2015.
- [Fotakis and Spirakis, 2008] Dimitris Fotakis and Paul G Spirakis. Cost-balancing tolls for atomic network congestion games. *Internet Mathematics*, 5(4):343–363, 2008.
- [Fotakis *et al.*, 2010] Dimitris Fotakis, George Karakostas, and Stavros G Kolliopoulos. On the existence of optimal taxes for network congestion games with heterogeneous users. In *SAGT'10*, pages 162–173. Springer, 2010.
- [Karakostas and Kolliopoulos, 2004] George Karakostas and Stavros G Kolliopoulos. Edge pricing of multicommodity networks for heterogeneous selfish users. In *FOCS'04*, pages 268–276, 2004.
- [Meir and Parkes, 2015a] Reshef Meir and David Parkes. Congestion games with distance-based strict uncertainty. In *Proceedings of the 29th AAAI Conference on Artificial Intelligence*, 2015.
- [Meir and Parkes, 2015b] Reshef Meir and David Parkes. Playing the wrong game: Smoothness bounds for congestion games with behavioral biases. In *NETECON'15*, pages 67–70, 2015.
- [Nikolova and Stier-Moses, 2015] Evdokia Nikolova and Nicolas E Stier-Moses. The burden of risk aversion in mean-risk selfish routing. In *ACM-EC'15*, pages 489–506. ACM, 2015.
- [Ordóñez and Stier-Moses, 2010] Fernando Ordóñez and Nicolás E Stier-Moses. Wardrop equilibria with risk-averse users. *Transportation Science*, 44(1):63–86, 2010.
- [Pigou, 1920] Arthur Cecil Pigou. *The economics of welfare*. Palgrave Macmillan, 1920.
- [Rosenthal, 1973] Robert W Rosenthal. A class of games possessing pure-strategy Nash equilibria. *International Journal of Game Theory*, 2(1):65–67, 1973.
- [Roughgarden and Tardos, 2007] Tim Roughgarden and Eva Tardos. Introduction to the inefficiency of equilibria. In Noam Nisan *et al.*, editor, *Algorithmic Game Theory*, pages 443–459. Cambridge University Press, 2007.
- [Roughgarden, 2007] Tim Roughgarden. Routing games. In Noam Nisan *et al.*, editor, *Algorithmic game theory*, pages 459–484. Cambridge University Press, 2007.
- [Roughgarden, 2009] Tim Roughgarden. Intrinsic robustness of the price of anarchy. In *STOC'09*, pages 513–522. ACM, 2009.
- [Sandholm, 2007] William H Sandholm. Pigouvian pricing and stochastic evolutionary implementation. *Journal of Economic Theory*, 132(1):367–382, 2007.
- [Singh, 2008] Chandramani Singh. Marginal cost pricing for atomic network congestion games. Technical report, Department of Electrical Communication Engineering, Indian Institute of Science, 2008.
- [Swamy, 2012] Chaitanya Swamy. The effectiveness of stackelberg strategies and tolls for network congestion games. *ACM Transactions on Algorithms*, 8(4):36, 2012.
- [Tumer and Agogino, 2006] Kagan Tumer and Adrian Agogino. Agent reward shaping for alleviating traffic congestion. In *Workshop on Agents in Traffic and Transportation*. Citeseer, 2006.
- [Yang and Zhang, 2008] Hai Yang and Xiaoning Zhang. Existence of anonymous link tolls for system optimum on networks with mixed equilibrium behaviors. *Transportation Research Part B: Methodological*, 42(2):99–112, 2008.

A Omitted proofs

Theorem 2. *Suppose that (G, \hat{G}) is $(\hat{\lambda}, \hat{\mu})$ -BS. Let σ be any equilibrium (pure, mixed, correlated, or coarse-correlated) of the game \hat{G} . Then $SC(G, \sigma) \leq \frac{\hat{\lambda}}{1-\hat{\mu}} SC(G, OPT(G))$.*

Proof. For a correlated profile σ we denote $SC(G, \sigma) = E_{\mathbf{a} \sim \sigma}[SC(G, \mathbf{a})]$.

By Def. 1, there is a profile \mathbf{a}' s.t. Eq. (2) holds for every profile \mathbf{a} .

It is sufficient to prove for a CCE σ . By definition of CCE, for any $i \in N, b_i \in A_i$,

$$E_{\mathbf{a} \sim \sigma}[\hat{C}_i(\mathbf{a})] \leq E_{\mathbf{a} \sim \sigma}[\hat{C}_i(\mathbf{a}_{-i}, b_i)]. \quad (3)$$

$$SC(G, \sigma) = E_{\mathbf{a} \sim \sigma}[SC(G, \mathbf{a})] \leq E_{\mathbf{a} \sim \sigma}[SC(G, \mathbf{a}')] \quad (4)$$

$$\begin{aligned} &+ \left(\sum_{i=1}^n E_{\mathbf{a} \sim \sigma}[\hat{C}_i(\mathbf{a}_{-i}, a'_i)] - E_{\mathbf{a} \sim \sigma}[\hat{C}_i(\mathbf{a}')] \right) \\ &= E_{\mathbf{a} \sim \sigma} \left[SC(G, \mathbf{a}') + \sum_{i=1}^n \hat{C}_i(\mathbf{a}_{-i}, a'_i) - \hat{C}_i(\mathbf{a}') \right] \quad (5) \end{aligned}$$

$$\begin{aligned} &= E_{\mathbf{a} \sim \sigma} \left[\sum_{i=1}^n \left(C_i(\mathbf{a}') + \hat{C}_i(\mathbf{a}_{-i}, a'_i) - \hat{C}_i(\mathbf{a}') \right) \right] \\ &\leq E_{\mathbf{a} \sim \sigma} \left[\hat{\lambda} SC(G, OPT(G)) + \hat{\mu} SC(G, \mathbf{a}') \right] \quad (6) \end{aligned}$$

$$= \hat{\lambda} SC(G, OPT(G)) + \hat{\mu} SC(G, \sigma), \quad (7)$$

where Inequality (4) follows from Eq. (3) with $b_i = a'_i$, (5)+(7) from linearity of expectation, and (6) from Eq. (2) applied for each \mathbf{a} . By rearranging terms, we get the bound in the theorem. \square