

Playing the Wrong Game: Bounding Externalities in Diverse Populations of Agents

Reshef Meir

Technion—Israel Institute of Technology
Haifa, Israel

David Parkes

Harvard University
Cambridge, MA, USA

ABSTRACT

The robustness of multiagent systems can be affected by mistakes or behavioral biases (e.g., risk-aversion, altruism, toll-sensitivity), with some agents playing the “wrong game.” This can change the set of equilibria, and may in turn harm or improve the social welfare of agents in the system. We are interested in bounding what we call the *biased price of anarchy* (BPoA) in populations with diverse agent behaviors, which is the ratio between welfare in the “wrong” equilibrium and optimal welfare. We study nonatomic routing games, and derive an externality bound that depends on a key topological parameter of the underlying network. We then prove two general BPoA bounds for games with diverse populations: one that relies on the network structure and the *average bias* of all agents in the population, and one that is independent of the structure but depends on the *maximal bias*. Both types of bounds can be combined with known results to derive concrete BPoA bounds for a variety of specific behaviors (e.g., varied levels of risk-aversion).

KEYWORDS

Price of Anarchy; Behavioral Game Theory; Selfish Routing

ACM Reference Format:

Reshef Meir and David Parkes. 2018. Playing the Wrong Game: Bounding Externalities in Diverse Populations of Agents. In *Proc. of the 17th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2018)*, Stockholm, Sweden, July 10–15, 2018, IFAAMAS, 9 pages.

1 INTRODUCTION

Game theory is founded on the assumption that agents are rational decision makers, i.e. maximizing their utility, and that groups of agents reach an equilibrium outcome. In many games there is some objective measure of welfare that can be accurately measured in terms of money, time, and so on. Utility on the other hand, is subjective. It is not always easy to identify an explicit utility function that an agent is trying to maximize, if such exists. Even when there are objective factors that affect agents’ utility, such as profit, effort, uncertainty, and temporal delays, various agents may weigh these factors differently or have subjective constraints and biases. Thus different agents may demonstrate different behavior even in the same situation.

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 take into account not just the expected congestion, but also the likelihood of an unexpected

delay, a heuristic safety margin and so on. Moreover, different commuters may have different levels of risk-aversion, or act upon different heuristics.

The implications of these subjective differences and biases on a multiagent system are two-fold. First, from the perspective of an outside observer who cares about a particular objective (say, total latency), the agents are playing the “wrong game.” This is true whether they are failing to optimize their true cost due to cognitive limitations and biases [18], or optimizing their true subjective costs that differ from those of the “objective” observer.

A second related issue, is that the behavior of one agent may exert significant negative externality on the utility of another. The extreme case is when some agents are adversarial, and act in a way that tries to *minimize* the utility of some other agents.

It is well known that even in games without biases or subjective utilities, there may be negative externalities, and that equilibrium outcomes may be suboptimal in terms of the utilitarian social welfare. This inefficiency is often quantified as the *Price of Anarchy* (PoA), which is the ratio between the social welfare at the worst equilibrium and the optimal social welfare. Biases may change the equilibria of the game, and thus have a large effect on the PoA. Externalities are widely studied in multiagent systems, but usually in the context of well-defined behavior such as maximizing utility in a game [28, 34] or finding an optimal or stable solution [6, 10, 11]. How can we hope to bound externalities in a diverse society of agents with assorted biases and behaviors? The answer will lie in a proper abstraction of such behaviors, but before that we will present our questions more formally.

Research goals. Suppose that in game G , each agent i has some true cost function c^i (in this paper we consider negative utilities, i.e. costs). Now, each agent i sees her cost as some other function \hat{c}^i , and thus we are interested in the equilibria of the *biased game* \hat{G} comprised of modified utilities or costs $\{\hat{c}^i\}_{i=1,2,\dots}$. We ask the following:

- Is there a bound to the negative externality that type j exerts on type i , that applies *regardless* of the subjective cost \hat{c}^j ? (Example: can we bound the equilibrium latency of type i commuters without any assumption on the risk-aversion level of others?)
- Is there a bound to the social cost (in the “objective game” G) of \hat{G} ’s worst equilibrium? (Example: can we bound the total latency in equilibrium, given all commuters’ subjective cost functions?)

We refer to the ratio between the latter measure and the optimal social cost in G as the *Biased Price of Anarchy* (BPoA), and note that it coincides with the PoA when $\hat{G} = G$.

While the first question is straight-forward, the second question may raise some conceptual debate. One may argue that since agents are acting so as to minimize their biased cost \hat{c}^i rather than

Proc. of the 17th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2018), M. Dastani, G. Sukthankar, E. André, S. Koenig (eds.), July 10–15, 2018, Stockholm, Sweden. © 2018 International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). All rights reserved.

c^i , this is the cost we should take into account when computing the social cost or welfare. We justify using G as the baseline for welfare as follows. First, the social cost may be the objective used by the system designer or analyst, as in our example above. A different analyst might care about a different goal, resulting in a different objective game G . Second, the social cost may be the sum of the agents' true utilities, while the agents are bounded-rational. For example they may be unaware of some roads [2] or the exact latency functions [41]; they may have uncertainty regarding actual congestion [42] or the overall amount of other agents [4]; and may assign wrong probabilities to rare outcomes [23].

Smoothness in routing games. Nonatomic routing games in the Wardrop model [43] are a good testbed for the questions above: they have very convenient theoretical properties, such as the uniqueness of equilibrium (up to identical utilities); equilibrium inefficiency and in particular the PoA is very well understood; and several biases have already been suggested and studied in the context of routing games (see Related Work). The *smoothness method* allows to leverage a property of the edge cost functions to obtain a tight upper bound on the Price of Anarchy that is independent of the network topology: if all cost functions are (λ, μ) -smooth for some parameters λ, μ , then the PoA is bounded by $\frac{\lambda}{1-\mu}$ [17, 36, 37]. Thus the smoothness of cost functions lets us *abstract away* the details of the game and prove PoA bounds on large classes of games.

1.1 Paper structure and contribution

After a short overview of nonatomic routing games, price of anarchy, and a measure of graph complexity called *serial-parallel width* (Section 2), we prove a tight bound on the negative externality in any routing game without any assumption on agents' behavior (Section 3). Our bound generalizes previous results from specific behaviors on series-parallel networks [2, 33] to arbitrary behaviors and networks (parametrized by their serial-parallel width).

Next, we adopt *smoothness* as an abstraction for general biases and behaviors. In Section 4 we extend the definition of a smooth cost function to $(\hat{\lambda}, \hat{\mu})$ -biased-smoothness that takes into account both the "true" and the "biased" cost functions. This approach follows similar extensions for specific, modified costs [9, 12], and we review recent smoothness bounds for such specific settings.

In Section 5, we consider games where agents have diverse biases, and use our results from Section 3 along with biased-smoothness to derive several bounds on the Biased Price of Anarchy. For example, for symmetric games over series-parallel networks (which have parallel-width of 1), we derive a BPoA bound in terms of the "average" biased-smoothness: $\text{BPoA} \leq O(1) \sum_i \frac{r_i}{r} \frac{\hat{\lambda}^i}{1-\hat{\mu}^i}$, where $\frac{r_i}{r}$ is the fraction of type i agents in the population. For arbitrary networks, we get a bi-criteria result that depends both on the average bias as above, and on the serial-parallel width of the underlying network.

For the special case of polynomial cost functions, we leverage (Section 6) known results on the PoA in heterogeneous unbiased games to derive a structure-independent BPoA bound. In contrast to our main results, this bound depends on the *worst* bias rather than on the average bias.

Omitted proofs, additional results and further discussion are available in the full version of this paper (arxiv.org/abs/1411.1751).

1.2 Related Work

The most-well studied, modified cost function comes about as a result of tolls, where the travel time plus the imposed toll can be thought of as a modified cost function. In this context, most papers focus on the objective of minimizing total latency [14, 20, 26], and on the design of optimal or practical toll schemes [9, 21, 25, 40]. Heterogeneous biases arise when different agents have different sensitivity to imposed tolls. We explain in the relevant sections how these papers technically relate to our work.

Chen et al. [12, 13] apply smoothness analysis to provide BPoA bounds for various games (including atomic congestion games) where agents are altruistic, i.e., part of their utility is derived from the social welfare. In the context of nonatomic routing games, their model is formally equivalent to toll-sensitivity (see Section 4.2).

Acemoglu et al. [2] study nonatomic congestion games where some agents are unaware of the existence of certain edges, which is equivalent to having a wrong cost function that assigns infinite costs to some edges. They prove that on directed series-parallel networks, such ignorance can only lead to a worse equilibrium than under true information, yet the *worst-case* PoA remains the same.

Another behavioral bias that has been studied in congestion games is risk aversion [3, 4, 30, 33, 35], which can often be written as biased cost functions (see Section 4.2 for details).


Closest to our work, a recent paper by Cole et al. [15] considers topological conditions on the network such that diversity will always weakly reduce the social cost. This is guaranteed when the network is series-parallel, which is highly related to a special case of our main theorem.

Finally, Babaioff et al. [5] consider congestion games where some of the agents are malicious. Babaioff et al. focused on the effect of a negligible amount of malicious agents, and showed examples where it can be either detrimental or (surprisingly) beneficial to the other agents, but without any upper bounds. Indeed, malicious behavior can be considered as another form of bias.

2 PRELIMINARIES

A *network* is a 2-terminal directed multigraph $G = \langle V, E, s, t \rangle$, where $s, t \in V$ are special vertices (*source* and *target*), and every edge $e \in E$ belongs to some simple $s - t$ path.

A network is *series-parallel* [19, 24] if it is either a single edge, or composed recursively by joining two series-parallel networks in series or in parallel. E.g., merging $\{s_1, s_2\}$ and $\{t_1, t_2\}$ in the following networks also results in a series-parallel network:



2.1 Nonatomic routing games

Following the definitions of Roughgarden [36] and Roughgarden and Tardos [38], a *nonatomic routing game* (NRG) is a tuple $G = \langle G, m, c, s, t, r \rangle$, where

- $G = \langle V, E \rangle$ is a directed multigraph;
- $m \in \mathbb{N}$ is the number of agent types;
- $c = (c_e)_{e \in E}$, where $c_e(x) \geq 0$ is the cost incurred when x agents use edge e ;
- $s, t \in V^m$, where (s_i, t_i) are the source and target nodes of type i agents;

- $\mathbf{r} = (r_i)_{i \leq m}$, where $r_i > 0$ is the total mass of type i agents. The total mass of agents of all types is $\sum_{i \leq m} r_i = r$, and we assume unless specified otherwise that $r = 1$.

We denote by $A_i \subseteq 2^E$ the set of all directed simple paths between the pair of nodes (s_i, t_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 simple paths. We assume that all cost functions mentioned in the paper (including biased costs mentioned later on) are non-decreasing, continuous, differentiable and semi-convex (i.e., $xc_e(x)$ is convex). Such cost functions are called *standard* [36].

Player-specific costs. A *nonatomic routing game with player-specific costs* (PNRG) is a tuple $\mathbf{G} = \langle G, m, (c^i)_{i \leq m}, s, t, \mathbf{r} \rangle$. The difference from a NRG is that agents of each type i experience a cost of $c_e^i(x)$ when x agents use edge e . We can have multiple types with the same source and target nodes to allow diversity of behavior. To avoid confusion, we refer to (s_i, t_i) (or A_i) as the *demand type* and to c^i as the *cost type*. Thus the type i specifies both the demand type and the cost type.

A PNRG is *symmetric* if all agents have the same demand type, i.e., $A_i = A$ for all i . A PNRG is a *resource selection game* (RSG) if G is a network of parallel links. That is, if the action of every agent is to select a single (s, t) edge.

Flows. A *flow* (or *action profile*) of a PNRG is a vector $\mathbf{f} \in \mathbb{R}_+^{|A| \times m}$, where $f_{p,i}$ is the amount of agents of type i that use path $p \in A_i$. In a valid flow, $\sum_{p \in A_i} f_{p,i} = r_i$ for all i . The total traffic on path $p \in A$ is denoted by $f_p = \sum_{i=1}^m f_{p,i}$. Similarly, the total traffic on edge $e \in E$ is denoted by $f_e = \sum_{p: e \in p} f_p$. Denote the *support of type i strategy* in flow \mathbf{f} by $P_i(\mathbf{f}) = \{p \in A : f_{p,i} > 0\}$. That is, all paths used by type i agents in flow \mathbf{f} .

The cost for an agent of type i in game \mathbf{G} , selecting a path $p \in A_i$ in flow \mathbf{f} , is $c^i(p, \mathbf{f}) = \sum_{e \in p} c_e^i(f_e)$.

Social cost. For an NRG (not player-specific), we denote by $SC_i(\mathbf{G}, \mathbf{f}) = \sum_{p \in A_i} f_{p,i} c(p, \mathbf{f})$ the cost experienced by type i agents in flow \mathbf{f} . By summing over all types, we get the social cost:

$$SC(\mathbf{G}, \mathbf{f}) = \sum_{i \leq m} SC_i(\mathbf{G}, \mathbf{f}) = \sum_{i \leq m} \sum_{p \in P_i(\mathbf{f})} f_{p,i} c(p, \mathbf{f}) = \sum_{e \in E} c_e(f_e) f_e$$

Thus the social cost only depends on the total traffic per edge. We denote by $\mathbf{f}^o(\mathbf{G}) \in \arg\min_{\mathbf{f}} SC(\mathbf{G}, \mathbf{f})$ some profile with minimal total cost, and $OPT(\mathbf{G}) = SC(\mathbf{G}, \mathbf{f}^o(\mathbf{G})) = \min_{\mathbf{f}} SC(\mathbf{G}, \mathbf{f})$.

Equilibrium. A flow \mathbf{f} for an PNRG is an *equilibrium* in game \mathbf{G} if for every agent type i , any used path $p \in P_i(\mathbf{f})$ and any $p' \in A_i$, we have $c^i(p, \mathbf{f}) \leq c^i(p', \mathbf{f})$. That is, if no agent can switch to a path with a lower cost. This provides the analogy of a Nash equilibrium for nonatomic games.

It is known that in any NRG there is at least one equilibrium, and that this can be reached by a simple best-response dynamic. Further, all equilibria have the same social cost and in every equilibrium all agents of type i experience the same cost [1, 8, 31, 38]. Player-specific NRGs are also guaranteed to have at least one equilibrium [39], however, equilibrium costs may not be unique, and best-response dynamics may not converge, except in special cases [22].

Affine routing games. In an *affine* NRG, all cost functions take the form of a linear function. That is, $c_e(x) = a_e x + b_e$ for some constants $a_e \geq 0, b_e \geq 0$. The social cost can be written as $SC(\mathbf{G}, \mathbf{f}) = \sum_{e \in E} a_e (f_e)^2 + b_e f_e$. *Pigou's example* is the special case of an affine RSG with two resources, where $c_1(x) = 1$ and $c_2(x) = ax$. We denote by $G_P(a)$ the instance where $c_2(x) = ax$ (see Fig. 1a).

The price of anarchy. Let $EQ(\mathbf{G})$ be the set of equilibria in game \mathbf{G} . The *price of anarchy* (PoA) of a game is the ratio between the social cost in the worst equilibrium in $EQ(\mathbf{G})$ and the optimal social cost [27]. Formally, $PoA(\mathbf{G}) = \sup_{\mathbf{f}^* \in EQ(\mathbf{G})} \frac{SC(\mathbf{G}, \mathbf{f}^*)}{OPT(\mathbf{G})}$. For example in affine NRGs, it is known that $PoA(\mathbf{G}) \leq \frac{4}{3}$, and this bound is attained by the Pigou example of $G_P(1)$ [38].

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

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

A NRG \mathbf{G} is (λ, μ) -smooth if all cost functions in \mathbf{G} are (λ, μ) -smooth. For any (λ, μ) -smooth NRG, $PoA(\mathbf{G}) \leq \frac{\lambda}{1-\mu}$ [17, 37]. 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 [16, 38]). For example, affine functions are $(1, \frac{1}{4})$ -smooth, which again entails a PoA bound of $\frac{4}{3}$.

2.2 Serial-Parallel Width

Consider a network $G = \langle V, E, s, t \rangle$.

Definition 2.1. A set of edges $S \subseteq E$ is *parallel* if there is some $S' \subseteq E$ s.t. $S \subseteq S'$, and S' is a minimal cut between s and t in the network G .

Definition 2.2. A set of edges $S \subseteq E$ is *serial* if there is a simple directed $s - t$ path p containing S .

Definition 2.3 (Serial-parallel Width). The *serial-parallel width* of a network, $SPW(\mathbf{G})$, is the size of the largest set $S \subseteq E$ that is both parallel and serial.

Intuitively, a serial-parallel width of k means there are at least k non-intersecting source-target paths, and some additional path that edge-intersects all of them.

Example 2.4. Consider the Braess network in Fig. 1b (ignore the costs). The minimal $s - t$ cuts are: $\{sa, sb\}$, $\{at, bt\}$, $\{sa, bt\}$ and $\{sb, ab, at\}$. Thus the set $\{sa, bt\}$ is both parallel and serial, which means $SPW(G_B) \geq 2$. The set $\{sa, at\}$ is serial but not parallel; and $\{sa, sb, ab\}$ is neither. In fact, the *only* parallel set of size greater than 2 is $\{sb, ab, at\}$, which is not serial, thus $SPW(G_B) < 3$. We conclude that the serial-parallel width of the Braess network is 2.

Definition 2.5. For any $k \geq 2$, we define the *k-serial-parallel graph* $G_{SP(k)} = \langle V, E, s, t \rangle$, where $V = \{s, t, a_2, \dots, a_k, b_1, \dots, b_{k-1}\}$, and $E = \bigcup_{i=2}^{k-1} \{(s, a_i), (a_i, b_i), (b_i, t), (b_i, a_{i+1})\} \cup \{(s, b_1), (a_k, t)\}$.

See Figure 2 for an example. The 2-serial-parallel network is the Braess network G_B . The serial-parallel width of the k -serial-parallel network is exactly k , where the parallel edges are $\{(s, b_1), (a_2, b_2), \dots, (a_{k-1}, b_{k-1}), (a_k, t)\}$. The network $G_{SP(k)}$ was used in [5] to derive examples of games with high Price of Malice, and we will use it later in a similar way. In a companion paper, Meir and Deligkas [29]




Figure 1: (1a) Pigou's example. (1b) An objective game G (a quadratic variation of the Braess paradox). (1c) The game \hat{G}^α is the same game played by pessimistic agents with parameter $\alpha = 3$. The biased costs of edges with fixed costs like sb do not change, but the biased cost on the edge sa for example is $\hat{c}_{sa}^\alpha(x) = c_{sa}(3x) = (3x)^2 = 9x^2$.

proved that the serial-parallel width of acyclic networks can be characterized completely, where $\text{SPW}(G) < k$ if and only if $G_{SP(k)}$ (or some small variants of which) is not embedded in G . For $k = 1$, this entails another characterization.

PROPOSITION 2.6 (MEIR AND DELIGKAS [29]). *Let G be an acyclic network, then $\text{SPW}(G) = 1$ if and only if G is series-parallel.*

3 BOUNDING EXTERNALITIES

Suppose we are given a game $G = \langle G, m, (c^j)_{j \leq m}, s, t, r \rangle$, and agents play some equilibrium f^* of G . For any single type i , all type i agents have the same cost in f^* . We denote this cost by $C^{i,*} = c^i(p, f^*)$, where $p \in P_i(f^*)$ is an arbitrary path used by a type i agent. Note that $r_i \cdot C^{i,*} = \text{SC}_i(G, f^*)$.

Our goal in this section is to bound $C^{i,*}$. A-priori, this may seem difficult, as we do not assume *anything* about the types of the other agents. However, we will show that the negative externality of the other types can be bounded using the structural parameters of the network G (the serial-parallel width), and even adversarial agents cannot be much worse for i than more type i agents.

Given a game G and type i , we define a new game G^i by setting both the cost type and the demand type of all agents in G to i . That is, $G^i = \langle G, 1, c_i, s_i, t_i, r \rangle$. We also define a game $G^{(k)} = \langle G, m, c, s, t, k \cdot r \rangle$, where the demand in G is multiplied by k . Finally, we set $G^{i,(k)} = (G^i)^{(k)}$, i.e., $G^{i,(k)}$ is a game where there all $k \cdot r$ agents are of type i . Note that $G^{i,(1)} = G^i$, and denote $C^i = C^{i,1}$.

Let g^* be an equilibrium of $G^{i,(k)}$. As $G^{i,(k)}$ is a symmetric game, all agents have the same cost in g^* . We denote this cost by $C^{i,(k)} = c^i(p, g^*)$, where $p \in \bar{P}(g^*)$ is an arbitrary used path.

Finally, let $G|_i = \langle V^i, E^i, s_i, t_i \rangle$ be the network obtained from $G = \langle V, E \rangle$, by eliminating all edges and vertices that are not in A_i .

The following bound is our main result in this section. Prop. 3.2 shows that the bound is tight.

THEOREM 3.1 (EXTERNALITY THEOREM). *Let G be any PNRG played on a network G . If $\text{SPW}(G|_i) \leq k$, then $C^{i,*} \leq C^{i,(k)}$.*

Before we prove the theorem, we explain its implications. For $k = 1$ (e.g., series-parallel networks), this means that there are no

negative externalities due to type differences—the cost of type i agents may only increase when all others are of the same type. We note that this result (along with Prop. 3.2) strictly generalizes Theorem 5 from [2], which is attained as a special case for $k = 1$ and specific demand types (namely, agents that ignore certain edges); and Theorem 5.7 in [33] which is attained as a special case for $k = 1$ and symmetric games. [15] considers a model similar to [33], and compare the social cost with diverse population versus uniform population with the same average bias. There is a tight technical connection between our results and their Theorem 1, see the full version for details. Theorem 3.1 above also implies that the Price of Malice [5] is 0 in series-parallel networks (this is since our theorems make no assumption on the behavior of agents of other types).

For larger values of k , Theorem 3.1 means that regardless of what the other agents are doing, the cost to the type i agents is never more than their equilibrium cost in a game where all agents are of type i , and their number is multiplied by k .

PROOF. Denote $\bar{P} = P_i(g^*)$ (the set of paths $p \in A_i$ s.t. $g_{p,i}^* > 0$). Assume towards a contradiction that $C^{i,*} > C^{i,(k)}$. This means that $c^i(p, f^*) > C^{i,(k)} = c^i(p, g^*)$ for any path $p \in \bar{P}$, as $c^i(p, f^*)$ is either $C^{i,*}$ or higher. For any path $p \in \bar{P}$,

$$\sum_{e \in p} c_e^i(f_e^*) = c^i(p, f^*) > c^i(p, g^*) = \sum_{e \in p} c_e^i(g_e^*),$$

thus there is an edge $e = e_p \in p$ s.t. $c_e^i(f_e^*) > c_e^i(g_e^*)$. Consider the set $E_{\bar{P}} = \{e_p | p \in \bar{P}\}$. Since c_e^i is monotone, this means $f_e^* > g_e^*$ for every edge $e \in E_{\bar{P}}$.

Consider the weighted directed graph $H = \langle V, F \rangle$, where $F = \{e \in p | p \in \bar{P}\}$, and the capacity (weight) of every edge $e \in F$ is g_e^* .

By construction, $E_{\bar{P}}$ is a cut between s_i and t_i in H . Let $\bar{E} \subseteq E_{\bar{P}}$ s.t. \bar{E} is a minimal cut (not necessarily of minimum size or minimum weight), then

$$\sum_{e \in \bar{E}} f_e^* > \sum_{e \in \bar{E}} g_e^* \geq k \cdot r. \quad (2)$$

The first inequality is since $f_e^* > g_e^*$ holds for every $e \in E_{\bar{P}}$ and thus for every $e \in \bar{E}$. The second inequality follows from the min-cut-max-flow theorem, since \bar{E} is some $s_i - t_i$ cut in H , and the flow between s_i, t_i is exactly $k \cdot r$ in any valid flow of $G^{i,(k)}$ (which by definition has a mass of $k \cdot r$ type i agents).

Now, as there are only r agents in $G = (G, m, (c^i)_{i \leq m}, s, t, r)$, then by Eq. (2) and the pigeon-hole principle there must be some agents in profile f^* using strictly more than k edges from the set \bar{E} . Choose some path used by such an agent and denote it by $p' \in A_i$. We define $S = p' \cap \bar{E}$, thus by our selection of p' , $|S| > k$.

Finally, all edges of S are contained in the minimal cut \bar{E} between s_i, t_i (thus S is parallel), and S is also contained by the simple path p' (thus S is serial). By definition of serial-parallel width, $\text{SPW}(G|_i) \geq |S| > k$, which is a contradiction to our assumption. \square

3.1 Lower bound

We say that a class of cost functions is *nontrivial* if it contains all constant functions, as well as at least one strictly increasing function.




Figure 2: The solid edges compose the network $G_{SP(k)}$. The edge labels are used in Prop. 3.2 and 5.4, where edges with the same label (e.g. c_{ab}) have the same cost function. The dashed (s, t) edge is not part of $G_{SP(k)}$ and is used only in Prop. 5.4.

Theorem 3.1 is tight in a very strong sense: the bound cannot be improved for *any network* and *any nontrivial class* of cost functions.

PROPOSITION 3.2. *Let any $k \geq 2$, any network G with $SPW(G) = k$, any nontrivial class of cost functions C , and any $\delta > 0$. There is a game with two types $G = \langle G, (c^1, c^2), s, t, r \rangle$, such that $c_e^1, c_e^2 \in C$ for all $e \in E$, and $C^{1,*} > C^{1,(k)} - \delta$.*

PROOF. We first prove for the k -Serial-Parallel network $G_{SP(k)}$. The demand type of both agents is (s, t) , thus all paths from s to t are valid strategies. We define edge costs as follows. Recall that C contains all constant functions, and at least one strictly increasing function c . For type $i = 1$, we set $c_{ab}^i(x) = c(x)$. The cost of all other edges is 0. For type $j = 2$, we set $c_{ab}^j(x) = c_{ba}^j(x) = 0$, $c_{sa}^j(x) = c_{bt}^j(x) = c(1)$. We set $r_i = \epsilon$ (ϵ will be determined later) and $r_j = 1 - \epsilon$.

Thus in the unique equilibrium g^* of game $G^{i,(k)}$, the agents (k in total) split evenly over all k short paths, and $C^{i,(k)} = c_{ab}^i(\frac{k}{k}) = c_{ab}^i(1) = c(1)$.

In the unique equilibrium f^* of game G , all type j agents take the long path $p = (s, b_1, a_2, b_2, \dots, b_{k-1}, a_k, t)$ since $c^j(p, f^*) = 0$. The type i agents split evenly over all k short paths, and thus each experiences a cost of $C^{i,*} = c_{ab}^i(r_j + \frac{r_i}{k}) = c(r_j + \frac{r_i}{k}) > c(r_j)$.

Let $\epsilon > 0$ such that $c(1 - \epsilon) > c(1) - \delta$. Since c is continuous and strictly increasing such ϵ must exist. Finally, $C^{i,*} > c(r_j) = c(1 - \epsilon) > c(1) - \delta = c(1) = C^{i,(k)} - \delta$, as required.

Now, consider an arbitrary network G with $SPW(G) \geq k$. Consider some serial-parallel set S of size k and the path p that contains it, and S will play the role of the ab edges above. Specifically, For type i agents we set $c_e^i \equiv c_{ab}^i \equiv c$ for any $e \in S$, and $c_e^i \equiv c \equiv 0$ for all other edges. For type j , we set $c_e^j \equiv 0$ for all $e \in p$, and $c_e^j \equiv c(1)$ for all other edges. It is easy to verify that $C^{i,*}, C^{i,(k)}$ are the same as for $G_{SP(k)}$. \square

4 BIASED PRICE OF ANARCHY

Given an NRG $G = \langle G, m, c, s, t, r \rangle$ and modified cost functions \hat{c}_e^i for every type $i \leq m$ and edge $e \in E$, we obtain a PNRG $\hat{G} = \langle G, m, (\hat{c}^i)_{i \leq m}, s, t, r \rangle$. We assume that agents act based on their modified cost functions, irrespective of whether this is a rational behavior or not. We refer to \hat{G} as the *biased game*, where every agent of type i experiences a cost of \hat{c}_e^i rather than c_e . Both games G and \hat{G} have a role in our model, and we often denote the overall setting as $\mathcal{G} = \langle G, \hat{G} \rangle$. The way we interpret \mathcal{G} is that agents play the biased game \hat{G} (and thus, it is the equilibria of \hat{G} that matter), whereas their true costs are according to game G . We say that \mathcal{G} is *uniform* if all agents in \hat{G} have the same cost type (same bias).

Biased Price of Anarchy. We measure the price of anarchy in a game with biased costs by comparing the equilibria of \hat{G} (denoted by f^*) to the optimum of G . Formally:

$$\text{BPoA}(\mathcal{G}) = \text{BPoA}(G, \hat{G}) = \sup_{f^* \in \text{EQ}(\hat{G})} \frac{SC(G, f^*)}{\text{OPT}(G)}. \quad (3)$$

In the uniform bias case where $\hat{c}^i = \hat{c}$ for all i , the game \hat{G} is just another NRG. Chen et al. [13] referred to the BPoA (when applied to altruism) as the *Robust PoA*. The *Price of Risk Aversion* [33] and the *Deviation Ratio* [26] are similar concepts to the BPoA, except they compare f^* to the equilibrium of the unbiased game G .

A simple example of a biased cost is induced by *pessimism*, which is one way to model risk-aversion [30]. Suppose that whenever faced with some expected load f_e on edge e , an agent takes a safety margin by playing as if the actual load is $\alpha \cdot f_e$ (for some fixed private parameter $\alpha \geq 1$). Such an agent will play as if every cost function c_e is replaced with a new cost function \hat{c}_e^α , where $\hat{c}_e^\alpha(x) = c_e(\alpha x)$ (see Fig. 1). We denote the (uniform) game where all agents play according to $(\hat{c}_e^\alpha)_{e \in E}$ by \hat{G}^α .

Example 4.1. In the optimal flow of the objective game G in Fig. 1, the amount of agents taking the long path is ~ 0.155 , and thus $\text{OPT}(G) \cong 1.23$. The equilibrium f^* of G is suboptimal (as all agents take the long path $s - a - b - t$), and $\text{PoA}(G) \cong \frac{2}{1.23} = 1.626$. In contrast, in the equilibrium f^* of \hat{G}^α for $\alpha = 3$ (Fig. 1c), agents divide equally among the two short paths, which leads to $\text{BPoA}(G, \hat{G}^\alpha) \cong \frac{1.25}{1.23} = 1.016$.

Interestingly, the bias of the agents on Example 4.1 somewhat mitigates the increase in social cost that is incurred by rational behavior.

4.1 Smoothness for Biased Costs

Our goal is to provide bounds on the biased Price of Anarchy for a given game with biased costs $\langle G, \hat{G} \rangle$. That each of c_e and \hat{c}_e^i are smooth is insufficient to provide such a bound.

Example 4.2. Consider a Pigou game with pessimistic agents $\mathcal{G}_\alpha = \langle G_\alpha, \hat{G}_\alpha^\alpha \rangle$ where $G_\alpha = G_P(2/\alpha)$ and thus $\hat{G}_\alpha^\alpha = G_P(2)$. For any α the equilibrium of \hat{G}_α^α (and thus of \mathcal{G}_α) is the same: $1/2$ of all agents use each resource. However as α increases, the optimal flow of G_α shifts more agents to the resource with variable cost, and the optimal social cost decreases to 0. Thus the gap between the

equilibrium cost and the optimal cost (the BPoA) goes to infinity with α even though both of $G_\alpha, \hat{G}_\alpha^\alpha$ are affine.

It is crucial, then, to extend the definition of smoothness to games with biased costs in a way that takes into account both c and \hat{c} . This technique has been applied before for specific modified costs, for example nonatomic games with restricted tolls [9] and atomic games with altruistic agents [12]. We provide a general extension.

Definition 4.3. Let $\lambda \geq 0, \mu < 1$. The function c is (λ, μ) -biased-smooth w.r.t. biased cost function \hat{c} , if for any $x, x' \in \mathbb{R}_+$,

$$c(x)x + \hat{c}(x)(x' - x) \leq \lambda c(x')x' + \mu c(x)x. \quad (4)$$

It is instructive to check the familiar case where there is no bias. Indeed, if $\hat{c} = c$, and c is (λ, μ) -smooth, then

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

and Eq. (4) collapses to “standard” smoothness (Eq. (1)).

Recall that the PoA of a (λ, μ) -smooth game is bounded by $\frac{\lambda}{1-\mu}$. This bound extends to games with biased costs that are biased-smooth and when all agents have the *same bias*. This was explicitly shown for specific biases, but we write down the general formulation for completeness.

THEOREM 4.4 (BONIFACI ET AL. [9]). Consider a uniform game with biased costs $\mathcal{G} = \langle G, \hat{G} \rangle$ where every cost function c_e is $(\hat{\lambda}, \hat{\mu})$ -biased smooth w.r.t. biased cost function \hat{c}_e . Let $\hat{f}^* \in EQ(\hat{G})$, and f any valid flow. Then $SC(G, \hat{f}^*) \leq \frac{\hat{\lambda}}{1-\hat{\mu}} SC(G, f)$.

In particular, $BPoA(G, \hat{G}) \leq \frac{\hat{\lambda}}{1-\hat{\mu}}$.

PROOF. The proof extends the standard proof of PoA bounds for nonatomic congestion games via smoothness arguments. In any equilibrium \hat{f}^* , the *variational inequality* $\sum_e \hat{c}_e(\hat{f}_e^*)\hat{f}_e^* \leq \sum_e \hat{c}_e(\hat{f}_e^*)f_e$ holds (see [17, 37]). Thus,

$$\begin{aligned} SC(G, \hat{f}^*) &\leq SC(G, \hat{f}^*) + \sum_e \hat{c}_e(\hat{f}_e^*)f_e - \sum_e \hat{c}_e(\hat{f}_e^*)\hat{f}_e^* \\ &\leq \sum_e [\hat{\lambda}c_e(f_e)f_e + \hat{\mu}c_e(\hat{f}_e^*)\hat{f}_e^*] = \hat{\lambda}SC(G, f) + \hat{\mu}SC(G, \hat{f}^*). \end{aligned}$$

We get the bound by reorganizing the terms. The only part that differs from the standard smoothness is the first inequality. \square

An alternative derivation of optimal tolls. We can also check that the extension provides the right result in regard to modified costs $\tilde{c}(x) = c(x) + c'(x)x$ (where $c'(x) = \frac{\partial c(x)}{\partial x}$) that represent optimal tolls and should lead to optimal play [7]. Let’s confirm this result via a biased-smoothness argument.

OBSERVATION 4.5. If we set modified cost $\tilde{c}(x) = c(x) + c'(x)x$, then cost function c is $(1, 0)$ -biased smooth w.r.t. to \tilde{c} .

This follows immediately from the convexity of $xc(x)$, and affirms via Theorem 4.4 that $BPoA(G, \hat{G}) = 1$ for any game. For example, adding optimal tolls to the game in Fig. 1b will set $\tilde{c}_{sa}(x) = \tilde{c}_{bt}(x) = 3x^2$, and in the only equilibrium, $SC(G, f^*) = 1.23 = OPT(G)$.

4.2 Bounds for Specific Biases

Theorem 4.4 lets us bound the BPoA of a population of agents with *uniform bias*, and in the next sections we will prove theorems that bound the BPoA of populations with *heterogeneous biases*. However, all these theorems require some explicit bounds on the biased smoothness parameters $\hat{\lambda}$ and $\hat{\mu}$ in a given scenario (i.e., for the specific bias and specific class of cost functions).

There are by now many such studied biases. We already introduced pessimism, which depends on a single parameter α , and where the biased cost is $\hat{c}_e^\alpha(x) = c_e(\alpha x)$. Toll-sensitive agents [14] have a single parameter β , and their biased cost is $\hat{c}_e^\beta(x) = c_e(x) + \beta xc_e'(x)$ (meaning that marginal tolls are imposed, but agents discount or over-weigh them by a factor of β). Altruist agents [13] have exactly the same biased cost function as tolls, except that now β should be interpreted as how much they care about hurting others. In the Mean-Var risk-aversion model [32, 33], $c_e(x)$ is a distribution over costs, and the biased cost is $\hat{c}_e^\gamma(x) = \mathbb{E}[c_e(x)] + \gamma VAR[c_e(x)]$, i.e., γ represents the sensitivity of agents to variance. We assume that $VAR[c_e(x)] \leq \tau$ for some constant τ . *Small error* means that c and \hat{c} are within a small multiplicative factor of $1 + \delta$ from one another (see Sec. 6 for details).

Figure 3 summarizes known biased smoothness bounds for routing games with *affine* cost functions (see the full version for bounds for more general classes of cost functions). Note that from every column in the table we can derive a BPoA bound via Theorem 4.4. For example, for any biased game $\langle G, \hat{G} \rangle$ where all cost functions in G are affine, and all agents in \hat{G} are pessimistic with parameter $\alpha \leq 2$, we have $BPoA(G, \hat{G}^\alpha) \leq \frac{4}{4-\alpha^2}$. This bound, which is tight, is illustrated in Fig. 3b. Note that for certain values of α , the BPoA is lower than the PoA, i.e., bias steers the society to better outcomes (we also saw this phenomenon in Example 4.1).

We note that BPoA bounds have also been derived independently using different techniques by Nikolova and Stier-Moses [33], by Meir and Parkes [30] and by Kleer et al. [26].

5 STRUCTURE-DEPENDENT BOUNDS

In this section, we derive a bound, tight up to a constant, on the equilibrium cost for agents participating in a game with heterogeneous biases. This bound depends on their own biased-smoothness parameters, as well as on the structure of the network. In some cases, this also allows us to derive BPoA bounds that depend on the *average bias*. Most importantly, our bounds in this section do not assume any specific class of cost functions, and work for any function and any bias, as long as biased smoothness holds.

5.1 Upper bound

Our primary question is whether we can get a bound on the social cost of any equilibrium of \mathcal{G} in terms of the smoothness parameters of all types, and the optimal social cost.

In the next theorem and corollaries, each c_e is (λ, μ) -smooth (as per Eq. (1)), each \hat{c}_e^i is (λ^i, μ^i) -smooth, and c_e is $(\hat{\lambda}^i, \hat{\mu}^i)$ -biased-smooth w.r.t. \hat{c}_e^i (as per Def. 4.3). Recall that $\frac{t_i}{r}$ is the fraction of agents of type i .

THEOREM 5.1 (SMOOTHNESS THEOREM). Consider any game with biased costs $\mathcal{G} = \langle G, \hat{G} \rangle$ where $SPW(G) \leq k$. Let \hat{f}^* be an equilibrium




Figure 3: Biased-smoothness bounds for affine cost functions. For a trivial bias of $\alpha = 1, \beta = 0, \gamma\tau = 0$, or $\delta = \hat{\delta} = 0$, we always get the familiar affine smoothness of $(\hat{\lambda} = 1, \hat{\mu} = \frac{1}{4})$. The bound on Altruism with $\beta \leq 1$ is due to Chen et al. [13]. The bound on small error is due to Prop. 6.2 applied to affine cost functions. The proofs for all other bounds are in the full version.

of \hat{G} . Then for any type i ,

$$SC_i(G, \hat{f}^*) \leq \frac{r_i}{r} \frac{\lambda \lambda^i \hat{\lambda}^i}{(1 - \sqrt{\mu})^2 (1 - \mu^i)(1 - \hat{\mu}^i)} \frac{1}{k} OPT(G^{i,(k)}).$$

In simple words, the theorem says that for every agent type, the equilibrium cost to all r_i agents of this type may increase (compared to the optimal cost) by a factor that only depends on the smoothness of cost functions, their own bias, and the structure of the network, but not on the biases of other types. This result is bi-criteria, as we compare to the optimum of a game with higher demand. Since costs are increasing functions, $\frac{1}{k} OPT(G^{i,(k)})$ is increasing in k .

LEMMA 5.2. For every $x > 0$, and every type i , it holds that:

$$(a) \hat{c}^i(x) \geq (1 - \hat{\mu}^i)c(x);^1 \text{ and } (b) \hat{c}^i(x) \leq \frac{\lambda \lambda^i}{(1 - \sqrt{\mu})^2} c(x).$$

PROOF SKETCH. We use the biased-smoothness definition: (a) with $x' = 0$; (b) with $x' = \sqrt{1/\mu} \cdot x$. \square

PROOF SKETCH OF THEOREM 5.1. Let \hat{g}^* be an equilibrium of $\hat{G}^{i,(k)}$. Define $\hat{C}^{i,*} = \hat{c}^i(p, \hat{f}^*)$ for some used path $p \in P_i(\hat{f}^*)$, and $\hat{C}^{i,(k)} = \hat{c}^i(p, \hat{g}^*)$ for some used path $p \in P_i(\hat{g}^*)$. In words, $\hat{C}^{i,*}, \hat{C}^{i,(k)}$ are the *perceived* costs to agents of type i in equilibria \hat{f}^* and \hat{g}^* , respectively. Note that they are not affected by which used path we choose.

The next steps are: to upper bound the ratio between $rk \cdot \hat{C}^{i,(k)} = SC(\hat{G}^{i,(k)}, \hat{g}^*)$ and $OPT(G^{i,(k)})$, using Lemma 5.2(b) on every used edge; and to lower bound $\hat{C}^{i,*}$ by the *true cost* $c(p, \hat{f}^*)$ of path p , using Lemma 5.2(a) on all edges of p . The most important part is the inequality $\hat{C}^{i,*} \leq \hat{C}^{i,(k)}$, which is due to Theorem 3.1 applied to the biased game \hat{G} .

Finally, recall that what we need to bound is social cost for all type i agents in the “real game” G , which is $SC_i(G, \hat{f}^*)$. Since each such type i agent has a cost of exactly $c(p, \hat{f}^*)$, we now chain all the previous inequalities to get the desired upper bound. \square

5.2 Implications

Theorem 5.1 has a number of useful corollaries. Some use the fact that for series-parallel networks $k = 1$ (Prop. 2.6).

COROLLARY 5.3. Consider any game with biased costs $\mathcal{G} = \langle G, \hat{G} \rangle$ where $SPW(G) \leq k$. Let \hat{f}^* be an equilibrium of \mathcal{G} . For any type i :

¹For some biases (e.g. tolls), we have a tighter inequality $\hat{c}^i(x) \geq c(x)$ from the definition. In such cases we can eliminate the factor $1 - \hat{\mu}^i$ in Theorem 5.1 and all following corollaries.

$$(1) \text{ if } G \text{ is affine, } SC_i(G, \hat{f}^*) \leq 4 \frac{r_i}{r} \frac{\lambda \lambda^i \hat{\lambda}^i}{(1 - \mu^i)(1 - \hat{\mu}^i)} k \cdot OPT(G^i);$$

$$(2) \text{ if type } i \text{ is unbiased, } SC_i(G, \hat{f}^*) \leq \frac{r_i}{r} \frac{\lambda}{1 - \mu} \frac{1}{k} OPT(G^{i,(k)}).$$

Moreover, if G is symmetric, then:

$$(3) SC(G, \hat{f}^*) \leq \left(\sum_{i \leq m} \frac{r_i}{r} \frac{\lambda \lambda^i \hat{\lambda}^i}{(1 - \sqrt{\mu})^2 (1 - \mu^i)(1 - \hat{\mu}^i)} \right) \frac{1}{k} OPT(G^{(k)});$$

$$(4) \text{ if } G \text{ is affine, } BPoA(\mathcal{G}) \leq O(1)k \sum_{i \leq m} \frac{r_i}{r} \frac{\lambda^i}{1 - \hat{\mu}^i};$$

$$(5) \text{ If } G \text{ is series-parallel, then } BPoA(\mathcal{G}) \leq O(1) \sum_{i \leq m} \frac{r_i}{r} \frac{\lambda^i}{1 - \hat{\mu}^i}.$$

These corollaries provide us with more explicit bounds. In particular: (1) for affine cost functions we can get rid of the bi-criteria result since the optimal social cost is linear in the demand (more generally, for polynomials of degree d the factor becomes $O(k^d)$). (2) entails that in a setting where most agents are unbiased and there is only a small fraction of adversarial agents (as in [5]), the damage to the unbiased agents is limited. (3)-(5) show that the equilibrium social cost in symmetric games depends on the *average bias* over all agents: each type of biased agents can only increase the social cost by a factor that is proportional to their mass (and also affected by the serial-parallel width of the network). Klee et al. [26] proved a result similar in spirit to our Corollary 5.3(5), but restricted to toll-sensitivity.

5.3 Lower Bounds

Our main theorems both rely on the restricted structure of the underlying network. Prop. 3.2 already shows the necessity of this restriction when some agents behave adversarially.

We next show that the neither smoothness parameters $(\hat{\lambda}_i, \hat{\mu}_i)$ nor the structural parameter $k = SPW(G)$ can be relaxed in Theorem 5.1 and in its corollaries.

PROPOSITION 5.4. Let $k \geq 2$, any smoothness parameters $(\hat{\lambda}^i, \hat{\mu}^i)$, and any $r_i < \frac{1}{2}$. There is an affine symmetric game with biases $\mathcal{G} = \langle G, \hat{G} \rangle$ with $SPW(G) = k$, such that in the unique equilibrium \hat{f}^* of \hat{G} , $SC_i(G, \hat{f}^*) \geq \Omega(1) \frac{\hat{\lambda}^i}{1 - \hat{\mu}^i} k \cdot SC_i(G, f^0)$.

PROOF SKETCH. See Fig. 2 for $G_{SP(k)}$ with one additional edge e_{st} . We define the real costs as follows: $c_{sa} = c_{bt} = c_{ba} \equiv 0, c_{st} \equiv \frac{q}{2}, c_{ab}(x) = x$, where $q = \frac{\hat{\lambda}^i}{1 - \hat{\mu}^i}$. In the flow of G where all agents split evenly, the cost is $C^0 = \frac{1}{k}$. In the biased game, $\hat{c}_{ab}^i(x) = qx$, and all other agents prefer the long path. Thus in equilibrium \hat{f}^* all type i agents use edge e_{st} for a cost of $\frac{q}{2} \geq \frac{\hat{\lambda}^i}{2(1 - \hat{\mu}^i)} \geq \frac{\hat{\lambda}^i}{2(1 - \hat{\mu}^i)} k \cdot C^0$. \square

Since we only used affine cost functions, the bounds in both Theorem 5.1 and Cor. 5.3(1) are tight up to a constant. Also, since the network used in the proof is embedded in $G_{SP(k+1)}$, and thus embedded in any acyclic network with $SPW(G) > k$ (due to [29]), the example in Prop. 5.4 can be constructed for *any acyclic network* G with $SPW(G) \geq k + 1$.

6 STRUCTURE-INDEPENDENT BOUNDS

In this section, we leverage known PoA bounds in routing games with player-specific costs [22] to obtain BPoA bounds in biased games with heterogeneous agents. These are bounds that do not depend on the network structure, but only on the cost functions and agents' biases.

The bounds in this section are typically worse than those in Section 5, as they depend on the worst bias rather than the average bias, and this dependency is polynomial rather than linear. Yet, these bounds hold regardless of network structure.

For two cost functions c, \hat{c} and $r > 0$, denote by $\Delta(c, \hat{c}, r) = \sup_{x \in [0, r]} \frac{c(x)}{\hat{c}(x)}$. Let $\Psi(G) = \max_{i,j \leq m} \max_{e \in E} \Delta(c_e^i, c_e^j, r)$.

THEOREM 6.1 (GAIRING ET AL. [22]). *Consider any PNRG $G = \langle G, (c^i)_{i \leq m}, r \rangle$ where c^i are polynomials of degree d for all i . Then $PoA(G) \leq (d + 1) \cdot \Psi(G)^d$.*²

Intuitively, the PoA is low if for each edge e and any x , all functions $(c_e^i(x))_{i \leq m}$ attain similar values. We can use this PoA bound to prove a similar BPoA bound that is independent of the network structure. Given a game with biased costs $\mathcal{G} = \langle G, \hat{G} \rangle$, let:

- $\bar{\Phi}(\mathcal{G}) = \max_{i \leq m} \max_{e \in E} \Delta(\hat{c}_e^i, c_e, r)$;
- $\Phi(\mathcal{G}) = \max_{i \leq m} \max_{e \in E} \Delta(c_e, \hat{c}_e^i, r)$;
- and $\Psi(\mathcal{G}) = \Psi(\hat{G})$.

We first observe that Δ is tightly related to biased-smoothness, as biased-smoothness bounds entail bounds on Δ , and vice versa.

PROPOSITION 6.2. *Consider c, \hat{c} such that c is (λ, μ) -smooth, $\Delta(c, \hat{c}, r) \leq 1 + \delta$, and $\Delta(\hat{c}, c, r) \leq 1 + \hat{\delta}$ for some $\delta, \hat{\delta} \geq 0$ and any $r > 0$. Then c is $\left((1 + \delta)\lambda, \frac{\mu + \hat{\delta}}{1 + \hat{\delta}}\right)$ -biased-smooth w.r.t. \hat{c} .*

PROPOSITION 6.3. *Consider c, \hat{c} such that c is (λ, μ) -smooth, and $(\hat{\lambda}, \hat{\mu})$ -biased smooth w.r.t. \hat{c} . Then for any $r > 0$ $\Delta(c, \hat{c}, r) \leq \frac{1}{1 - \hat{\mu}}$; $\Delta(\hat{c}, c, r) \leq \frac{\lambda}{(1 - \sqrt{\mu})^2} \hat{\lambda}$. Also, if c is a polynomial of degree at most d , $\Delta(\hat{c}, c, r) \leq (d + 1) \cdot e \cdot \hat{\lambda}$, where e is the natural logarithm base.*

Proposition 6.3 follows from Lemma 5.2 and from the smoothness of polynomial functions. Next, we derive a BPoA bound. Due to tightness results on the PoA [22], we cannot hope to significantly improve the bound in this approach.

THEOREM 6.4. *Consider any game $\mathcal{G} = \langle G, \hat{G} \rangle$ where all of \hat{c}_e^i are polynomials of degree at most d . Then $BPoA(\mathcal{G}) \leq (d + 1)(\Phi(\mathcal{G})\bar{\Phi}(\mathcal{G}))^{d+1}$.*

Finally, we can combine Proposition 6.3 and Theorem 6.4 to obtain a BPoA bound that depends only on the smoothness parameters of the “worst” type:

COROLLARY 6.5. *Consider any game with biased costs $\mathcal{G} = \langle G, \hat{G} \rangle$ where for all $i \leq m$ and $e \in E$: (a) c_e, \hat{c}_e^i are polynomials of degree at most d ; and (b) c_e is $(\hat{\lambda}, \hat{\mu})$ -biased smooth w.r.t. \hat{c}_e^i .*

Then $BPoA(\mathcal{G}) \leq (d + 1)^{d+2} e^{d+1} \left(\frac{\hat{\lambda}}{1 - \hat{\mu}}\right)^{d+1}$.

7 CONCLUSION AND FUTURE WORK

We have considered strategic settings in which participants are making routing decisions based on individually perceived costs. From the perspective of a system analyst who considers the objective costs, the agents' behavior deviates from optimal self-interested play. Whether these deviations come from a cognitive limitation, subjective preferences, or external influence such as tolls, it is important to understand how the equilibria of the game are affected.

Our work is the first to provide (Biased) PoA bounds for populations with heterogeneous arbitrary biases, and the first to consider heterogeneous biases in general networks. Our main results (Sec. 3 and 5) bound the equilibrium cost of each agent type separately, based only on their own biases and on the structure of the network. When considering the entire population, this entails a bi-criteria bound on the social cost (in some cases can be written as a Biased PoA bound) that depends on the *average* bias in the population. These results can help estimate the worst case effect when the exact structure of the population is not known in advance: one only has to know the smoothness bounds of each type and a bound on the quantity of each type.

Since in nonatomic routing games all mixed and correlated equilibria have the same social cost anyway [8], smoothness does not add more “robustness,” and is thus perhaps considered less interesting in such games. Our results show that smoothness is useful for a *different kind of robustness*, namely to heterogeneous biases.

An additional result (Sec. 6) relies on the PoA of the heterogeneous “wrong” game, and while it does not require the network to have a particular structure, it provides bounds stated in terms of the *worst* bias in the population, and only applies for polynomial cost functions. We can think of these results as bounding the negative externalities of each type *on all others*, whereas the results in Sections 3 and 5 bound the externalities that all others may inflict on agents of type i .

We hope that our results and techniques will inspire further progress in understanding both the effect of network topology and the effect of bounded rationality on equilibria.

Some challenges are to derive BPoA bounds that only depend on the average bias (i.e., not on the network structure); and bounds on the negative externality that become negligible for a small amount of adversarial agents.

Biased smoothness can also be applied to obtain robust PoA bounds in *atomic* congestion games and other normal form games (i.e. for mixed and correlated equilibria), by extending the standard smoothness definition as in [13]. An important open question is whether atomic games with heterogeneous biases can be similarly analyzed, by showing biased-smoothness independently for each type of agent.

ACKNOWLEDGMENTS

We thank Pieter Kleer, Igal Milchtaich, Nicolás Stier-Moses, Evdokia Nikolova and Éva Tardos for helpful discussions and comments.

²The tight bound (Thm. 4.10 in Gairing et al. [22]) is better by a factor of up to $(d + 1)$.

REFERENCES

- [1] Heayat Zakaie Aashtiani and Thomas L Magnanti. 1981. Equilibria on a congested transportation network. *SIAM Journal on Algebraic Discrete Methods* 2, 3 (1981), 213–226.
- [2] Daron Acemoglu, Ali Makhdomi, Azarakhsh Malekian, and Asuman Ozdaglar. 2016. Informational Braess' Paradox: The Effect of Information on Traffic Congestion. *arXiv preprint arXiv:1601.02039* (2016).
- [3] Haris Angelidakis, Dimitris Fotakis, and Thanasis Lianeas. 2013. Stochastic Congestion Games with Risk-Averse Players. In *Proceedings of the 6th International Symposium on Algorithmic Game Theory*. Springer, 86–97.
- [4] Itai Ashlagi, Dov Monderer, and Moshe Tennenholtz. 2009. Two-terminal routing games with unknown active players. *Artificial Intelligence* 173, 15 (2009), 1441–1455.
- [5] Moshe Babaioff, Robert Kleinberg, and Christos H Papadimitriou. 2007. Congestion games with malicious players. In *Proceedings of the 8th ACM conference on Electronic commerce*. ACM, 103–112.
- [6] Bikramjit Banerjee and Landon Kraemer. 2010. Coalition structure generation in multi-agent systems with mixed externalities. In *Proceedings of the 9th International Conference on Autonomous Agents and Multiagent Systems: volume 1-Volume 1*. International Foundation for Autonomous Agents and Multiagent Systems, 175–182.
- [7] M. Beckmann, B. McGuire, and C. Winsten. 1956. *Studies in the Economics of Transportation*. Yale University Press, New Haven.
- [8] Avrim Blum, Eyal Even-Dar, and Katrina Ligett. 2006. Routing without regret: On convergence to Nash equilibria of regret-minimizing algorithms in routing games. In *Proceedings of the twenty-fifth annual ACM symposium on Principles of distributed computing*. ACM, 45–52.
- [9] Vincenzo Bonifaci, Mahyar Salek, and Guido Schäfer. 2011. Efficiency of Restricted Tolls in Non-atomic Network Routing Games. In *Proceedings of the 4th International Symposium on Algorithmic Game Theory*. 302–313.
- [10] Simina Brânzei, Tomasz Michalak, Talal Rahwan, Kate Larson, and Nicholas R Jennings. 2013. Matchings with externalities and attitudes. In *Proceedings of the 2013 international conference on Autonomous agents and multi-agent systems*. International Foundation for Autonomous Agents and Multiagent Systems, 295–302.
- [11] Patrice Caire, Baptiste Alcalde, Leendert van der Torre, and Chattrakul Sombatheera. 2011. Conviviality measures. In *The 10th International Conference on Autonomous Agents and Multiagent Systems-Volume 2*. International Foundation for Autonomous Agents and Multiagent Systems, 895–902.
- [12] Po-An Chen, Bart De Keijzer, David Kempe, and Guido Schäfer. 2011. The robust price of anarchy of altruistic games. In *Proceedings of the 7th Workshop on Internet and Network Economics*. Springer, 383–390.
- [13] Po-An Chen, Bart De Keijzer, David Kempe, and Guido Schäfer. 2014. Altruism and its impact on the price of anarchy. *ACM Transactions on Economics and Computation* 2, 4 (2014), 17.
- [14] Richard Cole, Yevgeniy Dodis, and Tim Roughgarden. 2003. Pricing network edges for heterogeneous selfish users. In *Proceedings of the thirty-fifth annual ACM symposium on Theory of computing*. ACM, 521–530.
- [15] Richard Cole, Thanasis Lianeas, and Evdokia Nikolova. 2017. When Does Diversity of User Preferences Improve Outcomes in Selfish Routing? *arXiv preprint arXiv:1702.07806* (2017).
- [16] José R Correa, Andreas S Schulz, and Nicolás E Stier-Moses. 2004. Selfish routing in capacitated networks. *Mathematics of Operations Research* 29, 4 (2004), 961–976.
- [17] José R Correa, Andreas S Schulz, and Nicolás E Stier-Moses. 2008. A geometric approach to the price of anarchy in nonatomic congestion games. *Games and Economic Behavior* 64, 2 (2008), 457–469.
- [18] Giovanna Devetag and Massimo Warglien. 2008. Playing the wrong game: An experimental analysis of relational complexity and strategic misrepresentation. *Games and Economic Behavior* 62, 2 (2008), 364 – 382.
- [19] David Eppstein. 1992. Parallel recognition of series-parallel graphs. *Information and Computation* 98, 1 (1992), 41–55.
- [20] Lisa Fleischer, Kamal Jain, and Mohammad Mahdian. 2004. Tolls for heterogeneous selfish users in multicommodity networks and generalized congestion games. In *Foundations of Computer Science, 2004. Proceedings. 45th Annual IEEE Symposium on*. IEEE, 277–285.
- [21] Dimitris Fotakis, Dimitris Kalimeris, and Thanasis Lianeas. 2015. Improving Selfish Routing for Risk-Averse Players. In *International Conference on Web and Internet Economics*. Springer, 328–342.
- [22] Martin Gairing, Burkhard Monien, and Karsten Tiemann. 2011. Routing (un-) splittable flow in games with player-specific affine latency functions. *ACM Transactions on Algorithms (TALG)* 7, 3 (2011), 31.
- [23] Ralph Hertwig, Greg Barron, Elke U Weber, and Ido Erev. 2004. Decisions from experience and the effect of rare events in risky choice. *Psychological science* 15, 8 (2004), 534–539.
- [24] Andreas Jakob, Maciej Liśkiewicz, and Rüdiger Reischuk. 2006. Space efficient algorithms for directed series-parallel graphs. *Journal of Algorithms* 60, 2 (2006), 85–114.
- [25] Tomas Jelinek, Marcus Klaas, and Guido Schäfer. 2014. Computing Optimal Tolls with Arc Restrictions and Heterogeneous Players. In *STACS*. 433–444.
- [26] Pieter Kleer and Guido Schäfer. 2016. The impact of worst-case deviations in non-atomic network routing games. In *International Symposium on Algorithmic Game Theory*. Springer, 129–140.
- [27] Elias Koutsoupias and Christos Papadimitriou. 1999. Worst-case equilibria. In *Proceedings of the 16th annual conference on Theoretical aspects of computer science*. Springer-Verlag, 404–413.
- [28] Piotr Krysta, Tomasz Michalak, Tuomas Sandholm, and Michael Wooldridge. 2010. Combinatorial auctions with externalities. In *AAMAS'10*. 1471–1472.
- [29] Reshef Meir and Argyrios Deligkas. 2017. Directed graph embeddings. *arXiv preprint arXiv:1711.01806* (2017).
- [30] Reshef Meir and David Parkes. 2015. Congestion Games with Distance-Based Strict Uncertainty. In *Proceedings of the 29th AAAI Conference on Artificial Intelligence*.
- [31] Igal Milchtaich. 2000. Generic uniqueness of equilibrium in large crowding games. *Mathematics of Operations Research* 25, 3 (2000), 349–364.
- [32] Evdokia Nikolova and Nicolás E Stier-Moses. 2014. A mean-risk model for the traffic assignment problem with stochastic travel times. *Operations Research* 62, 2 (2014), 366–382.
- [33] Evdokia Nikolova and Nicolas E Stier-Moses. 2015. The burden of risk aversion in mean-risk selfish routing. In *Proceedings of the Sixteenth ACM Conference on Economics and Computation*. ACM, 489–506.
- [34] Jean Oh and Stephen F Smith. 2008. A few good agents: multi-agent social learning. In *Proceedings of the 7th international joint conference on Autonomous agents and multiagent systems-Volume 1*. International Foundation for Autonomous Agents and Multiagent Systems, 339–346.
- [35] Georgios Piliouras, Evdokia Nikolova, and Jeff S Shamma. 2016. Risk sensitivity of price of anarchy under uncertainty. *ACM Transactions on Economics and Computation (TEAC)* 5, 1 (2016), 5.
- [36] Tim Roughgarden. 2003. The price of anarchy is independent of the network topology. *J. Comput. System Sci.* 67, 2 (2003), 341–364.
- [37] Tim Roughgarden. 2009. Intrinsic robustness of the price of anarchy. In *Proceedings of the 41st annual ACM symposium on Theory of computing*. ACM, 513–522.
- [38] Tim Roughgarden and Éva Tardos. 2004. Bounding the inefficiency of equilibria in nonatomic congestion games. *Games and Economic Behavior* 47, 2 (2004), 389–403.
- [39] David Schmeidler. 1973. Equilibrium points of nonatomic games. *Journal of Statistical Physics* 7, 4 (1973), 295–300.
- [40] Guni Sharon, Josiah Hanna, Tarun Rambha, Michael Albert, Peter Stone, and Stephen D Boyles. 2016. Delta-Tolling: Adaptive Tolling for Optimizing Traffic Throughput. In *ATT@IJCAI*.
- [41] Moshe Tennenholtz and Aviv Zohar. 2009. Learning equilibria in repeated congestion games. In *Proceedings of The 8th International Conference on Autonomous Agents and Multiagent Systems-Volume 1*. International Foundation for Autonomous Agents and Multiagent Systems, 233–240.
- [42] Shoshana Vasserman, Michal Feldman, and Avinatan Hassidim. 2015. Implementing the Wisdom of Waze. In *IJCAI*. 660–666.
- [43] John Glen Wardrop. 1952. Road paper. Some theoretical aspects of road traffic research. In *ICE Proceedings: Engineering Divisions*, Vol. 1. 325–362.