When Smoothness is Not Enough: Toward Exact Quantification and Optimization of the Price of Anarchy

RAHUL CHANDAN, University of California, Santa Barbara
DARIO PACCAGNAN, Imperial College London
JASON R. MARDEN, University of California, Santa Barbara

The price of anarchy (PoA) is a popular metric for analyzing the inefficiency of self-interested decision making. Although its study is widespread, characterizing the PoA can be challenging. A commonly employed approach is based on the *smoothness framework*, which provides tight PoA values under the assumption that the system objective consists in the sum of the agents' individual welfares. Unfortunately, several important classes of problems do not satisfy this requirement (e.g., taxation in congestion games), and our first result demonstrates that the smoothness framework does *not* tightly characterize the PoA for such settings. Motivated by this observation, this work develops a framework that achieves two chief objectives: i) to tightly characterize the PoA for such scenarios, and ii) to do so through a tractable approach. As a direct consequence, the proposed framework recovers and generalizes many existing PoA results, and enables efficient computation of incentives that optimize the PoA. We conclude by highlighting the applicability of our contributions to incentive design in congestion games and utility design in distributed welfare games.

Additional Key Words and Phrases: game theory, multiagent systems, price of anarchy, optimal incentives

1 INTRODUCTION

The widespread proliferation of smartphones and other smart devices has led to a momentous shift in the operation of shared technological infrastructure like road-traffic networks, cloud computing and the power grid, where the local behaviours and interactions of individual decision makers are increasingly influencing the system wide performance. Although such performance could be improved if a central coordinator was able to dictate the choices of individual decision makers, this approach is often infeasible owing to the distributed and self-interested nature of the very same decision making process. Within these settings, wide ranging inefficiencies can severely degrade system performance, a phenomenon that is typically referred to as the *tragedy of the commons* [25] in economics and the social sciences.

In light of these growing challenges, this paper focuses on (i) characterizing the impact of self-interested decision making on system performance and (ii) deriving locally implementable mechanisms to help mitigate these inefficiencies. Before delving into our specific contributions, we begin with two motivating examples: incentive design in congestion games and distributed coordination of multiagent systems. Both these problem settings have been widely studied in the operations research and game theory literature as they capture the adverse effects of local decision making on system wide performance. Applications include routing in traffic and communication networks [18, 41], distributed resource allocation [24, 26], credit assignment in teams [28], among many others. The two problems detailed below apply to vastly different fields of research and yet are intimately related as they prompt the same set of questions: how can we characterize and optimize the system performance as measured by the *price of anarchy*?

For the interested reader, the authors provide a software package, available in both MATLAB® and Python, that implements the techniques described in this manuscript at https://github.com/rahul-chandan/resalloc-poa.

Authors' addresses: R. Chandan, Department of Electrical and Computer Engineering, University of California, Santa Barbara (e-mail: rchandan@ucsb.edu); D. Paccagnan, Department of Computing, Imperial College London (e-mail: d.paccagnan@imperial.ac.uk); J. R. Marden, Department of Electrical and Computer Engineering, University of California, Santa Barbara (e-mail: jrmarden@ece.ucsb.edu).

1.1 Motivating Example #1 : Incentive design in congestion games

A widely studied model for self-interested resource allocation problems is that of (atomic) congestion games [39]. A congestion game consists of a set of users $N = \{1, \dots, n\}$ sharing the use of a common set of resources $\mathcal{R}$, where each resource $r \in \mathcal{R}$ is associated with a congestion function $c_r : \{1, \dots, n\} \to \mathbb{R}$. The term $c_r(k)$ identifies the cost a user experiences for selecting resource r given that there are $1 \leq k \leq n$ users concurrently selecting resource r. Further, each user $i \in N$ is associated with a given action set $\mathcal{A}_i \subseteq 2^{\mathcal{R}}$ that meets the individual needs. Within the context of traffic routing, for example, an action $a_i \in \mathcal{A}_i$ describes a path in the network connecting the user's source to destination. Given an admissible allocation of users to resources $a = (a_1, a_2, \dots, a_n) \in \mathcal{A} = \mathcal{A}_1 \times \cdots \times \mathcal{A}_n$, the system cost describes the sum of the costs incurred by all users, i.e.,

$$C(a) = \sum_{i \in N} \sum_{r \in a_i} c_r(|a|_r), \tag{1}$$

where $|a|_r = |\{i \in N : r \in a_i\}|$ denotes the number of users selecting resource r in allocation a. In this setting, an optimal allocation of users to resources consists in $a^{\text{opt}} \in \arg\min_{a \in \mathcal{A}} C(a)$.

One of the fundamental challenges associated with the allocation of resources in this problem setting is that users are often modeled as selfish decision makers. Specifically, each user $i \in N$ independently chooses an action $a_i \in \mathcal{A}_i$ with the aim of minimizing the individual cost $J_i : \mathcal{A} \to \mathbb{R}$ incurred over the selected resources, i.e.,

$$J_i(a) = \sum_{r \in \mathcal{R}} c_r(|a|_r). \tag{2}$$

The resulting allocation is then suitably described as a pure Nash equilibrium $a^{\text{ne}} \in \mathcal{A}$ of the game, or a generalization thereof (e.g., mixed Nash, correlated/coarse correlated equilibria). Accordingly, there has been significant research seeking to quantify the quality of Nash equilibria relative to the optimal allocation. This is typically measured using the notion of *price of anarchy* [29], i.e., the worst case ratio $C(a^{\text{ne}})/C(a^{\text{opt}})$ across a family of games.

The analysis of the price of anarchy in congestion games has a rich history and clearly demonstrates the inefficiencies associated with self interested decision making [1, 3, 16, 43]. Given these inefficiencies, there is also significant research interest in the design of incentives that alter the users' experienced costs, thereby influencing the set of Nash equilibria and, thus, improving the price of anarchy [7, 9, 17, 24, 27, 35]. Within this context, each resource $r \in \mathcal{R}$ is commonly associated with an incentive function $\tau_r : \{1, \dots, n\} \to \mathbb{R}$, where $\tau_r(k)$ denotes the incentive imposed on resource r when there are k users selecting it.[1] As a result, user $i \in N$ experiences a cost accounting for both the congestion on the resources and the imposed incentives, i.e.,

$$J_i(a) = \sum_{r \in a_i} \left[c_r(|a|_r) + \tau_r(|a|_r) \right]. \tag{3}$$

It should be stressed that, while the incentives modify the users' cost functions, they do not alter the assessment of the system cost, which still takes the form in (1). Thus, unless all incentive are identically zero, the sum of the users' costs in (3) does not equal the system cost in (1).

The objective of a system operator is to design admissible incentives, i.e., functions $\{\tau_r\}_{r \in \mathcal{R}}$ that satisfy a viable constraint on monetary budget, to improve the quality of the resulting collective behaviour. While this objective may seem deceptively straightforward, it requires the following theoretical advances which are currently unresolved:

[1]We will use the terminology of "tax" and "rebate" when referring to incentives satisfying $\tau_r(k) \geq 0$ and $\tau_r(k) \leq 0$ for all r and k, respectively.

– **Characterization:** What is the price of anarchy for a given set of incentives $\{\tau_r\}_{r\in\mathcal{R}}$?

– **Optimization:** What are admissible incentives $\{\tau_r\}_{r\in\mathcal{R}}$ that optimize the price of anarchy?

1.2 Motivating example #2 : Distributed coordination of multiagent systems

Alternatively, in systems where the set of users is a group of autonomous, computer controlled entities, the users' decision making processes can be explicitly selected by the system designer. Resource allocation problems represent a particular class of such multiagent systems. In a resource allocation problem a set of n users, $N = \{1, \ldots, n\}$, must be allocated to a set of resources $\mathcal{R}$. Each user $i \in N$ has a corresponding set of permissible actions, $\mathcal{A}_i$. For a given allocation $a = (a_1, \ldots, a_n) \in \mathcal{A}_1 \times \cdots \times \mathcal{A}_n = \mathcal{A}$, we consider a system welfare with the following separable form:

$$W(a) = \sum_{r\in\mathcal{R}} W_r(|a|_r),$$

where we refer to $W_r : \{1, \ldots, n\} \to \mathbb{R}$ as the resource welfare function on resource r, and denote with $a^{\text{opt}} \in \arg\max_{a\in\mathcal{A}} W(a)$ an optimal allocation. Applications of this model are not limited to distributed sensing [23, 47], resource allocation [20, 24] and task assignment [2, 15].

Requirements for scalability and security of multiagent systems compounded with constraints on the users' communication and computation capabilities make centralized coordination undesirable or even impossible. Thus, there has recently been increased interest in the distributed coordination of such systems [13, 23, 24, 26], where the users determine their actions according to a prescribed local decision making process. A natural paradigm for the design of distributed coordination algorithms consists in i) the assignment of a local utility function to each user, and ii) the choice of a learning rule that specifies each user's decision making process in light of the perceived utility [31, 42]. The class of distributed welfare games [31] provides a framework for studying resource allocation problems under this lens. In this context, each user $i \in N$ is associated with a utility function $U_i : \mathcal{A} \to \mathbb{R}$ of the following form:

$$U_i(a) = \sum_{r\in a_i} F_r(|a|_r),$$

where the utility generating functions $F_r : \{1, \ldots, n\} \to \mathbb{R}$, $r \in \mathcal{R}$, are *subject to our design.* As such, the sum of the users' utilities need not be equal to the system welfare. Remarkably, the performance (formally, the approximation ratio) of many learning rules including best response and no-regret dynamics matches the price of anarchy of a corresponding game whereby each user in N is associated to the action set $\mathcal{A}_i$ and utility U_i, see [40]. Thus, in order to derive efficient coordination algorithms, we must address the following questions:

– **Characterization:** What is the price of anarchy of a given set of utility generating functions?

– **Optimization:** What are utility generating functions that optimize the price of anarchy?

1.3 Our contributions

The focus of this manuscript is on developing an exact, computationally efficient technique to address both the *characterization* and *optimization* questions highlighted above for a broad class of games that includes congestion games and distributed welfare games. The specific contributions associated with this paper are as follows:

(1) In Section 2, we propose a generalization of the smoothness framework introduced by Roughgarden [40]. We show that for *any* cost minimization game (or analagous welfare maximization

game), this new framework gives an improved bound on the price of anarchy when compared to the original smoothness framework (Proposition 2.2).

(2) Our second result focuses on a generalization of congestion games where the system-level objective is not necessarily equal to the sum of the users' individual welfares. Here, we demonstrate that our framework *tightly* characterizes the price of anarchy for any such game (Theorem 3.5). This is in contrast to the original smoothness framework which provides only an upper bound on the price of anarchy in these settings.
(3) Our third result shifts attention to the efficient characterization of the price of anarchy. In particular, the problem of characterizing the price of anarchy is transformed to the problem of computing an optimal set of parameters (λ, μ) over a given admissible space. In Theorem 3.5, we provide a tractable linear program that can characterize the exact price of anarchy for any set of generalized congestion games using our framework.
(4) Our fourth result focuses on optimizing the price of anarchy. Specifically, we show that the linear program for computing the price of anarchy can be modified to derive incentives that minimize the price of anarchy in generalized congestion games (Theorem 3.7).
(5) Lastly, in Section 5, we demonstrate the applicability of our methodology to the problems of incentive design in congestion games and utility design in distributed welfare games introduced in Sections 1.1 and 1.2. Additionally, we show how our approach recovers and unifies a variety of existing results on the price of anarchy.

All above results extend unchanged to mixed Nash, correlated and coarse correlated equilibria.

1.4 Related literature

The notion of price of anarchy was introduced by Koutsoupias and Papadimitriou in 1999 as a metric to quantify equilibrium performance [29]. While a number of works have initially derived bounds on such metric, the breakthrough in the analysis of the price of anarchy came with the introduction of smoothness style arguments in two studies on atomic congestion games with affine latency functions [3, 16]. The smoothness framework was later formalized and generalized by Roughgarden [40]. This approach has not only proven to be useful for characterizing the efficiency of many classes of equilibria [1, 10, 34], it has also been applied more broadly in problems including learning [22] and mechanism design [44]. The smoothness framework provides several advantages when deriving bounds on the price of anarchy: it is tight for well studied families of games; and, it consists of a standard set of linear inequalities that govern the price of anarchy bound. However, as we show in this manuscript, the original smoothness argument does not provide exact bounds on the price of anarchy in settings when the social welfare is not aligned with the system-level objective, i.e., $\sum_i J_i(a) \neq C(a)$, and thus provides an inaccurate approach to the problems of incentive design and utility design. The generalization of smoothness proposed in this manuscript resolves these deficiencies, while retaining the strengths of the original smoothness approach.

The notion of generalized smoothness presented in this work is most similar to the style of argument used by Gairing [23] to quantify the price of anarchy of covering problems. This work also builds upon the results of Paccagnan et al. [36], who provide a linear programming framework for characterizing and optimizing the efficiency of pure Nash equilibria in restricted classes of resource allocation games. Generalized smoothness permits a non-trivial extension of their framework: we are now able to construct linear programs for computing and optimizing the coarse-correlated equilibrium efficiency, relative to a broader class of problems that includes the class of atomic congestion games. For an in-depth study on optimal local incentive design within the class of atomic congestion games, we refer the interested reader to Paccagnan et al. [35].

Our derivation of a linear programming technique to compute upper and lower bounds on the price of anarchy is inspired by the *primal-dual approach* [6, 33]. The primal-dual approach was used in Nadav and Roughgarden [33] to understand when the smoothness bound from [40] is exact. It was then specialized to the class of weighted congestion games in [6], applied to bound the efficiency of approximate Nash equilibria in [5] and used to identify the best achievable price of anarchy for weighted polynomial congestion games with incentives based on the optimal allocation in [7]. It is important to note that these prior works also propose linear programming techniques for computing upper and lower bounds on the price of anarchy. However, as we discuss in Section 3.4, our technique provides an *exact* and *computationally tractable* characterization of the price of anarchy, while the techniques introduced in the above works are either inexact (i.e., the bounds do not always match) or are not computationally efficient (i.e., the complexity of computing the bounds grows exponentially in the number of users n) for the class of generalized congestion games considered in this paper. The improvements we obtain stem from the formalization of the generalized smoothness framework in addition to the use of a succinct parameterization that guarantees tightness of the price of anarchy bound. As the price of anarchy bound we obtain is exact, there is no further analysis required: our linear program automatically generates a worst case game instance (lower bound) and a matching generalized smoothness argument (upper bound). Furthermore, our framework can be modified to efficiently compute incentives and utilities that optimize the price of anarchy.

1.5 Outline

This article is organized as follows. Section 2 defines the class of games and the performance metrics that we consider throughout this paper, reviews the original notion of smoothness [40] and defines the novel generalized smoothness argument. Section 3 refines our study to the class of generalized congestion games, presents our results relating to the characterization of tight and tractable bounds on the price of anarchy using the primal-dual approach in conjunction with a novel game parameterization and the derivation of optimal incentives under this specialized game model. Section 4 presents analogous results for the welfare maximization problem setting without proof. Section 5 applies our theoretical results to the problems of incentive design in congestion games and utility design in distributed welfare games. Section 6 includes our conclusions and a brief discussion on future work.

2 GENERALIZED SMOOTHNESS IN COST MINIMIZATION GAMES

This section introduces the class of games and performance metrics used throughout this paper. We proceed to review the smoothness framework from Roughgaren [40] and highlight its limitations. We then introduce a revised framework, termed generalized smoothness, that alleviates these limitations and improves upon the efficiency guarantees provided by the original smoothness framework.

2.1 Cost minimization games

We consider the class of cost minimization problems in which there is a set of users $N = \{1, \ldots, n\}$, and where each user $i \in N$ is associated with a given action set $\mathcal{A}_i$ and a cost function $J_i : \mathcal{A} \to \mathbb{R}$. The system cost induced by an allocation $a = (a_1, \ldots, a_n) \in \mathcal{A} = \mathcal{A}_1 \times \cdots \times \mathcal{A}_n$ is measured by the function $C : \mathcal{A} \to \mathbb{R}_{>0}$, and an optimal allocation satisfies

$$a^{\text{opt}} \in \arg\min_{a \in \mathcal{A}} C(a). \tag{4}$$

We represent a cost minimization game as defined above as a tuple $G = (N, \mathcal{A}, C, \mathcal{J})$, where $\mathcal{J} = \{J_1, \ldots, J_n\}$. Note that the example highlighted in Section 1.1 represents a special class of cost minimization games, where the users' cost functions and the system cost are separable over a given set of shared resources.

The main focus of this work is on characterizing the degradation in system wide performance resulting from local decision making. To that end, we focus on the solution concept of (pure) Nash equilibrium as a model of the emergent behaviour in such systems. A Nash equilibrium is defined as any allocation $a^{\text{ne}} \in \mathcal{A}$ such that

$$J_i(a^{\text{ne}}) \leq J_i(a_i, a^{\text{ne}}_{-i}) \quad \forall a_i \in \mathcal{A}_i, \forall i \in N. \tag{5}$$

For a given game G, let $\text{NE}(G)$ denote the set of all allocations $a \in \mathcal{A}$ that satisfy Equation (5). Assuming the set $\text{NE}(G)$ is non-empty, we define the *price of anarchy* of the game G as

$$\text{PoA}(G) := \frac{\max_{a\in\text{NE}(G)} C(a)}{\min_{a\in\mathcal{A}} C(a)} \geq 1. \tag{6}$$

The price of anarchy represents the ratio between the costs of the worst-performing pure Nash equilibrium in the game G, and the optimal allocation. For a given class of cost minimization games $\mathcal{G}$, which may contain infinitely many game instances, we further define the price of anarchy as,

$$\text{PoA}(\mathcal{G}) := \sup_{G\in\mathcal{G}} \text{PoA}(G) \geq 1. \tag{7}$$

Note that a lower price of anarchy corresponds to an improvement in worst case equilibrium performance and $\text{PoA}(\mathcal{G}) = 1$ implies that all Nash equilibria of all games $G \in \mathcal{G}$ are optimal.

2.2 Smoothness in cost minimization games

The framework of (λ, μ)-smoothness, introduced in [40], is widely used in the existing literature aimed at characterizing the price of anarchy over various classes of games. A cost minimization game G is termed (λ, μ)-smooth if the following two conditions are met:

(i) For all $a \in \mathcal{A}$, we have $\sum_{i=1}^n J_i(a) \geq C(a)$;

(ii) For all $a, a' \in \mathcal{A}$, there exist $\lambda > 0$ and $\mu < 1$ such that

$$\sum_{i\in N} J_i(a'_i, a_{-i}) \leq \lambda C(a') + \mu C(a). \tag{8}$$

If a game G is (λ, μ)-smooth, then the price of anarchy of game G is upper bounded by

$$\text{PoA}(G) \leq \frac{\lambda}{1-\mu}.$$

Observe that if all the games in a class $\mathcal{G}$ are shown to be (λ, μ)-smooth, then the price of anarchy of the class $\text{PoA}(\mathcal{G})$ is also upper bounded by $\lambda/(1-\mu)$. We refer to the best upper bound obtainable using a smoothness argument on a given class of games $\mathcal{G}$ as the *robust price of anarchy*, i.e.,

$$\text{RPoA}(\mathcal{G}) := \inf_{\lambda>0,\mu<1} \left\{ \frac{\lambda}{1-\mu} \text{ s.t. Equation (8) holds } \forall G \in \mathcal{G} \right\}. \tag{9}$$

It is important to note that the robust price of anarchy represents only an upper bound on the price of anarchy, i.e., for any class of (λ, μ)-smooth games $\mathcal{G}$, it holds that $\text{PoA}(\mathcal{G}) \leq \text{RPoA}(\mathcal{G})$, where it could be that $\text{PoA}(\mathcal{G}) < \text{RPoA}(\mathcal{G})$.

2.3 Generalized smoothness in cost minimization games

In this section, we provide a generalization of the smoothness framework, termed *generalized smoothness*. We will then proceed to show how this new framework provides tighter efficiency bounds and covers a broader spectrum of problem settings than the original smoothness framework, defined in the previous section.

Definition 2.1 (Generalized smoothness). The cost minimization game G is (λ, μ)-generalized smooth if, for any two allocations $a, a' \in \mathcal{A}$, there exist $\lambda > 0$ and $\mu < 1$ satisfying,

$$\sum_{i=1}^{n} J_i(a_i', a_{-i}) - \sum_{i=1}^{n} J_i(a) + C(a) \leq \lambda C(a') + \mu C(a). \tag{10}$$

Note that we maintain the notation of (λ, μ) as in the original notion of smoothness for ease of comparison. In the specific case when $\sum_{i=1}^{n} J_i(a) = C(a)$ for all $a \in \mathcal{A}$, observe that the smoothness conditions in Equation (10) are equivalent to the original smoothness conditions in Equation (8). As with Equation (9), we define the *generalized price of anarchy* of a class of cost minimization games $\mathcal{G}$ as the best upper bound obtainable using a generalized smoothness argument, i.e.,

$$\text{GPoA}(\mathcal{G}) := \inf_{\lambda > 0, \mu < 1} \left\{ \frac{\lambda}{1 - \mu} \text{ s.t. Equation (10) holds } \forall G \in \mathcal{G} \right\}. \tag{11}$$

In our first result we show that (i) price of anarchy bounds under the generalized smoothness framework follow in the same way as the original smoothness framework without the restriction that $\sum_{i=1}^{n} J_i(a) \geq C(a)$ for all $a \in \mathcal{A}$ and (ii) the generalized smoothness framework provides stronger bounds on the price of anarchy than the original smoothness framework whenever both are defined.

Proposition 2.2. *For any (λ,μ)-generalized smooth game G, the following statements hold:*

(i) The price of anarchy of G is upper bounded as $\text{PoA}(G) \leq \lambda/(1 - \mu)$.

(ii) If the game G is (λ,μ)-smooth, then $\text{RPoA}(G) \geq \text{GPoA}(G) \geq \text{PoA}(G)$. *Furthermore, if* $\sum_{i=1}^{n} J_i(a) > C(a)$ *holds for all $a \in \mathcal{A}$, then* $\text{RPoA}(G) > \text{GPoA}(G) \geq \text{PoA}(G)$.

Proof. Proof. For the proof of statement (i), observe that, for all $a^{\text{ne}} \in \text{NE}(G)$ and $a^{\text{opt}} \in \mathcal{A}$,

$$C(a^{\text{ne}}) \leq \sum_{i=1}^{n} J_i(a_i^{\text{opt}}, a_{-i}^{\text{ne}}) - \sum_{i=1}^{n} J_i(a^{\text{ne}}) + C(a^{\text{ne}}) \leq \lambda\, C(a^{\text{opt}}) + \mu\, C(a^{\text{ne}}). \tag{12}$$

The inequalities hold by Equations (5) and (10), respectively. Rearranging gives the result.

The remainder of the proof focuses on statement (ii). Since the condition $\sum_{i=1}^{n} J_i(a) \geq C(a)$ for all $a \in \mathcal{A}$ implies that any pair of (λ, μ) satisfying Equation (8) necessarily satisfies Equation (10), we note that the generalized price of anarchy is less than or equal to the robust price of anarchy, i.e., $\text{RPoA}(G) \geq \text{GPoA}(G) \geq \text{PoA}(G)$.

Note that for any game $G = (N, \mathcal{A}, C, \mathcal{J})$ with $\sum_{i=1}^{n} J_i(a) > C(a)$ for all $a \in \mathcal{A}$ there must exist a uniform scaling factor $0 < \gamma < 1$ such that $\sum_{i=1}^{n} \gamma J_i(a) \geq C(a)$, but for which the price of anarchy remains the same, i.e., for $G' = (N, \mathcal{A}, C, \mathcal{J}')$ where $\mathcal{J}' = \{\gamma J_1, \ldots, \gamma J_n\}$, it holds that $\text{PoA}(G') = \text{PoA}(G)$. The price of anarchy remains the same despite the rescaling, because the inequalities in Equation (5) are unaffected by a positive scaling factor (i.e., $\text{NE}(G) = \text{NE}(G')$), and because the optimal cost remains unchanged since the scaling does not impact the system cost. Further, one can verify from Equation (8) that $\text{RPoA}(G) > \text{RPoA}(G')$, and thus $\text{RPoA}(G) > \text{RPoA}(G') \geq \text{PoA}(G') = \text{PoA}(G)$. Finally, we know that $\text{GPoA}(G')$ is less than or equal to $\text{RPoA}(G')$ and can verify from Equation (10) that $\text{GPoA}(G) = \text{GPoA}(G')$. Thus,

$$\text{RPoA}(G) > \text{RPoA}(G') \geq \text{GPoA}(G') = \text{GPoA}(G) \geq \text{PoA}(G).$$

$\square$

Further comparisons between the original notion of smoothness and generalized smoothness can be made, as summarized by the following observations. These observations are stated without proof for brevity, but can easily be verified by the reader.

– *Observation #1*: The price of anarchy and generalized price of anarchy are shift-, and scale-invariant, i.e., for any given $\gamma > 0$ and $(\delta_1, \dots, \delta_n) \in \mathbb{R}^n$,

$$\mathrm{PoA}((N, \mathcal{A}, C, \{J_i\}_{i=1}^n)) = \mathrm{PoA}((N, \mathcal{A}, C, \{\gamma J_i + \delta_i\}_{i=1}^n)),$$
$$\mathrm{GPoA}((N, \mathcal{A}, C, \{J_i\}_{i=1}^n)) = \mathrm{GPoA}((N, \mathcal{A}, C, \{\gamma J_i + \delta_i\}_{i=1}^n)).$$

Neither of these properties hold for the robust price of anarchy.

– *Observation #2*: The robust price of anarchy is optimized by budget-balanced user cost functions, i.e., $\sum_{i \in N} J_i(a) = C(a)$ for all $a \in \mathcal{A}$. In general, this does not hold for the price of anarchy and generalized price of anarchy.

– *Observation #3*: For a given cost minimization game G, we define an average coarse-correlated equilibrium as a probability distribution $\sigma \in \Delta(\mathcal{A})$ satisfying, for all $a' \in \mathcal{A}$,

$$\mathbb{E}_{a \sim \sigma}\left[\sum_{i=1}^{N} J_i(a)\right] := \sum_{a \in \mathcal{A}}\left[\sigma_a \sum_{i=1}^{N} J_i(a)\right] \leq \sum_{a \in \mathcal{A}}\left[\sigma_a \sum_{i=1}^{N} J_i(a_i, a'_{-i})\right], \tag{13}$$

where $\sigma_a \in [0, 1]$ is the probability associated with action $a \in \mathcal{A}$ in the distribution σ. Note that the set of average coarse correlated equilibria contains all of the game's pure Nash equilibria, mixed Nash equilibria, correlated equilibria and coarse-correlated equilibria [40]. The generalized price of anarchy tightly characterizes the average coarse correlated equilibrium performance of any cost minimization game G, and, thus, of any class of cost minimization games $\mathcal{G}$. The proof follows identically to the result by Nadav and Roughgarden [33] that proves this claim for the robust price of anarchy under an alternative definition of average coarse correlated equilibrium. The two equilibrium definitions match for games with $\sum_{i=1}^n J_i(a) = C(a)$.

3 GENERALIZED SMOOTHNESS IN GENERALIZED CONGESTION GAMES

The previous section introduced the framework of generalized smoothness and showed that the resulting generalized price of anarchy provides improved bounds on the price of anarchy when compared with the robust price of anarchy. However, deriving the generalized price of anarchy still requires solving for the optimal λ and μ given in Equation (11). In this section, we show that the optimal parameters λ and μ for a generalization of the well-studied class of congestion games can be computed as solutions of a tractable linear program. Furthermore, we demonstrate that the generalized price of anarchy tightly characterizes the price of anarchy for this important class of games, extending the canonical results on the robustness of the price of anarchy from Roughgarden [40] to the broader class of generalized congestion games that we present below.

3.1 Generalized congestion games

In this section, we consider a generalization of the congestion game framework that consists of a user set $N = \{1, \dots, n\}$ and a resource set $\mathcal{R}$. Given an allocation $a \in \mathcal{A}$, the system cost and user cost functions have the following separable structure:

$$C(a) = \sum_{r \in \mathcal{R}} C_r(|a|_r), \tag{14}$$

$$J_i(a_i, a_{-i}) = \sum_{r \in a_i} F_r(|a|_r), \tag{15}$$

where $C_r : \{0, 1, \dots, n\} \to \mathbb{R}_{\geq 0}$ and $F_r : \{1, \dots, n\} \to \mathbb{R}$ define the resource cost functions and cost generating functions, respectively. We will denote a congestion game by the tuple $G =$

$(N, \mathcal{R}, \mathcal{A}, \{C_r, F_r\}_{r\in\mathcal{R}})$. This game model covers many of the existing models studied in the game theoretic literature, including congestion games [39].

Example 3.1 (Congestion Games). In congestion games, each resource $r \in \mathcal{R}$ is associated with a congestion function $c_r : \{1, \dots, n\} \to \mathbb{R}_{\geq 0}$. Here, the resource cost and cost generating functions are $C_r(k) = k \cdot c_r(k)$ and $F_r(k) = c_r(k)$ for any $k \geq 1$. Note that $C_r(k) = k \cdot F_r(k)$ for this case, hence the definitions of smoothness and generalized smoothness coincide.

Example 3.2 (Congestion Games with Incentives). When incentives are introduced into the congestion game setup, each resource $r \in \mathcal{R}$ is also associated with an incentive function $\tau_r : \{1, \dots, n\} \to \mathbb{R}$. For this class of games, the resource cost and cost generating functions take on the form where $C_r(k) = k \cdot c_r(k)$ and $F_r(k) = c_r(k) + \tau_r(k)$ for any $k \geq 1$. Note that for the case when $\tau_r(k) > 0$ for all r (i.e., taxes), then $C_r(k) < k \cdot F_r(k)$ and the generalized price of anarchy provides a strictly closer bound on the price of anarchy than the robust price of anarchy, by Proposition 2.2. Furthermore, when $\tau_r(k) < 0$ for all r (i.e., rebates), then $C_r(k) > k \cdot F_r(k)$ and the original smoothness framework is inadmissible, by definition.

3.2 Tight price of anarchy for generalized congestion games

Our goal in this section is to characterize the price of anarchy for a given set of generalized congestion games $\mathcal{G}$. We begin by defining our set of games.

Definition 3.3. A generalized congestion game $G = (N, \mathcal{R}, \mathcal{A}, \{C_r, F_r\}_{r\in\mathcal{R}})$ is generated from basis function pairs $\{C^j, F^j\}$, $j = 1, \dots, m$, if there exists a set of coefficients $\alpha_r^1, \dots \alpha_r^m \geq 0$ such that $C_r(k) = \sum_{j=1}^m \alpha_r^j \cdot C^j(k)$ and $F_r(k) = \sum_{j=1}^m \alpha_r^j \cdot F^j(k)$ for any $k \in \{1, \dots, n\}$ and any $r \in \mathcal{R}$.

We let $\mathcal{G}$ denote the set of all generalized congestion games with a maximum of n users that can be generated from given basis function pairs $\{C^j, F^j\}$, $j = 1, \dots, m$. The following example demonstrates how a limited set of basis function pairs can actually model a diverse set of games.

Example 3.4 (Affine and Polynomial Congestion Games). One commonly studied class of congestion games is affine congestion games, where each resource $r \in \mathcal{R}$ is associated with a cost generating function $F_r(k) = a_r \cdot k + b_r$ and resource cost function $C_r(k) = k \cdot (a_r \cdot k + b_r) = k \cdot F_r(k)$ for any $k \geq 1$, where $a_r, b_r \geq 0$. Observe that all admissible function pairs $\{C_r, F_r\}$ can be represented as linear combinations of the basis function pairs $\{C^1, F^1\}, \{C^2, F^2\}$ where $\{C^1(k), F^1(k)\} = \{k, 1\}$ (case where $a_r = 0$ and $b_r = 1$) and $\{C^2(k), F^2(k)\} = \{k^2, k\}$ (case where $a_r = 1$ and $b_r = 0$). Similarly, the function pairs $\{C_r, F_r\}$ of any polynomial congestion game of degree $d \geq 1$, i.e., where each resource is associated with a cost generating function of the form $F_r(k) = \sum_{j=1}^{d+1} \alpha_r^j \cdot k^{j-1}$ such that $\alpha_r^1, \dots, \alpha_r^{d+1} \geq 0$ and $C_r(k) = k \cdot F_r(k)$, can be represented as linear combinations of $d+1$ basis function pairs in the same fashion as in the affine case.

The following theorem provides our main contribution pertaining to the price of anarchy in generalized congestion games. Throughout, we define $C^j(0) = F^j(0) = F^j(n+1) = 0$, $j = 1, \dots, m$, for ease of notation and without loss of generality. Additionally, $\mathcal{I}_{\mathcal{R}}(n)$ is defined as the set of all triplets $(x, y, z) \in \{0, 1, \dots, n\}^3$ that satisfy: (i) $1 \leq x + y - z \leq n$ and $z \leq \min\{x, y\}$; and, (ii) $x + y - z = n$ or $(x - z)(y - z)z = 0$. The structure of the set $\mathcal{I}_{\mathcal{R}}(n)$ comes from our game parameterization and will be fully addressed in the proof of Theorem 3.5.

Theorem 3.5. *Let $\mathcal{G}$ denote the set of all generalized congestion games with a maximum of n users generated from basis function pairs $\{C^j, F^j\}$, $j = 1, \dots, m$, and let ρ^{opt} be the optimal value of the* *following (tractable) linear program:*

$$\rho^{\text{opt}} = \underset{\nu \in \mathbb{R}_{\geq 0}, \rho \in \mathbb{R}}{\textit{maximize}} \quad \rho$$
$$\textit{subject to:} \quad C^j(y) - \rho C^j(x) + \nu[(x-z)F^j(x) - (y-z)F^j(x+1)] \geq 0, \qquad \forall j = 1, \dots, m, \quad \forall (x,y,z) \in \mathcal{I}_{\mathcal{R}}(n). \tag{16}$$

Then, it holds that $\text{PoA}(\mathcal{G}) = \text{GPoA}(\mathcal{G}) = 1/\rho^{\text{opt}}$.

There are two significant findings associated with Theorem 3.5. First, observe that the generalized price of anarchy achieves a tight bound on the price of anarchy for any set of generalized congestion games. Therefore, there is no loss in characterizing the price of anarchy using the generalized smoothness bound. Second, the price of anarchy associated with a set of generalized congestion games $\mathcal{G}$ can be characterized by means of a tractable linear program that scales linearly in its complexity with the number of basis function pairs, m, and quadratically with the number of users n. Thus, there are computationally efficient mechanisms for characterizing the price of anarchy in a given set of congestion games.

3.3 Proof of Theorem 3.5

The following informal outline of the proof for Theorem 3.5 is directly followed by the formal proof, which follows a similar structure:
– *Step 1* : We define our game parameterization, which represents any generalized congestion game $G \in \mathcal{G}$ with $O(mn^3)$ parameters $\theta(x,y,z,j) \geq 0$ corresponding with basis pairs $\{(C^j, F^j)\}$, $j = 1, \dots, m$, and triplets $x, y, z \in \{0, \dots, n\}$ such that $1 \leq x + y - z \leq n$ and $z \leq \min\{x, y\}$.
– *Step 2* : For any class of generalized congestion games $\mathcal{G}$, we observe that an upper bound on the generalized price of anarchy can be computed as a fractional program with $m \times |\mathcal{I}_{\mathcal{R}}(n)|$ constraints under the game parameterization presented in Step 1.
– *Step 3* : Following a change of variables, we observe that the linear program in Equation (16) is equivalent to the fractional program from Step 2. We then provide a game $G \in \mathcal{G}$ with price of anarchy equal to the upper bound on the generalized price of anarchy, implying that $\text{PoA}(G) \geq \text{GPoA}(\mathcal{G})$. Since $\text{PoA}(G) \leq \text{PoA}(\mathcal{G}) \leq \text{GPoA}(\mathcal{G})$, it must then be that $\text{PoA}(G) = \text{PoA}(\mathcal{G}) = \text{GPoA}(\mathcal{G})$ for any class of generalized congestion games $\mathcal{G}$, concluding the proof.

Proof of Theorem 3.5. The proof is shown in three steps, corresponding with the informal outline:

– *Step 1* : For a given game $G \in \mathcal{G}$, our game parameterization is defined as follows for allocations $a, a' \in \mathcal{A}$: For every resource $r \in \mathcal{R}$, we define integers $x_r, y_r, z_r \geq 0$ where $x_r = |a|_r$ is the number of users that select r in a, $y_r = |a'|_r$ is the number of users that select r in a' and $z_r = |\{i \in N \text{ s.t. } r \in a_i\} \cap \{i \in N \text{ s.t. } r \in a'_i\}|$ is the number of users that select r in both a and a'. Note that $1 \leq x_r + y_r - z_r \leq n$ and $z_r \leq \min\{x_r, y_r\}$ must hold for all $r \in \mathcal{R}$. For all $x, y, z \geq 0$ such that $1 \leq x + y - z \leq n$ and $z \leq \min\{x, y\}$, and all $j = 1, \dots, m$, we define the parameters

$$\theta(x,y,z,j) = \sum_{r \in \mathcal{R}(x,y,z)} \alpha_r^j, \tag{17}$$

where $\mathcal{R}(x,y,z) = \{r \in \mathcal{R} \text{ s.t. } (x_r, y_r, z_r) = (x,y,z)\}$. and $\alpha_r^j \geq 0$, $j = 1, \dots, m$, are the coefficients in the basis representation of the resource cost function C_r and cost generating function F_r. Although the parameterization into values $\theta(x,y,z,j) \geq 0$ is of size $O(mn^3)$, we show in Step 2 that only $O(mn^2)$ parameters are needed in the computation of the price of anarchy.

– *Step 2* : For any generalized congestion game $G \in \mathcal{G}$, we denote an optimal allocation as a^{opt}, and a Nash equilibrium as a^{ne}, i.e. $a^{\text{ne}} \in \text{NE}(G)$ such that $\text{PoA}(G) \geq C(a^{\text{ne}})/C(a^{\text{opt}})$. We observe that using the above definitions of (x_r, y_r, z_r) for $a = a^{\text{ne}}$ and $a' = a^{\text{opt}}$, it follows that

$$\sum_{i=1}^{n} J_i(a_i^{\text{opt}}, a_{-i}^{\text{ne}}) = \sum_{r \in \mathcal{R}} (y_r - z_r) F_r(x_r + 1) + z_r F_r(x_r).$$

Informally, if a user $i \in N$ selects a given resource $r \in \mathcal{R}$ in both a_i^{ne} and a_i^{opt}, then by deviating from a_i^{ne} to a_i^{opt}, the user does not add to the load on r, i.e., $|a_i^{\text{opt}}, a_{-i}^{\text{ne}}|_r = |a^{\text{ne}}|_r = x_r$. However, if $r \in a_i^{\text{opt}}$ and $r \notin a_i^{\text{ne}}$, then $|a_i^{\text{opt}}, a_{-i}^{\text{ne}}|_r = |a^{\text{ne}}|_r + 1 = x_r + 1$.

Recall that for all $r \in \mathcal{R}$, it must hold that $z_r \leq \min\{x_r, y_r\}$, and $1 \leq x_r + y_r - z_r \leq n$. We define the set of triplets $\mathcal{I}(n) \subseteq \{0, 1, \ldots, n\}^3$ as

$$\mathcal{I}(n) := \{(x, y, z) \in \mathbb{N}^3 \text{ s.t. } 1 \leq x + y - z \leq n \text{ and } z \leq \min\{x, y\}\},$$

and $\gamma(\mathcal{G})$ as the value of the following fractional program:

$$\begin{aligned} \gamma(\mathcal{G}) := \inf_{\lambda > 0, \mu < 1} \quad & \frac{\lambda}{1 - \mu} \\ \text{s.t.} \quad & (z - x) F^j(x) + (y - z) F^j(x + 1) + C^j(x) \leq \lambda C^j(y) + \mu C^j(x), \\ & \qquad \forall j = 1, \ldots, m, \quad \forall (x, y, z) \in \mathcal{I}(n). \end{aligned} \tag{18}$$

Observe that, by Equation (17), the generalized smoothness condition in Equation (10) can be rewritten within the context of generalized congestion games as

$$\begin{aligned} & \sum_{(x,y,z) \in \mathcal{I}(n)} \sum_{j=1}^{m} \left[(x - z) F^j(x) - (y - z) F^j(x + 1) + C^j(x) \right] \theta(x, y, z, j) \\ \leq & \sum_{(x,y,z) \in \mathcal{I}(n)} \sum_{j=1}^{m} \left[\lambda C^j(y) + \mu C^j(x) \right] \theta(x, y, z, j) \end{aligned}$$

It must then hold that for any pair (λ, μ) in the feasible set of the fractional program in Equation (18), all games $G \in \mathcal{G}$ are (λ, μ)-generalized smooth, i.e., $\gamma(\mathcal{G}) \geq \text{GPoA}(\mathcal{G})$. This is because the generalized smoothness condition for generalized congestion games can be expressed as a weighted sum with positive coefficients over a subset of the constraints in Equation (18).

To conclude Step 2 of the proof, we show that it is sufficient to define $\gamma(\mathcal{G})$ in Equation (18) over the reduced set of constraints corresponding to $j \in \{1, \ldots, m\}$ and triplets in $\mathcal{I}_{\mathcal{R}}(n) \subseteq \mathcal{I}(n)$, where

$$\mathcal{I}_{\mathcal{R}}(n) := \{(x, y, z) \in \mathcal{I}(n) \text{ s.t. } x + y - z = n \text{ or } (x - z)(y - z)z = 0\}.$$

For each $j \in \{1, \ldots, m\}$ and any $(x, y, z) \in \mathcal{I}(n)$, observe that the constraint in Equation (18) is equivalent to $yF^j(x+1) - xF^j(x) + z[F^j(x) - F^j(x+1)] \leq \lambda C^j(y) + (\mu - 1)C^j(x)$. If $F^j(x+1) \geq F^j(x)$, the strictest condition on λ and μ corresponds to the lowest value of z. Thus, $z = \max\{0, x + y - n\}$, and either $(x - z)(y - z)z = 0$ or $x + y - z = n$. Otherwise, if $F^j(x + 1) < F^j(x)$, then the largest value of z is strictest, i.e., $z = \min\{x, y\}$ which satisfies $(x - z)(y - z)z = 0$.

– *Step 3* : In order to derive the game instances with price of anarchy matching $\gamma(\mathcal{G})$, it is convenient to perform the following change of variables: $\nu(\lambda, \mu) := 1/\lambda$ and $\rho(\lambda, \mu) := (1 - \mu)/\lambda$. For ease of notation, we will refer to the new variables simply as ν and ρ, respectively, i.e., $\nu = \nu(\lambda, \mu)$ and $\rho = \rho(\lambda, \mu)$. For each $j \in \{1, \ldots, m\}$ and each $(x, y, z) \in \mathcal{I}_{\mathcal{R}}(n)$, it is straightforward to verify that the constraints in Equation (18) can be rewritten in terms of ν and ρ as

$$C^j(y) - \rho C^j(x) + \nu[(x - z) F^j(x) - (y - z) F^j(x + 1)] \geq 0.$$

Thus, the value $\gamma(\mathcal{G})$ must be equal to $1/\rho^{\text{opt}}$, where ρ^{opt} is the value of the following linear program:

$$\begin{aligned}\rho^{\text{opt}} = \underset{\nu\in\mathbb{R}_{\geq 0},\rho\in\mathbb{R}}{\text{maximize}} \quad & \rho \\ \text{subject to:} \quad & C^j(y) - \rho C^j(x) + \nu[(x-z)F^j(x) - (y-z)F^j(x+1)] \geq 0, \\ & \forall j = 1,\dots,m, \quad \forall (x,y,z) \in \mathcal{I}_{\mathcal{R}}(n).\end{aligned} \tag{19}$$

It is important to note here that while $\gamma(\mathcal{G})$ is the infimum of a fractional program (see, e.g., Equation (18)), the value ρ^{opt} can be computed as a maximum because the feasible set is bounded and closed. Firstly, since $\gamma(\mathcal{G})$ is an upper bound on the price of anarchy, its inverse (i.e., ρ) must be in the bounded and closed interval $[0,1]$. Additionally, one can verify that ν is not only bounded from below by 0, but also from above by the quantity

$$\begin{aligned}\bar{\nu} := \min_{j\in\{1,\dots,m\}} \quad \underset{(x,y,z)\in\mathcal{I}_{\mathcal{R}}(n)}{\text{minimize}} \quad & \frac{C^j(y)}{(y-z)F^j(x+1) - (x-z)F^j(x)} \\ \text{subject to:} \quad & (x-z)F^j(x) - (y-z)F^j(x+1) < 0 \text{ and } C^j(x) = 0,\end{aligned} \tag{20}$$

which comes from the constraints in Equation (19) corresponding to triplets $(x,y,z) \in \mathcal{I}_{\mathcal{R}}(n)$ such that $C^j(x) = 0$ and $(x-z)F^j(x) - (y-z)F^j(x+1) < 0$. Such a value must exist, as we assume $C^j(0) = 0$. One can verify that any $j \in \{1,\dots,m\}$ and $(x,y,z) \in \mathcal{I}_{\mathcal{R}}(n)$ such that $C^j(x) = 0$ and $(x-z)F^j(x) - (y-z)F^j(x+1) \geq 0$ correspond to constraints that are satisfied trivially in Equation (19) since $\nu \geq 0$, by definition, and $C^j(y) \geq 0$ for all $y = 0,1,\dots,n$, by assumption.

We denote with $\mathcal{H}^j(x,y,z)$ the halfplane of (ν,ρ) values that satisfy the constraint corresponding to $j \in \{1,\dots,m\}$ and $(x,y,z) \in \mathcal{I}_{\mathcal{R}}(n)$, i.e.,

$$\mathcal{H}^j(x,y,z) := \left\{(\nu,\rho) \in \mathbb{R}_{\geq 0} \times \mathbb{R} \text{ s.t. } \rho \leq \frac{C^j(y)}{C^j(x)} + \frac{1}{C^j(x)}\nu\left[(x-z)F^j(x) - (y-z)F^j(x+1)\right]\right\}.$$

The set of feasible (ν,ρ) is the intersection of these $m \times |\mathcal{I}_{\mathcal{R}}(n)|$ halfplanes. Since the objective is to maximize ρ, any solution $(\nu^{\text{opt}},\rho^{\text{opt}})$ to the linear program in Equation (19) must be on the (upper) boundary of the feasible set. We argue below that a solution $(\nu^{\text{opt}},\rho^{\text{opt}})$ can only exist in one of the three following scenarios: (1) at the intersection of two halfplanes' boundaries, where one halfplane has boundary line with positive slope, and the other has boundary line with nonpositive slope; (2) on a halfplane boundary line with positive slope at $\nu^{\text{opt}} = \bar{\nu}$; or (3) at $(\nu^{\text{opt}},\rho^{\text{opt}}) = (0,0)$.

We denote with $\partial\mathcal{H}^j(x,y,z)$ the boundary line of the halfplane $\mathcal{H}^j(x,y,z)$, i.e.,

$$\partial\mathcal{H}^j(x,y,z) := \left\{(\nu,\rho) \in \mathbb{R}_{\geq 0} \times \mathbb{R} \text{ s.t. } \rho = \frac{C^j(y)}{C^j(x)} + \frac{1}{C^j(x)}\nu\left[(x-z)F^j(x) - (y-z)F^j(x+1)\right]\right\}.$$

Observe that the boundary lines of halfplanes corresponding to the choice $y = z = 0$ have ρ-intercept equal to zero and slope $xF^j(x)/C^j(x)$. If $F^j(x) \leq 0$ for any $j \in \{1,\dots,m\}$ and $x \in \{1,\dots,n\}$, then an optimal pair (ν,ρ) is trivially at the origin, i.e., $(\nu^{\text{opt}},\rho^{\text{opt}}) = (0,0)$ (i.e., scenario (3) above). Note that the ρ-intercept of any halfplane boundary cannot be below 0, as we only consider cost functions such that $C^j(k) \geq 0$ for all k and all j. Otherwise, the maximum value of ρ occurs at the intersection of a boundary line with positive slope and a boundary line with nonpositive slope (i.e., scenario (1) above) or on a boundary line with positive slope at $\nu = \bar{\nu}$ (i.e., scenario (2) above). We illustrate this reasoning in Fig. 1.

Observe that for Scenarios (1) and (2), the pair $(\nu^{\text{opt}},\rho^{\text{opt}})$ is at the intersection of two boundary lines, which we denote as $\partial\mathcal{H}^j(x,y,z)$ and $\partial\mathcal{H}^{j'}(x',y',z')$. The parameters $j, j' \in \{1,\dots,m\}$ and

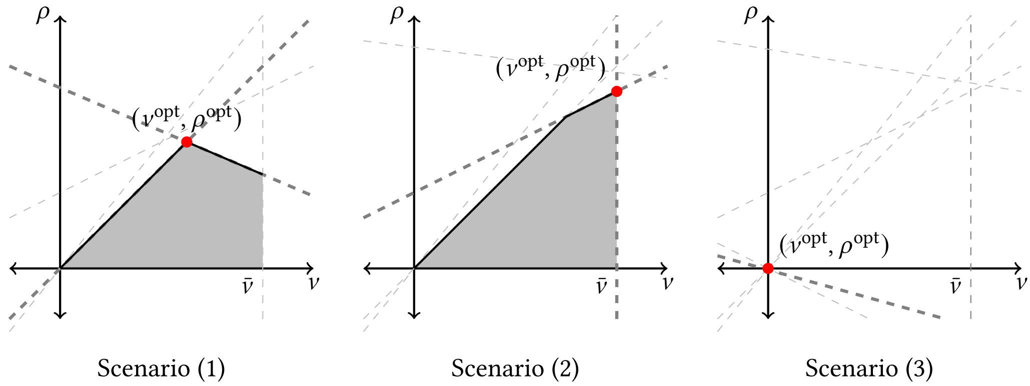

Fig. 1. *The three different scenarios in which optimal solutions* $(\nu^{\text{opt}}, \rho^{\text{opt}})$ *to Equation* (19) *can exist.* We illustrate the reasoning behind each of the three scenarios for optimal solutions $(\nu^{\text{opt}}, \rho^{\text{opt}})$ to the linear program in Equation (19). Since the objective of Equation (19) is to maximize ρ, the optimal values will be at the (upper) boundary of the feasible set, illustrated with a solid, bolded line in each of the examples above. Additionally, the optimal solution $(\nu^{\text{opt}}, \rho^{\text{opt}})$ is marked by a solid, red dot in the illustrations above. In Scenario (1), on the left, $(\nu^{\text{opt}}, \rho^{\text{opt}})$ lie on the intersection of a boundary line with positive slope and a boundary line with nonpositive slope. In Scenario (2), centre, $(\nu^{\text{opt}}, \rho^{\text{opt}})$ lie on the intersection of a boundary line with positive slope at $\nu = \bar{\nu}$, which is defined in Equation (20). In Scenario (3), on the right, there exists a halfplane boundary line with nonpositive slope and ρ-intercept equal to zero, and so $(\nu^{\text{opt}}, \rho^{\text{opt}}) = (0, 0)$. Using the parameters corresponding to the halfplanes on which the pair $(\nu^{\text{opt}}, \rho^{\text{opt}})$ lays, we can construct games $G \in \mathcal{G}$ with $\text{PoA}(G) = 1/\rho^{\text{opt}}$ in each of these scenarios.

$(x, y, z), (x', y', z') \in \mathcal{I}_{\mathcal{R}}(n)$ satisfy the following:

$$\begin{aligned} \nu^{\text{opt}}\left[(x - z)F^j(x) - (y - z)F^j(x + 1)\right] &= \rho^{\text{opt}} C^j(x) - C^j(y), \\ \nu^{\text{opt}}\left[(x' - z')F^{j'}(x') - (y' - z')F^{j'}(x' + 1)\right] &= \rho^{\text{opt}} C^{j'}(x') - C^{j'}(y'), \end{aligned} \tag{21}$$

because $(\nu^{\text{opt}}, \rho^{\text{opt}})$ is on both boundary lines. Further, there exists $\eta \in [0, 1]$ such that

$$\eta\left[(x - z)F^j(x) - (y - z)F^j(x + 1)\right] + (1 - \eta)\left[(x' - z')F^{j'}(x') - (y' - z')F^{j'}(x' + 1)\right] = 0. \tag{22}$$

Equation (22) holds in Scenario (1) because one of the boundary lines has positive slope, i.e., $(x-z)F^j(x)-(y-z)F^j(x+1) > 0$, while the other has nonpositive slope, and in Scenario (3) because one boundary line has positive slope while the other is the vertical line $\nu = \bar{\nu}$ which corresponds to a particular choice of $j \in \{1, \dots, m\}$ and $(x, y, z) \in \mathcal{I}_{\mathcal{R}}(n)$ such that $(x-z)F^j(x)-(y-z)F^j(x+1) < 0$ by Equation (20).

Next, for the parameters $j, j' \in \{1, \dots, m\}$, $(x, y, z), (x', y', z') \in \mathcal{I}_{\mathcal{R}}(n)$, and $\eta \in [0, 1]$ obtained above, we construct a game instance $G \in \mathcal{G}$ such that $\text{PoA}(G) = 1/\rho^{\text{opt}}$. Let $\mathcal{R}_1 = \{r_1, \dots, r_n\}$ and $\mathcal{R}_2 = \{r_{n+1}, \dots, r_{2n}\}$ denote two disjoint cycles of resources. Every resource $r \in \mathcal{R}_1$ has cost function $C_r(k) = \eta C^j(k)$, and cost generating function $F_r(k) = \eta F^j(k)$ for all k. Meanwhile, every $r \in \mathcal{R}_2$ has cost function $C_r(k) = (1 - \eta)C^{j'}(k)$, and cost generating function $F_r(k) = (1 - \eta)F^{j'}(k)$ for all k. We define the user set $N = \{1, \dots, n\}$, where each user $i \in N$ has action set $\mathcal{A}_i = \{a_i^{\text{ne}}, a_i^{\text{opt}}\}$. In action a_i^{ne}, the user i selects x consecutive resources in $\mathcal{R}_1$ starting with r_i, i.e. $\{r_i, r_{(i \bmod n)+1}, \dots, r_{((i+x-2) \bmod n)+1}\}$, and x' consecutive resources in $\mathcal{R}_2$ starting with resource r_{n+i}. In a_i^{opt}, user i selects y consecutive resources in $\mathcal{R}_1$ ending with resource $r_{((i+z-2) \bmod n)+1}$, i.e. $\{r_{((i+z-y-1) \bmod n)+1}, \dots, r_{((i+z-2) \bmod n)+1}\}$, and y' consecutive resources in $\mathcal{R}_2$ ending with resource

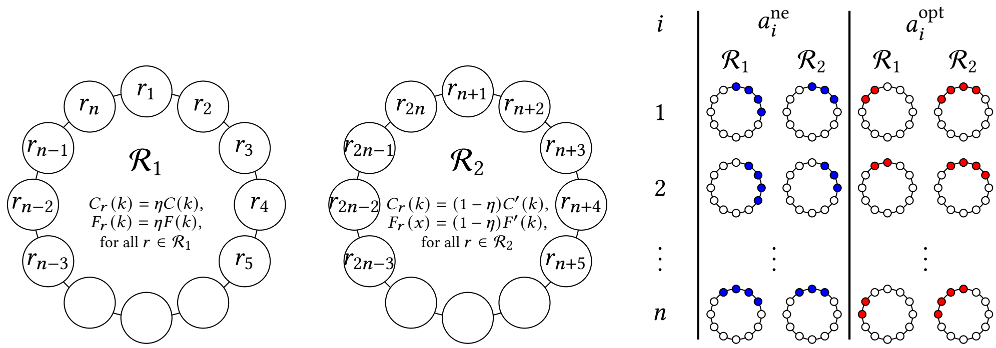

Fig. 2. *The game instance construction G consisting of n users, and two disjoint cycles $\mathcal{R}_1$ and $\mathcal{R}_2$, as described in the proof of Theorem 3.5, Step 2 for Scenarios (1) and (2).* Consider the set of games $\mathcal{G}$, where n is the maximum number of users and $\mathcal{Z}$ is a set of basis functions pairs, and suppose that $(\nu^{\mathrm{opt}}, \rho^{\mathrm{opt}})$ satisfy the conditions of Scenarios (1) or (2). Further, suppose that the parameters for which Equation (21) and Equation (22) hold are $C, F, C', F' \in \mathcal{Z}$, $(x, y, z) = (4, 2, 0)$, $(x', y', z') = (3, 4, 2) \in \mathcal{I}_{\mathcal{R}}(n)$ and $\eta \in [0, 1]$. In the above figure, we illustrate the game $G \in \mathcal{G}$ such that $\mathrm{PoA}(G) = \mathrm{PoA}(G') = 1/\rho^{\mathrm{opt}}$ according to the reasoning for constructing game instances in Scenarios (1) and (2). Observe that each resource $r \in \mathcal{R}_1$ has $C_r(k) = \eta C(k)$, and $F_r(k) = \eta F(k)$, whereas each resource $r \in \mathcal{R}_2$ has $C_r(k) = (1-\eta)C'(k)$, and $F_r(x) = (1-\eta)F'(k)$, for all $k \in \{1, \ldots, n\}$. Each user $i \in N$ has two actions a_i^{ne} and a_i^{opt}, as defined in the table on the right. Observe that every resource in $\mathcal{R}_1$ is selected by 4 users in the allocation $a^{\mathrm{ne}} = (a_1^{\mathrm{ne}}, \ldots, a_n^{\mathrm{ne}})$, and 3 users in $a^{\mathrm{opt}} = (a_1^{\mathrm{opt}}, \ldots, a_n^{\mathrm{opt}})$, where no user $i \in N$ has a common resource between its actions a_i^{ne} and a_i^{opt}, i.e., $x_r = 4 = x$, $y_r = 3 = y$, and $z_r = 0 = z$ for all $r \in \mathcal{R}_1$. Similarly, $x_r = 3 = x'$, $y_r = 4 = y'$, and $z_r = 2 = z'$, for each resource $r \in \mathcal{R}_2$.

$r_{n+((i+z'-2) \bmod n)+1}$. We provide an illustration of this game construction in Fig. 2. Observe that $a^{\mathrm{ne}} = (a_1^{\mathrm{ne}}, \ldots, a_n^{\mathrm{ne}})$ satisfies the conditions for a Nash equilibrium,

$$\begin{aligned} J_i(a^{\mathrm{ne}}) &= \eta x F^j(x) + (1-\eta) x' F^{j'}(x') \\ &= \eta [z F^j(x) + (y-z) F^j(x+1)] + (1-\eta)[z' F^{j'}(x') + (y'-z') F^{j'}(x'+1)] = J_i(a_i^{\mathrm{opt}}, a_{-i}^{\mathrm{ne}}), \end{aligned}$$

which holds by Equation (22). Then, by the above equality and Equation (21),

$$\begin{aligned} 0 &= \sum_{i=1}^{n} J_i(a_i^{\mathrm{opt}}, a_{-i}^{\mathrm{ne}}) - \sum_{i=1}^{n} J_i(a^{\mathrm{ne}}) \\ &= \frac{1}{\nu^{\mathrm{opt}}} \left[n \cdot \eta \left[\rho^{\mathrm{opt}} C^j(x) - C^j(y) \right] + n \cdot (1-\eta) \left[\rho^{\mathrm{opt}} C^{j'}(x') + C^{j'}(y') \right] \right] \\ &= \frac{1}{\nu^{\mathrm{opt}}} \left[\rho^{\mathrm{opt}} C(a^{\mathrm{ne}}) - C(a^{\mathrm{opt}}) \right], \end{aligned}$$

where $a^{\mathrm{opt}} = (a_1^{\mathrm{opt}}, \ldots, a_n^{\mathrm{opt}})$. Thus, $\mathrm{PoA}(G) = 1/\rho^{\mathrm{opt}}$.

For Scenario (3), observe that $\rho^{\mathrm{opt}} = 0$, and so $1/\rho^{\mathrm{opt}}$ is unbounded. Recall that, in this scenario, there exist $j \in \{1, \ldots, m\}$ and $x \in \{1, \ldots, n\}$ such that $F^j(x) \leq 0$. We use the basis function pair $\{C^j, F^j\}$ to construct a game G with unbounded price of anarchy. Consider a game instance with x users and resource set $\mathcal{R} = \{r_1, r_2\}$, where $x \in \{1, \ldots, n\}$ is the value that minimizes the function $F(x)$, i.e., $F^j(x) = \min_{k \in \{1, \ldots, n\}} F^j(k) \leq 0$. Every user $i \in \{1, \ldots, x\}$ has action set $\mathcal{A}_i = \{\{r_1\}, \{r_2\}\}$. The resource r_1 has cost function $C_r(k) = \eta C^j(k)$ and cost generating function $F_r(k) = \eta F^j(k)$ for all k. Similarly, the resource r_2 has cost function $C_r(k) = (1-\eta) C^j(k)$ and cost generating function $F_r(k) = (1 - \eta)F(k)$. It is straightforward to verify that, for η approaching 0 from above, the allocation in which all users select r_1 is an equilibrium and the price of anarchy is unbounded. □

3.4 Comparison to Existing Literature

There has been a significant amount of research focused on characterizing the price of anarchy in congestion games. Accordingly, in this section, we position the results of Theorem 3.5 in the broader context of smoothness [40] and the primal-dual approach [6, 33]. First, it is important to recognize that both the smoothness and generalized smoothness frameworks can be written as linear programs for any family of games $\mathcal{G}$. For example, observe that the generalized price of anarchy satisfies $\mathrm{GPoA}(\mathcal{G}) = 1/\rho^{\mathrm{opt}}$, where

$$\begin{aligned} \rho^{\mathrm{opt}} = \underset{\nu \in \mathbb{R}_{\geq 0}, \rho \in \mathbb{R}}{\text{maximize}} \quad & \rho \\ \text{subject to:} \quad & C(a') - \rho C(a) + \nu \left[\sum_{i=1}^{n} J_i(a) - \sum_{i=1}^{n} J_i(a_i', a_{-i}) \right] \geq 0, \quad \forall a, a' \in \mathcal{A}, \forall G \in \mathcal{G}, \end{aligned} \tag{23}$$

which follows from Equations (10) and (11) for the change of variables $\nu = 1/\lambda$ and $\rho = (1 - \mu)/\lambda$. Although the price of anarchy bound obtained using the above linear program is the best achievable following a generalized smoothness argument, computing such a bound for a family of games is intractable as there may be exponentially many constraints, even for modest values of the maximum number of users n. The novelty of the result in Theorem 3.5 is in identifying a game parameterization for any set of generalized congestion games such that the number of linear program constraints only grows linearly in m and quadratically in n while verifying and preserving the tightness of the generalized price of anarchy bound.

Our result is inspired by several previous works, most notably Bilò [6] and Roughgarden [40], that introduce game parameterizations to reduce the complexity of smoothness bounds. We note that many of the bounds proposed in these previous works remain intractable, as the number of linear program constraints grows exponentially in n. Nonetheless, [6] provides a tractable linear program for deriving upper bounds on the price of anarchy that has two decision variables and $\mathcal{O}(n^2)$ constraints. Here, we demonstrate that upper bounds computed using the tractable linear program in [6] are not tight, even for affine congestion games with $n = 2$ users.

Example 3.6. For the set of affine congestion games $\mathcal{G}$ with a maximum of n users, Bilò [6] proposes the following linear program for computing an upper bound γ^{opt} on the price of anarchy:

$$\begin{aligned} \gamma^{\mathrm{opt}} = \underset{\kappa \in \mathbb{R}_{\geq 0}, \gamma \in \mathbb{R}}{\text{maximize}} \quad & \gamma \\ \text{subject to:} \quad & \gamma y^2 - x^2 + \kappa[x^2 - (x+1)y] \geq 0, \quad \forall x, y \in \{0, 1, \ldots, n\}. \end{aligned} \tag{24}$$

We observe that solving the linear programs in Equations (16) and (24) for the set of affine congestion games $\mathcal{G}$ with $n \leq 2$ users yields $\mathrm{PoA}(\mathcal{G}) = 2$ and $\mathrm{PoA}(\mathcal{G}) \leq 2.5$, respectively.[2] It then holds that upper bounds on the price of anarchy derived from the tractable linear program in Reference [6] are not tight, as $\mathrm{PoA}(\mathcal{G}) = 2 < 2.5$ in this example.

Observe that the linear program in Equation (16) closely resembles the linear program in Equation (24). In fact, these two linear programs are identical in structure as they are both tractable reductions of the linear program in Equation (23). They only differ in the parameterization of

[2]One can verify that the solution to the linear program in Equation (16) is $(\nu^{\mathrm{opt}}, \rho^{\mathrm{opt}}) = (0.5, 0.5)$, while the solution to the linear program in Equation (24) is $(\kappa^{\mathrm{opt}}, \gamma^{\mathrm{opt}}) = (1.5, 2.5)$.

the constraint set. In this respect, the game parameterization we identify in this work is critical in retaining tightness of the generalized price of anarchy using an extremely modest number of constraints. In contrast, though the parameterization used in [6] has a comparable number of constraints, we observed in Example 3.6 that it loses tightness.

3.5 Optimizing the price of anarchy

The previous section focused on how to characterize the price of anarchy in any set of generalized congestion games. In this section, we shift our focus to the derivation of cost generating functions that optimize the price of anarchy. That is, given a set of resource cost functions $C^1, \ldots, C^m$, what is the corresponding set of cost generating functions $F^1, \ldots, F^m$ that minimizes the resulting price of anarchy $\mathrm{PoA}(\mathcal{G})$? Recall from the introduction that this line of questioning is relevant to the problem of incentive design given in Section 1.1, when the price of anarchy is the performance bound of interest.

The following theorem provides a tractable and scalable methodology for computing the set of cost generating functions that minimize the price of anarchy.

Theorem 3.7. *Let $C^1, \ldots, C^m$ denote a set of resource cost functions defined for n users, and let $(F^{\mathrm{opt},j}, \rho^{\mathrm{opt},j})$, $j = 1, \ldots, m$, be solutions to the following m linear programs:*

$$\begin{aligned} &\underset{F \in \mathbb{R}^n, \rho \in \mathbb{R}}{\textit{maximize}} \quad \rho \\ &\textit{subject to:} \quad C^j(y) - \rho C^j(x) + (x-z)F(x) - (y-z)F(x+1) \geq 0, \quad \forall (x,y,z) \in \mathcal{I}_{\mathcal{R}}(n). \end{aligned} \tag{25}$$

Then the cost generating functions $F^{\mathrm{opt},1}, \ldots, F^{\mathrm{opt},m}$ minimize the price of anarchy and the price of anarchy corresponding to basis function pairs $\{C^j, F^{\mathrm{opt},j}\}$, $j = 1, \ldots, m$, satisfies

$$\mathrm{PoA}(\mathcal{G}) = \max_{j \in \{1,\ldots,m\}} \frac{1}{\rho^{\mathrm{opt},j}}.$$

Theorem 3.7 states that we can derive cost generating functions $F^{\mathrm{opt},1}, \ldots, F^{\mathrm{opt},m}$ that minimize the price of anarchy by solving m independent linear progams, where each $F^{\mathrm{opt},j}$ can be derived using only information about the corresponding resource cost function C^j. Accordingly, the price of anarchy of this optimized system corresponds to the worst price of anarchy associated with any single pair $\{C^j, F^{\mathrm{opt},j}\}$, i.e.,

$$\mathrm{PoA}(\mathcal{G}) = \max_{j \in \{1,\ldots,m\}} \mathrm{PoA}(\mathcal{G}^j),$$

where $\mathcal{G}^j \subseteq \mathcal{G}$ represents the set of generalized congestion games induced by n and the basis function pair $\{C^j, F^{\mathrm{opt},j}\}$. Observe that this statement is not true in general for an arbitrary set of basis function pairs, i.e., there exist sets of basis function pairs $\{C^j, F^j\}$, $j = 1, \ldots, m$, such that[3]

$$\mathrm{PoA}(\mathcal{G}) > \max_{j \in \{1,\ldots,m\}} \mathrm{PoA}(\mathcal{G}^j).$$

However, when we restrict our attention to optimal cost generating functions for each C^j, the above strict inequality holds with equality. This is the key observation in the proof of Theorem 3.7.

Proof of Theorem 3.7. For each $j \in \{1, \ldots, m\}$, the function $F^{\mathrm{opt},j}$ maximizes $\rho^{\mathrm{opt},j}$ by the following reasoning: For each resource cost function C^j, we wish to find the function $F^{\mathrm{opt},j}$ that

[3]For example, consider the set of generalized congestion games $\mathcal{G}$ induced by $n = 3$, and $\{\{C^1, F^1\}, \{C^2, F^2\}\}$, where $\{C^1(k), F^1(k)\} = \{k^2, k\}$ and $\{C^2, F^2\} = \{k, k\}$ for all $k = 1, \ldots, n$. Using the linear program in Equation (16), we get $\mathrm{PoA}(\mathcal{G}^1)) = 2.5$, $\mathrm{PoA}(\mathcal{G}^2)) = 2.0$, and $\mathrm{PoA}(\mathcal{G}) = 2.6$. For this particular choice of $\mathcal{G}$, observe that $\mathrm{PoA}(\mathcal{G}) > \max_{j \in \{1,\ldots,m\}} \mathrm{PoA}(\mathcal{G}^j)$.

maximizes ρ in Equation (16). Finding such a function is equivalent to finding the solution to

$$(F^{\text{opt},j}, \nu^{\text{opt},j}, \rho^{\text{opt},j}) \in \arg \max_{F\in\mathbb{R}^n, \nu\in\mathbb{R}_{\geq 0}, \rho\in\mathbb{R}} \quad \rho$$
$$\text{s.t.} \quad C^j(y) - \rho C^j(x) + \nu[(x-z)F(x) - (y-z)F(x+1)] \geq 0, \quad \forall (x,y,z) \in \mathcal{I}_{\mathcal{R}}(n).$$

It is important to note that an optimal function $F^{\text{opt},j}$ must exist since the above program is feasible for $F^j(k) = 0$, $k = 1, \ldots, n$, $\nu = 1$ and $\rho \leq \min_{x,y} C^j(y)/C^j(x)$, and is bounded since any pair $\{C^j, F^{\text{opt},j}\}$ generates a set of games $\mathcal{G}^j$ so $\rho^{\text{opt},j} = 1/\text{PoA}(\mathcal{G}^j) \in [0,1]$ must hold by Theorem 3.5.

To obtain a linear program, we combine the decision variables ν and F in $\tilde{F}(k) := \nu F(k)$ to get

$$(\tilde{F}^j_{\text{opt}}, \tilde{\rho}^j_{\text{opt}}) \in \arg \max_{F\in\mathbb{R}^n, \rho\in\mathbb{R}} \quad \rho$$
$$\text{s.t.} \quad C^j(y) - \rho C^j(x) + (x-z)F(x) - (y-z)F(x+1) \geq 0, \quad \forall (x,y,z) \in \mathcal{I}_{\mathcal{R}}(n).$$

Note that $\tilde{F}^j_{\text{opt}} \in \mathbb{R}^n$ must be feasible as $\tilde{F}^{\text{opt},j}(k) = \nu^{\text{opt},j} F^{\text{opt},j}(k)$, and we know that $F^{\text{opt},j} \in \mathbb{R}^n$ exists. Further, $\tilde{\rho}^{\text{opt},j} = \rho^{\text{opt},j}$, as equilibrium conditions are invariant to scaling of F.

For the set of generalized congestion games $\mathcal{G}$ induced by n and basis function pairs $\{C^j, \tilde{F}^{\text{opt},j}\}$, $j = 1, \ldots, m$, and the set of games $\mathcal{G}^j$ induced by n and the basis function pair $\{C^j, \tilde{F}^{\text{opt},j}\}$, it holds that $\text{PoA}(\mathcal{G}) \geq \max_{j\in\{1,\ldots,m\}} \text{PoA}(\mathcal{G}^j)$. We conclude by proving that the converse also holds, i.e.,

$$\text{PoA}(\mathcal{G}) \leq \max_{j\in\{1,\ldots,m\}} \text{PoA}(\mathcal{G}^j).$$

Simply note that the values $(\nu, \rho) = (1, \rho^{\text{opt}})$ are feasible in the linear program in Equation (16) for the function pairs $\{C^j, \tilde{F}^{\text{opt},j}\}$, $j = 1, \ldots, m$, where $\rho^{\text{opt}} := \min_j \rho^{\text{opt},j}$. This implies that $\text{PoA}(\mathcal{G}) \leq 1/\rho^{\text{opt}}$. Observing that $1/\rho^{\text{opt}} = \max_{j\in\{1,\ldots,m\}} \text{PoA}(\mathcal{G}^j)$ concludes the proof.

□

4 GENERALIZED SMOOTHNESS IN WELFARE MAXIMIZATION GAMES

Although the primary focus of this paper is on cost minimization settings, many of the results that we obtain can be analogously derived for welfare maximization problems. A welfare maximization problem consists of a set $N = \{1, \ldots, n\}$ of users, where each user $i \in N$ is associated with a finite action set $\mathcal{A}_i$. The global objective is to *maximize* the system's welfare, which is measured by the welfare function $W : \mathcal{A} \to \mathbb{R}_{>0}$, i.e. we wish to find the allocation $a^{\text{opt}} \in \mathcal{A}$, such that $a^{\text{opt}} \in \arg\max_{a\in\mathcal{A}} W(a)$. As with cost minimization problems, we consider a game-theoretic model where each user $i \in N$ is associated with a utility function $U_i : \mathcal{A} \to \mathbb{R}$, which it uses to evaluate its own actions against the collective actions of the other users. A *welfare maximization game* is a tuple $G = (N, \mathcal{A}, W, \{U_i\})$.

Given a welfare maximization game G, a pure Nash equilibrium is defined as an allocation $a^{\text{ne}} \in \mathcal{A}$ such that $U_i(a^{\text{ne}}) \geq U_i(a_i, a^{\text{ne}}_{-i})$ for all $a \in \mathcal{A}_i$, and all $i \in N$. The price of anarchy in welfare maximization games is defined similarly to Equation (6) and Equation (7),[4]

$$\text{PoA}(G) := \frac{\max_{a\in\mathcal{A}} W(a)}{\min_{a\in\text{NE}(G)} W(a)} \geq 1, \quad \text{PoA}(\mathcal{G}) := \sup_{G\in\mathcal{G}} \text{PoA}(G) \geq 1,$$

where a lower value of the price of anarchy corresponds to an improvement in performance.

[4]For consistency with the previous sections, we opt to define the price of anarchy in welfare maximization games as the ratio between the welfare at the optimal allocation and the system welfare at the worst performing Nash equilibrium, in contrast with previous works [23, 40]. This is achieved by inverting the ratio, i.e., defining the price of anarchy as the worst case ratio between the welfare at optimum, and the welfare at the equilibria in $\text{NE}(G)$. By adopting this formalism, we retain the overall objective of *minimizing* the system's price of anarchy.

4.1 Generalized smoothness in welfare maximization games

We begin with the definition of generalized smoothness in welfare maximization games and then provide the analogue of Proposition 2.2.

Definition 4.1. The welfare maximization game G is (λ, μ)-generalized smooth if, for any two allocations $a, a' \in \mathcal{A}$, there exist $\lambda > 0$ and $\mu > -1$ satisfying,

$$\sum_{i=1}^{n} U_i(a_i', a_{-i}) - \sum_{i=1}^{n} U_i(a) + W(a) \geq \lambda W(a') - \mu W(a). \tag{26}$$

Proposition 4.2. *The price of anarchy of a (λ, μ)-generalized smooth welfare maximization game G is upper bounded as,*

$$\mathrm{PoA}(G) \leq \frac{1+\mu}{\lambda}.$$

We define the generalized price of anarchy of a set of welfare maximization games $\mathcal{G}$ as

$$\mathrm{GPoA}(\mathcal{G}) := \inf_{\lambda>0, \mu>-1} \left\{ \frac{1+\mu}{\lambda} \text{ s.t. Equation (26) holds } \forall G \in \mathcal{G} \right\}. \tag{27}$$

As with cost minimization games, all efficiency guarantees also extend to average coarse-correlated equilibria (as in Observation #3) and there are also provable advantages of generalized smoothness over the original smoothness framework in terms of characterizing price of anarchy bounds (as in Proposition 2.2). We do not explicitly state or prove these parallel results to avoid redundancy.

4.2 Generalized smoothness in distributed welfare games

In this section, we consider *distributed welfare games* [31] as described in Section 1.1, which are the welfare maximization analogue to generalized congestion games. Games in this class feature a set of users $N = \{1, \ldots, n\}$ and a finite set of resources $\mathcal{R}$. The system welfare and user utility functions are defined as

$$W(a) = \sum_{r \in \mathcal{R}} W_r(|a|_r), \qquad U_i(a) = \sum_{r \in a_i} F_r(|a|_r),$$

where, for each $r \in \mathcal{R}$, $W_r : \{0, 1, \ldots, n\} \to \mathbb{R}_{\geq 0}$ and $F_r : \{1, \ldots, n\} \to \mathbb{R}_{\geq 0}$ are the resource welfare function and utility generating function, respectively. For the remainder of this section, given basis function pairs $\{W^j, F^j\}$, $j = 1, \ldots, m$, we define the set of local welfare maximization games $\mathcal{G}$ in the same fashion as for generalized congestion games given in Section 3. Distributed welfare games have been used to model several problems of interest as in [4, 13, 23, 28, 31].

The following theorem provides the analagous results derived for generalized congestion games to the domain of distributed welfare games. As before, we define $W^j(0) = F^j(0) = F^j(n+1) = 0$, for $j = 1, \ldots, m$, for ease of notation. Theorem 4.3 is stated without proof as the reasoning follows almost identically to the proofs of Theorems 3.5 and 3.7.

Theorem 4.3. *Let $\mathcal{G}$ denote the set of all distributed welfare games with a maximum of n users generated from basis function pairs $\{W^j, F^j\}$, $j = 1, \ldots, m$. The following statements hold true:*

(i) The price of anarchy and the generalized price of anarchy satisfy $\mathrm{PoA}(\mathcal{G}) = \mathrm{GPoA}(\mathcal{G})$.

(ii) Let ρ^{opt} be the value of the following linear program:

$$\begin{aligned} \rho^{\mathrm{opt}} = \min_{\nu \in \mathbb{R}_{\geq 0}, \rho \in \mathbb{R}} \quad & \rho \\ \textit{s.t.} \quad & W^j(y) - \rho W^j(x) + \nu \left[(x-z)F^j(x) - (y-z)F^j(x+1) \right] \leq 0 \\ & \qquad \forall j = 1, \ldots, m, \quad \forall (x, y, z) \in \mathcal{I}_{\mathcal{R}}(n), \end{aligned} \tag{28}$$

Then, it holds that $\mathrm{PoA}(\mathcal{G}) = \rho^{\mathrm{opt}}$.

(iii) Let the parameters $(F^{\mathrm{opt},j}, \rho^{\mathrm{opt},j})$, $j = 1, \ldots, m$, *be solutions to the following* m *linear programs:*

$$(F^{\mathrm{opt},j}, \rho^{\mathrm{opt},j}) \in \arg \min_{F \in \mathbb{R}^n, \rho \in \mathbb{R}} \rho \quad \textit{subject to:}$$
$$W^j(y) - \rho W^j(x) + (x-z)F(x) - (y-z)F(x+1) \le 0, \quad \forall (x,y,z) \in \mathcal{I}_{\mathcal{R}}(n). \tag{29}$$

Then, the utility generating functions $F^{\mathrm{opt},1}, \ldots, F^{\mathrm{opt},m}$ *minimize the price of anarchy, and the price of anarchy corresponding to basis function pairs* $\{W^j, F^{\mathrm{opt},j}\}$, $j = 1, \ldots, m$, *satisfies*

$$\mathrm{PoA}(\mathcal{G}) = \max_{j \in \{1,\ldots,m\}} \rho^{\mathrm{opt},j}.$$

5 ILLUSTRATIVE EXAMPLES

In the introduction, we motivated our study by considering two seemingly distinct problems: incentive design in congestion games and utility design in distributed welfare games. In this section, we utilize these same classes of problems (among others) to demonstrate the breadth of our approach. For an in-depth study discussing the application of the machinery derived here to the design of incentives in congestion games, we refer to Paccagnan et al. [35].

5.1 Price of anarchy in congestion games and their variants

Theorem 3.5 allows to determine the exact price of anarchy for any game that can be cast as generalized congestion game. In this section we illustrate the applicability of this result by i) recovering/extending classical findings on the price of anarchy of congestion games, ii) computing the efficiency of marginal cost incentives, iii) providing novel price of anarchy results for perception-parametrized congestion games.

i) Congestion games. Congestion games constitute a subclass of problems to which Theorem 3.5 applies. This follows readily upon letting $C_r(k) = k \cdot c_r(k)$, $F_r(k) = c_r(k)$ for $k = 1, \ldots, n$ and $r \in \mathcal{R}$ in Equations (14) and (15), where $c_r(\cdot)$ describes the original resource congestion. Hence, we are able to compute their price of anarchy by simply solving the linear program in Theorem 3.5. As a special case, we recover well-known price of anarchy results [1, 3, 16] for polynomial congestion games of maximum degree d, i.e., congestion games where the resource congestion is obtained by non-negative linear combinations of monomials $1, x, \ldots, x^d$. Although the bounds provided in these works are exact (for large n), the authors had to combine traditional smoothness arguments with nontrivial worst-case game constructions.[5] In contrast, the linear program in Theorem 3.5 provides exact price of anarchy values for all n, can be solved efficiently (featuring only two decision variables and $(d+1)O(n^2)$ constraints), and does not require ad-hoc worst-case constructions. We solve such program as a function of the number of users n and the maximum degree d, reporting the results in Fig. 3. To the best of our knowledge, this is the first characterization of the dependence of the price of anarchy in polynomial congestion games on the number of users n. Remarkably, we note (table in Fig. 3), that the price of anarchy values for $n = 5$ users exactly match their corresponding asymptotic values ($n \to \infty$) from [1, 3, 16], suggesting that very small instances are sufficient to produce highly inefficient equilibria.

Finally, we remark that the machinery developed here can be used to characterize the price of anarchy for a variety of congestion functions often employed in the literature. This includes the well-studied Bureau of Public Roads (BPR) function [45] where $c_r(x) = T_r \cdot \left[1 + 0.15 \cdot \left(\frac{x}{K_r}\right)^4\right]$, and $T_r \ge 0$, $K_r \in \mathbb{N}_{\ge 1}$ are the free flow congestion and capacity of road r. Solving the corresponding linear

[5]Limited to this settings, the smoothness and generalized smoothness inequalities coincide since the system cost equals the sum of the users' costs when no incentives are employed.

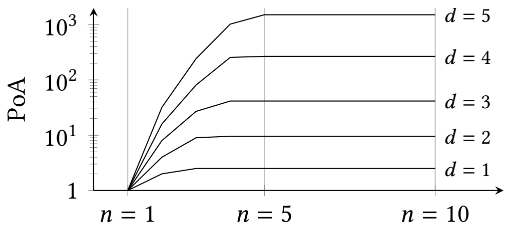

d	PoA, $n = 5$ (Theorem 3.5) PoA, $n \longrightarrow \infty$ ([1, 3, 16])
1	2.50
2	9.58
3	41.54
4	267.64
5	1513.57

Fig. 3. Evolution of the price of anarchy in polynomial congestion games of order $d = 1, \ldots, 5$ as a function of the number of users (left). These values were obtained by solving the corresponding linear program in Theorem 3.5. Observe that the price of anarchy plateaus at $n = 5$, matching the asymptotic bounds ($n \longrightarrow \infty$) previously obtained in the literature [1, 3, 16]. This suggests that small instances are sufficient to produce highly inefficient equilibria.

program in Theorem 3.5, one obtains a price of anarchy of approximately 36.09 for $n = 50$ users and $K_r \in \{1, \ldots, 50\}$. This highlights that, although BPR functions are polynomials of order $d = 4$, their special structure allows significant reductions in the price of anarchy (from 267.64 to 36.09).

ii) Marginal cost incentives in congestion games. Marginal cost incentives have been repeatedly proposed to improve the performance of Nash equilibria in congestion games, e.g., [30, 37]. In the nonatomic variant of this model (whereby users are treated as divisible entities), these incentives guarantee optimal equilibrium efficiency, i.e., their price of anarchy is exactly 1. In the classical atomic setting, marginal cost incentives take the form

$$\tau_r(k) = (k-1)[c_r(k) - c_r(k-1)],$$

allowing for the deployment of our framework to compute their efficiency. This follows readily upon letting $C_r(k) = k \cdot c_r(k)$ and $F_r(k) = k \cdot c_r(k) - (k-1) \cdot c_r(k-1)$ for all k and r in Equations (14) and (15). Thus, using the linear program in Theorem 3.5, we compute the corresponding price of anarchy for polynomial congestion games of order $d = 1, \ldots, 5$ with $n = 100$ (Column 3 in Section 5.2). Perhaps surprisingly, while marginal cost incentives promote optimal performance in the nonatomic settings, their use in the atomic model significantly deteriorates the system's efficiency, with a price of anarchy greater than that experienced when no incentives are used (Columns 2, 3 in Section 5.2).

iii) Perception-parametrized congestion games. The *perception-parameterized* congestion game model was proposed by Kleer and Schäfer [27] to unify the notions of risk sensitivity [8, 38], and altruism [11, 14] in affine congestion games. In this model, the system and user costs are

$$C(a) = \sum_{r \in \mathcal{R}} |a|_r \cdot c_r(1 + \sigma(|a|_r - 1)), \qquad J_i(a) = \sum_{r \in a_i} c_r(1 + \gamma(|a|_r - 1)),$$

where $\sigma, \gamma \geq 0$ are fixed parameters and $c_r(x)$ is an affine function. Among other parameterizations, $\sigma = \gamma = 1$ models "classical" congestion games as in Example 3.1 and $\sigma = 1, \gamma \geq 1$ models congestion games with altruistic users, whereas $\sigma = \gamma \geq 0$ describes congestion games in which each user $i \in N$ participates in the game with probability $p_i = \sigma = \gamma$ [19]. Note that, for given $\sigma, \gamma \geq 0$, the corresponding class of perception-parameterized congestion games is covered by our framework. To see this, it suffices to set $C_r(k) = k \cdot c_r(1 + \sigma(k-1))$, $F_r(k) = c_r(1 + \gamma(k-1))$ for all k and r in Equations (14) and (15) to cover the case of affine resource costs as well as more general cases. Thus, evaluating the price of anarchy of perception-parameterized congestion games – which remains an open problem even in the affine case – is equivalent to solving the linear program in Theorem 3.5.

In Fig. 4, we plot the solution for the affine case, $n = 20$ users and $\sigma, \gamma \in [0, 2]$. Not only do we recover the exact asymptotic bounds of [27] where applicable (region enclosed by white line), but we also provide a complete characterization of the price of anarchy for all $\sigma, \gamma \in [0, 2]$.

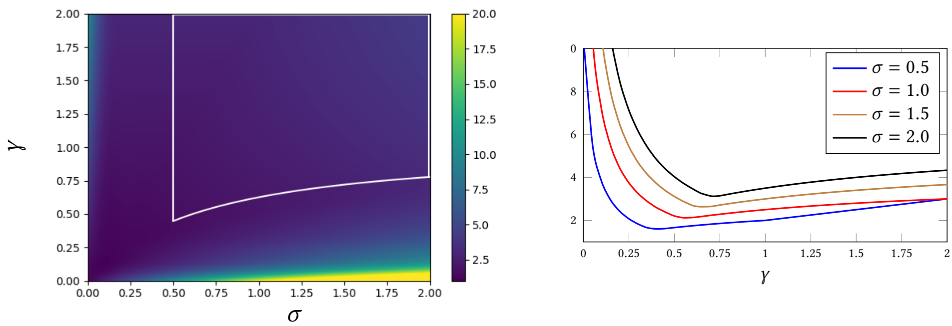

Fig. 4. Exact price of anarchy for perception-parameterized affine congestion games with $n = 20$ users and $\sigma, \gamma \in [0, 2]$, computed via Theorem 3.5 (left). Corresponding values for fixed $\sigma \in \{0.5, 1, 1.5, 2\}$ (right). Kleer and Schäfer [27] give asymptotic values limited to the region enclosed in the white perimeter, which we recover exactly and generalize.

5.2 Optimal local incentives in congestion games

In this section we consider *local* incentives, i.e, incentives that map each resource r of the game to an incentive function τ_r by leveraging solely information on the corresponding congestion function c_r. Whilst previous works also consider incentives that utilize global information, e.g., [7, 9], incentives based solely on local information have a number of advantages including limited informational requirements, scalability, ability to accommodate resources that are dynamically added or removed, and robustness against a number of variations. To ease presentation, we will refer to local incentives simply as incentives.

As illustrated in the previous section, Theorem 3.5 allows us to evaluate the price of anarchy of commonly studied classes of games, e.g., congestion games and generalization thereof. In contrast, we observe that Theorem 3.7 does *not* directly provide a machinery for the design of optimal incentives. To see this observe that, while Theorem 3.7 allows to determine the best *linear* incentive, optimal incentives might very well not satisfy this structural property.[6] Surprisingly, Paccagnan et al. [35] recently showed that the best linear incentive is *optimal* (i.e., its performance can not be improved, even by a nonlinear incentive). Building upon this result, we are then guaranteed that the incentives derived in Theorem 3.7 are the best possible. Thus, we solve the linear program derived in Theorem 3.7 and report the optimal values of the price of anarchy for polynomial congestion games of degree $d = 1, \ldots, 5$ and $n = 100$ in Column 5 of Section 5.2. We observe that the achieved price of anarchy is significantly lower (better) than the setting without incentives (Column 2).

We conclude the section by highlighting that our framework can also accommodate commonly-studied constraints on the set of admissible incentives. For example, fixed incentives (i.e., incentives that are constant in the congestion) can be studied by imposing $\tau_r(k) = \tau_r$, which corresponds to

[6] Given a set of congestion games where each resource cost is obtained by the linear combination $c_r(k) = \sum_{j=1}^{m} \alpha_r^j \cdot c^j(k)$, we say that an incentive T is *linear* if it satisfies $T(c_r) = \sum_{j=1}^{m} \alpha_r^j \cdot T(c^j)$ (i.e., if it is obtained by a linear combination of the incentives $\{T(c^j)\}_{j=1}^{m}$ using the *same* coefficients that define c_r).

substituting $F^j(k) = c^j(k) + \tau^j$, into the linear program in Theorem 3.5. Including $\sigma^j = \nu \cdot \tau^j$ as decision variables, we obtain a linear program with $m+2$ decision variables and $O(mn^2)$ constraints for computing the optimal fixed incentives. We report the corresponding optimal price of anarchy for polynomial congestion games of degree $d = 1, \ldots, 5$ in Column 6 of Section 5.2, and observe that such simple incentives already provide a good improvement upon the setting without incentives.

Table 1. *Price of anarchy in polynomial congestion games.* Price of anarchy for polynomial congestion games with degree $d = 1, \ldots, 5$ and $n = 100$. The second column contains the asymptotic values ($n \to \infty$) without incentives [1, 3, 16]. In the third column, we report the values corresponding to the use of marginal cost incentives (computed through Theorem 3.5). The fourth and fifth columns feature optimal price of anarchy values for general and fixed local incentives, respectively (computed through Theorems 3.5 and 3.7).

d	No Incentive	Marginal Cost	Optimal Local Incentives	Optimal Fixed Incentives
1	2.50	3.00	2.012	2.15
2	9.58	13.00	5.101	5.33
3	41.54	57.36	15.551	18.36
4	267.64	391.00	55.452	89.41
5	1513.57	2124.21	220.401	469.74

5.3 Optimal utility design in distributed welfare games

The preceding discussion showcases how the machinery we developed can be used to compute and optimize the price of anarchy in well-studied classes of cost minimization games. While a similar approach can be followed also for welfare maximization problems (and in fact, recent results build on top of this work to derive state-of-the-art approximation algorithms with explicit guarantees [4, 13, 35, 46]), we purposely decide to take a different perspective in this section. Specifically, we aim at demonstrating the *robustness* of the proposed approach and the *quality* of the corresponding results.

Toward this goal, rather than fixing a specific set of welfare functions, we consider a general setting, whereby $W_r : \{1, \ldots, n\} \to \mathbb{R}$ merely satisfies two properties: non-decreasingness, i.e., $W_r(k+1) \geq W_r(k)$, and concavity (or diminishing returns property), i.e., $W_r(k+1) - W_r(k) \leq W_r(k) - W_r(k-1)$, for all k and r. Observe that these properties are commonly encountered in application areas including vehicle-target assignment problems [15, 32], multiwinner elections [21] and sensor coverage [23, 47]. For any given welfare function satisfying non-decreasingness and concavity, we compare the performance (price of anarchy) obtained by *optimal utilities* against that of the commonly-advocated-for *identical interest* design, whereby each user's utility coincides with the system welfare, i.e., $U_i(a) = W(a)$. We do so for 10^5 unique resource welfare functions W_r, randomly generated by sorting 10 independently values from a uniform distribution over $[0, 1]$ from largest to smallest and setting $W_r(k)$ to be the sum over the first k sorted values. Nondecreasingness and concavity of W_r follow readily. For each generated resource welfare, we determine the corresponding optimal price of anarchy and the price of anarchy of the identical interest design through the solution of the linear programs in (29) and (28), respectively.[7] The left

[7]Instead of determining the price of anarchy of the identical interest utilities directly, we can compute the price of anarchy of the marginal contribution utilities (taking the form $U_i(a) = W(a) - W(\emptyset_i, a_{-i})$), as these two values coincide. To see this note that the underlying set of Nash equilibria remains the same under either utility as, for any $a, a' \in \mathcal{A}$, since $W(a) - W(a'_i, a_{-i}) \geq 0 \iff W(a) - W(\emptyset_i, a_{-i}) - [W(a'_i, a_{-i}) - W(\emptyset_i, a_{-i})] \geq 0$. Here, $W(\emptyset_i, a_{-i})$ denotes the system welfare when user i selects no action and the remaining users select their action in a.

panel in Fig. 5 depicts the resulting empirical distribution of the price of anarchy values, while the right reports the ratio between the price of anarchy in the identical interest and optimal settings (this ratio is never lower than one, as expected). We conclude by observing that, whilst the identical interest design may initially appear to be an intuitive and appealing option, Fig. 5 clearly highlights that strictly better performance can be readily achieved using the machinery developed here.

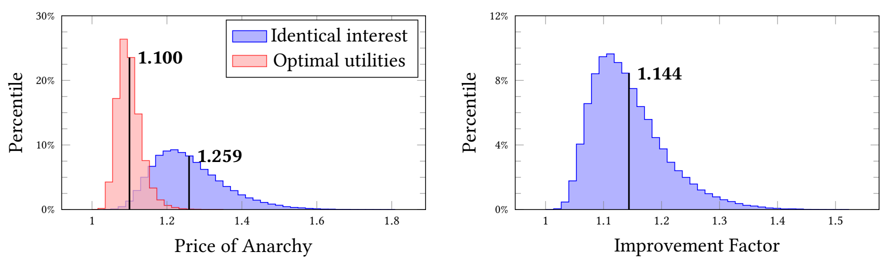

Fig. 5. *Price of anarchy for identical interest and optimal utilities in distributed welfare games.* Left: Empirical distribution of the price of anarchy with optimal utilities vs. identical interest design. Right: Improvement factor when using the optimal design in place of the identical interest design. The mean of each distribution is indicated by a bisecting solid black line. Observe that the optimal utilities offer significant improvement over the identical interest utilities, reducing the price of anarchy by a factor of approximately 1.144 on average. The values in the above charts were obtained using the linear programs in Theorem 4.3 for nondecreasing, concave resource welfare functions.

6 CONCLUSIONS AND FUTURE WORK

Though well-studied, the price of anarchy can still be difficult to compute as ad-hoc approaches are often needed. As a result, the design of incentives that optimize this metric is even more challenging, with only few results available. Motivated by this observation, our work provides a framework achieving two fundamental goals: to tightly characterize and optimize the price of anarchy through a computationally tracable approach.

Toward this end, we first introduced the notion of *generalized smoothness*, which we showed always produces tighter or equal price of anarchy bounds compared to the original smoothness approach. We proved that such bounds are *exact* for generalized congestion and local welfare maximization games, unlike those obtained through a simple smoothness argument. Additionally, we showed that the problems of computing and optimizing the price of anarchy can be posed (and solved) as tractable linear programs, when considering these broad problem classes. Finally, we demonstrated the ease of applicability, strength and breadth of our approach by recovering and generalizing existing results on the computation of the price of anarchy, as well as by tackling the problems of incentive design in congestion games and utility design in distributed welfare games. In this regard, the list of illustrative example provided in Section 5 is certainly non-exhaustive.

Overall, we feel that the proposed approach has significant potential, especially since it can be used as a "black box" to compute and optimize the exact price of anarchy in many problems of interest. The linear programs derived here can be used, for example, as "computational companions" to support the analytical study of the price of anarchy, e.g, by providing evidence, or disproving certain conjectures. For this reason, we compliment our work with a software package that implements the techniques and linear programs derived here in the hope that they can be of help for new research to come.[8]

[8]https://github.com/rahul-chandan/resalloc-poa

We conclude observing that the price of anarchy represents but one of many metrics for measuring an algorithm's performance. Nevertheless, we believe that the techniques introduced here can be suitably extended to analyze different metrics (e.g., the price of stability) and to understand whether optimizing for the price of anarchy has any unintended consequences on them.

ACKNOWLEDGMENTS

A preliminary version of this work appeared in [12]. This work is supported by ONR grants #N00014-17-1-2060 and #N00014-20-1-2359, NSF grant #ECCS-1638214, and SNSF grant #P2EZP2_181618.

REFERENCES

[1] Aland, S., Dumrauf, D., Gairing, M., Monien, B., and Schoppmann, F. Exact price of anarchy for polynomial congestion games. *SIAM Journal on Computing 40*, 5 (2011), 1211–1233.

[2] Augustine, J., Chen, N., Elkind, E., Fanelli, A., Gravin, N., and Shiryaev, D. Dynamics of profit-sharing games. *Internet Mathematics 11*, 1 (2015), 1–22.

[3] Awerbuch, B., Azar, Y., and Epstein, A. The price of routing unsplittable flow. *SIAM Journal on Computing 42*, 1 (2013), 160–177.

[4] Barman, S., Fawzi, O., and Fermé, P. Tight approximation guarantees for concave coverage problems. *arXiv preprint arXiv:2010.00970* (2020).

[5] Bilò, V. On the robustness of the approximate price of anarchy in generalized congestion games. In *International Symposium on Algorithmic Game Theory* (2016), Springer, pp. 93–104.

[6] Bilò, V. A unifying tool for bounding the quality of non-cooperative solutions in weighted congestion games. *Theory of Computing Systems 62*, 5 (2018), 1288–1317.

[7] Bilò, V., and Vinci, C. Dynamic taxes for polynomial congestion games. *ACM Transactions on Economics and Computation (TEAC) 7*, 3 (2019), 1–36.

[8] Caragiannis, I., Fanelli, A., Gravin, N., and Skopalik, A. Computing approximate pure nash equilibria in congestion games. *ACM SIGecom Exchanges 11*, 1 (2012), 26–29.

[9] Caragiannis, I., Kaklamanis, C., and Kanellopoulos, P. Taxes for linear atomic congestion games. *ACM Transactions on Algorithms (TALG) 7*, 1 (2010), 1–31.

[10] Caragiannis, I., Kaklamanis, C., Kanellopoulos, P., Kyropoulou, M., Lucier, B., Leme, R. P., and Tardos, É. Bounding the inefficiency of outcomes in generalized second price auctions. *J. Econ. Theory 156* (2015).

[11] Caragiannis, I., Kaklamanis, C., Kanellopoulos, P., Kyropoulou, M., and Papaioannou, E. The impact of altruism on the efficiency of atomic congestion games. In *International Symposium on Trustworthy Global Computing* (2010), Springer, pp. 172–188.

[12] Chandan, R., Paccagnan, D., and Marden, J. R. When smoothness is not enough: Toward exact quantification and optimization of the price-of-anarchy. In *2019 IEEE 58th Conference on Decision and Control (CDC)* (2019), IEEE, pp. 4041–4046.

[13] Chandan, R., Paccagnan, D., and Marden, J. R. Tractable mechanisms for computing near-optimal utility functions. In *Proceedings of the 20th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)* (2021), IFAAMAS, p. to appear.

[14] Chen, P.-A., Keijzer, B. D., Kempe, D., and Schäfer, G. Altruism and its impact on the price of anarchy. *ACM Transactions on Economics and Computation (TEAC) 2*, 4 (2014), 1–45.

[15] Chopra, S., Notarstefano, G., Rice, M., and Egerstedt, M. A distributed version of the hungarian method for multirobot assignment. *IEEE Transactions on Robotics 33*, 4 (2017), 932–947.

[16] Christodoulou, G., and Koutsoupias, E. The price of anarchy of finite congestion games. In *Proceedings of the thirty-seventh annual ACM symposium on Theory of computing* (2005), ACM, pp. 67–73.

[17] Christodoulou, G., Koutsoupias, E., and Nanavati, A. Coordination mechanisms. In *International Colloquium on Automata, Languages, and Programming* (2004), Springer, pp. 345–357.

[18] Colini-Baldeschi, R., Cominetti, R., Mertikopoulos, P., and Scarsini, M. When is selfish routing bad? the price of anarchy in light and heavy traffic. *Operations Research 68*, 2 (2020), 411–434.

[19] Cominetti, R., Scarsini, M., Schröder, M., and Stier-Moses, N. Price of anarchy in bernoulli congestion games with affine costs. *arXiv e-prints* (2019), arXiv–1903.

[20] Dai, H., Huang, Y., and Yang, L. Game theoretic max-logit learning approaches for joint base station selection and resource allocation in heterogeneous networks. *IEEE Journal on Selected Areas in Communications 33*, 6 (2015), 1068–1081.

[21] Dudycz, S., Manurangsi, P., Marcinkowski, J., and Sornat, K. Tight approximation for proportional approval voting. *Proceedings of IJCAI'20* (2020), 276–282.
[22] Foster, D. J., Li, Z., Lykouris, T., Sridharan, K., and Tardos, E. Learning in games: Robustness of fast convergence. In *Advances in Neural Information Processing Systems* (2016).
[23] Gairing, M. Covering games: Approximation through non-cooperation. In *International Workshop on Internet and Network Economics* (2009), Springer, pp. 184–195.
[24] Gkatzelis, V., Kollias, K., and Roughgarden, T. Optimal cost-sharing in general resource selection games. *Operations Research 64*, 6 (2016), 1230–1238.
[25] Hardin, G. The tragedy of the commons. *Science 162*, 3859 (1968), 1243–1248.
[26] Harks, T., and Miller, K. The worst-case efficiency of cost sharing methods in resource allocation games. *Operations research 59*, 6 (2011), 1491–1503.
[27] Kleer, P., and Schäfer, G. Tight inefficiency bounds for perception-parameterized affine congestion games. *Theoretical Computer Science 754* (2019), 65–87.
[28] Kleinberg, J., and Oren, S. Mechanisms for (mis) allocating scientific credit. In *Proceedings of the forty-third annual ACM symposium on Theory of computing* (2011), ACM, pp. 529–538.
[29] Koutsoupias, E., and Papadimitriou, C. Worst-case equilibria. In *Annual Symposium on Theoretical Aspects of Computer Science* (1999), Springer, pp. 404–413.
[30] Maillé, P., and Stier-Moses, N. E. Eliciting coordination with rebates. *Transportation Science 43*, 4 (2009), 473–492.
[31] Marden, J. R., and Wierman, A. Distributed welfare games. *Operations Research 61*, 1 (2013), 155–168.
[32] Murphey, R. A. Target-based weapon target assignment problems. In *Nonlinear assignment problems*. Springer, 2000, pp. 39–53.
[33] Nadav, U., and Roughgarden, T. The limits of smoothness: A primal-dual framework for price of anarchy bounds. In *International Workshop on Internet and Network Economics* (2010), Springer, pp. 319–326.
[34] Nisan, N., Roughgarden, T., Tardos, E., and Vazirani, V. V. *Algorithmic game theory*. Cambridge university press, 2007.
[35] Paccagnan, D., Chandan, R., Ferguson, B. L., and Marden, J. R. Optimal taxes in atomic congestion games. *ACM Transactions on Economics and Computation (to appear), arXiv preprint arXiv:1911.09806v4* (2021).
[36] Paccagnan, D., Chandan, R., and Marden, J. R. Utility design for distributed resource allocation—part i: Characterizing and optimizing the exact price of anarchy. *IEEE Transactions on Automatic Control 65*, 11 (2019), 4616–4631.
[37] Pigou, A. *The economics of welfare*. Macmillan, 1920.
[38] Piliouras, G., Nikolova, E., and Shamma, J. S. Risk sensitivity of price of anarchy under uncertainty. *ACM Transactions on Economics and Computation (TEAC) 5*, 1 (2016), 1–27.
[39] Rosenthal, R. W. A class of games possessing pure-strategy nash equilibria. *International Journal of Game Theory 2*, 1 (1973), 65–67.
[40] Roughgarden, T. Intrinsic robustness of the price of anarchy. *Journal of the ACM (JACM) 62*, 5 (2015), 32.
[41] Scarsini, M., Schröder, M., and Tomala, T. Dynamic atomic congestion games with seasonal flows. *Operations Research 66*, 2 (2018), 327–339.
[42] Shamma, J. S. *Cooperative control of distributed multi-agent systems*. Wiley Online Library, 2007.
[43] Suri, S., Tóth, C. D., and Zhou, Y. Selfish load balancing and atomic congestion games. *Algorithmica 47*, 1 (2007), 79–96.
[44] Syrgkanis, V., and Tardos, E. Composable and efficient mechanisms. In *Proceedings of the forty-fifth annual ACM symposium on Theory of computing* (2013), ACM.
[45] United States Bureau of Public Roads. *Traffic assignment manual for application with a large, high speed computer*, vol. 2. US Department of Commerce, Bureau of Public Roads, 1964.
[46] Vijayalakshmi, V. R., and Skopalik, A. Improving approximate pure nash equilibria in congestion games. In *International Conference on Web and Internet Economics* (2020), Springer, pp. 280–294.
[47] Zhu, M., and Martinez, S. Distributed coverage games for energy-aware mobile sensor networks. *SIAM Journal on Control and Optimization 51*, 1 (2013), 1–27.