Playing the Wrong Game

Citation

Meir, Reshef, and David Parkes. 2015. "Playing the Wrong Game." SIGMETRICS Perform. Eval. Rev. 43 (3) (November 19): 67–70. doi:10.1145/2847220.2847242.

Published Version

doi:10.1145/2847220.2847242

Permanent link

<http://nrs.harvard.edu/urn-3:HUL.InstRepos:30403700>

Terms of Use

This article was downloaded from Harvard University's DASH repository, and is made available under the terms and conditions applicable to Open Access Policy Articles, as set forth at <http://nrs.harvard.edu/urn-3:HUL.InstRepos:dash.current.terms-of-use#OAP>

Share Your Story

The Harvard community has made this article openly available.
Please share how this access benefits you. [Submit a story](#).

[Accessibility](#)

Playing the Wrong Game: Smoothness Bounds for Congestion Games with Behavioral Biases

Reshef Meir and David Parkes

ABSTRACT

In many situations a player may act so as to maximize a perceived utility that is not exactly her utility function, but rather some other, biased, utility. Examples of such biased utility functions are common in behavioral economics, and include risk attitudes, altruism, present-bias and so on. When analyzing a game, one may ask how inefficiency, measured by the Price of Anarchy (PoA) is affected by the perceived utilities.

The smoothness method [16, 15] naturally extends to games with such perceived utilities or costs, regardless of the game or the behavioral bias. We show that such biased-smoothness is broadly applicable in the context of nonatomic congestion games. First, we show that on series-parallel networks we can use smoothness to yield PoA bounds even for diverse populations with different biases. Second, we identify various classes of cost functions and biases that are smooth, thereby substantially improving some recent results from the literature.

1. INTRODUCTION

Game theory is founded on the assumption that players are rational decision makers, i.e. maximizing their utility, and that groups of agents reach an equilibrium outcome. However agents, either human or automated, often have bounded resources that prevent them from finding the optimal response in every situation. Further, human decision makers are prone to various cognitive and behavioral biases, such as risk-aversion, loss-aversion, tendency to focus on short-term utility (present-bias) and so on.

As a concrete example, commuters may have some information on the expected congestion at each route via traffic reports or a cellphone app. However they also know that this information is inaccurate, and a risk-averse driver might try to minimize some combination of the expected latency and the variance rather than the expectation alone. An even more heuristic pessimistic approach (e.g. if the driver does not know the distribution of latency), is to simply add a “safety margin” of 20% to the congestion at each road. A less pessimistic commuter may only add 5% or 10%.

The implications of these limitations and biases on game playing can be similarly described: in essence, the players are playing the “wrong game,” by either applying some simple heuristics, or optimizing a different utility function from

the one in the game specification [6]. Moreover, different agents may have different perspective on the game.

Fortunately, as long as we can incorporate the biases that affect agents into suitably modified utility (or cost) functions, we can still use traditional equilibrium concepts. More formally, suppose that in the real underlying game G each agent i has some utility function u_i . Now, each player i sees her utility as some other function $\hat{u}_i$, and thus we are interested in the equilibria of the *biased game* $\hat{G}$ comprised of modified utilities $\{\hat{u}_i\}$.

Our starting point in this paper is the smoothness framework [15] in nonatomic congestion games. The smoothness framework connects a particular property of the edge cost functions (say, $(1, \frac{1}{4})$ -smoothness of affine functions), with a tight upper bound on the price of anarchy (PoA), this bound also being independent of the network topology [13].

Smoothness for modified costs has also been considered; e.g. by Bonifaci et al. [3] in the context of taxes that can also be viewed as a perturbation on utilities. The definition of smoothness naturally extends to account for bias, and the modified definition (which contains both c and $\hat{c}$) guarantees a PoA bound in a similar manner to standard smoothness PoA analysis. Moreover, it is easy to see that while [3] considered a particular modified cost function, the same biased-smoothness approach works for any combination of cost function c and bias $\hat{c}$. PoA bounds have been similarly attained via smoothness for games with altruistic players [4].¹

Thus the PoA is not only robust to different notions of equilibrium, but also to players that are not playing exactly by the game specification. However this brings about two important challenges: First, it is unlikely that all of the participants in a game have exactly the same bias (see example on risk-aversion). Thus we want to combine smoothness results of several different types to attain a single PoA bound. Second, can we find broad classes of games and behavioral biases for which smoothness applies?

In this extended abstract we provide positive results in both of the above directions. We first show (in Section 3) that under certain conditions on the network structure and the cost functions, the PoA is close to the average of the PoA of all participating agents, weighted by their frequency. In contrast, without such a structural restriction, even a small fraction of “bad” agents may inflict unbounded damage on the society. Our second technical contribution is to show how biased-smoothness can be proved for several classes of

¹For more related literature, see the full version of this paper [9].

cost functions (general, convex, affine and quadratic) with particular behavioral biases, namely different levels of tax sensitivity, pessimism [8], or risk-aversion [12, 11]. We do this in Sections 4.1, 4.2 and 4.3.

Our results significantly improve the upper bounds of Meir and Parkes [8] and Nikolova and Stier-Moses [11], while using much simpler proofs.

2. PRELIMINARIES

Nonatomic routing games. Following the definitions of Roughgarden and Tardos [16], a *nonatomic routing game* (NRG) is a tuple $G = \langle V, E, m, \mathbf{c}, \mathbf{u}, \mathbf{v}, \mathbf{n} \rangle$, where

- (V, E) is a directed graph;
- $m \in \mathbb{N}$ is the number of agent types;
- $\mathbf{c} = (c_e)_{e \in E}$, $c_e(x) \geq 0$ is the cost (or *latency*) incurred when x agents use edge e ;
- $\mathbf{u}, \mathbf{v} \in V^m$, where (u_i, v_i) are the source and target nodes of type i agents;
- $\mathbf{n} \in \mathbb{R}_+^m$, where $n_i \in \mathbb{R}_+$ is the total mass of type i agents. The total mass of all agents is $\sum_{i \leq m} n_i = n$.

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 agents of type i . We denote by $A = \cup_i A_i$ the set of all directed source-target paths. We assume that the costs c_e are non-decreasing, continuous, differentiable and semi-convex (i.e., $xc_e(x)$ is convex). Such cost functions are called *standard* [13]. An example of standard cost functions are *affine costs* where $c(x) = ax + b$ for some $a, b \in \mathbb{R}_+$.

A *nonatomic routing game with player-specific costs* is a tuple $G = \langle V, E, m, (\hat{\mathbf{c}}^i)_{i \leq m}, \mathbf{u}, \mathbf{v}, \mathbf{n} \rangle$. It is essentially the same as the NRG definition above, except that agents of each type i experience a cost of $\hat{c}_e^i(x)$.

A *state* $\mathbf{s}$ describes the paths selected by all agents. Due to space constraints we omit the formal definitions, but we can think of state $\mathbf{s}$ as a vector, where s_e is the total traffic through edge $e \in E$.

The *social cost* in a profile $\mathbf{s}$ in game G is $SC(G, \mathbf{s}) = \sum_{e \in E} s_e c_e(s_e)$.

A state $\mathbf{s}$ for an NRG is an *equilibrium* in game G if no agent can gain by selecting a different path. It is well known that all equilibria (including correlated equilibria) have the same social cost; and in every equilibrium all agents of type i experience the same cost [10, 16, 2].

Optimal taxation. Let G^* be a modification of G , where we replace every $c_e(x)$ with $c_e^*(x) = c_e(x) + xc_e'(x)$ ($c'(x)$ is a shorthand for $\frac{\partial c(x)}{\partial x}$). Then, $SC(G, \mathbf{s})$ becomes a *potential function* of G , and thus every equilibrium has minimum social cost (see [1, 14]). We can think of c^* as the original cost plus an *optimal tax*, s.t. the incentives of all agents are completely aligned with those of the society.

The price of anarchy. Let $EQ(G)$ be the set of equilibria in game G . The *price of anarchy* (PoA) of a game is the ratio between the social cost in the worst equilibrium in $EQ(G)$ and the optimal social cost [7]. Since all equilibria have the same cost, we can write $\text{PoA}(G) = \frac{SC(\mathbf{s}^*)}{SC(G, \text{OPT}(G))}$, where $\mathbf{s}^*$

is any equilibrium of G . E.g. in affine NRGs, it is known that $\text{PoA}(G) \leq \frac{4}{3}$, and this bound tight [16].

Smoothness. A cost function c is (λ, μ) -smooth for $\lambda \geq 0, \mu < 1$ if for any $x, x' \geq 0$, it holds that

$$c(x)x' \leq \lambda c(x')x' + \mu c(x)x. \quad (1)$$

A game G is (λ, μ) -smooth if all cost functions in G are (λ, μ) -smooth. For any (λ, μ) -smooth game, $\text{PoA}(G) \leq \frac{\lambda}{1-\mu}$ [5, 15]. Moreover, w.l.o.g. $\lambda = 1$ (that is, for any class of cost functions there is an optimal pair $(1, \mu)$ for some μ). For example, affine routing games are $(1, \frac{1}{4})$ -smooth [15], providing an alternative derivation of the PoA bound of $\frac{4}{3}$.

2.1 Introducing Biased Costs

Given a cost function c , we denote by $\hat{c}$ a modified cost function. For a class C , we denote by $\hat{C}$ where the modification is determined according to a mapping from cost functions to cost functions. Given a game G and modified cost functions $\hat{C}$, we get a new game $\hat{G}$, called the *biased game*. This modified game is identical to G , except any cost function c_e is replaced with $\hat{c}_e$, generated according to a mapping.

DEFINITION 1. A nonatomic routing game with biased costs is a tuple $\mathcal{G} = \langle G, (\hat{C}^i)_{i \leq m} \rangle$, where all players of type i play as if the game is $\hat{G}^i = \langle V, E, m, \hat{\mathbf{c}}^i, \mathbf{u}, \mathbf{v}, \mathbf{n} \rangle$.

We denote the (player-specific) NRG where each type i has costs $\hat{\mathbf{c}}^i = (\hat{c}_e^i)_e$ by $\hat{G}$. Note that at least in the simple case where $\hat{\mathbf{c}}^i = \hat{\mathbf{c}}$ for all i (all agents have the same bias), the game $\hat{G}$ is just another NRG. Thus $\hat{G}$ and $\mathcal{G}$ have a potential function with a unique locally minimal value.

One way to interpret $\mathcal{G}$ is that players play the wrong game $\hat{G}$, whereas their true costs are according to the underlying game G .

Biased Price of Anarchy. Given the above, we measure the price of anarchy in a game with biased costs by comparing the equilibria of $\hat{G}$ to the optimum of G . Formally:

$$\text{PoA}(G, (\hat{C}^i)_{i \leq m}) = \sup_{\mathbf{s} \in EQ(\hat{G})} \frac{SC(G, \mathbf{s})}{SC(G, \text{OPT}(G))}. \quad (2)$$

An equilibrium of $\hat{G}$, and thus of $\mathcal{G}$, exists even with diverse types, although it may not be a potential game. This follows from general existence results on nonatomic games with convex strategy spaces and continuous utilities [17].

2.2 Smoothness for Biased Costs

Our goal is to provide bounds on the biased Price of Anarchy for a given game $\langle G, (\hat{C}^i)_i \rangle$. The fact that each of $\hat{C}^i$ and $\hat{C}^i$ are smooth is insufficient to provide such a bound, since $\hat{c}^i$ and c may be completely unrelated.

In the remainder of this section we extend the definition of smoothness to games with biased costs, in a way that takes into account both c and $\hat{c}$. This technique is not new and has been applied before for specific modified costs, for example nonatomic games with restricted taxes [3] and atomic games with altruistic players [4]. The extension to general biases is essentially the same, but is given here for completeness, and since it is required for our new results.

DEFINITION 2. The class of functions C is (λ, μ) -biased-smooth (w.r.t. biased cost function $\hat{c}$) if for any cost function $c \in C$ and any $x, x' \in \mathbb{R}_+$,

$$c(x)x + \hat{c}(x)(x' - x) \leq \lambda c(x')x' + \mu c(x)x. \quad (3)$$

As a first sanity check, we test how biased smoothness behaves in the trivial case where there is no bias. Indeed, if $\hat{c} = c$ for all $c \in C$, and C is (λ, μ) -smooth, then:

$$c(x)x + \hat{c}(x)(x' - x) = c(x)x' \leq \lambda c(x')x' + \mu c(x)x.$$

That is, Eq. (3) collapses to the standard definition of smoothness from Eq. (1).

Recall that the PoA of a (λ, μ) -smooth game is bounded by $\frac{\lambda}{1-\mu}$. A key observation is that this bound extends to games with biased costs when all agents have the same bias.

PROPOSITION 1. Consider a game G where all cost functions are from C , and C is $(\hat{\lambda}, \hat{\mu})$ -biased smooth w.r.t. biased costs $\hat{C}$. Let $\mathbf{s}$ be any equilibrium of the game $\hat{G}$, and $\mathbf{s}^*$ any valid state.² Then $SC(G, \mathbf{s}) \leq \frac{\hat{\lambda}}{1-\hat{\mu}} SC(G, \mathbf{s}^*)$.

Proposition 1 was proved by Bonifaci et al. [3] for a particular modified cost (truncated taxes), but the proof works the same for any pair of c and $\hat{c}$. In fact, the proof is a minor variation of the standard smoothness argument, e.g. from [5, 15]. A price of anarchy bound of $\frac{\hat{\lambda}}{1-\hat{\mu}}$ follows as an immediate corollary.

An alternative derivation of optimal taxation. Given the above, we can now do a second sanity check of the concept of biased smoothness. It is known that the modified costs $c^*(x) = c(x) + c'(x)x$ should lead to an optimal play [1]. We want to get this result via a biased-smoothness argument.

LEMMA 1. Set $\hat{c}(x) = c^*(x) = c(x) + c'(x)x$, then c is $(1, 0)$ -biased smooth w.r.t. $\hat{c}$.

PROOF. Let $h(x) = xc(x)$. By convexity of h ,

$$c(x)x + \hat{c}(x)(x' - x) = h(x) + h'(x)(x' - x) \leq h(x') = x'c(x'),$$

affirming that the price of anarchy of $\langle G, C^* \rangle$ is 1. $\square$

3. DIVERSE POPULATION

Consider a game $\mathcal{G} = \langle G, (\hat{C}^i)_i \rangle$, where the modified cost function is obtained for each agent type via a distinct, type-specific mapping. Suppose we can show biased smoothness of $(\hat{\lambda}^i, \hat{\mu}^i)$ for each of the uniform games $\langle G, \hat{C}^i \rangle$ (i.e., games that are identical to $\mathcal{G}$, except that all agents have the same bias). The primary question is whether we can get a bound on $\text{PoA}(\mathcal{G})$ in terms of $(\hat{\lambda}^i, \hat{\mu}^i)_{i \leq m}$, that is, the smoothness parameters for each type.

Unfortunately, without further restrictions on the game, attaining such a bound is impossible, as even a small fraction of “bad” agents (with poor biased-smoothness guarantees) may significantly increase the PoA (see full version of this paper for details).

We circumvent this difficulty by considering networks that are series-parallel. A *directed series-parallel (DSP) graph* is an acyclic directed graph (V, E, u, v) with a source u and a target v , that is composed recursively by merging two DSPs

²We use $\hat{\lambda}, \hat{\mu}$ to specify the biased smoothness parameters, and λ, μ for “standard” smoothness.

C	general		convex
β	$\beta \leq 1$	$\beta \geq 1$	$\beta \geq 1$
$\hat{\lambda}$	1	β	$1 + (\beta - 1)\mu$
$\hat{\mu}$	$(1 - \beta)\mu$	0	0
PoA upper bound	$\frac{1}{1 - (1 - \beta)\mu}$	β	$1 + (\beta - 1)\mu$

Table 1: Biased-smoothness bounds under tax-sensitivity for cost functions that are $(1, \mu)$ -smooth.

C	affine		quadratic	
β	$\beta \leq 1$	$\beta \geq 1$	$\beta \leq 1$	$\beta \geq 1$
$\hat{\lambda}$	1	$\frac{(1+\beta)^2}{4\beta}$	1	$1 + \frac{2\beta-2}{3\sqrt{3}}$
$\hat{\mu}$	$\frac{(\beta+1)^2}{4} - \beta$	0	$\frac{2\beta'^3}{3\sqrt{3}} - \beta'^2 + 1$	0
PoA	$\frac{1}{(\beta+1) - \frac{1}{4}(\beta+1)^2}$	$\frac{(1+\beta)^2}{4\beta}$	$\frac{1}{\beta'^2 - \frac{2\beta'^3}{3\sqrt{3}}}$	$\frac{(1+2\beta)^3}{27\beta^2}$

Table 2: Biased-smoothness bounds for tax-sensitivity, affine and quadratic cost functions ($\beta' = \sqrt{1 + 2\beta}$). All bounds are tight.

either in a serial manner (source to target) or in a parallel manner (merge sources and merge targets). The basic graph is a single directed edge $u - v$.

Our main result is the following bound. For a proof and more details see the full version.

THEOREM 1. Suppose that $\mathcal{G} = \langle G, (\hat{C}^i)_{i \leq m} \rangle$ is a game over a DSP network with convex cost functions. Then $\text{PoA}(\mathcal{G}) = \sum_i \frac{n_i}{n} \frac{\lambda^i \hat{\lambda}^i}{(1-2\mu)(1-\mu^i)(1-\hat{\mu}^i)} = O(\sum_i \frac{n_i}{n} \frac{\lambda^i}{1-\hat{\mu}^i})$.

In the above theorem, the class C is (λ, μ) -smooth, each class C^i is (λ^i, μ^i) -biased smooth, and C is $(\hat{\lambda}^i, \hat{\mu}^i)$ -biased-smooth w.r.t. each $\hat{C}^i$. Recall that $\frac{n_i}{n}$ is the fraction of agents of type i .

To get some feel of this bound for specific classes, suppose for example that both of C and $\hat{C}^i$ are affine (as in Section 4.2). We get $\text{PoA}(\mathcal{G}) \leq \frac{8}{3} \sum_i \frac{n_i}{n} \frac{\lambda^i}{(1-\hat{\mu}^i)}$.

4. SMOOTHNESS FOR COMMON BIASES

In this section, we assume that all agents have the same bias, but make no assumptions on the network.

4.1 Tax-Sensitive Agents

Suppose that the center imposes the theoretically optimal tax of $c'_e(x)x$ on every edge e . However for an agent with tax sensitivity $\beta \geq 0$, $\hat{c}^\beta(x) = c(x) + \beta c'(x)x$, and the monetary part of the cost is adjusted by a factor of β (say, due to different wealth). Due to space constraints, we omit the full propositions. Rather, we provide the smoothness parameters and PoA bounds as a function of β for various classes of cost functions in a tabular form (Tables 1 and 2).

4.2 Pessimist Agents

Suppose that agents are pessimistic, in the sense that they play according to a congestion amount that is larger by a factor of $r > 1$ than the true congestion [8], that is, $\hat{c}^r(x) = c(rx)$. For affine cost functions, it holds that

$$\hat{c}^r(x) = rax + b = ax + b + (r-1)ax = c(x) + (r-1)c'(x)x,$$

and pessimism with factor $r \geq 1$ coincides with tax-sensitivity of $\beta = r - 1$. As a result, we can immediately apply the

upper bounds for tax-sensitive agents to pessimistic agents with affine cost functions, thereby improving the previous upper bounds from [8]:

The double red lines are the tight bounds on the PoA in NRGs with affine costs, for agents with tax-sensitivity or pessimism (from Table 2). For reference, we also show the upper bounds derived in Table 1 for arbitrary cost functions (in light gray); using convexity only (thick green line). The upper bound from Meir and Parkes [2015] is in dashed line.

4.3 Risk-averse agents in the Mean-Var Model

Following [11], consider an arbitrary cost function $c(x)$ and an arbitrary distribution $\epsilon(x)$ s.t. $\text{var}(\epsilon(x)) \leq \kappa c(x)$ for all $x > 0$. Denote $v(x) = \text{var}(\epsilon(x)) \geq 0$, and suppose that $c(x)$ is $(1, \mu)$ -smooth for some $\mu < 1$. Define the biased cost as $\hat{c}^\gamma(x) = c(x) + \gamma v(x)$.

Nikolova and Stier-Moses bounded the “Price of Risk Aversion” (PRA), which is the ratio between the social welfare in the biased equilibrium and that in the non-biased equilibrium. Their main result is that the PRA is upper bounded by $1 + \kappa\gamma\eta$, where η is a parameter that depends on the network and may be as large as the number of vertices. In particular this leads to a bound of $\text{PoA}(\mathcal{G}) \leq (1 + \gamma\kappa\eta)^{\frac{1}{1-\mu}}$ for any class of $(1, \mu)$ cost functions under risk aversion.

We show the following:

PROPOSITION 2. *For $\gamma \geq 0$, $c(x)$ is $(1 + \gamma\kappa, \mu)$ -biased-smooth w.r.t. $\hat{c}^\gamma(x)$. Thus,*

$$\text{PoA}(\mathcal{G}) \leq (1 + \gamma\kappa)^{\frac{1}{1-\mu}}.$$

Our result improves not just the PoA bound that follows from [11], but also their PRA bound: Since the unbiased equilibrium is at least as costly as the optimal state, it follows that the PRA is at most the (biased) PoA, and is thus upper bounded by $(1 + \gamma\kappa)^{\frac{1}{1-\mu}}$. Crucially, μ is fixed and *does not depend* on the network structure. Thus our results show the factor η is in fact redundant when considering either the PoA or the PRA.

5. DISCUSSION

We have considered strategic settings in which participants are playing the wrong game, and perceiving utilities in some inaccurate or biased form. Whether these modified utilities, and thus deviations from rational play, come from a cognitive limitation, a behavioral bias, or a different perception in regard to taxes or other payments, it is important to understand how the equilibria of the game are affected.

Biased smoothness is a tool that enables PoA analysis under such modified utilities. Our work is the first to provide PoA bounds on populations with arbitrary biases, diverse

across agents, in nonatomic congestion games. The analysis framework reduces the problem to that of analyzing the smoothness parameters of each payoff and behavioral type of agent. Perhaps surprisingly, some biases (e.g. moderate pessimism) leads to better equilibria and lower PoA than under no bias at all. The fact that for specific biases we get tight PoA bounds (Sections 4.1, 4.2), and significantly improve bounds derived in the traditional way (Section 4.3) can be seen as an indication that our definition of biased smoothness (borrowed from Bonifaci et al. [3]) is the “correct” one.

We emphasize that all of our PoA bounds in Section 4 hold regardless of the network topology. If we further assume that the network is series-parallel, then we can combine these bounds for assorted populations using Theorem 1.

6. REFERENCES

- [1] M. Beckmann, B. McGuire, and C. Winsten. *Studies in the Economics of Transportation*. Yale University Press, New Haven, 1956.
- [2] A. Blum, E. Even-Dar, and K. Ligett. Routing without regret: On convergence to nash equilibria of regret-minimizing algorithms in routing games. In *PODC’06*, pages 45–52. ACM, 2006.
- [3] V. Bonifaci, M. Salek, and G. Schäfer. Efficiency of restricted tolls in non-atomic network routing games. *Algorithmic Game Theory*, pages 302–313, 2011.
- [4] P.-A. Chen, B. De Keijzer, D. Kempe, and G. Schäfer. The robust price of anarchy of altruistic games. In *WINE’11*, pages 383–390. Springer, 2011.
- [5] J. R. Correa, A. S. Schulz, and N. E. Stier-Moses. A geometric approach to the price of anarchy in nonatomic congestion games. *Games and Economic Behavior*, 64(2):457–469, 2008.
- [6] G. Devetag and M. Warglien. Playing the wrong game: An experimental analysis of relational complexity and strategic misrepresentation. *Games and Economic Behavior*, 62(2):364 – 382, 2008.
- [7] E. Koutsoupias and C. Papadimitriou. Worst-case equilibria. *Computer science review*, 3(2):65–69, 2009.
- [8] R. Meir and D. Parkes. Congestion games with distance-based strict uncertainty. In *AAAI’15*, 2015.
- [9] R. Meir and D. Parkes. Playing the wrong game: Smoothness bounds for congestion games with behavioral biases. *arXiv preprint arXiv:1411.1751*, 2015.
- [10] I. Milchtaich. Generic uniqueness of equilibrium in large crowding games. *Mathematics of Operations Research*, 25(3):349–364, 2000.
- [11] E. Nikolova and N. Stier-Moses. The burden of risk aversion in mean-risk selfish routing. In *EC’15*, 2015. To appear.
- [12] E. Nikolova and N. E. Stier-Moses. A mean-risk model for the traffic assignment problem with stochastic travel times. *Operations Research*, 62(2):366–382, 2014.
- [13] T. Roughgarden. The price of anarchy is independent of the network topology. *Journal of Computer and System Sciences*, 67(2):341–364, 2003.
- [14] T. Roughgarden. Routing games. In N. et al., editor, *Algorithmic game theory*. Cambridge University Press New York, 2007.
- [15] T. Roughgarden. Intrinsic robustness of the price of anarchy. In *STOC’09*, pages 513–522. ACM, 2009.
- [16] T. Roughgarden and É. Tardos. Bounding the inefficiency of equilibria in nonatomic congestion games. *Games and Economic Behavior*, 47(2):389–403, 2004.
- [17] D. Schmeidler. Equilibrium points of nonatomic games. *Journal of Statistical Physics*, 7(4):295–300, 1973.